Validated Profile: University Vertical

June 27, 2025

Solution overview

This document aims to provide guidance and act as a validated reference for deploying university networks using Catalyst Center and Cisco Software-Defined Access (SD-Access) solutions.

The education industry is undergoing significant transformations, including the adoption of smart campus technologies, automation, hybrid learning environments, and secure distance learning. The exponential increase in connected endpoints, driven by students and faculty bringing personal devices to campus, poses unique challenges. Additionally, university students and faculty often travel globally and require immediate access to research materials.

Like other industries, educational networks demand advanced network services, seamless mobility, high availability, and efficient management. However, university networks have unique requirements, including enhanced security for residential services, support for wireless-heavy infrastructures, and robust wireless mobility.

This document highlights key considerations tailored to address the specific needs of the education sector.

Scope

This guide serves as a comprehensive roadmap for understanding the challenges faced by university networks, exploring common use cases, and demonstrating how Cisco SD-Access can address these needs. While it does not delve into detailed configuration steps, it provides valuable insights to support the development of an effective university network strategy.

Traditional network versus Cisco SD-Access

This section provides an overview of the key differences between traditional network and Cisco SD-Access.

Traditional networks:

- Traditional networks require network devices to be configured manually.
- They often require a separate overlay network for segmentation.
- Security policies are typically enforced between network boundaries.
- Scaling the network can be complex and time-consuming.
- Troubleshooting is often reactive and requires manual intervention.
- Limited visibility into network traffic and application performance.

Cisco SD-Access:

- SD-Access automates network provisioning and management through intent-based automation.
- It simplifies network design by carrying Virtual Network (VN) and Security Group Tag (SGT) information in the VXLAN overlay while using a single underlay network for both connectivity and segmentation.
- Security policies are applied dynamically based on user and device identity.
- SD-Access scales more easily through automation and centralized control.
- Troubleshooting is proactive with network-wide visibility and analytics.
- SD-Access provides detailed insights into network traffic and application performance.

In summary, Cisco SD-Access offers a more streamlined and flexible approach compared to traditional networks, with centralized management, improved scalability, and enhanced security features.

Challenges in traditional networks

Today there are many challenges in managing the network, because of manual configuration and fragmented tool offerings. Manual operations are slow and error prone. Issues are exacerbated because of a constantly changing environment. The growth of users and different device types makes it more complex to configure and maintain a consistent user policy across the network.

- Network deployment challenges:

  Setup or deployment of a single network switch can take several hours due to scheduling requirements and the need to work with different infrastructure groups. In some cases, deploying a batch of switches can take several weeks.

- Network security challenges:

  Security is a critical component of managing modern networks. Organizations need to protect resources and make changes efficiently in response to real-time needs. In traditional networks, it can be challenging to track VLANs, access control lists (ACLs), and IP addresses to ensure optimal policy and security compliance.

- Wireless and wired network challenges:

  Disparate networks are common in many organizations, because different systems are managed by different departments. Typically, the main IT network is operated separately from building management systems, security systems, and other production systems. This leads to duplication of network hardware procurement and inconsistency in management practices.

- Network operations challenges:

  IT teams often contend with outdated change management tools, difficulty in maintaining productivity, and slow issue resolution.

Advantages of Cisco SD-Access

Cisco SD-Access is designed to address the demands of rapid digitization. The core philosophy of the Cisco SD-Access architecture revolves around policy-based automation, enabling secure user and device segmentation across both wired and wireless connectivity.

Automation and simplicity boost productivity, allowing IT staff to innovate quickly and lead the industry in digital transformation, thereby enhancing operational effectiveness. A consistent segmentation framework aligned with business policies, regardless of transport medium (wired or wireless), is crucial for core effectiveness.

Cisco SD-Access provides technological advantages, including:

- Simplified operations:

  Simplifies network operations by providing a single, intuitive interface for managing the entire infrastructure, reducing complexity and operational overhead.

- Automation:

  Automates routine network operations such as configuration, provisioning, and management. This reduces the risk of human error and increases efficiency. Catalyst Center streamlines the deployment, minimizing the need for interaction with Command Line Interfaces (CLI).

- Agility:

  Network operations become more agile and align with business requirements by minimizing manual configuration steps.

- Security:

  Provides enhanced security and segmentation through Virtual Networks (VNs) and Security Group Tags (SGTs). SD-Access provides a strong framework for securing and managing complex enterprise networks through macro-segmentation with Virtual Networks (VNs), and micro-segmentation with Security Group Tags (SGTs).

- Consistent policies for wired and wireless:

  Extends segmentation, visibility, and policy from wired to wireless networks. Distributed wireless termination scales network throughput while centralizing management and troubleshooting.

- Support for business analytics:

  Aggregates analytics and telemetry information into a single platform, aiding business decisions and facilitating growth or diversification planning.

University network overview

For guidance and recommendations on constructing a new greenfield deployment of the Cisco SD-Access fabric tailored to the challenges and use cases of a university network, see the sections for an in-depth exploration of SD-Access fabric components. Discover the benefits Cisco SD-Access solutions provide in addressing the unique requirements and challenges of the education sector.

Traditional networks can be managed using Cisco Prime Infrastructure or Catalyst Center. Catalyst Center offers advanced automation, monitoring, and telemetry capabilities for both traditional networks and SD-Access environments. If you are currently managing a network with Cisco Prime Infrastructure and plan to migrate to Catalyst Center, see Cisco Prime Infrastructure to Cisco Catalyst Center Migration.

To transition existing Cisco Catalyst legacy networks to a Cisco SD-Access fabric, see Migration to Cisco SD-Access, which outlines options for migrating existing networks with both wired and wireless endpoints.

Cisco Catalyst Center

Catalyst Center (formerly known as Cisco DNA Center) is a centralized network management and orchestration platform designed to simplify network operations and management. It provides a single dashboard to manage and monitor your network infrastructure, including switches, routers, and wireless access points (APs).

Using Catalyst Center, network administrators can do various tasks, including:

- Automate network provisioning:

  Easily deploy network devices and services using automated workflows, reducing the time and effort required for configuration.

- Monitor network health:

  Gain visibility into the entire network, including device status, traffic patterns, and performance metrics, to quickly identify and resolve issues.

- Implement security policies:

  Define and enforce security policies across the network, ensuring compliance and protecting against threats.

- Manage software updates:

  Simplify the process of updating device software and firmware, ensuring that network devices are up to date with the latest features and security patches.

- Troubleshoot network problems:

  Use built-in tools and analytics to diagnose and resolve network issues quickly, minimizing downtime and disruption.

Overall, Catalyst Center helps organizations streamline network operations, improve efficiency, and enhance security, making it an essential tool for managing modern network infrastructures.

The Catalyst Center platform is available in various form factors, including physical and virtual appliances. For details, see these resources:

- Cisco Catalyst Center Data Sheet (for supported platform and scale).
- Cisco Catalyst Center Installation Guide

Cisco Identity Service Engine

Cisco Identity Services Engine (ISE) is a security policy management and control platform. It automates and simplifies access control and security compliance for wired, wireless, and VPN connectivity. Cisco ISE offers secure access to network resources, enforces security policies, and delivers comprehensive visibility into network access.

The key features of Cisco ISE include:

- Policy-based access control:

  Define and enforce policies based on user roles, device types, and other contextual information.

- Authentication and authorization:

  Support for various authentication methods (for example, 802.1X, MAB, web authentication) and enables dynamic authorization based on changing conditions.

- Endpoint compliance:

  Assess the compliance of endpoints with security policies and enforce remediation actions, if necessary.

- Guest access:

  Provide secure guest access to the network with customizable guest portals and sponsor approval workflows.

- Bring Your Own Device (BYOD) support:

  Enable secure BYOD initiatives with device onboarding and policy enforcement.

- Integration and Ecosystem:

  Integrate with other security and networking technologies through APIs and partner ecosystem.

- Visibility and reporting:

  Gain insights into network access and security posture through comprehensive reporting and analytics.

Cisco ISE is a critical component of Cisco's security and network access control portfolio, providing organizations with a centralized and scalable solution to address their security and access control needs. Cisco Identity Services Engine (ISE) supports both standalone and distributed deployment models. Multiple distributed nodes can be deployed collectively to enhance failover resiliency and scalability. Minimally, a basic two-node ISE deployment is recommended for SD-Access single-site deployments, with each ISE node running all services (personas) for redundancy.

For more details, see:

- Cisco Identity Services Engine Administrator Guides
- Performance and Scalability Guide for Cisco Identity Services Engine

Cisco SD-Access fabric

Cisco SD-Access fabric is a networking architecture that uses software-defined networking (SDN) concepts to automate network provisioning, segmentation, and policy enforcement. It aims to simplify network operations, enhance security, and improve user experiences in modern digital workplaces.

Key features of Cisco SD-Access fabric include:

- Unified wired and wireless automation:

  One of the standouts features of SD-Access is its ability to integrate wired and wireless networks into a single, automated management framework.

- Network segmentation:

  Divides the network into virtual segments based on user and device identity, enabling granular control over access and security policies.

- Centralized policy management:

  Policies are defined centrally and enforced consistently across the entire network, reducing the risk of misconfiguration and policy conflicts.

- ISE:

  Provides authentication and authorization services, ensuring that only authorized users and devices can access the network.

- Catalyst Center:

  Serves as the management and orchestration platform for SD-Access, providing a single pane of glass for network management and troubleshooting.

- Scalability:

  Supports large-scale deployments, enabling organizations to easily scale their networks as their needs expand.

- Enhanced security:

  Improves network security by dynamically segmenting the network and enforcing security policies based on user and device identity.

Overall, Cisco SD-Access fabric aims to simplify network management, improve security, and enhance scalability, making it an attractive option for organizations looking to modernize their network infrastructure.

Fabric architecture overview

Cisco SD-Access fabric architecture is designed to simplify network operations, enhance security, and improve user experiences. It is based on the principles of software-defined networking (SDN) and incorporates various components to achieve these goals:

- Underlay network:

  The physical network infrastructure that provides basic connectivity between devices. It typically consists of switches, routers, and cables.

- Overlay network:

  A logical network built on top of the underlay network that provides virtualized connectivity between devices. It enables network segmentation and policy enforcement without the need for physical reconfiguration.

- Control plane:

  Manages the overall operation of the network, including routing, forwarding, and policy enforcement. It is typically implemented using a centralized controller, such as Catalyst Center.

- Data plane:

  Handles the actual forwarding of data packets within the network. It is implemented on network devices, such as switches and routers, and operates based on the instructions provided by the control plane.

- Policy plane:

  Defines and enforces network policies, such as access control and segmentation. It ensures that network resources are used efficiently and securely.

- Management plane:

  Provides tools and interfaces for managing and monitoring the network. It includes features such as configuration management, monitoring, and troubleshooting.

Overall, Cisco SD-Access fabric architecture offers a comprehensive solution for modernizing network infrastructure, providing scalability, security, and automation capabilities to meet the evolving needs of digital businesses.

Network architecture

Fabric technology supports the SD-Access architecture on campus, enabling the use of VNs (overlay networks) running on a physical network (underlay network) to create alternative topologies for connecting devices. In SD-Access, the user-defined overlay networks are provisioned as virtual routing and forwarding (VRF) instances that provide separation of routing tables.

Fabric roles

A fabric role is an SD-Access software construct running on physical hardware. These software constructs are designed with modularity and flexibility in mind. For example, a device can run either a single role or multiple roles. Care should be taken to provision SD-Access fabric roles in alignment with the underlying network architecture, ensuring a distributed function approach. Separating roles across different devices provides the highest level of availability, resilience, deterministic convergence, and scalability.

The SD-Access fabric roles include:

- Control plane node
- Border node
- Edge node
- Intermediate node
- Fabric wireless controller
- Fabric-mode APs

Control plane node

SD-Access fabric control plane node combines LISP map-server functionalities and map-resolver functionalities on a single node. It maintains a database that tracks all endpoints within the fabric site, mapping them to fabric nodes. This design separates an endpoint's IP or MAC address from its physical location (nearest router), ensuring efficient network operations.

Key functions of the control plane node:

- Host Tracking Database (HTDB):

  Acts as a central repository for EID-to-RLOC bindings, where the routing locator (RLOC) is the loopback zero IP address of a fabric node. It functions similarly to a traditional LISP site, storing endpoint registrations.

- Endpoint Identifier (EID):

  Identifies endpoint devices using MAC, IPv4, or IPv6 addresses in the SD-Access network.

- Map server:

  Receives endpoint registrations, associates them with their corresponding RLOCs, and updates the HTDB accordingly.

- Map resolver:

  Responds to queries from fabric devices, providing EID-to-RLOC mappings from the HTDB. This allows devices to determine the appropriate fabric node for forwarding traffic.

Border node

SD-Access fabric border node serves as the gateway between a fabric site and external networks, handling network virtualization interworking and the propagation of SGTs beyond the fabric.

Key functions of border nodes:

- EID subnet advertisement:

  Uses the Border Gateway Protocol (BGP) to advertise endpoint prefixes outside the fabric, ensuring return traffic is directed correctly.

- Fabric site exit point:

  Functions as the default gateway for edge nodes using LISP Proxy Tunnel Router (PxTR). Internal border nodes can register known subnets with the control plane node.

- Network virtualization extension:

  Extends segmentation beyond the fabric using VRF-lite and VRF-aware routing protocols.

- Policy mapping:

  Maintains SGT information outside the fabric via SGT Exchange Protocol (SXP) or inline tagging in Cisco metadata.

- VXLAN encapsulation and de-encapsulation:

  Converts external traffic into VXLAN for the fabric and removes VXLAN for outgoing traffic, acting as a bridge between the fabric and non-fabric networks.

Edge node

SD-Access fabric edge nodes function like access layer switches in a traditional campus LAN. They operate based on ingress and egress tunnel routers (xTR) in LISP and must be deployed using a Layer 3 routed access design. These edge nodes perform several key functions:

- Endpoint registration:

  Each edge node maintains a LISP control plane session with all control plane nodes. When an endpoint is detected, it is added to a local database called the EID-table. The edge node then sends a LISP map-register message to update the control plane's HTDB (Host Tracking Database).

- Anycast Layer 3 gateway:

  All edge nodes sharing the same EID subnet use a common IP and MAC address for seamless mobility and optimal forwarding. The anycast gateway is implemented as a Switched Virtual Interface (SVI) with a uniform MAC address across all edge nodes in the fabric.

- Layer 2 bridging:

  Edge nodes handle Layer 2 traffic for endpoints within the same VLAN. They determine whether to bridge or route packets and use VXLAN Layer 2 VNIs (equivalent to VLANs) to bridge traffic to the correct destination. If traffic needs to exit the fabric, a Layer 2 border node is used.

- User-to-VN mapping:

  Endpoints are assigned to VNs by associating them with VLANs linked to an SVI and VRF. This mapping ensures fabric segmentation at both the Layer 2 and Layer 3 LISP VNIs, even at the control plane level.

- AAA authentication:

  Edge nodes can statically or dynamically assign endpoints to VLANs using 802.1X authentication. Acting as a Network Access Device (NAD), they collect authentication credentials, send them to an authentication server, and enforce access policies.

- VXLAN encapsulation and de-Encapsulation:

  When an edge node receives traffic from an endpoint (directly connected, via an extended node, or through an AP), it encapsulates it in VXLAN and forwards it across the fabric. Depending on the destination, the traffic is sent to another edge node or a border node. When encapsulated traffic arrives at an edge node, it is de-encapsulated and delivered to the endpoint. This mechanism enables endpoint mobility, allowing devices to move between edge nodes without changing their IP addresses.

Intermediate node

Intermediate nodes are part of the Layer 3 network used for interconnections among devices operating in fabric roles, such as the connections between border nodes and edge nodes. These interconnections are established in the global routing table on the devices and are collectively known as the underlay network. For example, in a three-tier campus deployment where the core switches are provisioned as border nodes and the access switches as edge nodes, the distribution switches function as the intermediate nodes.

Intermediate nodes do not require VXLAN encapsulation/de-encapsulation, LISP control plane messaging, or SGT awareness. Their primary function is to provide IP reachability and physical connectivity, while also supporting the increased maximum transmission unit (MTU) to accommodate larger IP packets encapsulated with fabric VXLAN information. Essentially, intermediate nodes route and transport IP traffic between devices operating in fabric roles.

Fabric wireless controller

Both fabric wireless controllers and nonfabric wireless controllers provide AP image and configuration management, client session management, and mobility services. Fabric wireless controllers offer additional services for fabric integration, such as registering MAC addresses of wireless clients into the HTDB of the fabric control plane nodes during wireless client join events and supplying fabric edge node RLOC-association updates to the HTDB during client roam events. Fabric integration with a wireless controller occurs on a per-SSID basis. Fabric-enabled SSID traffic is tunneled by the AP using VXLAN encapsulation to the fabric edge node, while centrally switched SSID traffic is tunneled by the AP using the Control and Provisioning of Wireless Access Points (CAPWAP) protocol to the wireless controller. Thus, the wireless controller can operate in a hybrid or mixed mode, where some SSIDs are fabric-enabled while others are centrally switched.

- Traditional vs. SD-Access data handling:

  In a traditional Cisco Unified Wireless network or nonfabric deployment, both control traffic and data traffic are tunneled back to the wireless controller using CAPWAP. From a CAPWAP control plane perspective, AP management traffic is generally lightweight, while client data traffic is the larger bandwidth consumer. Wireless standards have enabled progressively larger data rates for wireless clients, resulting in more client data being tunneled to the wireless controller. This requires a larger wireless controller with multiple high-bandwidth interfaces to support the increase in client traffic.

  In nonfabric wireless deployments, wired and wireless traffic have different enforcement points in the network. The wireless controller addresses quality of service and security when bridging the wireless traffic onto the wired network. For wired traffic, enforcement occurs at the first-hop access layer switch. This paradigm shifts entirely with SD-Access Wireless. In SD-Access Wireless, the CAPWAP tunnels between the wireless controllers and APs are used only for control traffic. Data traffic from wireless endpoints is tunneled to the first-hop fabric edge node, where security and policy can be applied in the same manner as for wired traffic.

- Network connectivity and wireless controller placement:

  Typically, fabric wireless controllers connect to a shared services network through a distribution block or data center network that is located outside the fabric and fabric border, with the wireless controller management IP address existing in the global routing table. For wireless APs to establish a CAPWAP tunnel for wireless controller management, the APs must be in a VN with access to this external device. This means that the APs are deployed in the global routing table, and the wireless controller's management subnet or specific prefix must be present in the Global Routing Table (GRT) within the fabric site.

  In the SD-Access solution, Cisco Catalyst Center configures wireless APs to reside within an overlay VN named INFRA_VN, which maps to the global routing table. This setup eliminates the need for route

leaking or fusion routing (a multi-VRF device selectively sharing routing information) to establish connectivity between the wireless controllers and the APs. Each fabric site must have a wireless controller unique to that site. Most deployments place the wireless controller within the local fabric site itself, rather than across a WAN, due to latency requirements for local mode APs.

- Latency requirements and deployment considerations:

  Fabric APs operate in local mode, which requires a Round-Trip Time (RTT) of 20 ms or less between the AP and the wireless controller. This typically means that the wireless controller is deployed in the same physical site as the APs. However, if this latency requirement is met through dedicated dark fiber or other very low-latency circuits between physical sites, and the wireless controllers are deployed physically elsewhere, such as in a centralized data center, the wireless controllers and APs can be in different physical locations.

  This deployment type, where fabric APs are located separately from their fabric wireless controllers, is commonly used in metro area networks and SD-Access for Distributed Campus environments. APs should not be deployed over WAN or other high-latency circuits from their wireless controllers in an SD-Access network. Maintaining a maximum RTT of 20 ms between these devices is crucial for performance.

Fabric mode APs

Fabric mode APs are Cisco Wi-Fi 7 (802.11be), Wi-Fi 6 (802.11ax) and 802.11ac Wave 2 APs associated with the fabric wireless controller that have been configured with one or more fabric-enabled SSIDs. These fabric mode APs continue to support the same wireless media services as traditional APs, such as applying Application Visibility and Control (AVC), Quality of Service (QoS), and other wireless policies. Fabric APs establish a CAPWAP control plane tunnel to the fabric wireless controller and join as local-mode APs. They must be directly connected to the fabric edge node or extended node switch within the fabric site. For their data plane, fabric APs establish a VXLAN tunnel to their first-hop fabric edge switch, where wireless client traffic is terminated and placed on the wired network.

Fabric APs are considered special case wired hosts. Edge nodes use the Cisco Discovery Protocol to recognize APs as these wired hosts, apply specific port configurations, and assign the APs to a unique overlay network called INFRA_VN. As wired hosts, APs have a dedicated EID space and are registered with the control plane node. This EID space is associated with the predefined INFRA_VN overlay network in the Cisco Catalyst Center UI. It is a common EID space (prefix space) and VN for all fabric APs within a fabric site. The assignment to this overlay VN simplifies management by using a single subnet to cover the AP infrastructure within a fabric site.

Figure 1. Key components involved in an SD-Access fabric deployment and component positions within an SD-Access network

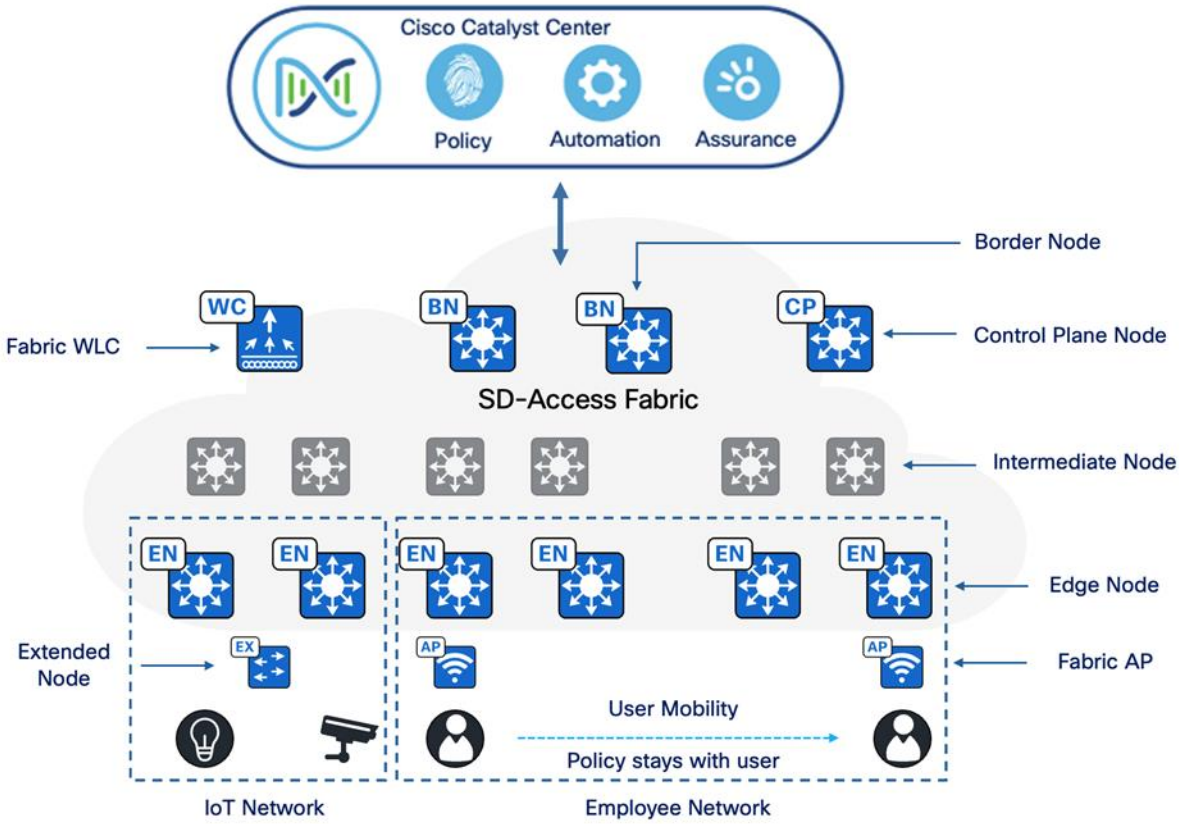

Fabric in a Box

Fabric In a Box (FIAB) integrates all the functionalities of a traditional SD-Access network such as border node, control plane node, and edge node into a single physical device. This device can be a single switch, a switch with hardware stacking capabilities, or part of a StackWise Virtual deployment.

The benefits of FIAB include:

- Simplicity
- Cost-effectiveness
- Faster deployment
- Ideal for branches and small-sized deployments

For more details, see the StackWise Virtual White Paper.

Extended nodes

SD-Access Extended Nodes enable the extension of the enterprise network to non-carpeted areas. Extended nodes provide a Layer 2 port extension to a fabric edge node while ensuring segmentation and applying group-based policies to the connected endpoints. Using Extended Nodes, organizations can extend the benefits of SD-Access such as enhanced security, simplified management, and consistent policy application to a broader range of devices and endpoints within their network.

For more details, see extended node design.

Fabric wireless controllers and APs

wireless controllers and traditional wireless controllers manage AP images and configurations, handle client sessions, and offer mobility services. Fabric wireless controllers additionally support fabric integration by registering MAC addresses of wireless clients into the host tracking database of fabric control plane nodes.

Fabric-mode APs that are Cisco Wi-Fi 7 (802.11be), Wi-Fi 6 (802.11ax) and 802.11ac Wave 2 APs associated with the fabric wireless controller, are configured with one or more fabric-enabled SSIDs. These fabric-mode APs retain support for wireless media services such as AVC, quality of service (QoS), and other wireless policies that are traditional APs.

For more details, see SD-Access Wireless Design and Deployment Guide.

SD-Access embedded wireless

For distributed branches and small campuses, wireless controller functionality can be achieved without a hardware wireless controller through the Cisco Catalyst 9800 Embedded Wireless Controller, available as a software package for Catalyst 9000 Series switches.

The Catalyst 9800 Embedded Wireless Controller supports SD-Access deployments in three topologies:

- Cisco Catalyst 9000 Series switches function as co-located border and control plane.
- Cisco Catalyst 9000 Series switches function as an edge node when the border and control plane node are on a routing platform.
- Cisco Catalyst 9000 Series switches functioning as fabric consolidation.

Figure 2. SD-Access supported topologies

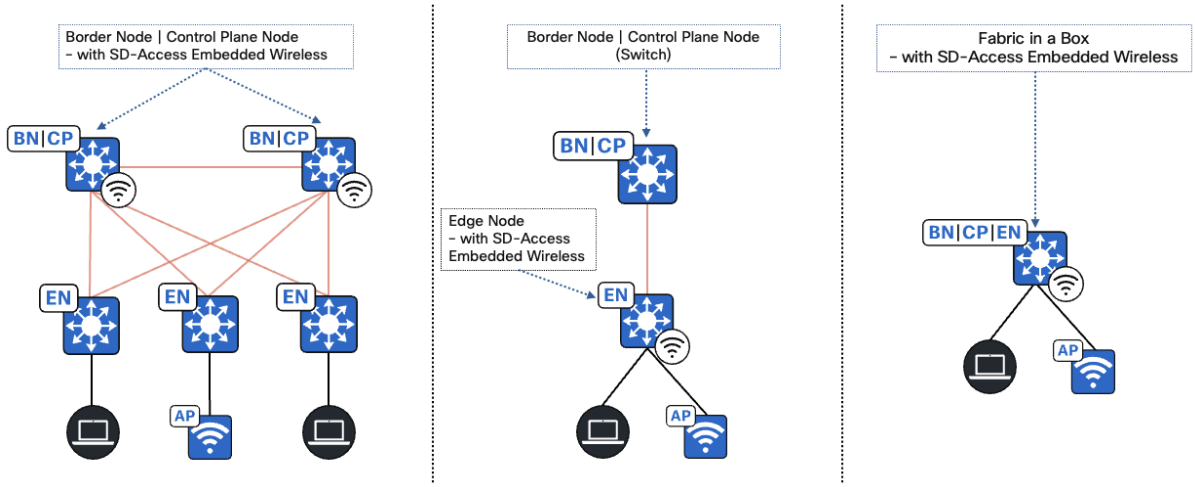

Transits

Transits can connect multiple fabric sites or link a fabric site to non-fabric domains such as a data center or the internet. Transits are a Cisco SD-Access construct that defines how Catalyst Center automates the border node configuration for connections between fabric sites or between a fabric site and an external domain. The two types of transits, include:

- IP-based transit:

  With IP-Based Transits, the Fabric VXLAN header is removed, leaving the original native IP packet. Once in native IP form, packets are forwarded using traditional routing and switching protocols between Fabric Sites. Unlike an SD-Access Transit, an IP-Based Transit is provisioned with a VRF-Lite connection to an upstream peer device. IP-Based Transits typically connect to a data center, WAN, or the Internet. Use an IP-Based Transit to connect to shared services using a VRF-A ware Peer.

- SD-Access transit:

  An SD-Access Transit uses VXLAN encapsulation and does not rely on a VRF-Lite connection to an upstream peer. Similar to IP-Based Transits, packets are forwarded using traditional routing and switching protocols between Fabric Sites. However, unlike IP-Based Transits, an SD-Access Transit is an overlay that operates on top of a W AN/MAN network, much like SD-W AN and DMVPN.

For more details about Cisco SD-Access components and architecture, see Cisco SD-Access.

IP-based transit and SD-Access transit comparison

IP-based transit:

- Leverages existing IP infrastructure:

  Uses traditional IP-based routing protocols to connect fabric sites.

- Requires VRF remapping:

  VRFs and Security Group Tags (SGTs) require to be remapped between sites, adding complexity.

- Supports existing IP networks:

  This approach is ideal if you already have an established IP-based WAN infrastructure.

- Offers flexibility:

  Provides more flexibility in terms of routing protocols and traffic engineering options.

SD-Access transit:

- Is the native SD-Access fabric:

  Uses LISP, VXLAN, and CTS for inter-site communication.

- Preserves SGTs:

  Maintains SGTs across fabric sites, enhancing security and policy enforcement.

- Centralizes control:

  Uses a domain-wide Control Plane node for simplified management.

- Requires dedicated infrastructure:

  Requires additional infrastructure for the SD-Access transit control plane.

Cisco Catalyst 9000 series switches

Cisco Catalyst 9000 series switches offer more flexible and highly scalable design options. Switches supported in different fabric roles offer secure, fast, and reliable connectivity to users and endpoints within the network.

For more details, see Catalyst 9000 switches.

Cisco Catalyst wireless controller and AP

Cisco Catalyst 9800 series wireless controllers and APs provide seamless network management and deployment in both on-premises and cloud for wireless clients.

See the data sheet for Catalyst 9800 and Catalyst 9100 devices:

- Cisco Catalyst 9800 Series
- Cisco Catalyst 9100 Series
- Cisco Access Point and Wireless Controller Selector

Compatibility matrix

Catalyst Center provides coverage for Cisco enterprise switching, routing, and mobility products.

For a complete list of supported Cisco products, see the compatibility matrix:

- Cisco Catalyst Center Compatibility Matrix
- Cisco SD-Access Compatibility Matrix

Profile deployment

This section provides design guidance for the education sector, emphasizing the requirements and the use of Cisco SD-Access to create a network that is simple, secure, and flexible. This section discusses topologies, use cases, and solutions tailored to meet standard deployment options for universities while also addressing the specific themes and requirements.

Solution topology

Figure 3. Topology solution for large-scale university deployments

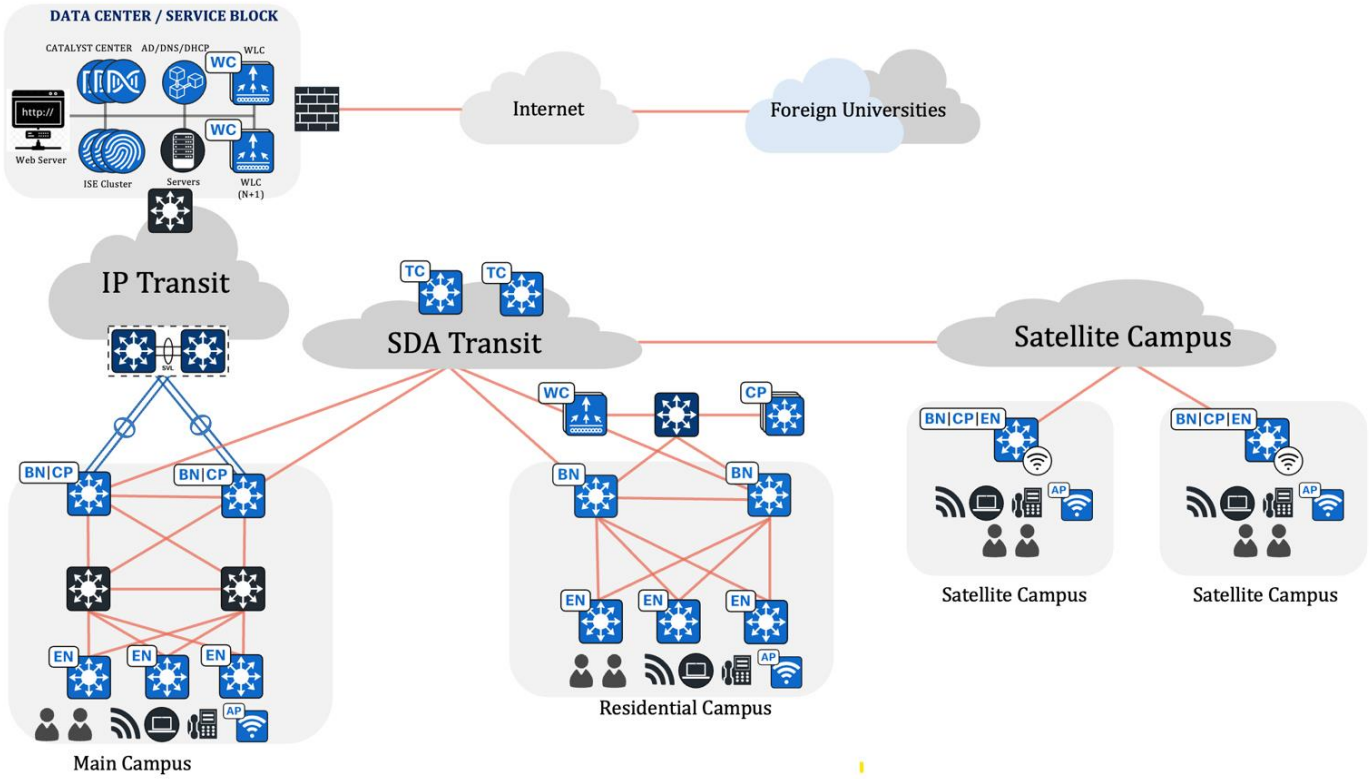

University business outcomes and challenges

A modern network infrastructure tailored to specific business outcomes empowers universities to remain competitive, drive innovation, and meet the evolving expectations of students, faculty, and the broader academic community. As universities increasingly rely on advanced and dependable network infrastructures to achieve their educational and operational goals. However, some of the universities' challenges and potential outcomes include:

- Security
- Compliance
- Operational
- Financial
- Experience

Security

For universities, enhancing security measures, mitigating risks, and ensuring regulatory compliance can be accomplished by implementing robust security protocols, conducting regular risk assessments, and adhering to relevant industry regulations and standards. Given the open and dynamic nature of university environments, which must balance accessibility for students, faculty, and researchers with safeguarding sensitive data, universities are particularly vulnerable to security threats. If not properly managed, malicious actors can exploit vulnerabilities, leading to substantial financial and reputational harm.

Operational

For universities, maintaining network uptime is crucial for seamless operations and the achievement of business objectives. Given the mission-critical nature of university networks, the primary goal is to approach 100% availability. Achieving five-nines availability (99.999% uptime) significantly advances this objective, permitting only 5 minutes and 16 seconds of downtime annually. Uninterrupted services are vital for ensuring the productivity of students and researchers, as well as supporting overall institutional success. By leveraging automation, monitoring, load balancing, and failover mechanisms, universities can attain or even exceed the five-nines availability benchmark.

Financial

Implementing a modern network infrastructure provides substantial financial advantages for universities by optimizing resource utilization, reducing costs, and creating opportunities for innovative revenue generation. For example, automating deployments across multiple campuses helps streamline expenses and increase efficiency. By embracing modern network solutions, universities can enhance operational effectiveness and service delivery while positioning themselves for long-term financial growth in an increasingly digital and competitive educational environment.

Experience

Optimize user and application experiences by leveraging modern technologies that enable critical business capabilities. While security, compliance, and availability are vital, a network with inconsistent or slow Quality of Service (QoS) can negatively impact user satisfaction and productivity. In time-sensitive environments, where delays can be critical, ensuring low latency and reliable QoS is essential to meet institutional demands effectively.

Solutions to university business outcomes

This section outlines solutions to help achieve the business outcomes defined for a university network deployment.

Security challenges

The education sector encounters substantial security challenges due to its intricate and ever-evolving environments. These include expanded attack surfaces, data breaches, insider threats, regulatory compliance requirements, advanced cyberattacks, and securing remote work. The SD-Access framework effectively tackles these issues with a robust suite of tools and capabilities:

- Macrosegmentation
- Microsegmentation
- Policy enforcement models
- Group-Based Policy Analytics
- AI endpoint analytics
- Endpoint security with zero-trust solution
- Isolation of guest users

Macrosegmentation

For university networks, assign different VRF instances to network endpoints, such as students, monitoring devices, and guests, to implement a recommended segmentation strategy. SD-Access provides the ability to macrosegment endpoints into distinct VRFs, which can be configured within the network using Catalyst Center.

Examples demonstrating the implementation of VNs include:

- INFRA VN:

  This VN is exclusively for APs, classic extended nodes, and policy extended nodes to ensure connectivity. It is mapped to the global routing table.

- Student VN:

  Use this VN for regular student access, ensuring secure and segregated connectivity for all internal users.

- Guest VN:

  This VN provides internet access to visitors and guests while ensuring they cannot access the internal network.

- Monitoring VN:

  Use this VN to dedicate network monitoring and management devices, ensuring they are isolated from regular user traffic.

- Resident VN:

  Use this VN for resident users, ensuring secure and segregated connectivity for all internal users.

By implementing VNs in an SD-Access network, a university can effectively segment and secure diverse types of traffic, enhancing overall network performance and security.

Microsegmentation

Microsegmentation simplifies the provisioning and management of network access control by using security groups to classify traffic and enforce policies, allowing for more granular security within SD-Access VNs.

Typically, within a single VN, you should further segment by grouping employees based on their department or placing devices such as printers in different security groups. Traditionally, this was done by placing groups in different subnets enforced by IP ACLs. However, Cisco SD-Access provides the flexibility of microsegmentation, allowing the use of the same subnet with a user and endpoint-centric approach. Dynamic authorization assigns different SGTs based on authentication credentials and Security-Group Access Control Lists (SGACLs) enforces these SGT-based rules.

When users connect to the network, they are authenticated using methods such as 802.1X and MAC authentication bypass (MAB). Network authorization then classifies the user's traffic using information such as identity, LDAP group membership, location, and access type. This classification information is propagated to a network device that enforces the dynamically downloaded policy, determining whether the traffic should be allowed or denied.

For more details, see the Software-Defined Access Macro Segmentation Deployment Guide.

Figure 4. Example illustrating both macrosegmentation and microsegmentation

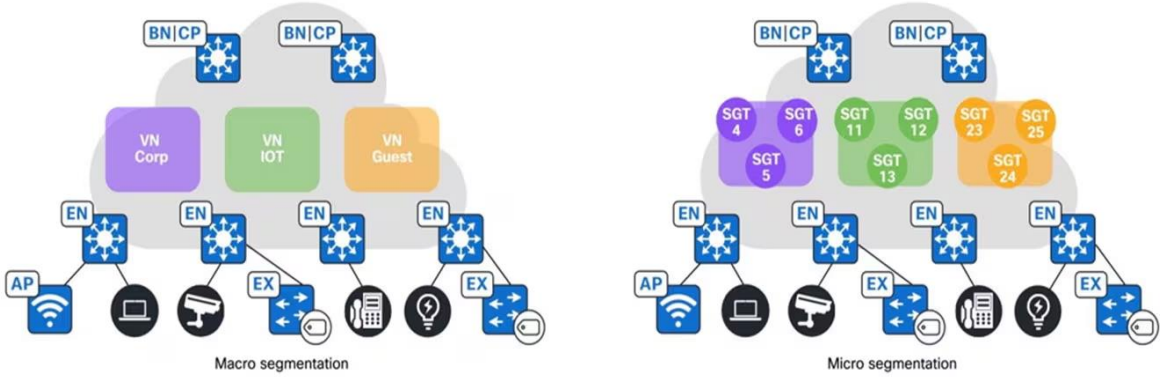

Policy enforcement models

Cisco TrustSec is a security solution designed to simplify network access provisioning and management while enforcing security policies across an organization. It enables comprehensive segmentation and access control based on roles and policies rather than traditional IP-based methods, enhancing security and operational efficiency across wired and wireless environments.

In computing and network security enforcement, policy enforcement models generally fall into two categories:

- Deny-list model (default permit IP):

  The default action permits IP traffic, and any restrictions must be explicitly configured using Security Group Access Lists (SGACLs). Use this model when there is an incomplete understanding of traffic flows within the network. It is relatively easy to implement.

- Allow-list model (default deny IP):

The default action denies IP traffic, so the required traffic must be explicitly permitted using SGACLs. Use this model when the customer has a good understanding of traffic flows within the network. This requires a detailed study of the control plane traffic, as it can block all traffic upon activation.

For more details, see Cisco ISE TrustSec Allow-List Model (Default Deny IP) with SD-Access.

Group-Based Policy Analytics

High-profile cyber-attack news is driving universities to move beyond perimeter security and implement internal network segmentation. However, the lack of visibility into user and device behavior within the network makes it difficult to create effective segmentation policies. Businesses are seeking solutions to navigate this complex landscape.

Cisco offers a solution on Catalyst Center that addresses these challenges by providing Group-Based Policy Analytics (GBPA). GBPA provides capabilities including:

- Discover and visualize group interactions:

  GBPA analyzes network traffic flows to identify how different network groups communicate, such as departments, functions, and so on.

- Identify communication patterns:

  GBPA pinpoints the specific ports and protocols used by different groups, providing granular insights into network behavior.

- Simplify policy creation:

  GBPA streamlines the process of building effective security policies to control communication between groups based on the discovered information.

Figure 5. GBPA information sources to create a holistic view of your network

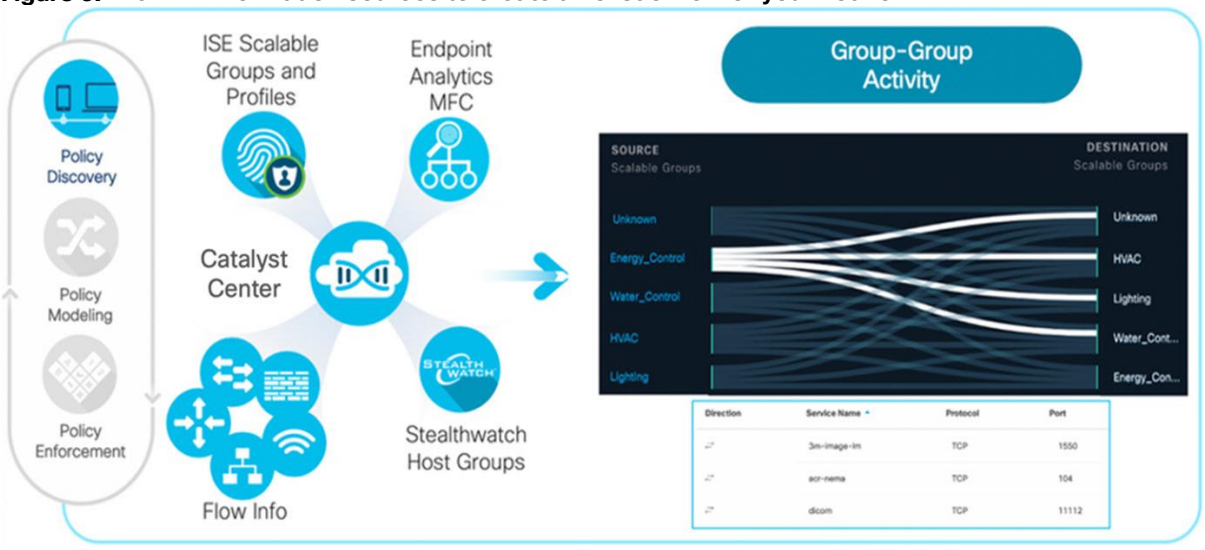

GBPA leverages information from these sources to create a holistic view of your network:

- Cisco ISE:

  When integrated with ISE, GBPA learns about network groups defined as Scalable Groups (SGTs) and Profile Groups, which categorize different types of connected devices.

- Endpoint Analytics:

  Endpoint Analytics leverages machine learning and multifactor classification to reduce unidentified devices on the network and provides more accurate profile groups for segmentation.

- Cisco Secure Network Analytics (Optional):

  Integration with Cisco Secure Network Analytics (SNA) allows GBPA to learn about Host Groups identified by SNA, further enriching network visibility.

- NetFlow data integration:

  GBPA incorporates NetFlow data from network devices to provide context for group information. This combined data is then visualized through graphs and tables, enabling administrators to clearly understand network behavior based on group interactions.

GBPA empowers network administrators with network discovery, visualization, and the tools to analyze security policy requirements. This comprehensive approach leads to the creation of more effective and targeted security policies for today's dynamic threat landscape.

AI Endpoint Analytics

Cisco AI Endpoint Analytics, next-generation endpoint visibility solution, provides deeper insights from your network and IT ecosystem, making all endpoints visible and searchable. It detects and reduces the number of unknown endpoints in your enterprise using the following techniques:

- Deep packet inspection (DPI):

  Gathers deeper endpoint context by scanning and understanding applications and communication protocols for IT, Building Automation, and Healthcare endpoints.

- Machine learning (ML):

  Intuitively groups endpoints with common attributes and helps IT administrators label them. These unique labels are then anonymously shared with other organizations as suggestions, assisting in reducing the number of unknown endpoints and grouping them based on new labels.

- Integrations with Cisco and third-party products:

  Provides additional network and non-network context to profile endpoints.

In summary, Cisco AI Endpoint Analytics addresses a critical challenge faced by many customers when implementing security policies: overcoming the lack of endpoint visibility, with high fidelity. It is available in Catalyst Center Release 2.1.2.x and higher as a new application. Customers with a subscription level of Cisco Catalyst Advantage and higher will have access to Cisco AI Endpoint Analytics. This technology primer will explore Cisco AI Endpoint Analytics and the benefits it offers to Cisco customers.

For more details, see these resources:

- Cisco SD-Access AI Endpoint Analytics
- Cisco Catalyst Center Guide - AI Endpoint Analytics

Endpoint security with zero-trust solutions

Endpoint security with zero-trust solutions in SD-Access is a comprehensive approach to network security that aims to protect endpoints, such as laptops, smart phones, and IoT devices, within an SD-Access environment. Applying zero-trust principles means that no device or user is automatically trusted, even if they are inside the network perimeter. Before granting access to network resources, each device is verified and authenticated.

The Cisco SD-Access zero-trust security solution provides the capability to automate network access policies using these features:

- Endpoint visibility:

You can identify and group endpoints. You can map their interactions through traffic flow analysis and define access policies.

- Trust monitoring:

  You can continuously monitor the endpoint behavior, scan for vulnerabilities, verify trustworthiness for continued access, and isolate rogue or compromised endpoints.

- Network segmentation:

  You can enforce group-based access policies and secure network through multilevel segmentation.

Cisco SD-Access can enforce the secure onboarding of network devices such as APs and switches using IEEE 802.1x mechanisms. This protects the network from unauthorized device attachment by maintaining closed authentication on all edge node access ports. Switches onboarded securely using closed authentication are called supplicant-based extended nodes (SBENs).

SBENs are provisioned as policy extended nodes by Catalyst Center to have a supplicant with EAP-TLS authentication on their uplink to the edge node. The EAP-TLS certificate is provisioned by Catalyst Center using the Catalyst Center Certificate Authority (CA). After successful onboarding, access to the port is purely based on authentication status. If the device or port goes down, the authentication session is cleared, and traffic is not allowed on the port. When the port comes back, it goes through dot1x authentication to regain access to the Cisco SD-Access network.

Secure AP onboarding authorizes the AP on a closed authentication port, allowing limited access to DHCP/DNS and Catalyst Center for the PnP workflow. The PnP workflow on Catalyst Center enhances to enable a dot1x supplicant on the AP, which the AP uses to authenticate with Cisco ISE.

For more details about SBEN, see the “Steps to Configure Supplicant-Based Extended Nodes” section in the Cisco Catalyst Center User Guide.

Isolation of guest users

Guest wireless isolation is a crucial security feature that ensures guest users remain completely separated from university networks while providing controlled access to the internet. To address this requirement, SD-Access introduces the Multisite Remote Border (MSRB) solution.

This solution allows traffic from a VN spanning multiple dispersed sites to be aggregated at a central location, known as the anchor site. Instead of configuring separate subnets for each site, the anchor site leverages a single, shared subnet for guest VN. By implementing a centralized and streamlined subnet structure, VN anchors simplify guest service deployments across multiple locations while maintaining consistent and secure traffic segmentation in university environments.

Enhancing security with external gateways

In SD-Access, a default gateway is present on all edge nodes for each subnet in a VN within a given fabric site. Traffic destined for a remote subnet is processed by the default gateway on the edge node and then routed to the appropriate destination.

In many networks, the default gateway needs to be on an external firewall rather than on the local edge node. Firewall traffic inspection is a common security and compliance requirement in such networks. By enabling the gateway outside of the fabric functionality, the default gateway is not provisioned on the edge nodes. Instead, it can be provisioned on an external device, such as a firewall, allowing traffic to be inspected before reaching its destination.

Seamless and secure campus network access

The increasing use of personal mobile devices and the growing dependence on cloud-based services have reshaped how universities handle technology and network access. Students and faculty now expect uninterrupted access to university resources from their own devices, regardless of location.

This evolution has transformed network security, access control, and resource management. Traditional approaches that rely on institution-owned devices and restricted networks are becoming outdated. To keep pace with these changes, universities are embracing BYOD strategies to provide flexibility while maintaining data security and regulatory compliance.

By integrating BYOD with SD-Access, universities can create a secure, adaptable network environment. SD-Access enforces access controls, protects sensitive data through network segmentation, and automates device onboarding, reducing the burden on IT teams. This approach strengthens security, improves operational efficiency, and fosters a more connected, technology-driven learning experience.

Compliance regulations

Compliance regulations are the rules and standards that institutions must follow to operate lawfully within a specific domain or jurisdiction. Advanced technologies assist in ensuring compliance by automating regulatory processes, improving data security, and enabling real-time monitoring and reporting to meet regulatory requirements effectively.

Staying compliant with industry regulations can be a complex task. Cisco SD-Access offers several features that can simplify this process:

- Role-Based Access Control
- Audit logs
- Configuration compliance
- Configuration drift

Role-Based Access Control

Role-Based Access Control (RBAC) in Catalyst Center provides a way to control access to features and operations based on the roles of individual users within the organization. RBAC helps enforce the principle of least privilege, ensuring that users have access only to the resources necessary for their roles. Catalyst Center supports the flexibility to assign permissions to users based on either a local or external RADIUS/TACACS database. You can assign roles to users and also grant access to specific applications within Catalyst Center.

Audit logs

Audit logs are a record of events or actions that have occurred within the Catalyst Center application. These logs typically include details such as who performed the action, what action was taken, and when it occurred. Audit logs are important for security and compliance purposes, as they help administrators track changes made to the network infrastructure, identify potential security breaches, and ensure that users are following proper procedures. By reviewing audit logs, administrators can gain insight into the activities within the Catalyst Center application and take appropriate actions as needed.

For more details, see View audit logs.

Configuration compliance

Compliance helps in identifying any intent deviation or out-of-band changes in the network that may be injected or reconfigured without affecting the original content. A network administrator can conveniently identify devices in Catalyst Center that do not meet compliance requirements for the various aspects of compliance, such as software images, PSIRT, network profiles, and so on.

You can automate compliance checks or do on demand compliance checks using these schedule options:

- Automated compliance check:

  Uses the latest data collected from devices in Catalyst Center. This compliance check listens to the traps and notifications from various services, such as inventory and SWIM, to assess data.

- Manual compliance check:

  Let's you manually trigger the compliance in Catalyst Center.

- Scheduled compliance check:

  A scheduled compliance job runs every day at 11:00 pm and triggers the compliance check for devices that have not undergone a compliance check in the past seven days.

Catalyst Center currently supports the following types of compliance checks:

- Flag compliance errors when running configuration on network devices differs from the startup configuration view that Catalyst Center has for the device.
- Software image compliance flag to indicate if the golden image is not running on network devices.
- Flag fabric compliance errors if the configurations deployed by the SD-Access Fabric workflows were tampered with, breaching out-of-band PSIRT compliance, to alert network administrators to existing vulnerabilities in the network.
- Network compliance alerts if the devices are not running configuration per the intent called out for the given site in Catalyst Center.

For more details, see Compliance User Guide.

Operational efficiency

Operational efficiency is crucial for universities, as it directly impacts productivity, cost management, and the quality of services provided to students and staff. By optimizing workflows and leveraging digital transformation, universities can maximize faculty and staff output, enhance administrative processes, and strengthen their reputation and institutional value. SD-Access addresses critical aspects of operational efficiency, including:

- HA
- System resiliency
- Reports
- Efficient troubleshooting

High availability

High availability (HA) is a critical component that ensures systems and applications remain operational and accessible to users with minimal interruption, even during technical setbacks such as hardware failures or software glitches. An overview for achieving HA for the components includes:

- Disaster recovery
- Resilient network architecture

- Fallback segments

Disaster Recovery

Universities have a low tolerance for disruptions in critical systems, including management, control, or data operations. The Catalyst Center ensures both intracluster and intercluster resiliency to safeguard operations. Its Disaster Recovery framework consists of three key components: the primary site, the recovery site, and the witness site. At any time, the primary and recovery sites operate in active or standby roles. The active site oversees network management, while the standby site maintains a continuously updated replica of the active site's data and services. In the event of an active site failure, the Catalyst Center automatically initiates a failover, transitioning the standby site into the active role to ensure continuity.

For more details, see Implement Disaster Recovery.

Resilient network architecture

Resilient network architecture in Cisco SD-Access is designed to ensure a highly available and reliable infrastructure, allowing critical university services to remain operational even during disruptions.

- Similar to the Virtual Switching System (VSS), StackWise Virtual (SVL) simplifies Layer 2 operations by combining two physical switches into a single logical switch at the control and management plane levels. This eliminates the need for spanning tree protocols and first-hop redundancy configurations, streamlining network management.
- By implementing Layer 3 routed access, the boundary between Layer 2 and Layer 3 shifts from the distribution layer to the access layer. This shift reduces the reliance on the distribution and collapsed core layers to manage Layer 2 adjacency and redundancy.

In university networks, traditional resilience techniques such as stacking and StackWise Virtual are complemented by strategies to safeguard regional hubs and campus headquarters against building-level failures. These measures ensure uninterrupted connectivity to data centers for critical academic and administrative applications.

Cisco SD-Access offers a flexible deployment model, enabling fabric borders to span multiple physical sites while integrating them seamlessly within a single fabric site, as illustrated in Figure 6:

Figure 6. Buildings 1 through 4 belong to the same fabric site with the colocated border nodes and control plane nodes located in different buildings

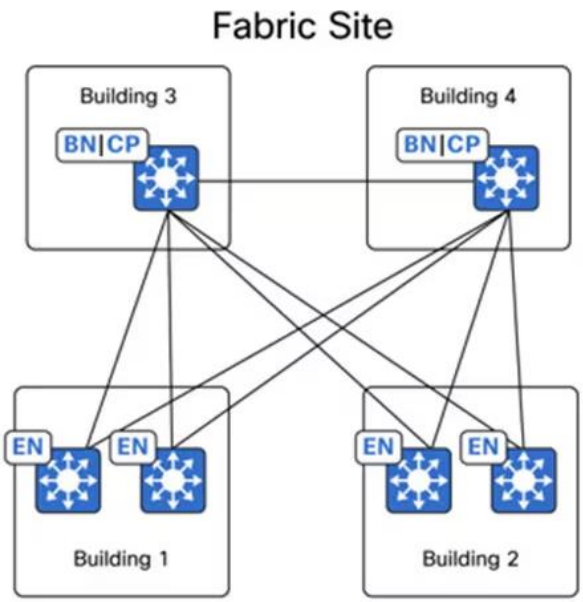

Cisco SD-Access offers the flexibility to designate priorities to these border node deployments. This allows for the prioritization of a border node or its exclusive use as the active border for traffic. In the event of a building failure, the border node in the alternate building can seamlessly assume all traffic from the edge nodes.

Fallback segments

In Cisco SD-Access, there is support for the Critical VLAN feature, which ensures that endpoints maintain a minimum level of network connectivity even when they lose connectivity to their ISE server due to outages like a WAN outage.

For wired clients (not applicable to wireless clients) that have already been onboarded, if the connection to the ISE Policy Service Node (PSN) is lost, the system pauses periodic re-authorization to prevent disruptions in the authentication path from affecting the data plane. For clients that have not yet been onboarded, the Critical VLAN feature assigns them to a specific VLAN if connectivity to ISE is lost, providing them with limited network access.

These Critical VLANs can use microsegmentation to enforce policies in the absence of ISE, but to achieve this, assign a security group during the provisioning of the anycast gateway for the critical VLAN such as VLAN-SGT mapping and configure the appropriate policy matrix to be downloaded onto the switches.

In summary, Critical VLAN in SD-Access ensures that even when devices cannot authenticate properly, they are not entirely disconnected from the network but are given limited access for remediation and troubleshooting purposes.

System resiliency

To ensure system resiliency, it is important to implement high availability and redundancy solutions for critical components of the network infrastructure. An overview about achieving this for these components includes:

- Catalyst Center HA
- ISE HA

- Cisco Wireless LAN Controller Redundancy

Catalyst Center HA

Catalyst Center HA is a feature designed to minimize downtime and increase network resilience. It achieves this by ensuring that critical services remain available in the event of hardware or software failures. HA in Catalyst Center typically involves deploying redundant hardware and software configurations to provide seamless failover and continuous operation. This helps organizations maintain network stability and reliability, even during unexpected events.

For more details, see Cisco Catalyst Center High Availability Guide.

ISE HA

Cisco ISE can be deployed in two main configurations: Standalone and Distributed.

- Standalone deployment:

  In a standalone deployment, a single ISE node serves all the necessary functions, including administration, policy services, and monitoring. This configuration is suitable for smaller networks where a single node can handle the workload and redundancy is not a critical requirement.

- Distributed deployment:

  In a distributed deployment, ISE nodes are distributed across multiple physical or virtual machines to provide scalability, redundancy, and high availability. This configuration is suitable for larger networks where scalability and redundancy are important.

Each deployment option has its own advantages and is chosen based on the specific requirements of the network in terms of scalability, redundancy, and performance. To support failover and to improve performance, you can set up a deployment with multiple Cisco ISE nodes in a distributed fashion.

For more details, see " Distributed Deployment Scenarios" in the Cisco Identity Services Engine Installation Guide.

Cisco Wireless Controller redundancy

Cisco Wireless Controller redundancy is essential for maintaining continuous wireless network services. In an HA pair setup, two wireless controllers are configured as a pair. One wireless controller functions as the primary (active) controller, managing all wireless clients and traffic, while the other serves as the secondary (standby) controller. The secondary controller stays synchronized with the primary controller's configuration and state.

If the primary controller encounters an issue, the secondary controller seamlessly takes over, ensuring uninterrupted wireless service. This redundancy feature significantly improves the reliability of wireless networks, providing failover capabilities in the event of wireless controller hardware or software failures. Consequently, users experience minimal disruption and maintain connectivity to the wireless network.

For more details, see Cisco Catalyst 9800 Series Wireless Controllers HA SSO Deployment Guide.

Reports

The Catalyst Center Reports feature provides a comprehensive suite of tools for deriving actionable insights into your network's operational efficiency. This feature enables data generation in multiple formats, with flexible scheduling and configuration options, allowing for tailored customization to meet your specific operational needs.

The Reports feature supports various use cases that include:

- Capacity planning:

  Understanding device utilization within your network.
- Pattern change analysis:

  Tracking changes in usage patterns, including clients, devices, bands, and applications.
- Operational reporting:

  Reviewing reports on network operations, such as upgrade completions and provisioning failures.
- Network health assessment:

  Evaluating the overall health of your network through detailed reports.

By leveraging Catalyst Center's reporting capabilities, you can significantly enhance your network's operational efficiency, ensuring a smooth-running, high-performing network environment.

For more details, see the Cisco Catalyst Center Platform User Guide.

Efficient troubleshooting

Efficient troubleshooting is a critical aspect of maintaining uninterrupted operations within a university IT infrastructure. The Catalyst Center offers advanced debugging features tailored to meet these requirements effectively. These tools enable IT administrators to swiftly identify, diagnose, and resolve Catalyst Center-related issues, ensuring optimal performance of the university network systems. The tools help with troubleshooting, including:

Validation Tool:
Before Catalyst Center version 2.3.5.x, the Audit and Upgrade Readiness Analyzer (AURA) tool was used to evaluate upgrade readiness. With the implementation of the restricted shell in version 2.3.5.x, most AURA upgrade checks are now integrated into the Catalyst Center. The Validation Tool assesses both the Catalyst Center appliance hardware and its connected external systems, identifying potential issues before they impact the university network.

For more details, see these links:

- Validate Cisco DNA Center Upgrade Readiness
- Use the Validation Tool

Catalyst Center Validation Tool provide invaluable support for university network administrators. These tools enable proactive maintenance, efficient troubleshooting, and improved network stability, significantly enhancing the operational efficiency of the university IT services.

Financial efficiency

Reducing operational expenses and increasing revenue are key priorities for universities striving for financial sustainability. By automating the deployment and monitoring of large-scale, multisite networks, universities can significantly lower operational costs, streamline administrative processes, and maintain efficient IT operations. This approach allows for the management of complex campus networks with minimal manual intervention, enhancing overall efficiency and enabling better allocation of resources. Some strategies adopted by universities to achieve financial efficiency include:

- Automation and monitoring
- IP Address Management (IPAM) integration
- IT Service Management (ITSM) integration
- SD-Access extension

Automation and monitoring

Automation and monitoring are essential components of modern IT infrastructure management. Automation can include tasks such as software deployment, configuration management, system provisioning, and workflow orchestration. By automating repetitive and time-consuming tasks, organizations can improve efficiency, reduce errors, and free up human resources to focus on more strategic activities. Monitoring, on the other hand, involves continuously observing and analyzing the performance and health of IT systems, networks, applications, and services. Below is an overview of how to implement these strategies for these components:

- LAN automation
- Plug and Play and Return Material Authorization
- Software Image Management
- Intelligent Capture
- Assurance and visibility

LAN automation

LAN automation in Catalyst Center is a feature designed to simplify the deployment and management of network infrastructure by automating the configuration and provisioning of network devices. This automation reduces the complexity and potential for errors associated with manual configuration, resulting in more efficient and reliable network operations.

Cisco LAN automation provides these key benefits:

- Zero-touch provisioning:

  Network devices are dynamically discovered, onboarded, and automated from their factory-default state to fully integrated in the network.

- End-to-end topology:

  Dynamic discovery of new network systems and their physical connectivity can be modelled and programmed. These new systems can be automated with Layer 3 IP addressing and routing protocols to dynamically build end-to-end routing topologies.

- Resilience:

  LAN automation integrates system and network configuration parameters that optimize forwarding topologies and redundancy. LAN automation enables system-level redundancy and automates best practices to enable best-in-class resiliency during planned or unplanned network outages.

- Security:

  Cisco-recommended network access and infrastructure protection parameters are automated, providing security from the initial deployment.

- Compliance:

  LAN automation helps eliminate human errors, misconfiguration, and inconsistent rules and settings that drain IT resources. During new system onboarding, LAN automation provides compliance across the network infrastructure by automating globally managed parameters from Catalyst Center.

For more details, see the Cisco Catalyst Center SD-Access LAN Automation Deployment Guide.

Plug and Play and Return Material Authorization

Catalyst Center features Plug and Play (PnP) functionality, which simplifies the deployment of Cisco Catalyst switches, routers, and wireless APs. With PnP, network administrators can easily onboard new

devices to the network without the need for manual configuration. Devices with PnP capability can automatically download the required software image and configuration from a PnP server, such as Catalyst Center, making the deployment process faster and more efficient.

Catalyst Center provides support for the Return Material Authorization (RMA) processes. In case of hardware failure or replacement, the RMA feature allows administrators to easily manage the return and replacement of faulty devices. This includes generating RMA requests, tracking the status of RMAs, and managing the replacement process through a centralized interface. Overall, the PnP and RMA features in Catalyst Center help streamline device deployment and replacement processes, reducing complexity and enhancing network management efficiency.

For more details, see the Network Device Onboarding for Cisco Catalyst Center Deployment Guide.

Software Image Management

The Catalyst Center Software Image Management (SWIM) feature simplifies and automates the process of managing software images across Catalyst switches, routers, and wireless devices in the network. Network administrators who wish to automate the upgrade of a Catalyst 9000 series switch at a branch or campus can use the Catalyst Center SWIM solution.

Catalyst Center stores all unique software images according to image type and version for the devices in your network. It allows you to view, import, and delete software images and push them to your network's devices. The software upgrade can be optimized by decoupling software distribution and activation to minimize downtime within the maintenance window. Overall, SWIM enhances operational efficiency, reduces downtime, and helps ensure network security and compliance by simplifying and automating the management of software images across Catalyst devices.

For more details, see the SWIM Deployment Guide.

Intelligent Capture

Catalyst Center Intelligent Capture (iCAP) is a powerful feature designed to enhance network troubleshooting and performance monitoring. It leverages advanced analytics and machine learning to provide deep insights into network traffic and client behaviors. iCap provides support for a direct communication link between Catalyst Center and APs, so each of the APs can communicate with Catalyst Center directly. Using this channel, Catalyst Center can receive packet capture (PCAP) data, AP and client statistics, and spectrum data. With the direct link from the AP to Catalyst Center through gRPC, iCap allows you to access data from APs that is not available from wireless controllers.

For more details, see the Cisco Intelligent Capture Deployment Guide.

Assurance and visibility

Catalyst Center manages your network by automating network devices and services but also provides network assurance and analytic capabilities. Catalyst Center collects telemetry from network devices, Cisco ISE, users/endpoints, applications, and other integrations across the network. Catalyst Center Network Analytics correlates data from various sources to help administrators or operators to offer comprehensive network insights into:

- Device 360 and Client 360:

  Provides the ability to view device or client connectivity, which includes information on topology, throughput, and latency from various times and different applications.

- Network time travel:

Provides the ability to go back in time and see the cause of a network issue.

- Application experience:

  Provides unparalleled visibility and performance control on the applications critical to your core business on a per-user basis.

- Network analytics:

  Provides recommendations for corrective actions for found issues in the network. These actions can involve guided remediation, where the engine specifies steps for a network administrator to do.

For details, see Cisco Catalyst Assurance.

IP Address Management integration

IP Address Management (IPAM) integration in Catalyst Center streamlines the process of managing IP addresses within a network. This integration provides a centralized platform to automate and simplify IP address allocation, tracking, and management. In SD-Access deployments, IPAM integration provides Catalyst Center access to existing IP address scopes. When configuring new IP address pools in Catalyst Center, it automatically updates the IPAM server, reducing the IP address management tasks.

Two third-party integration modules are included in Catalyst Center, one for IPAM provider Infoblox and one for Bluecat. Other IPAM providers may be configured for use with Catalyst Center by providing an IPAM provider REST API service that meets the Catalyst Center IPAM provider specification.

For more details, see Configure an IP Address Manager.

IT Service Management integration

IT Service Management (ITSM) refers to the implementation and management of quality IT services that meet the needs of a business. ServiceNow is a popular ITSM platform that provides a suite of applications to help organizations automate and streamline their IT services.

Catalyst Center and ServiceNow integration supports these capabilities:

- Integrating Catalyst Center into ITSM processes of incident, event, change, and problem management.
- Integrating Catalyst Center into ITSM approval and preapproval chains.
- Integrating Catalyst Center with formal change and maintenance window schedules.

The scope of the integration is mainly to check your network for assurance and maintenance issues and for events requiring software image updates for compliance, security, or any other operational triggers. Details about these issues are then published to an ITSM (ServiceNow) system or any REST endpoint.

For more details, see the Cisco Catalyst Center ITSM Integration Guide.

SD-Access extension

SD-Access extension is a critical capability that allows organizations to extend the reach of their SD-Access fabric, ensuring consistent policy enforcement, enhanced security, simplified management, and improved network performance across a broader range of environments and devices.

An extended node connects to the SD-Access in Layer 2 mode, facilitating the connection of IoT endpoints but does not support fabric technology. Using Catalyst Center, the extended node can be onboarded from a factory reset state through the PnP method, enabling security controls on the extended network and enforcing fabric policies for endpoints connected to the extended node.

To implement SD-Access extension, deploy extended nodes, which are available in three different types:

- Extended Node (EX):

  Extended Node is a Layer 2 switch that connects to a fabric edge node in a Cisco SD-Access network. It provides connectivity for IoT endpoints and other devices that do not support full SD-Access capabilities. Extended nodes are typically managed and configured through a centralized controller like Catalyst Center. They rely on the fabric edge for advanced network functions like LISP, VXLAN, and SGACL enforcement.

- Policy Extended Node (PEN):

  Policy Extended Node is a specific type of extended node that offers additional capabilities. It can perform 802.1X/MAB authentication, dynamically assign VLANs and SGTs to endpoints, and enforce SGACLs. This type of node provides a more granular level of policy control compared to a standard extended node, allowing for more flexible network segmentation and security.

- Supplicant-Based Extended Node (SBEN):

  Supplicant-Based Extended Node is an extended node that undergoes a stricter onboarding process. It requires an IEEE 802.1X supplicant configuration and completes a full authentication and authorization process before being allowed into the SD-Access network. This approach enhances security by ensuring that only authorized devices can access the network. SBENs are often used in environments with heightened security requirements.

Key points:

- Extended Nodes provide connectivity for endpoints that cannot directly participate in SD-Access.
- Policy Extended Nodes offer enhanced policy enforcement capabilities.
- Supplicant-Based Extended Nodes implement stricter security measures through 802.1X authentication.

Figure 7. An enterprise network using an Extended Node (EX), a Policy Extended Node (PEN), and a

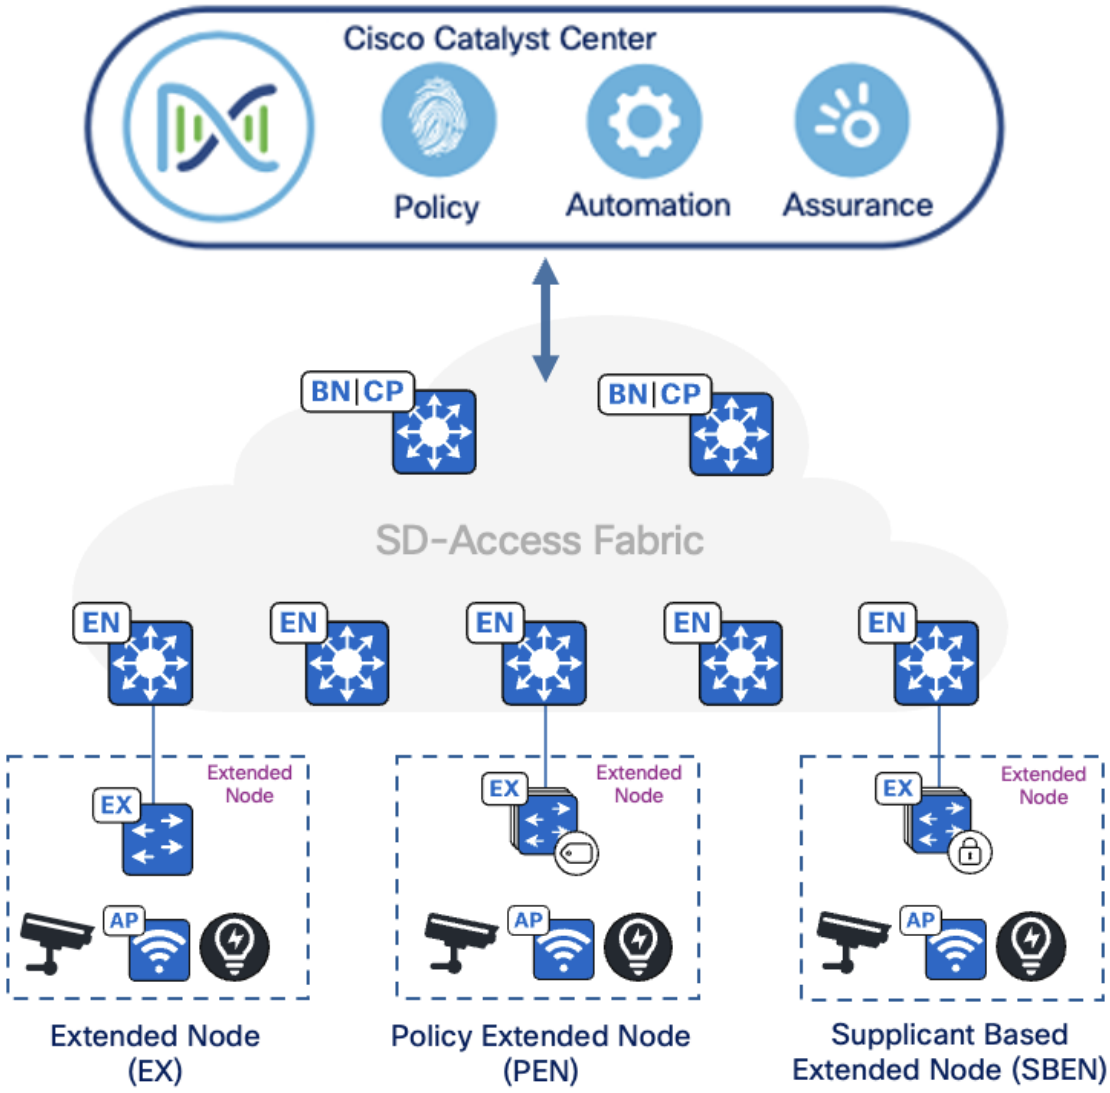

For more details, see:

- Cisco SD-Access Solution Design Guide
- Connected Communities Infrastructure - General Solution Design Guide

Experience improvement

Enhancing user and customer experiences through strategic use of modern technologies involves prioritizing Quality of Service (QoS), leveraging application visibility, and implementing video streaming, particularly in environments where performance directly impacts business operations and customer satisfaction. In today's competitive landscape, prioritizing QoS is not merely an option but a necessity for delivering exceptional user and customer experiences.

An overview of strategies for enhancing these areas, includes:

- QoS
- Application Visibility
- Video streaming across sites

QoS

QoS refers to the network capability to prioritize or differentiate service for selected types of network traffic. Configuring QoS ensures that network resources are utilized efficiently while meeting business objectives, such as ensuring enterprise-grade voice quality or delivering a high Quality of Experience (QoE) for video. Catalyst Center facilitates QoS configuration in your network through application policies.

These policies include core parameters, including:

- Application sets:

  Groups of applications with similar network traffic requirements. Each application set is categorized into business relevance groups (business relevant, default, or business irrelevant) that determine the priority of their traffic. QoS parameters for each group are defined according to Cisco Validated Design (CVD), and adjustments can be made to align with specific business goals.

- Site scope:

  Defines the scope to which an application policy applies. For example, a wired policy applies to all wired devices within the specified site scope, while a wireless policy applies to devices using a specific Service Set Identifier (SSID) within the defined scope.

Application visibility

Application visibility is a feature that allows network administrators to see which applications are running on their network, monitor their performance, and understand how network resources are being utilized. This is crucial for maintaining optimal network performance, ensuring security, and improving the user experience.

Catalyst Center empowers you to manage and gain insights into applications traversing your network. This includes identifying built-in applications, custom applications, and categorizing network traffic. The application visibility service, hosted as an application stack within Catalyst Center, enables the Controller-Based Application Recognition (CBAR) function on a specific device to classify thousands of networks, home-grown applications, and network traffic.

Application visibility is achieved through a combination of deep packet inspection, flow analysis, and application recognition technologies, providing a comprehensive view of network activity and application performance. By implementing CBAR, organizations can ensure that their critical applications work optimally, enhancing overall productivity and user satisfaction.

You can install these packages:

- Application policy:

  Allows you to automate QoS policies across LAN, WAN, and wireless within your campus and branch.

- Application registry:

  Allows you to view, manage, and create applications and application sets.

- Application visibility service:

  Provides application classification using Network-Based Application Recognition (NBAR) and CBAR techniques.

Video streaming across sites

Universities often need to conduct regular educational sessions, such as live lectures, seminars, and workshops, across multiple campuses or departments. These sessions typically involve live video broadcasts for lectures, demonstrations, and interactive Q&A discussions. To ensure efficient distribution

of video content to all campuses simultaneously without overwhelming the network, universities can utilize multicast technology.

Multicast data streams can originate from various sources, such as regional data centers or centralized academic servers. Cisco SD-Access architecture enables seamless, end-to-end multicast data traffic to flow efficiently within the university's network, whether it's local or global. SD-Access supports both headend replication and native multicast modes, offering the flexibility to designate Multicast Rendezvous Points (RP) within or outside the SD-Access fabric.

SD-Access supports two transport methods for forwarding multicast traffic: one uses the overlay, known as Head-End Replication, and the other uses the underlay, referred to as Native Multicast.

- Head-End Replication:

  Head-End Replication (or Ingress Replication) is performed either by the multicast first-hop router (FHR) when the multicast source is in the fabric overlay, or by the border nodes when the source is outside the fabric site.

- Native Multicast:

  Native Multicast eliminates the need for the ingress fabric node to perform unicast replication. Instead, the entire underlay network, including intermediate nodes, handles the replication. To support Native Multicast, the first-hop routers (FHRs), last-hop routers (LHRs), and all intermediate network infrastructure components must be multicast-enabled.

By integrating multicast technology within the SD-Access framework, the university can effectively deliver large-scale training sessions, fostering better communication and learning across all campuses.

Simplifying academic mobility

In today’s academic world, global collaboration is essential for students, faculty, and researchers. Universities increasingly focus on interdisciplinary work, partnerships, and international projects, making seamless communication and resource sharing crucial. However, accessing Wi-Fi at different institutions often requires temporary guest accounts, manual configurations, or IT support, leading to inconvenience and security risks.

Eduroam solves this problem by providing secure, hassle-free wireless access in over 100 countries and thousands of institutions. Users can connect instantly using their home institution’s credentials, eliminating the need for separate logins or complex setups. This ensures uninterrupted internet access, supporting academic and research activities while promoting global mobility. Universities benefit from improved security, reduced IT workload, and a connected academic community that fosters knowledge sharing and innovation.

By leveraging SD-Access, universities can efficiently incorporate Eduroam into their existing network, ensuring consistent and secure internet access for students, staff, and faculty—both on campus and at any participating Eduroam institution worldwide.

Seamless service discovery

In large-scale networks such as universities, enterprises, and multi-campus environments, finding and accessing networked devices and services, like printers, Apple TVs, and shared resources, can be complex. Traditional Bonjour service discovery is designed for small, local networks and relies on multicast communication, which does not scale effectively across VLANs or subnets. As a result, users struggle to discover and connect to services beyond their immediate network segment.

Wide Area Bonjour addresses this challenge by extending service discovery across different VLANs, subnets, and even geographically dispersed locations. By integrating with SD-Access and other network architectures, it enables seamless access to shared resources while ensuring efficient service advertisement, strong security, and optimal network performance. This enhances the user experience and supports modern, large-scale networking environments.

Network Deployment Options

These sections outline the deployment options for a university network.

Fabric site reference models

In university SD-Access deployments, different site templates are used for various campus locations, such as academic buildings, research centers, dormitories, or administrative offices. The primary design challenge is to assess the existing network and determine how to integrate SD-Access fabric sites into these areas. Standardizing designs into reference models simplifies this process.

These templates provide a structured approach to campus network design, categorizing sites based on factors such as network size, endpoint capacity, and architectural complexity. The guidelines provided in this section offer general recommendations and do not necessarily represent the maximum scale or performance limits of devices in a given reference model.

Site reference models include:

- FIAB site:

  Designed for small university buildings or specialized labs, supporting fewer than 1000 endpoints and 50 APs. Fabric in a Box provides resilience through switch stacking or StackWise Virtual. The border, control plane, edge, and wireless functions are colocated on a single redundant switching platform.

- Small site:

  Suitable for a single academic building or administrative office, supporting fewer than 10,000 endpoints and 500 APs. The border and control plane functions are colocated on one or two devices, while a separate wireless controller is deployed, optionally with high availability (HA).

- Medium site:

  Designed for a university building with multiple wiring closets or a group of smaller buildings, supporting fewer than 50,000 endpoints and 2,500 APs. The border and control plane functions may be colocated or provisioned on separate devices, and a dedicated wireless controller operates with HA.

- Large site:

  Intended for large university campuses or research institutions, supporting up to 100,000 endpoints and 10,000 APs. The border functions are distributed separately from the control plane on redundant devices, with multiple wireless controllers configured for HA.

Each fabric site includes a supporting set of control plane nodes, edge nodes, border nodes, and wireless controllers, sized appropriately from the listed categories. ISE PSNs are also distributed across the sites to meet survivability requirements.

Note: These reference models offer valuable guidance. Adjustments may be necessary based on your specific network requirements and constraints. To ensure optimal deployment of your SD-Access fabric, we recommend that you consult with a network design professional.

For details on fabric site reference models, see the Cisco SD-Access Solution Design Guide.

FIAB site reference model

The FIAB site reference model is designed for smaller university campuses or remote sites, typically supporting fewer than 1,000 endpoints. At the core of this design is a switch stack or StackWise Virtual, performing all three fabric roles: control plane node, border node, and edge node. In such deployments, SD-Access embedded wireless is commonly used to provide site-local wireless controller functionality.

Additionally, depending on the WAN or Internet circuit and latency, the site may include an ISE PSN for policy enforcement and authentication.

See Table 1 for general design guidelines for similar campus site sizes. These figures serve as recommendations and may not reflect exact limits for specific devices used in a deployment of this scale.

Table 1. FIAB site guideline (limits may be different)

Network element	Scale
Endpoints, target fewer than	1000
Control plane nodes	1
External border nodes	1
APs, target fewer than	50

Figure 8. Physical topology – FIAB site design

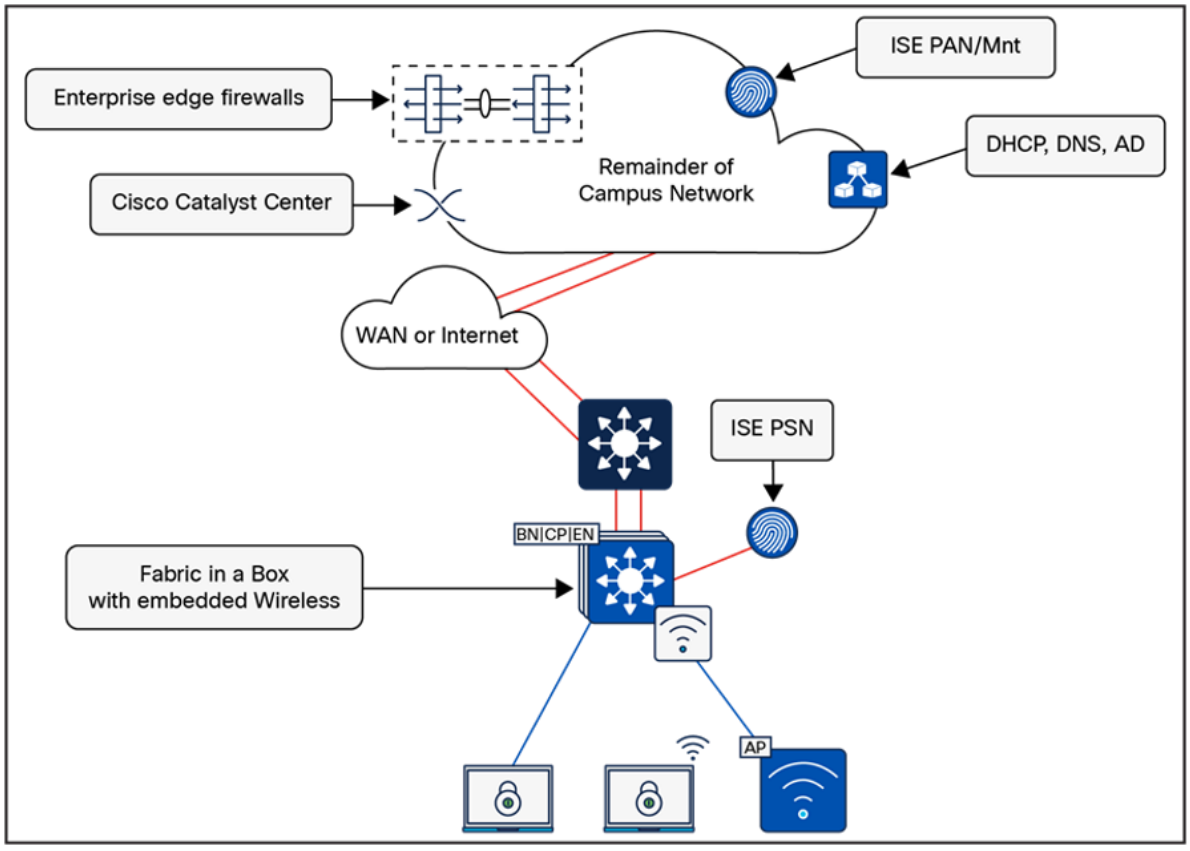

FIAB site consideration:

- Due to the smaller number of endpoints and the resulting lower impact, high availability and site survivability are not common requirements for a FIAB design. As with all reference designs, site-local services such as DHCP, DNS, wireless controllers, and ISE can enhance resiliency and survivability, though this increases complexity and requires additional equipment, such as a services block.

- High availability in this design is achieved through StackWise-480 or StackWise Virtual, both of which combine multiple physical switches into a single logical switch. If a chassis-based switch is used, redundancy is ensured through dual supervisors and power supplies.
- Wireless controllers can be deployed as physical units directly connected to the FIAB or as an embedded Catalyst 9800 controller. When using the embedded Catalyst 9800 with a switch stack or redundant supervisor, AP and client Stateful Switchover (SSO) are provided automatically.

Small site reference model

The small site reference model is designed for a single academic building or administrative office with a centralized wiring closet, typically supporting up to 4,000 endpoints and 100 APs. The physical network is usually a two-tier design, with a collapsed core/distribution layer and an access layer.

See Table 2 for general design guidelines for similar site sizes. These numbers serve as reference points and may not represent exact limits for specific devices used in such designs.

Table 2. Small site guidelines (limits may be different)

Network element	Scale
Endpoints, target fewer than	10,000
Fabric nodes, target fewer than	100
Control plane nodes	2
External border nodes	2
APs, target fewer than	500

Figure 9. Physical topology – Small site reference design

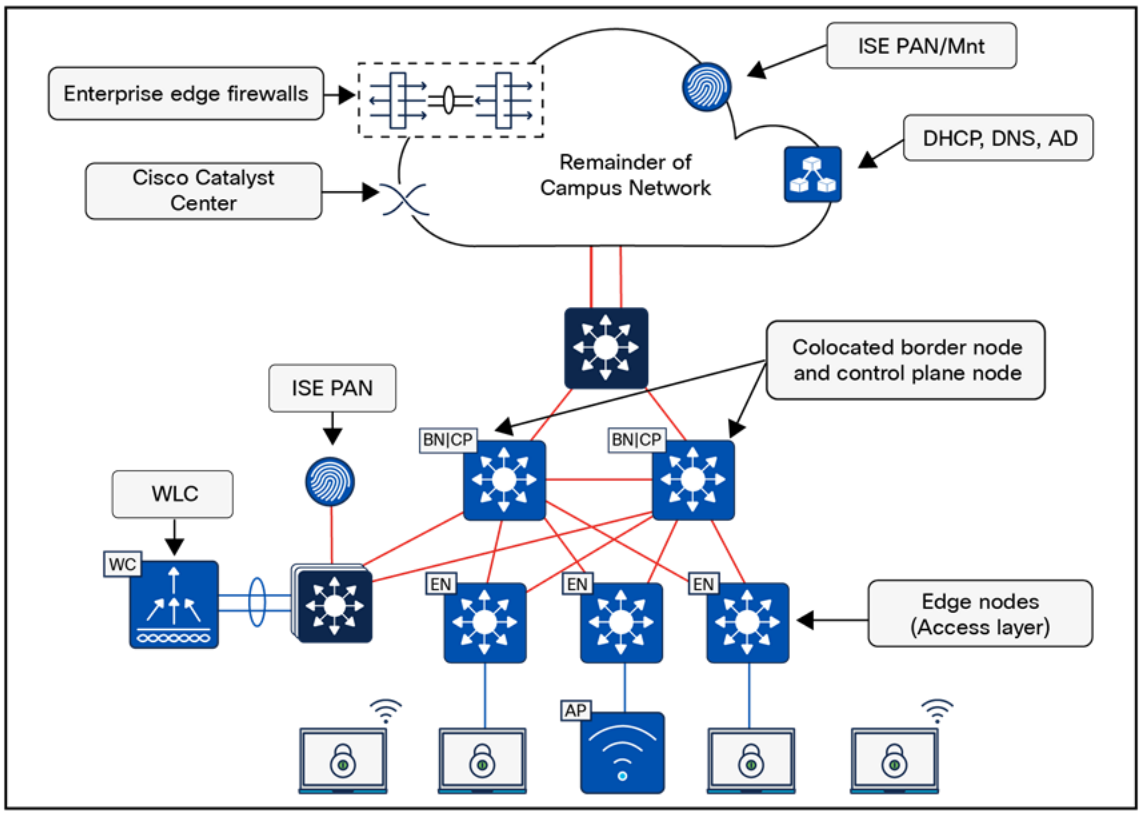

Small site considerations:

- For smaller deployments, an SD-Access fabric site is typically implemented using a two-tier design. In a small site, high availability is achieved by colocating the border and control plane node functionalities on the collapsed core switches and deploying them as a pair. To ensure resiliency and provide alternative forwarding paths in both the overlay and underlay, the collapsed core switches should be directly connected via a crosslink.
- The client and AP count necessitate the use of dedicated wireless controllers. To establish highly available links for the wireless controllers through physical connectivity, a services block is deployed. The wireless controllers connect to the services block switch via Layer 2 port channels, ensuring redundant interfaces. The services block, which consists of either a switch stack or StackWise Virtual, connects to both collapsed core switches through Layer 3 routed links. If DHCP, DNS, and other shared services are site-local, the services block may be deployed as a VRF-aware peer.

Medium site reference model

The medium site reference model is designed for university campuses with multiple buildings or a single building housing multiple wiring closets, supporting up to 50,000 endpoints. The physical network typically follows a three-tier architecture, consisting of core, distribution, and access layers. The border and control plane node functions can be colocated or deployed on separate devices.

See Table 3 for general design guidelines for sites of this scale. These numbers serve as recommendations and may not represent exact limits for specific network devices. To support the maximum endpoint capacity, a large Cisco Catalyst Center appliance is required, and in some cases, an extra-large appliance may be necessary.

Table 3. Medium site guidelines (limits may be different)

Network element	Scale
Endpoints, target fewer than	50,000
Fabric nodes, target fewer than	500
Control plane nodes (limit of 2 for FEW/SD-Access Wireless)	2 to 6
External border nodes	2
APs, target fewer than	2500

Figure 10. Physical topology – Medium site reference design

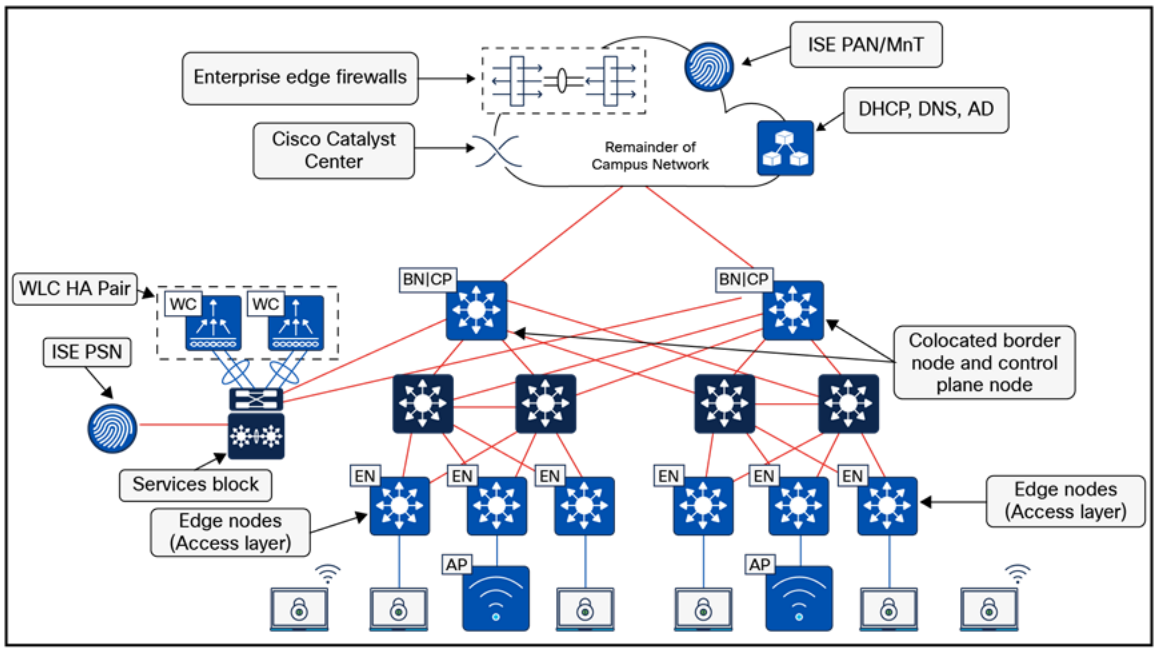

Medium site considerations:

- In a medium site, for both resiliency and alternative forwarding paths in the overlay and underlay, all devices within a given layer, except for the access layer, should be crosslinked to each other. Multiple distribution blocks do not need to be cross-connected to each other, but they should be cross-connected to all distribution switches within a block. If dedicated control plane nodes are used, they are typically connected to the core switches to ensure high availability for any edge node across various distribution blocks. For optimal forwarding and redundancy, they should be connected to both core switches. If interfaces and fiber are available, they may also be crosslinked to each other, though this is not a strict requirement.
- Physical wireless controllers should be deployed to support the wireless user scale. To enable high availability, a wireless controller HA-SSO pair is deployed with redundant physical connectivity to a services block using Layer 2 port channels. The services block is typically implemented with fixed-configuration switches operating as StackWise Virtual and connected to the core through Layer 3 routed links. This services block may function as a VRF-aware peer if DHCP, DNS, and other shared services are site-local.

Large site reference model

The large site reference model is designed for a single building with multiple wiring closets or a campus with multiple buildings. The physical network typically follows a three-tier architecture core, distribution, and access layers, and is capable of supporting up to 100,000 endpoints.

See Table 4 for general design guidelines for sites of this scale. These numbers serve as reference points and may not correspond to specific device limitations within a given design. Supporting the maximum endpoint capacity requires, at minimum, an extra-large Cisco Catalyst Center appliance and may necessitate a three-node cluster of extra-large Catalyst Center appliances. The Cisco Catalyst Center data sheet provides details on the scalability of various networking infrastructure devices used to implement an SD-Access fabric site.

Table 4. Large site guidelines (limits may be different)

Network element	Scale
Endpoints, target fewer than	100,000
Fabric nodes, target fewer than	1200
Control plane nodes	2 to 6
Border nodes (2 as internal and 2 as external) **In highly exceptional design scenarios, there may be multiple pairs of internal border nodes*.	2 to 4*
IP pools, target fewer than	1000
APs, target fewer than	10,000

Figure 11. Physical topology – Large site reference design

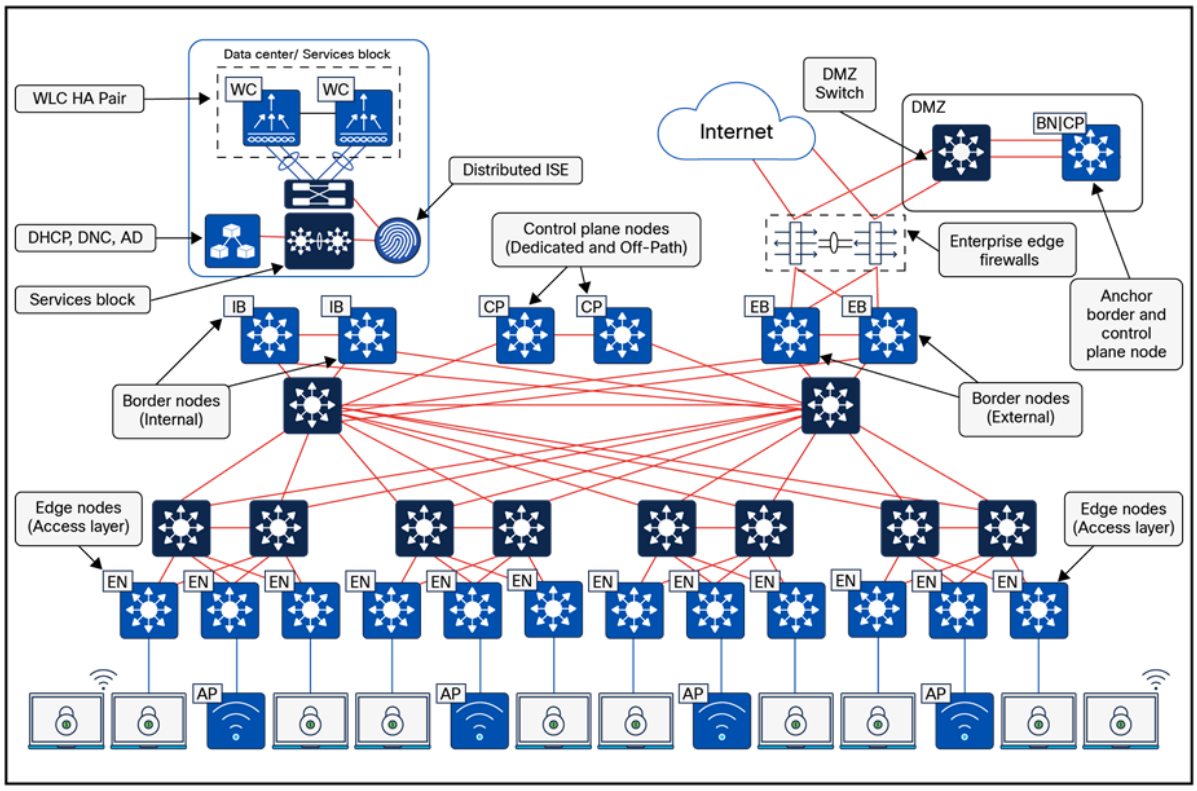

Large site considerations:

- Cisco Catalyst Center and the primary ISE Policy Administration Node are typically deployed at a large site location.
- Control plane nodes and border nodes should be dedicated devices deployed in redundant pairs. Dedicated control plane nodes should connect to each core switch to ensure resiliency and provide redundant forwarding paths. If interfaces and fiber are available, crosslinking the control plane nodes is recommended, though not required, as it offers an additional underlay forwarding path.
- One or more wireless controller HA-SSO pairs are deployed with redundant physical connectivity to a services block using Layer 2 port channels. The services block is typically part of the on-premises data center network.
- Dedicated internal border nodes are sometimes used to connect the fabric site to the data center core, while dedicated external border nodes are used to connect the site to the MAN, WAN, and internet. We recommend that you deploy the least number of border nodes that meet network design requirements, because as the quantity of border nodes increases so does the SD-Access administrative effort and routing complexity. Dedicated redundant routing infrastructure and firewalls are used to connect this site to external resources, and border nodes fully mesh to this infrastructure and to each other.
- The large site may contain the demilitarized zone (DMZ) where the anchored fabric border and control plane nodes for guest wireless are deployed.

SD-Access for distributed campus reference model

In a university environment, SD-Access for distributed campus is a solution that connects multiple independent fabric sites while maintaining security policy constructs (VRFs and SGTs) across these sites. Control plane signaling via the LISP protocol, along with fabric VXLAN encapsulation, is used between fabric sites. This ensures the preservation of macrosegmentation and microsegmentation policy constructs,

VRFs and SGTs, respectively. As a result, the network remains address-agnostic, enforcing end-to-end policies based on group membership.

In Figure 12, each fabric site is connected via a Metro Ethernet (Metro-E) private circuit. This deployment represents a large enterprise campus with multiple dispersed buildings within the same geographic area, where each building operates as an independent fabric site. The border nodes connected to this circuit are configured as external borders, colocated with a control plane node. IGP peering occurs across the circuit to establish IP reachability between the loopback interfaces (RLOCs) of the devices. The Metro-E circuit functions as the SD-Access transit between fabric sites.

The headquarters has direct internet access, while Fabric site 1 connects to the data center, where shared services are hosted. Internal border nodes at Fabric Site 1 import (register) data center prefixes into the overlay space, enabling VNs in each fabric site to access these services. For internet-bound traffic, packets are forwarded back to the headquarters, where they pass through a common security stack before egressing to the internet.

Transit control plane nodes deploy in a separate area, accessible through the SD-Access transit Metro-E network, but they are not in the direct forwarding path.

Figure 12. Physical topology – SD-Access for Distributed Campus reference design

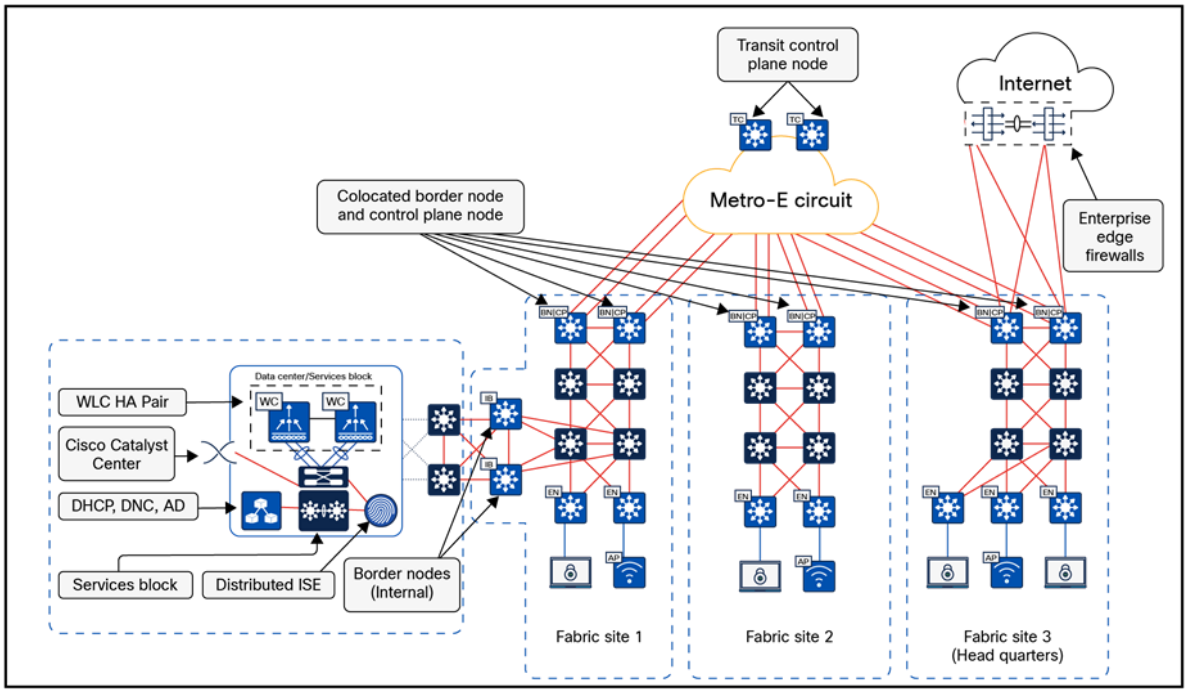

Distributed campus considerations:

- The core components enabling the distributed campus solution are the SD-Access transit and the transit control plane nodes. These architectural constructs are used exclusively in distributed campus deployments. The SD-Access transit serves as the physical network connection between fabric sites within the same city, on the same WAN, or across buildings in a large enterprise campus.

Wireless design

Designing wireless solutions within Cisco SD-Access framework in a university setting involves configuring and integrating multiple components to ensure seamless operation and management. Universities implementing an SD-Access fabric for their wired network have two options for incorporating wireless access:

- SD-Access Wireless Architecture
- Cisco Unified Wireless Network Wireless Over-the-Top

SD-Access Wireless Architecture

Cisco SD-Access provides a unique differentiator by integrating the wireless control plane with the overlay control plane of the wired world. Cisco SD-Access wireless offers a centralized control and management plane via the wireless controller with a distributed data plane providing the best of both worlds – centralized and distributed wireless designs.

The wireless controller integrates with the control plane node, registering endpoints as they are onboarded and updating their location as they roam. This is the first instance where there is synergy between the wireless and the wired control planes. This unique integration of wired and wireless brings several benefits to network users and the operations teams that support them:

- Simplification:

  Enables networks to have a single subnet for both wired and wireless clients.
- Consistency of policy:

  Extends wired policies to wireless traffic, with both enforced at the edge node.
- Improved performance:

  Removes the requirement for any form of anchoring using wireless roams on Layer 2.
- Distributed data plane:

  Enables higher overall wireless throughput compared to centrally switched wireless architectures.

Figure 13. Example of control plane and data plane traffic flow in Cisco SD-Access wireless

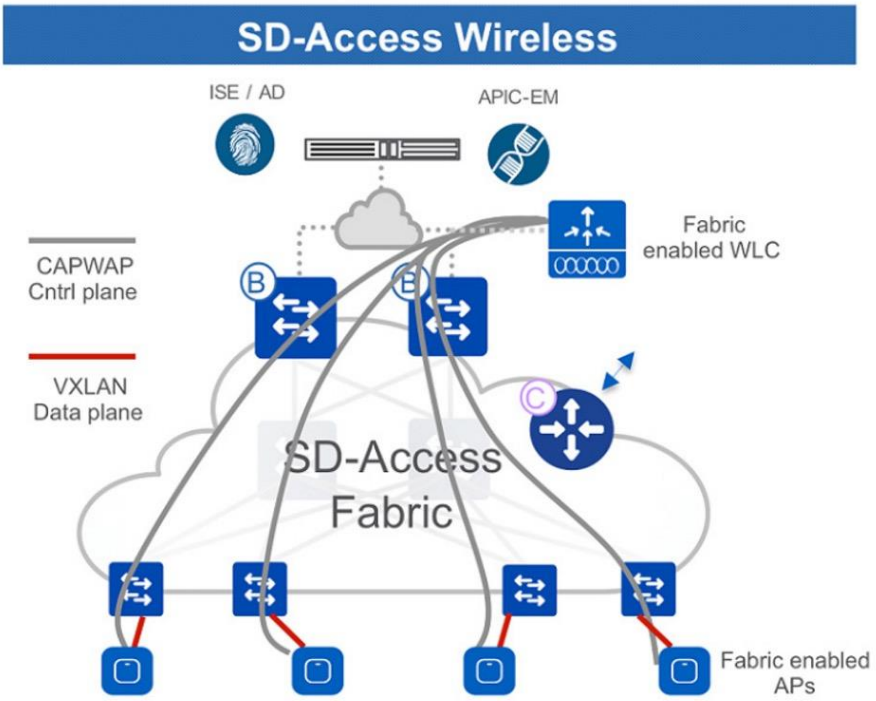

Cisco Unified Wireless Network Wireless Over-the-Top

Cisco SD-Access offers the flexibility to support a centralized wireless deployment known as wireless Over-the-Top (OTT). This support is crucial for several scenarios, such as:

- Existing Cisco wireless controllers and APs that are not SD-Access wireless-capable.
- Presence of third-party wireless devices in the network.
- Asymmetric migration pace between wired and wireless networks.

In wireless OTT deployments, wireless control, management, and data plane traffic travel through the fabric in a CAPWAP tunnel between the APs and wireless controller. This CAPWAP tunnel leverages the Cisco SD-Access fabric as a transport medium. While other vendor's wireless equipment may use different tunnelling protocols, the concept of using the SD-Access fabric as a transport remains the same.

Figure 14. Example of the flow of control plane and data plane traffic in wireless OTT:

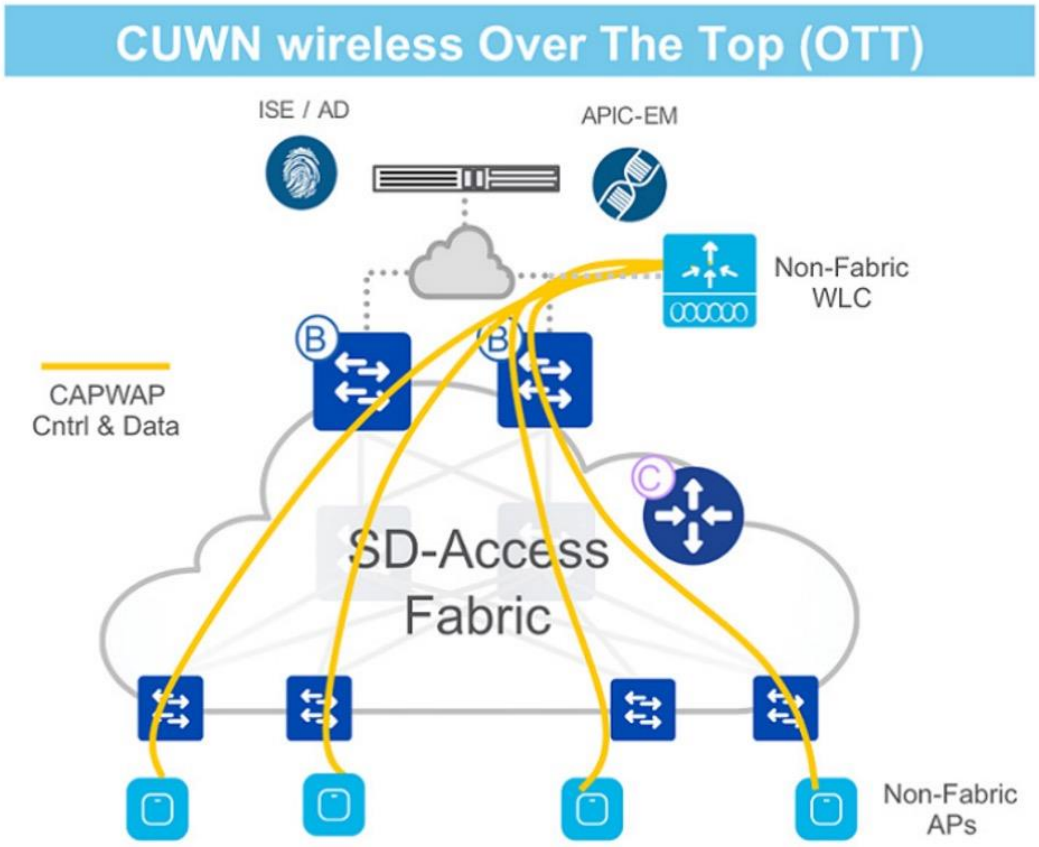

For more details, see Cisco SD-Access Wireless Design and Cisco Wireless Design and Deployment Guide.

Distributed site design

Cisco SD-Access for distributed campus is a solution designed for metro-area connectivity, linking multiple independent fabric sites while maintaining consistent security policies, such as VRF and SGT, across these sites. Although SD-Access has supported multisite environments for some time, there has not been a simple, automated method to synchronize policies between sites. Previously, at each site's fabric border node, fabric packets were de-encapsulated into native IP. While policy extension between sites was possible, it required a manual process, relied on SXP for policy propagation, and involved complex configurations of IP to SGT bindings within ISE.

With SD-Access transit for distributed campus, SXP is no longer needed, configurations are automated, and the complex mappings are simplified. This solution enables intersite communication with consistent, end-to-end automation and policy across the metro network.

SD-Access transit for distributed campus uses control plane signaling from the LISP protocol and maintains VXLAN encapsulation of packets between fabric sites. This preserves the macro and microsegmentation policy constructs of VRFs and SGTs, respectively, between fabric sites. The original Ethernet header of the packet is retained to enable the Layer-2 overlay service of SD-Access wireless. The result is a network that is address-agnostic because policy is maintained through group membership.

For more details, see Cisco SD-Access Distributed Campus Deployment Guide and Cisco SD-Access Distributed Campus Design.

MSRB

MSRB centralizes the routing of untrusted traffic within the fabric network to a designated location such as a firewall or DMZ. For example, in a scenario where a guest VN spans multiple sites, all guest traffic can be directed through a remote border located at the DMZ, effectively isolating it from enterprise traffic.

In a multisite network deployment, a designated MSRB manages traffic to and from a specific VN extended across multiple sites. This configuration enables the deployment of a VN across multiple fabric sites while maintaining a unified subnet across these locations. Consistently maintaining subnets across multiple fabric sites helps optimize IP address utilization. It establishes a centralized entry and exit point for that VN, providing several advantages:

- Centralized control:

  You can designate a common border switch, called the anchor border, to handle all traffic for a particular VN across various sites. This simplifies management and policy enforcement.

- Subnet consistency:

  MSRB enables you to use the same subnet for the VN across all sites. This eliminates the need to manage different subnets at each location, saving IP address space and simplifying configuration.

- Traffic isolation:

  MSRB is particularly useful for isolating untrusted traffic, such as guest Wi-Fi. All guest traffic across different sites can be tunneled to a central location, like a DMZ, for security purposes.

Here are some common terms that are used in the context of an MSRB:

- Anchor VN:

  A VN that exists across multiple fabric sites in a network. The associated IP subnet and segment are common across these multiple sites.

- Anchor site:

  The fabric site that hosts the common border and control plane for an anchor VN. anchor site handles the ingress and egress traffic for the anchor VN.

- Anchoring sites:

  Fabric sites other than the anchor site where the anchor VN is deployed.

- Anchor border node or MSRB:

  The fabric border node at the anchor site that provides the ingress and egress location for traffic to and from the anchor VN.

- Anchor control plane node:

  The fabric control plane node at the anchor site that accepts registrations and responds to requests for endpoints in the anchor VN.

In essence, MSRB simplifies network management, enhances security for isolated traffic, and optimizes IP address usage in Cisco SD-Access deployments with multiple sites.

Figure 15. Example of an MSRB deployment

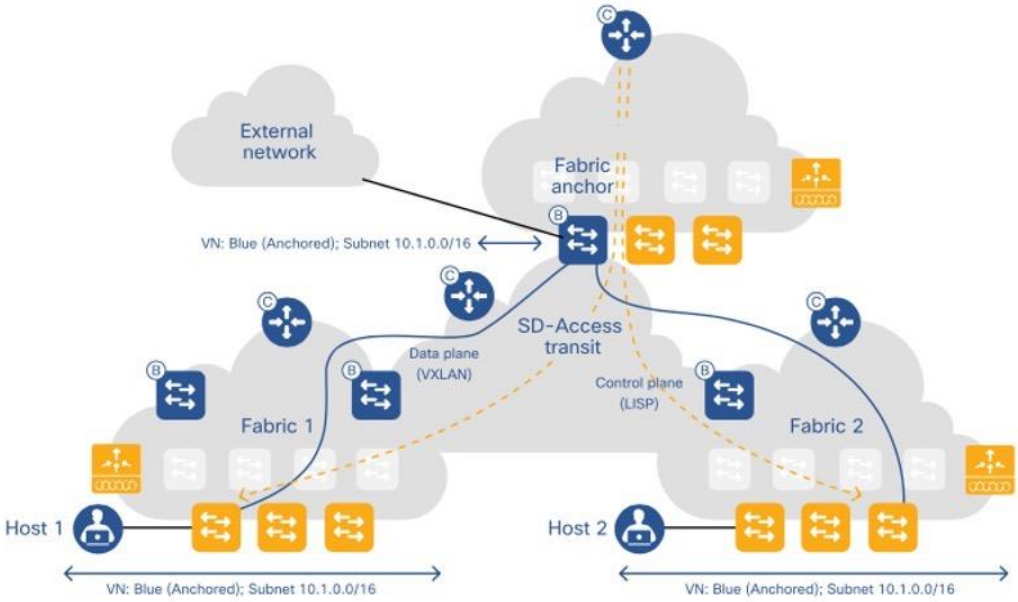

For more details regarding MSRB, see LISP VXLAN Fabric Configuration Guide.

Note: It is crucial to consider the MTU across the entire path to accommodate the additional 50-byte VXLAN header overhead. This is particularly important as the reachability of the anchor site border node may involve traversing multiple IP networks.

LISP Publish and Subscribe Design

LISP Publish and Subscribe (Pub/Sub) model is a significant enhancement to traditional LISP architecture. It streamlines the distribution of endpoint location information across the network, ensuring that all nodes receive timely and accurate data. With its efficiency, scalability, and ability to manage dynamic environments, the LISP Pub/Sub model is a crucial component in modern, large-scale network designs.

LISP Pub/Sub design eliminates the need for an additional protocol to register the LISP site registration table to control plane nodes in the fabric. LISP Pub/Sub feature is fully automated through Catalyst Center, which simplifies the deployment of an SD-Access fabric and removes the need for manual routing configuration.

LISP Pub/Sub architecture is a building block for other features and capabilities, such as:

- LISP Dynamic Default Border Node
- LISP Backup Internet
- LISP Affinity-ID
- LISP Extranet

LISP Pub/Sub uses a publish and subscribe model for routing information. Edge nodes subscribe to the default route, which includes the next-hop IP addresses of both border nodes. If a border node loses its upstream connection (and BGP peering), the default route is removed from the routing table for the affected VNs. The border node then updates the control plane to signal that it can no longer serve as the default route. As a result, the control plane informs all edge nodes subscribed to the default route, ensuring they stop using the failed route and instead rely on the default route toward the remaining active border node. This approach eliminates the need for BGP peering per VRF/VN between border nodes to maintain routing redundancy, thereby reducing manual configuration.

Deployment Considerations:

- LISP/BGP fabric sites and LISP Pub/Sub fabric sites cannot co-exist with the same SD-Access Transit Control Plane Nodes.
- Migration from one to another is not supported yet.
- LISP Pub/Sub is recommended for new network implementations.

Migration to Cisco SD-Access

Migrating to Cisco SD-Access involves a comprehensive approach encompassing assessment, design, implementation, and ongoing optimization. Leveraging Catalyst Center for automation and management ensures a streamlined and efficient migration process, ultimately resulting in a more secure, scalable, and manageable network environment.

Before beginning the migration of the existing network to Cisco SD-Access, consider these aspects:

- Network:

  MTU, network topology, IP addressing for underlay and overlay, and location of shared services.
- Policy:

  Existing policy definition and enforcement points, VNs, and SGTs.
- Hardware Platform:

  Switches, routers, wireless controllers, and APs that support SD-Access.
- Software Platform:

  Catalyst Center, ISE, network data platform.
- Scale of Deployment:

  Scale of hardware platforms on their role in the SD-Access architecture.
- Existing Network Design:

  Layer 2 access or routed access.

These are the primary approaches to migrate an existing network to SD-Access:

- Parallel approach

  An SD-Access network is built alongside the existing brownfield network. Switches are migrated from the brownfield network to the SD-Access network by physically patching cables. This approach simplifies change management and rollback procedures. However, setting up the parallel network requires additional rack space, power, and cabling infrastructure beyond what the brownfield network currently uses.
- Incremental approach

  This strategy involves migrating traditional switches from the brownfield network and converting them into SD-Access fabric edge nodes. The Layer 2 border handoff, discussed in the next section, facilitates this gradual migration. This approach is suitable for networks with existing equipment capable of supporting SD-Access or facing environmental constraints such as limited space and power.

For complete guidance and different options to migrate existing traditional networks to Cisco SD-Access, see the “Migration to Cisco SD-Access” chapter in the Cisco Software-Defined Access for Industry Verticals.

Validated solution use cases

These sections outline key use cases validated for university networks, serving as trusted templates to help institutions design and build their IT infrastructure. These solutions are rigorously tested and tailored to address the unique requirements of academic environments, ensuring reliability, scalability, and optimal performance.

Day-zero and day-1 network bring up use cases

- Deploy Catalyst Center and integrate it with external servers for a greenfield campus.
- Devices are onboarded using LAN Automation and PnP.
- Automate and simplify network device and fabric provisioning.
- Monitor inventory and manage network devices using Catalyst Center.
- Integrate with Cisco ISE for authentication and authorization of device and client.
- Manage and deploy wireless controllers and APs using Catalyst Center.
- Manage network settings for multiple sites using Cisco Catalyst for shared services.
- Deploy SD-Access multisite campus and manage traffic across campus.

Day-*n* network operations use cases

- Upgrade multiple devices, such as switches, routers, and wireless controllers using Catalyst Center.
- Onboard new floors to existing fabric sites.
- Onboard new fabric nodes with wired and wireless clients.
- Replace brownfield APs from Wave2 to 11 Ax.
- Add small campus using FIAB with an embedded wireless controller.
- Perform day-*n* credential changes such as, device password changes followed by device provision.

Solution use case scenarios

These use cases were implemented for this University Vertical profile, based on the topology illustrated.

Intent-based networking using Catalyst Center for university networks

University network administrators can achieve goals with Catalyst Center, including:

- Design a global network hierarchy and set both global and site-specific network configurations to optimize performance across the entire university.
- Automatically provision devices to streamline network deployment and management.
- Deploy the main campus network with dual borders and dual control plane nodes for redundancy and scalability, ensuring reliable connectivity as the university grows.
- Easily expand the network at both the main campus and satellite campus locations by onboarding new devices, such as:
    - Fabric edge switches using zero-touch PnP LAN automation or leverage existing IP/MPLS infrastructure for underlay reachability.
    - Classic-extended nodes into the fabric to support IoT devices with zero-touch plug-and-play.
    - Policy-extended nodes into the fabric for direct SGT support and enhanced traffic enforcement.
- Connect distributed campus sites using Cisco SD-Access transit, enabling efficient access to shared datacenter and internet services across all campuses.

This approach helps universities ensure a flexible, scalable, and secure network environment while minimizing manual configuration efforts.

Multitier security to protect sensitive university data

University administrators can implement the following measures to safeguard sensitive institutional data:

- Segment students, faculty, guests, IoT devices, and campus equipment into appropriate logical networks to restrict the lateral movement of threats within the university network.
- Utilize closed authentication mechanisms such as 802.1X (dot1x) or MAC Authentication Bypass (MAB) for both wired and wireless endpoints to prevent unauthorized access to university resources.
- Apply trusted Certificate Authority (CA) FQDN-based certificates through Catalyst Center to strengthen the network's security framework.
- Create user groups, categorize users and endpoints based on their identities, and define group-based policies to regulate traffic interactions between different groups.
- Monitor activities in Catalyst Center through comprehensive audit logs, which track system events, including what occurred, when and where it happened, and the users involved.
- Establish granular role-based access permissions for Catalyst Center users, ensuring administrators, faculty, and IT staff have access appropriate to their roles.

This multitiered security framework enhances the protection of university data and ensures proper access control across the campus network.

Cisco AI Endpoint Analytics

University administrators can effectively manage unknown endpoints and enhance endpoint security by implementing these:

- Configure Cisco AI Endpoint Analytics to identify spoofed endpoints based on classification and behavioral model learning tailored to the university network environment.

- Set up Catalyst Center to detect and flag compromised endpoints within the university's network. A comprehensive trust score is calculated for each endpoint based on the configured influencing parameters.
- Monitor the trust score and identify the type of threat associated with each endpoint. Take corrective actions or quarantine the malicious endpoint to protect the university network.

This ensures that university networks are continuously monitored for security threats, with proactive steps taken to maintain a safe and trusted environment.

Service and network resiliency

University administrators can implement these resiliency measures to ensure reliable service delivery and minimize network disruption:

- Deploy a redundant network infrastructure with features like dual Cisco SD-Access borders, dual control plane nodes, border and edge stacking, and dual transit control planes. This ensures rapid failover and minimal service interruption during network failures.
- Deploy Catalyst Center in a three-node HA cluster to ensure continuous service delivery and eliminate the need for manual intervention during node or service failures.
- Implement a distributed Cisco ISE deployment with PAN, PSN, M&T and pxGrid service failover for enhanced reliability and minimized service disruption in case of failures.
- Regularly back up Catalyst Center and device configurations to ensure quick recovery of previous configuration in case of unforeseen incidents.

These measures help ensure a stable and resilient network environment for university services.

Simplified network management

University administrators can leverage Catalyst Center for simplified network management by using these capabilities:

- Catalyst Center can be used to manage all network devices, including details like IP addresses and software versions, simplifying monitoring and management.
- VNs can be created and managed to ensure secure and organized network segmentation for different university departments or locations.
- SGTs and access policies can be applied to control traffic within a VN, ensuring better security and performance based on user roles.
- Catalyst Center SWIM can be used to upgrade network devices to a standard image, ensuring consistency and reducing administrative tasks.
- VLAN configurations can be optimized for efficient network operation and scalability as the university network grows.
- Automate the setup of new network devices with zero-touch provisioning (ZTP), saving time and reducing human error.
- Use templates to create and deploy network configurations, ensuring consistency and speeding up device setup or configuration changes.

These capabilities simplify network management, allowing university IT teams to maintain a reliable, secure, and scalable network environment.

Operations and maintenance with assurance and analytics

University administrators can leverage Catalyst Center Assurance and Analytics to optimize network operations by doing these tasks:

- Continuously track the performance of the university network, including wired and wireless infrastructure, to detect and resolve issues such as link failures, AP downtimes, or device malfunctions.
- Monitor the health and performance of both wired and wireless clients to identify and address client connectivity issues, ensuring smooth onboarding for students, faculty, and staff.
- Manage the network scalability by tracking up to 100,000 concurrent endpoints and 250,000 transient endpoints, which is crucial for large campuses with many users and devices.
- Identify unauthorized or rogue APs in the network and generate reports to take immediate corrective action, enhancing network security.
- Gain insights into application performance and end-user experience, ensuring that critical educational and administrative services run smoothly with minimal downtime.
- Ensure compliance with university security policies by monitoring and managing network devices and promptly addressing any vulnerabilities or rogue devices.
- Create customized dashboards to visualize key performance indicators (KPIs) and gain actionable insights into network performance.

This approach ensures that the university's network remains robust, secure, and responsive to the needs of its students, faculty, and staff.

Education roaming

Using eduroam, university administrators can enhance collaboration and connectivity globally by implementing eduroam with this configuration:

- Set up the eduroam Wi-Fi SSID with Enterprise security enabled and 802.1X authentication using Catalyst Center.
- Configure Cisco ISE to manage authentication requests from external eduroam servers and forward authentication requests for external users to the eduroam server for validation.
- Enable external users from other universities to connect to the eduroam SSID on the local campus and gain secure network access upon successful authentication.
- Allow traveling users from the local university to connect to eduroam SSIDs at foreign universities and access their network securely.

This configuration helps universities provide seamless, secure network access for both local and visiting users, fostering international collaboration and mobility.

For more details, see Configuring eduroam on Cisco Identity Services Engine (ISE).

Cisco Wide Area Bonjour

University administrators can leverage Cisco Wide Area Bonjour to enable service discovery across the university network by implementing these functions:

- Configure SDG-Agents and network information to enable efficient service routing across the university network.
- Set up service Policy-Based ZeroConf Services management and distribute Common Bonjour Services from end-user devices through the Cisco Wide Area Bonjour application.
- Enable university network users to access Bonjour services such as printing and screen sharing across Layer 3 domains.
- Use the Cisco Wide Area Bonjour dashboard to monitor SDG-Agent statistics, including service counts per subdomain and the operational status of policies.

This approach simplifies service discovery, enhancing user experience across different university campuses or network segments.

For more details, see the Cisco DNA Service for Bonjour Quick Configuration Guide.

BYOD

University administrators can use Cisco ISE to ensure secure BYOD access by implementing these functions:

- Configure Cisco ISE to grant privileged access to personal devices on the university network.
- Provisioning Enterprise SSID for BYOD endpoints using Catalyst Center.
- Allow students and faculty members to connect their personal devices to the university network and obtain privileged access.
- Using the My Devices Portal in Cisco ISE for convenient management and control of BYOD endpoints.
- Enforce authentication methods like 802.1X or WebAuth to ensure only authorized devices connect to the university network.
- Segregate BYOD devices from sensitive systems to protect university resources.
- Automatically identify device types and apply appropriate policies using device profiling.

These measures provide a seamless yet secure method for managing BYOD devices on university campuses while maintaining the integrity and security of the academic network.

For more details, see the Cisco ISE BYOD Prescriptive Deployment Guide.

Guest services with MSRB

University administrators can efficiently manage guest services by implementing these functions:

- Configure the guest VN and extend it across multiple university sites. This VN can be anchored at a remote border in a central anchor site.
- Isolate guest traffic within the guest VN and tunnel the traffic to the anchor borders for secure internet access.
- Set up a common guest subnet to onboard guest users seamlessly across both anchor and inherited fabric sites.
- Provision a Guest SSID with Captive Web Authentication (CWA) using Catalyst Center, ensuring that the same Guest SSID is consistent across all university sites.
- Implement an anchor site with dual anchor borders and control planes in physically distinct locations, ensuring redundancy in case of network failure.
- Monitor the trust score and threat type of connected devices, taking corrective action or quarantining malicious endpoints as necessary.

This approach ensures secure, scalable, and reliable guest access across all university campuses, enhancing the overall campus network experience.

Layer 2 termination outside the fabric

University network administrators can do these tasks:

- Configure the required Layer 2 VN using Catalyst Center.
- Configure the Layer 2 handoff on the dedicated Layer 2 border node.
- Terminate traffic outside the fabric (typically on a firewall) for traffic inspection.

This configuration enables efficient traffic management within the university network by ensuring proper routing and inspection of Layer 2 traffic outside the fabric.

Key solution notes

This section outlines the key technical details of the solution validation for the University Vertical profile, providing university administrators with a comprehensive reference for in-depth implementation understanding.

Eduroam wireless network access for universities

Eduroam is a global wireless network service that provides researchers, faculty, and students with network access when visiting other educational institutions. By connecting to the eduroam Wi-Fi network, users can access resources and the internet using their home institution credentials, no matter where they are in the world. The authorization for internet access and other network resources is managed by the visited institution.

The eduroam service uses IEEE 802.1X for authentication and operates through a hierarchical system of participating RADIUS servers. Both the home and foreign institution RADIUS servers must be part of the eduroam network for proper functioning.

The eduroam Wi-Fi SSID is secured with Enterprise Security and provisioned using Catalyst Center. To implement eduroam, configuration is required on both Cisco ISE and Catalyst Center. The eduroam policy settings and external RADIUS server configurations are managed within Cisco ISE.

Figure 16 and Figure 17 illustrate the configuration of the eduroam external servers.

Figure 16. External Radius Server Configuration - 1

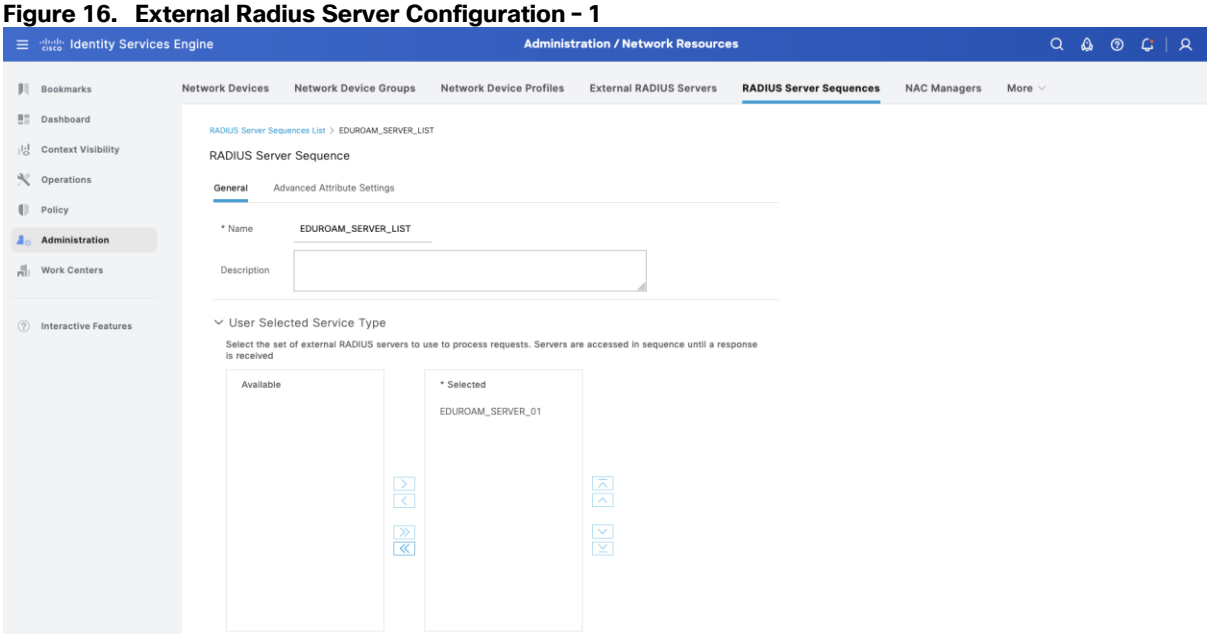

Figure 17. External Radius Server Configuration – 2

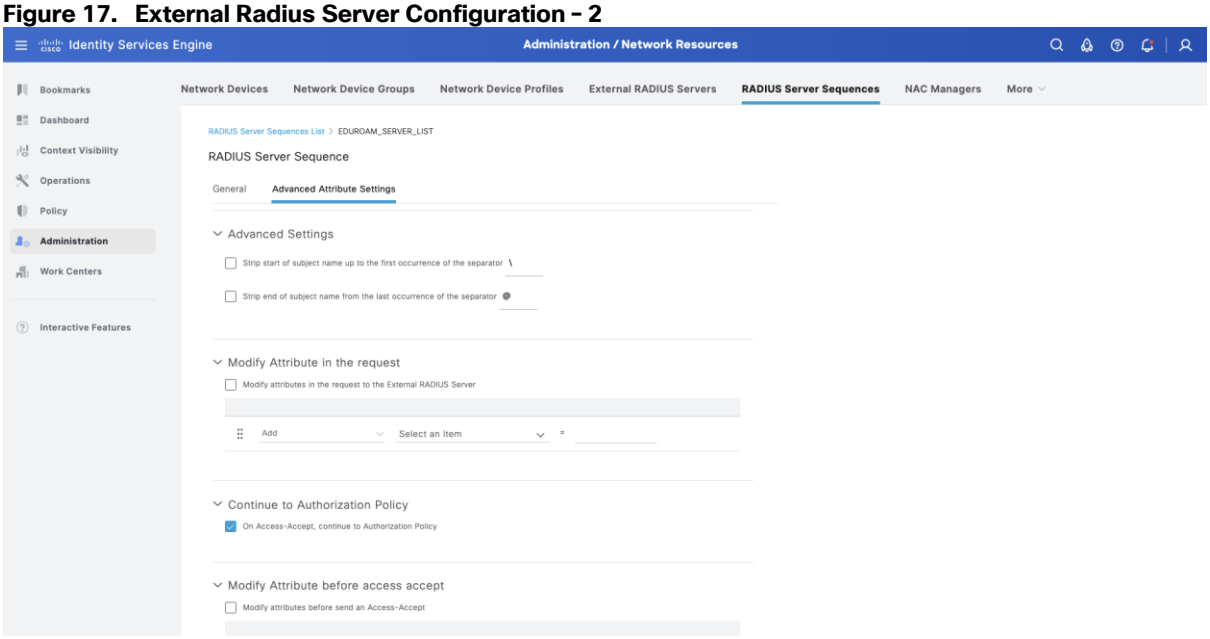

Figure 18 illustrates the configuration of the eduroam policy set.

Figure 18. Eduroam Policy Set configuration on Cisco ISE

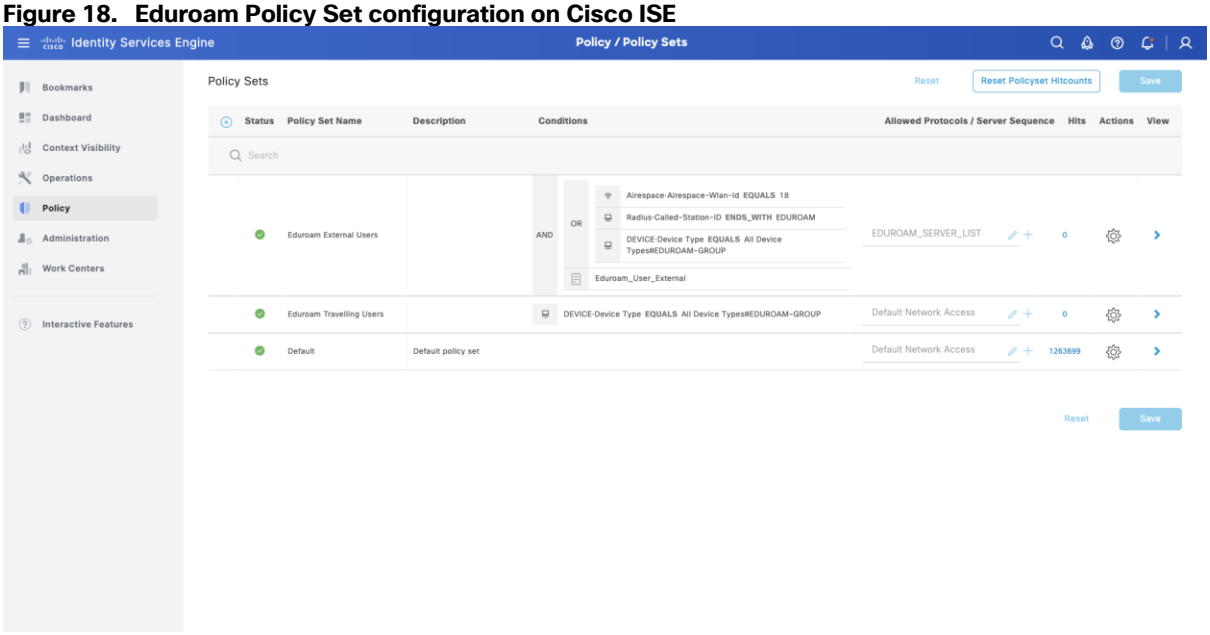

Figure 19 shows the eduroam SSID configuration on Catalyst Center.

Figure 19. Eduroam SSID Configuration on Catalyst Center

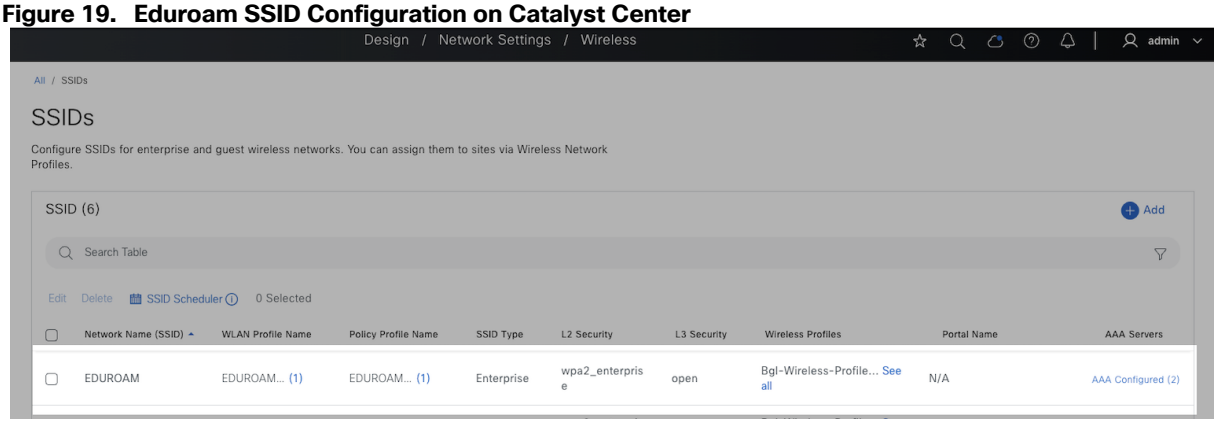

Eduroam supports two primary use cases for university students and faculty:

1. External user

   In this scenario, a student or faculty member from an external university visits a local campus. The user connects to the eduroam SSID and uses their home institution credentials for authentication. The request is forwarded to the eduroam servers, which then route it to the user's home institution. After the authentication is successful, the user is granted access to the network.

   Figure 20 shows the authentication flow.

Figure 20. External user authentication flow

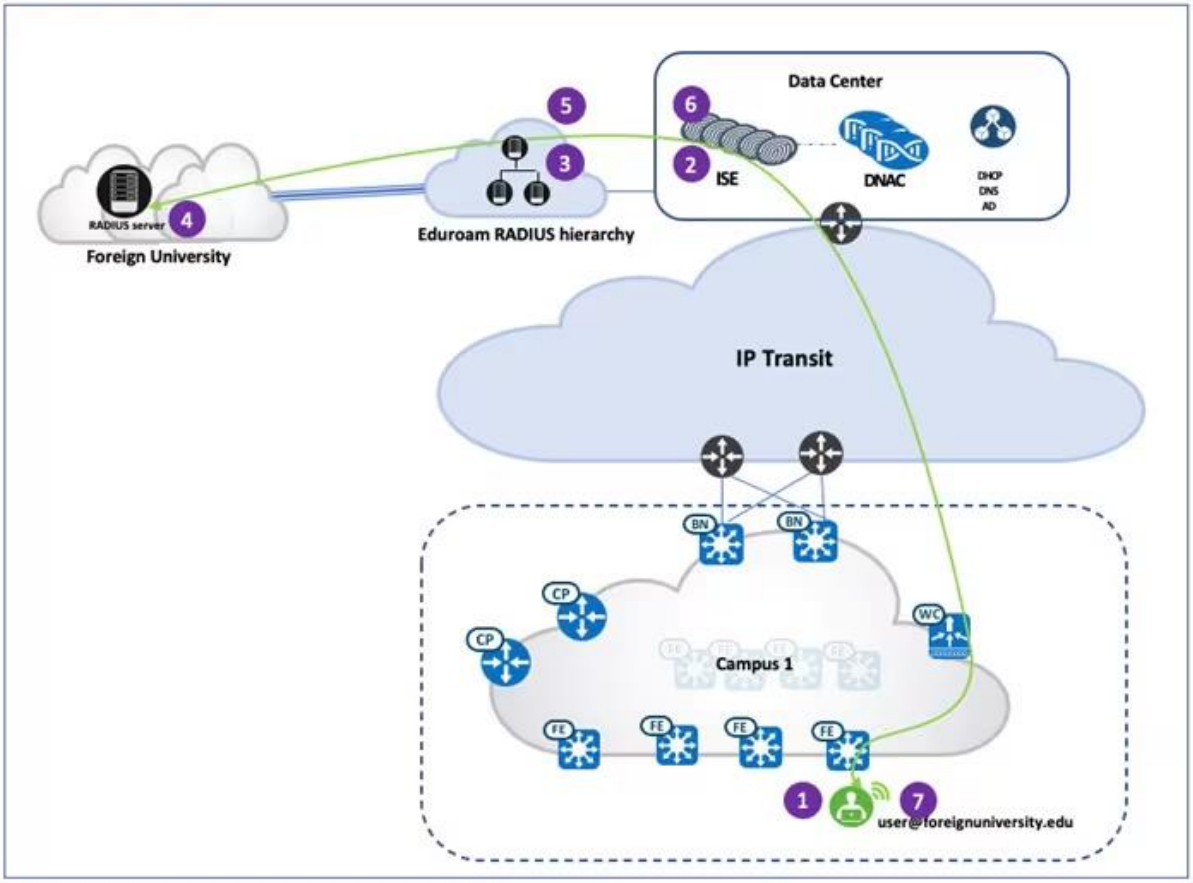

2. Traveling user

In this case, a student or faculty member from the home institution is physically located at a foreign university campus. When the user connects to the eduroam SSID at the foreign campus, the authentication request is sent to the eduroam RADIUS server, which then forwards the request to the user's home institution's RADIUS server. If authentication is successful, the home institution's RADIUS server sends an "Access Accept" response to the eduroam server, which then forwards the response to the foreign university. The user is authenticated and granted network access at the foreign university.

Figure 21 shows the authentication flow.

Figure 21. Traveling user authentication flow

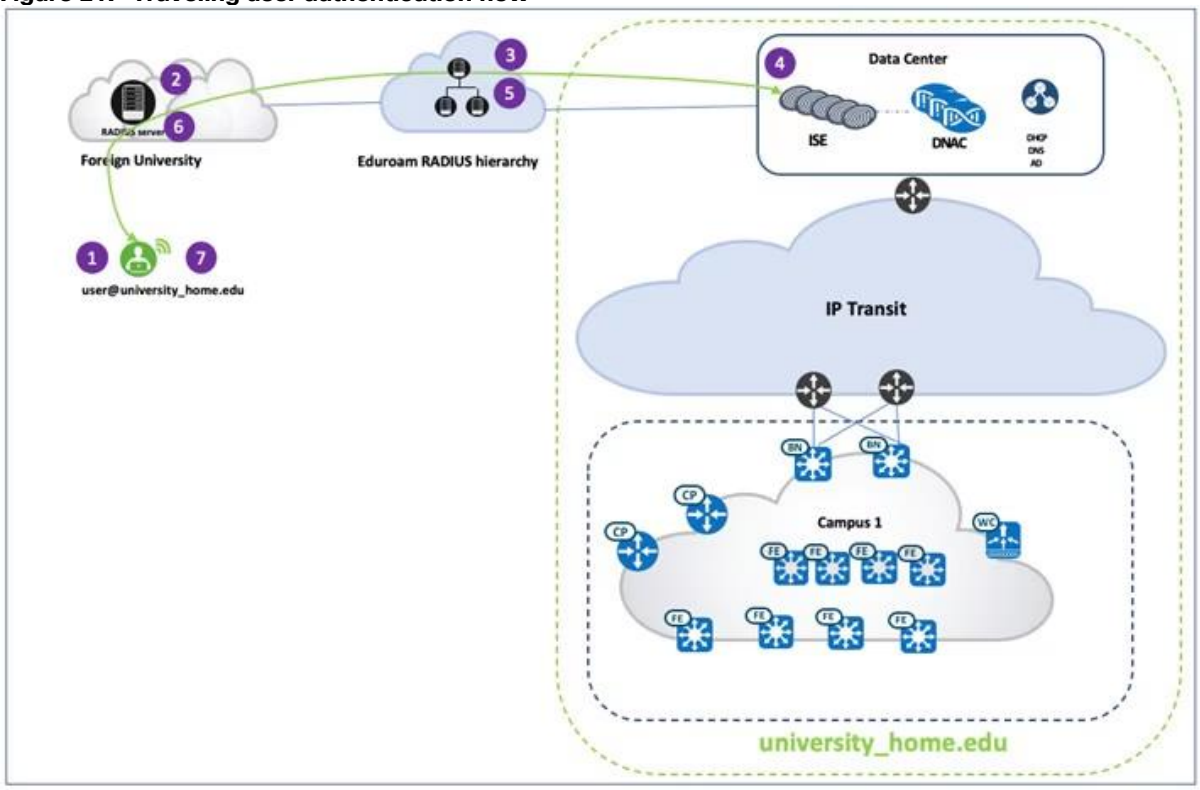

This structure ensures seamless and secure network access for university users, whether they are visiting or traveling abroad.

For more details, see Configuring eduroam on Cisco Identity Services Engine (ISE).

Guest services with MSRB in university deployments

MSRB is enabled on a per-VN basis. For an anchored VN, all edges in the anchoring sites use the anchor border and control plane nodes for data plane and control communication. Wireless controllers in the anchoring sites communicate with the anchor control plane node for wireless endpoint registration specific to the anchored VN.

Since anchor border reachability may traverse multiple IP networks, special attention must be paid to the MTU across the entire path to accommodate the VXLAN header overhead of 50 bytes. An anchored VN is configured to use the anchor site. After a guest endpoint joins the guest SSID and successfully completes Central Web Authentication using Cisco ISE, it is associated with the anchored guest VN. Guest traffic is tunneled to the anchor site border node and egresses to the Internet through a firewall.

Additionally, while guest traffic is VXLAN-encapsulated and passes through the fabric, the first hop or gateway for guest traffic can be outside the fabric and connected at Layer 2 to the MSRB. If required by design, such a device can be a firewall for inspection purposes.

Figure 22 and Figure 23 shows MSRB enabled in the Catalyst Center GUI.

Figure 22. Guest VN for anchor and anchoring sites

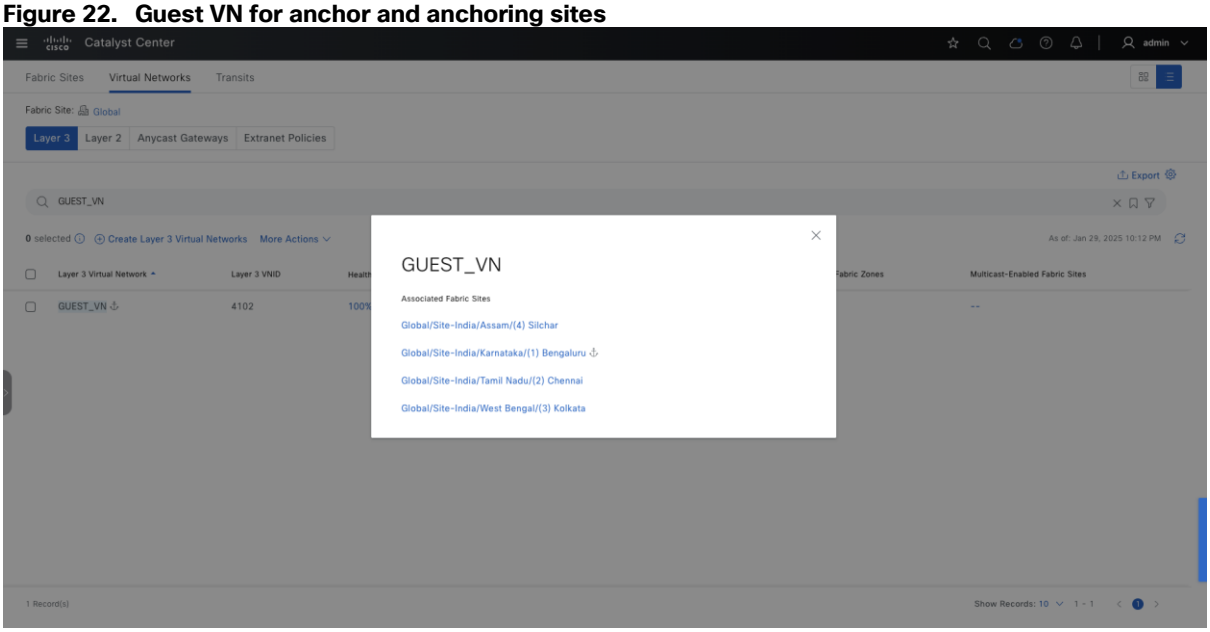

Figure 23. Same IP subnet is provisioned across multiple sites

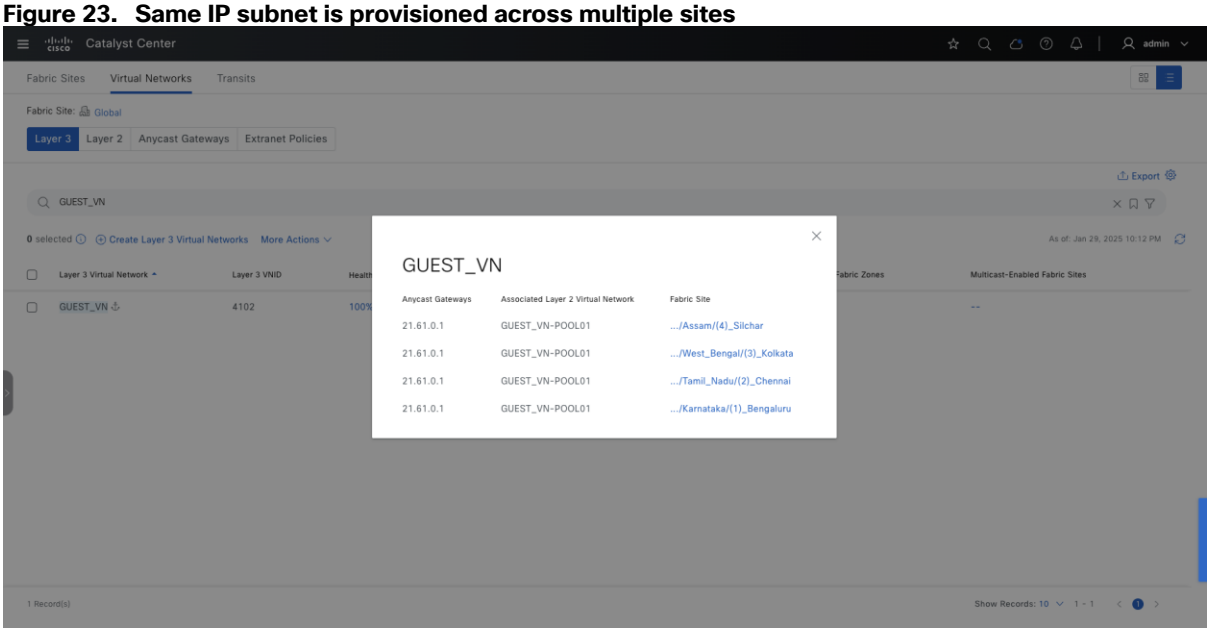

Figure 24 illustrates the traffic flow of anchored guest traffic.

Figure 24. Guest traffic flow implemented with an anchor VN

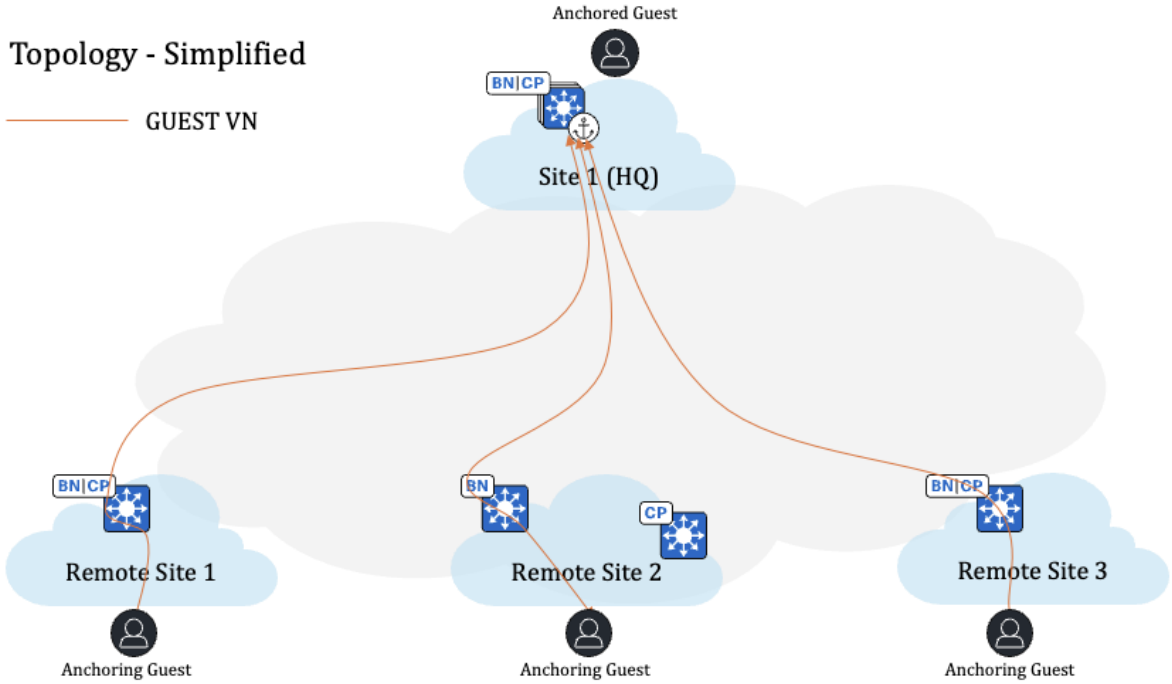

This MSRB solution, accessible through the Catalyst Center GUI, simplifies guest service management across multiple university sites, enhances security, and ensures a seamless experience for both students and visitors.

AI endpoint analytics for university networks

Universities need to manage a large number of users and their devices. With BYOD policies, each student or faculty member typically connects two to three devices to the network. As the network scales, security challenges also increase. Modern security threats often exploit vulnerable entry points to gain access to valuable institutional data. Identifying and tracking all devices on a network is time-consuming and complex.

The Cisco AI Endpoint Analytics feature addresses this challenge by identifying devices based on type, manufacturer, model, OS type, communication protocols, and ports using passive network telemetry monitoring and deep packet inspection. It allows administrators to create profiling rules to classify devices based on these attributes. Combined with machine learning, Catalyst Center can detect spoofed endpoints and assist administrators in taking appropriate action.

Cisco AI Endpoint Analytics is an additional application that runs within Catalyst Center. To deploy it:

- Download and install the application from the Catalyst Center catalog server.
- Enable it in Catalyst Center System Settings.
- Ensure Catalyst Center connects to the cloud to download the latest endpoint analytics models.

When installed, Cisco AI Endpoint Analytics can be accessed from the Catalyst Center home page by clicking the menu icon and choosing policy. Cisco AI Endpoint Analytics uses multiple methods to detect malicious endpoints, including:

- Changes in Profile Labels
- NAT mode detection

- Concurrent MAC address analysis
- Posture and authentication method analysis
- Machine learning–based anomaly detection

Each endpoint is assigned an overall trust score, which is a weighted average of multiple risk factors. A lower trust score indicates a higher risk, helping administrators identify and respond to potential threats.

For example, Cisco AI Endpoint Analytics can detect if a student's laptop attempts to spoof a library printer to gain unauthorized access to the university network.

Figure 25. Cisco AI Endpoint Analytics Trust Score details on Catalyst Center

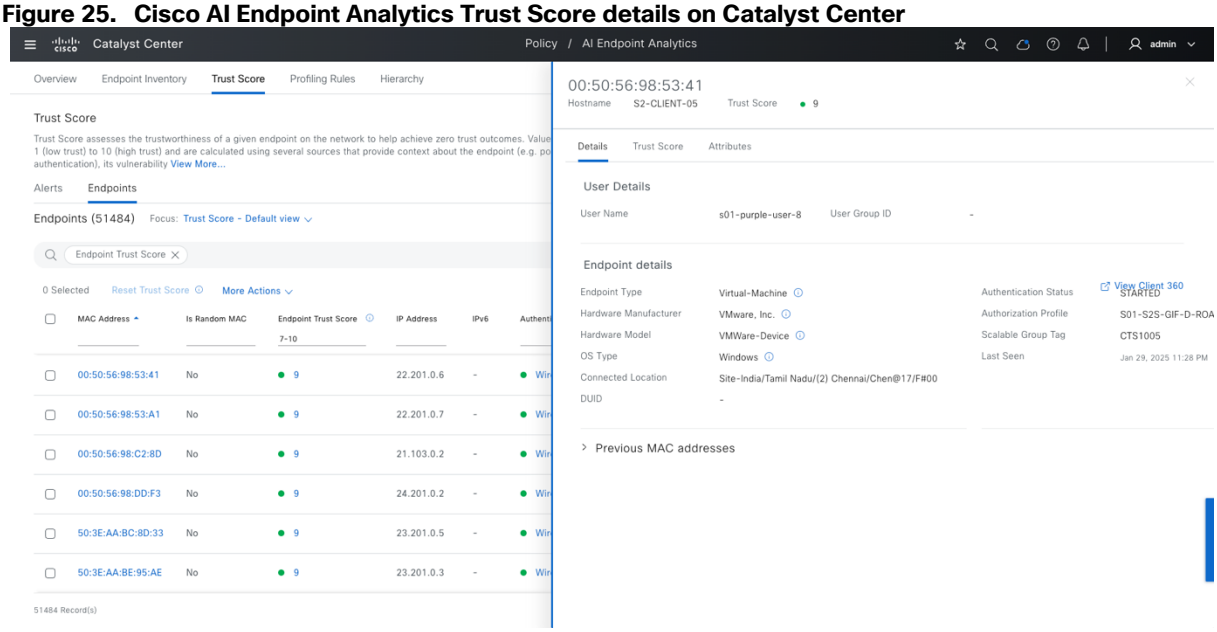

With AI-driven endpoint visibility, automated security policy enforcement, and real-time threat detection, Cisco AI Endpoint Analytics helps universities maintain a secure, scalable, and efficient IT infrastructure, protecting critical academic and research data while ensuring a seamless digital experience for students and faculty.

Layer 2 traffic termination outside the university network fabric

Universities may require Layer 2–level traffic inspection for specific network segments. To meet this requirement, all traffic within these segments must have its first hop outside the university's network fabric. This setup is implemented using a combination of Catalyst Center and a dedicated Layer 2 border node. While traffic remains VXLAN-encapsulated as it traverses the fabric, the initial gateway is located outside the fabric.

To configure this, the required network or SSID is deployed via Catalyst Center. Additionally, a Layer 2 handoff is configured at the border of the fabric site for the corresponding VN. The firewall or Layer 2 termination point is positioned at the receiving end of this handoff. The firewall is assigned an IP address within the same subnet as the network. The DHCP server responsible for assigning IP addresses must designate the Gateway IP as the firewall's Layer 2 termination IP instead of the fabric anycast gateway IP.

With these modifications, the firewall IP functions as a client within the VN. When a device obtains an IP address and initiates communication, the first hop (Layer 2 termination point outside the fabric) is resolved through Layer 2 LISP, allowing successful communication with external networks. This configuration enables advanced Layer 2 traffic inspection outside the university's network fabric.

Figure 26 and Figure 27 illustrate the configuration of Layer 2 VN and the Layer 2 handoff on the border node.

Figure 26. Layer 2 VN configuration using the Catalyst Center UI

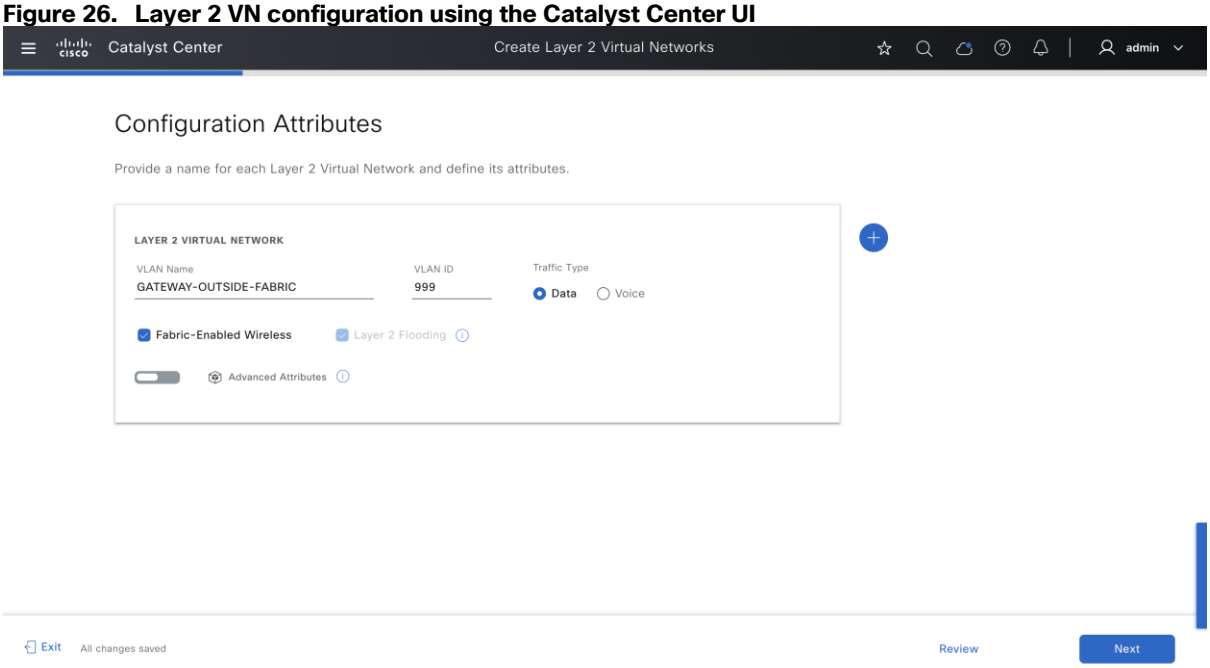

Figure 27. Performing Layer 2 handoff on the dedicated Layer 2 border node

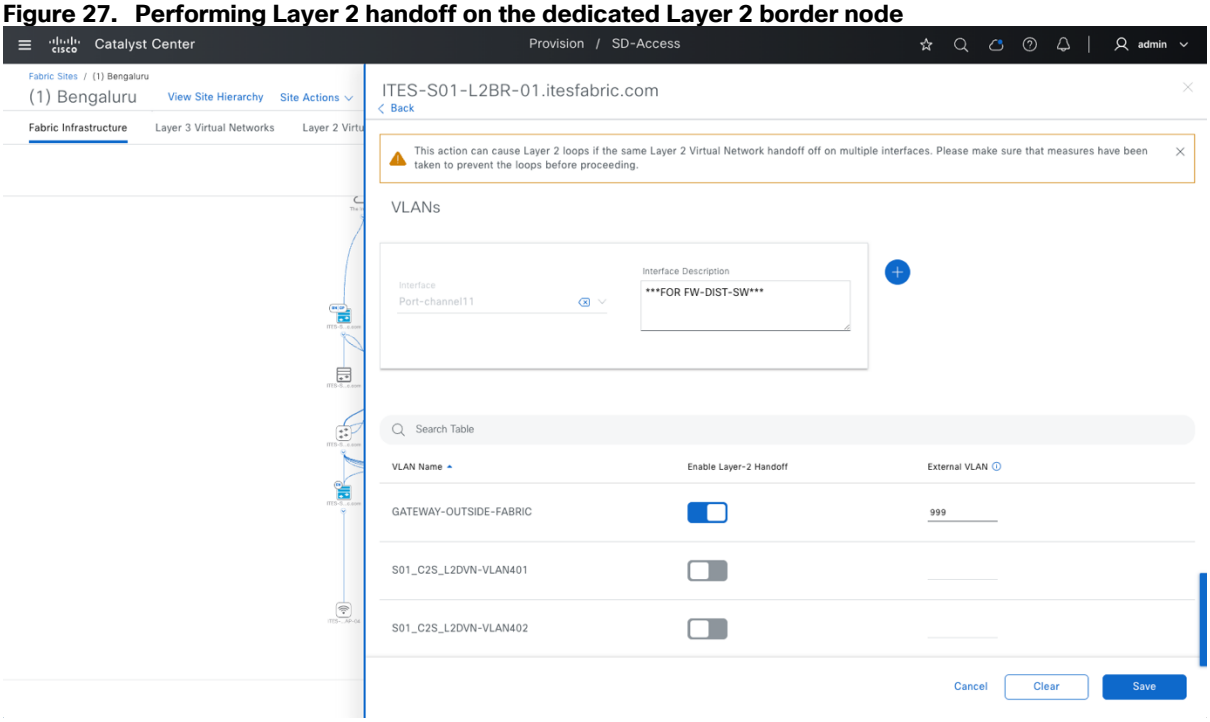

Wide Area Bonjour for university networks

In recent years, the number of clients capable of screen sharing has increased significantly, with major vendors supporting this feature on laptops and mobile devices. This growth has led to a rise in multicast sources, straining network table sizes. To manage this, multicast environments must be controlled by regulating source and receiver devices.

Two key methods can help:

- Designing a resilient multicast network with Rendezvous Points and ACLs to limit stream subscriptions.
- Leveraging Cisco Catalyst Center Wide Area Bonjour (WAB) to control and scale traffic.

Since clients operate in virtual networks, their multicast traffic is carried within the VXLAN overlay. Deploying Wide Area Bonjour requires understanding its scalability and the role of the Service Discovery Gateway (SDG), which facilitates communication with mDNS clients. Cisco SD-Access supports Bonjour services across fabric-based networks, benefiting universities deploying Catalyst 9000 Series Switches. VRF-aware Wide Area Bonjour enables secure, segmented mDNS service discovery without Layer 2 flooding, enhancing performance and scalability.

A key use case in universities involves fabric edge switches providing printer services to wired and wireless users across multiple locations. This allows students and faculty to print remotely without needing to be physically present near the printer.

Figure 28. Cisco SD-Access wired and wireless network design

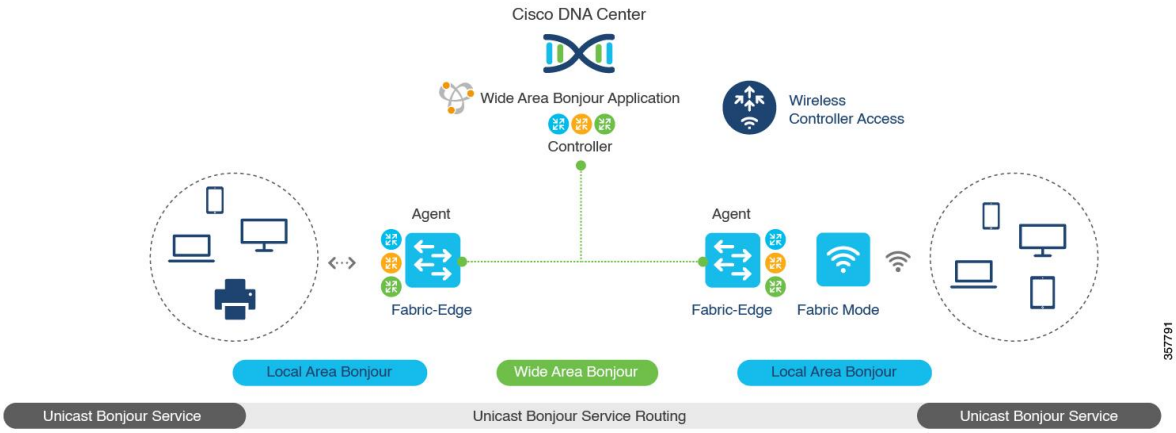

Setting Up Cisco Wide Area Bonjour in Catalyst Center:

- Wide Area Bonjour application:

  Cisco Wide Area Bonjour is a non-default application in Catalyst Center. Download and install the application from the Cisco Catalog Server. After successfully installing the Cisco Wide Area Bonjour application, you can access it from the Catalyst Center home page by clicking the menu icon and then selecting Tools.

- Manual or template based configuration:

  Cisco Wide Area Bonjour application on Catalyst Center does not push configurations to the Service Discovery Gateway (SDG) Agent switches or service peer devices. The SDG Agents and service peers must be configured either manually or through templates created using the Template Editor in Catalyst Center.

- Global service filter:

First step is to implement global service filters, which permit the Cisco Wide Area Bonjour application to dynamically discover and distribute service information between trusted Cisco Catalyst SDG Agent switches across IP networks.

- Create service filter:

  From Cisco Wide Area Bonjour application, select the service domain where the administrator wants to add the service filter and select the service types to permit announcements and queries. The administrator can edit, enable, or disable this setting after it is created.

- Configure source SDG Agents:

  From the Cisco Wide Area Bonjour application, select the SDG Agent and VLAN that announces the services. Administrators have the option to enable or disable the services for an IPv4 or IPv6 network.

- Configure Query SDG Agent:

  From the Cisco Wide Area Bonjour application, select the SDG Agent and VLAN that receives queries for the services (Printer). Administrators have the option to enable or disable the services for an IPv4 or IPv6 network.

Figures 29 to 30 showcase the Cisco Wide Area Bonjour dashboard, SDG Agents, and service instance in Catalyst Center.

Figure 29. Cisco Wide Area Bonjour dashboard in Catalyst Center

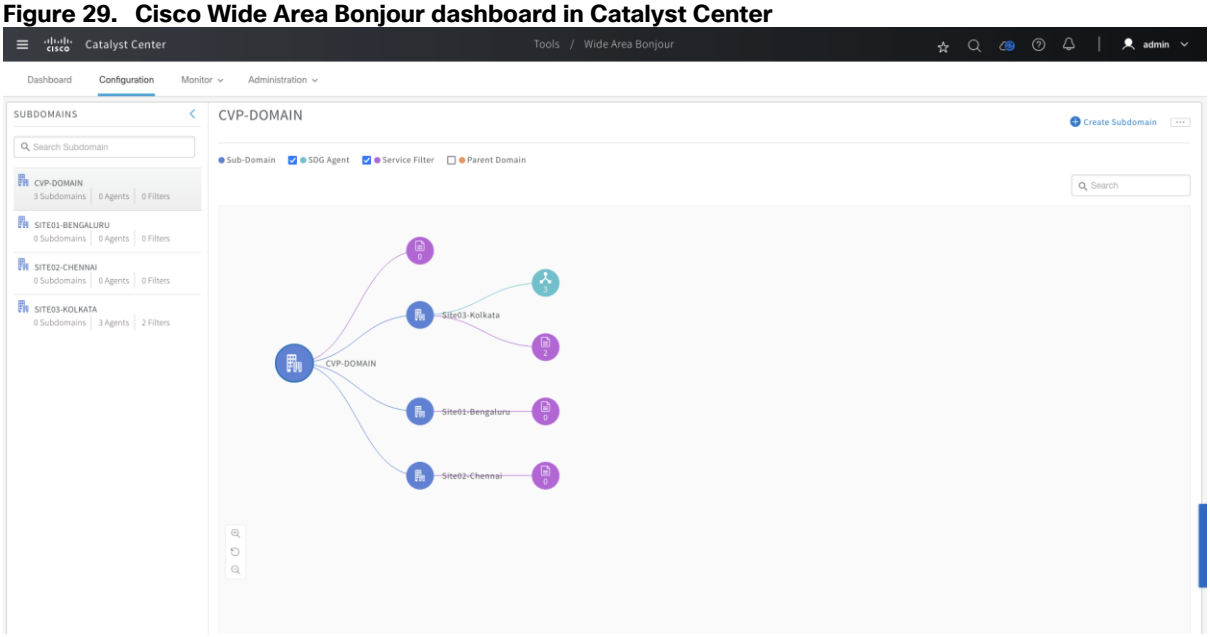

Figure 30. SDG agent details in Catalyst Center

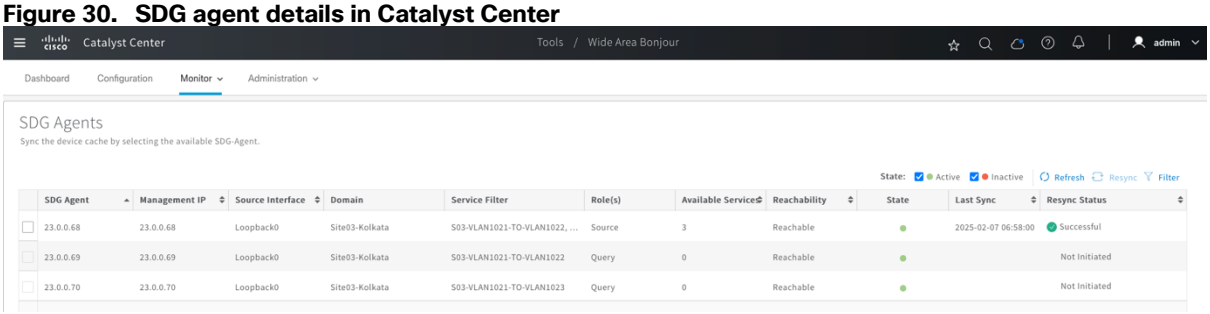

Figure 31. Service instance details in Catalyst Center

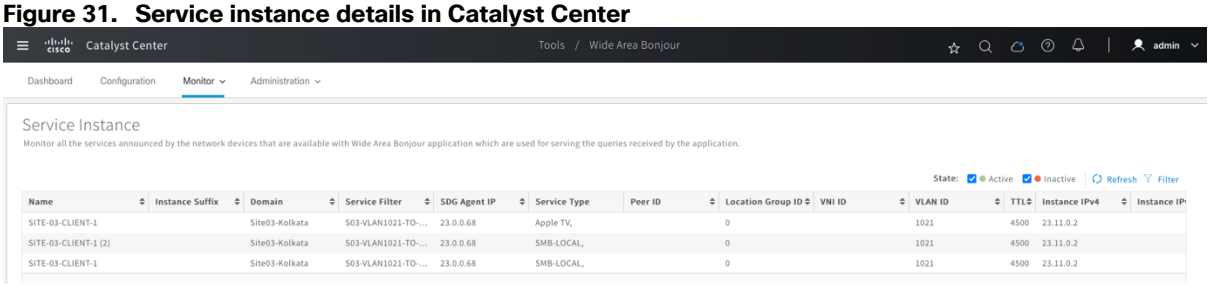

Note:
You need to enable Global Wireless Multicast Mode in the Cisco Catalyst 9800 Series wireless controller. By default, the Cisco wireless controller and APs prevent forwarding Layer 2 or Layer 3 multicast frames between wireless and wired network infrastructure.

If the service filter status appears green, it indicates that the policy is active. When a user's laptop connects remotely from VLAN-B, it can detect and access the printer services available in VLAN-A.

BYOD in a university network

Cisco ISE enables BYOD functionality for universities, allowing students and faculty members to securely connect their personal devices to the campus network. Users can onboard their devices through native supplicant provisioning or by registering them through the My Devices portal. University administrators can ensure that only compliant and secure devices—free from jailbreaking or rooting—gain access to the network. The system provides visibility into users, devices, and applications, ensuring that only authorized devices meeting security policies can connect.

With Single-SSID BYOD, when a user connects to a secure university WLAN, their device is onboarded. Upon automatic reconnection, the device gains full network access on the same WLAN.

Key components of a Cisco ISE BYOD configuration in a university environment include:

- Client Provisioning Policy:

  Defines the BYOD profile based on device type or user group, including certificate templates, SSID names, and proxy settings.

- Authentication and Authorization Policy:

  Determines the user portal experience during onboarding, authentication methods, and required network policies.

- Endpoint onboarding:

  Devices initiate the onboarding process via the BYOD portal, generate digital certificates, and configure network profiles. For Windows devices, Cisco ISE leverages the Network Setup Assistant (NSA) (also known as the Supplicant Provisioning Wizard) to streamline onboarding. Administrators must periodically update NSA to ensure compatibility with newer operating systems.

- Posture policy:

  Posture is a Cisco ISE service that allows you to check the compliance, also known as posture, of endpoints before allowing them to connect to your network. A posture agent, such as the Cisco ISE AnyConnect Posture Agent runs on the endpoint. The Client Provisioning service ensures that endpoints receive the appropriate posture agent. After the endpoint is brought up to compliance and successfully onboarded, the portal notifies the user that they now have full access. Users can open their browser and navigate to other destinations, while Cisco ISE registers the end point as a BYOD device. Figure 32 figure shows the client provisioning workflow.

Figure 32 illustrates a BYOD onboarding process on a Windows 10 laptop.

Figure 32. BYOD onboarding workflow

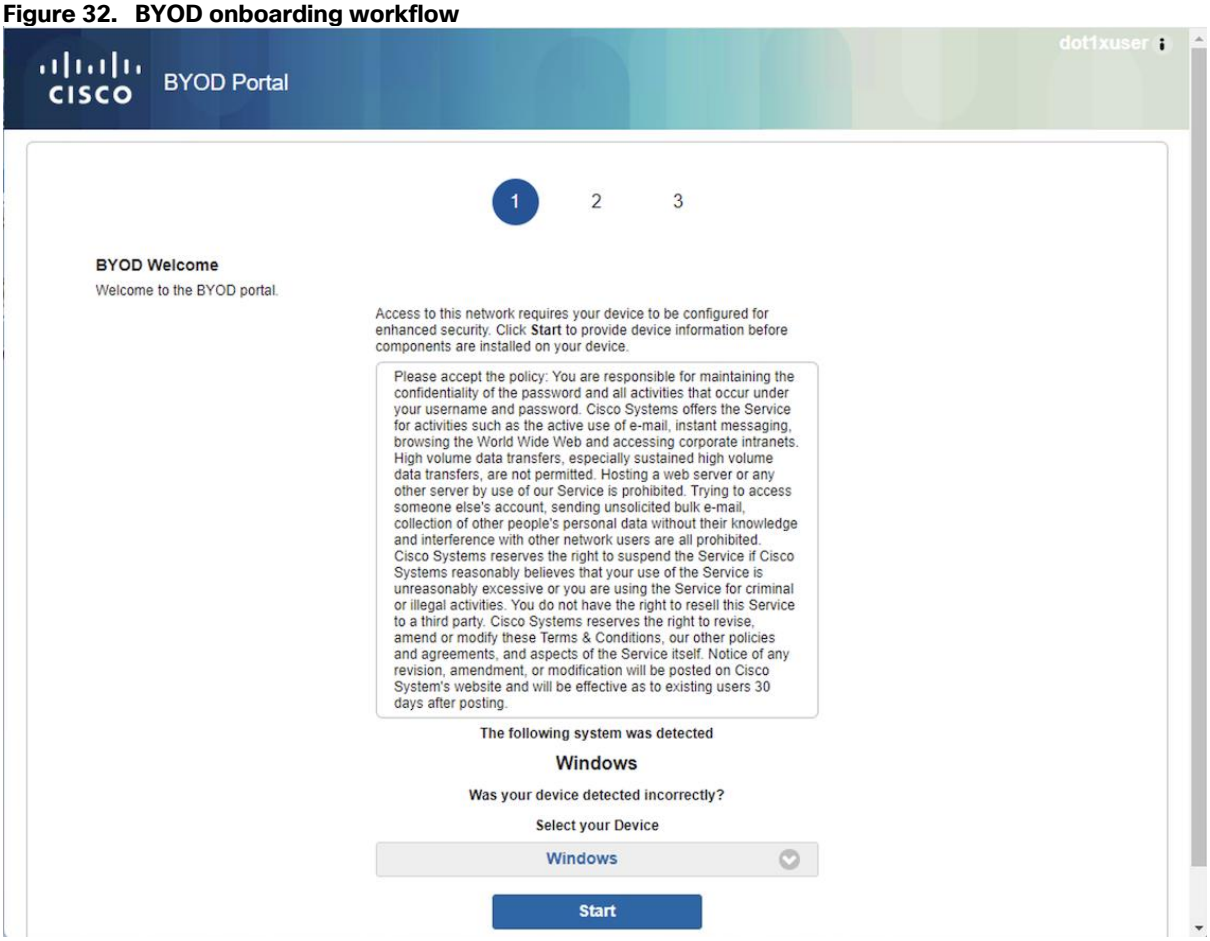

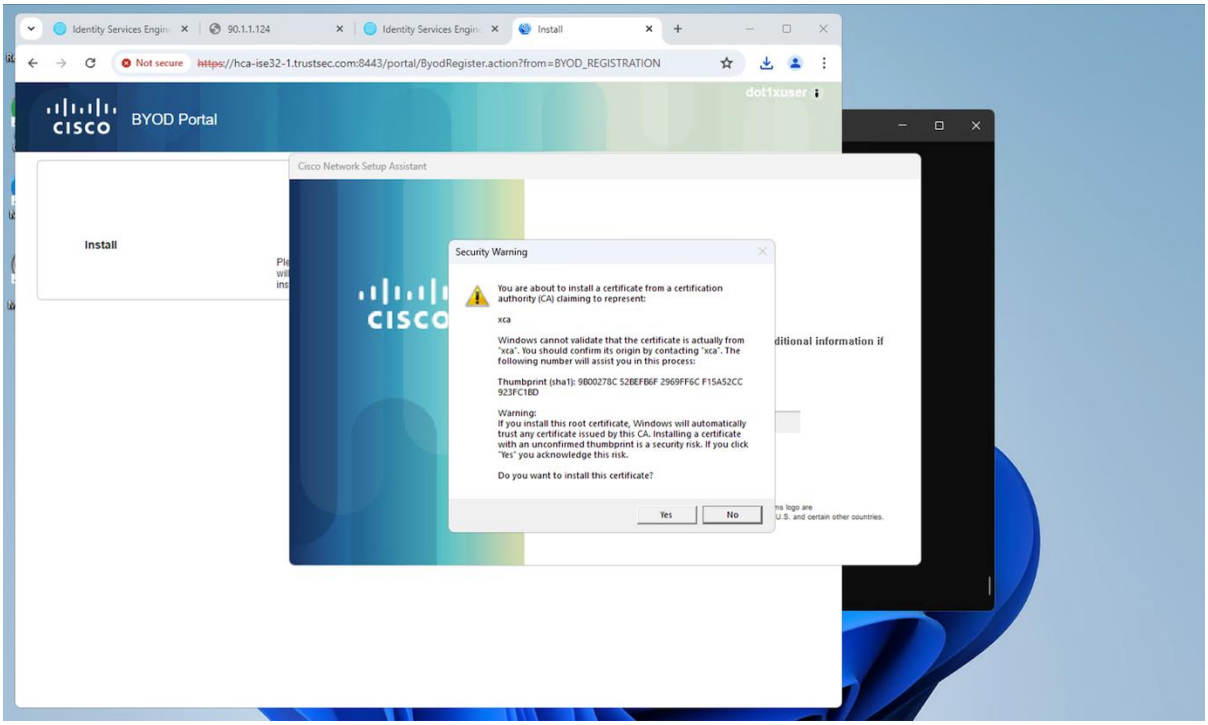

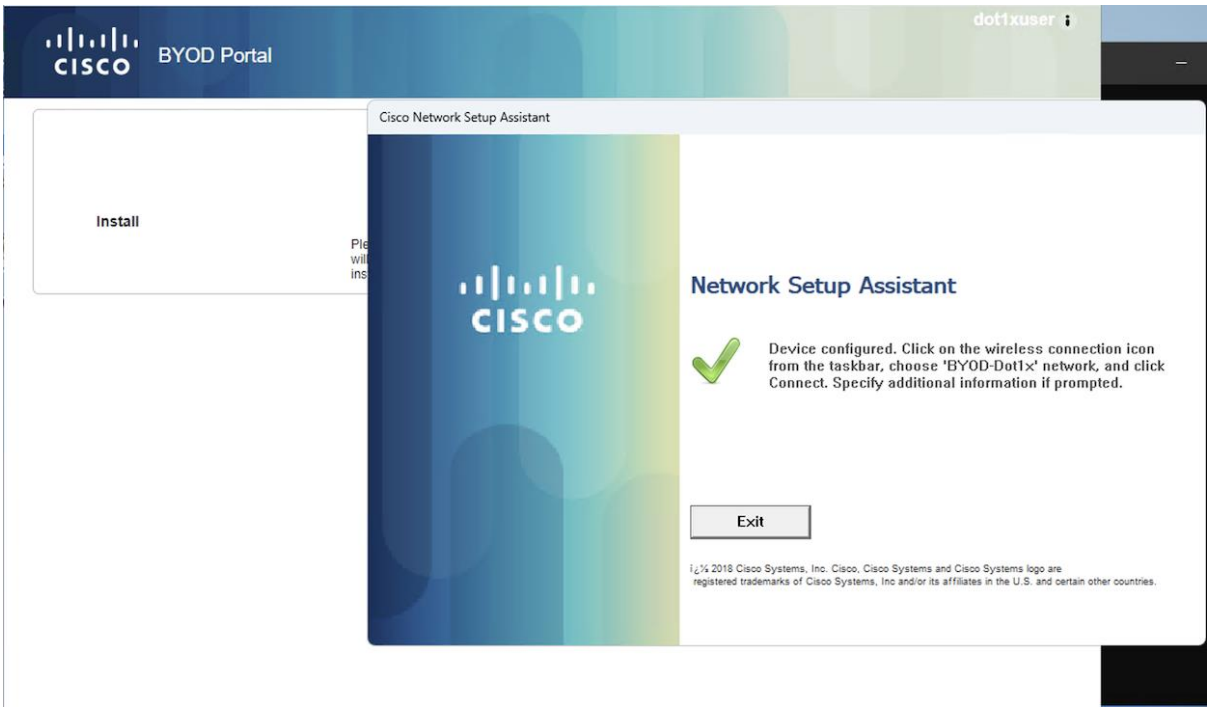

Figure 33 illustrates ISE Authentication of the BYOD devices.

Figure 33. BYOD device authentication

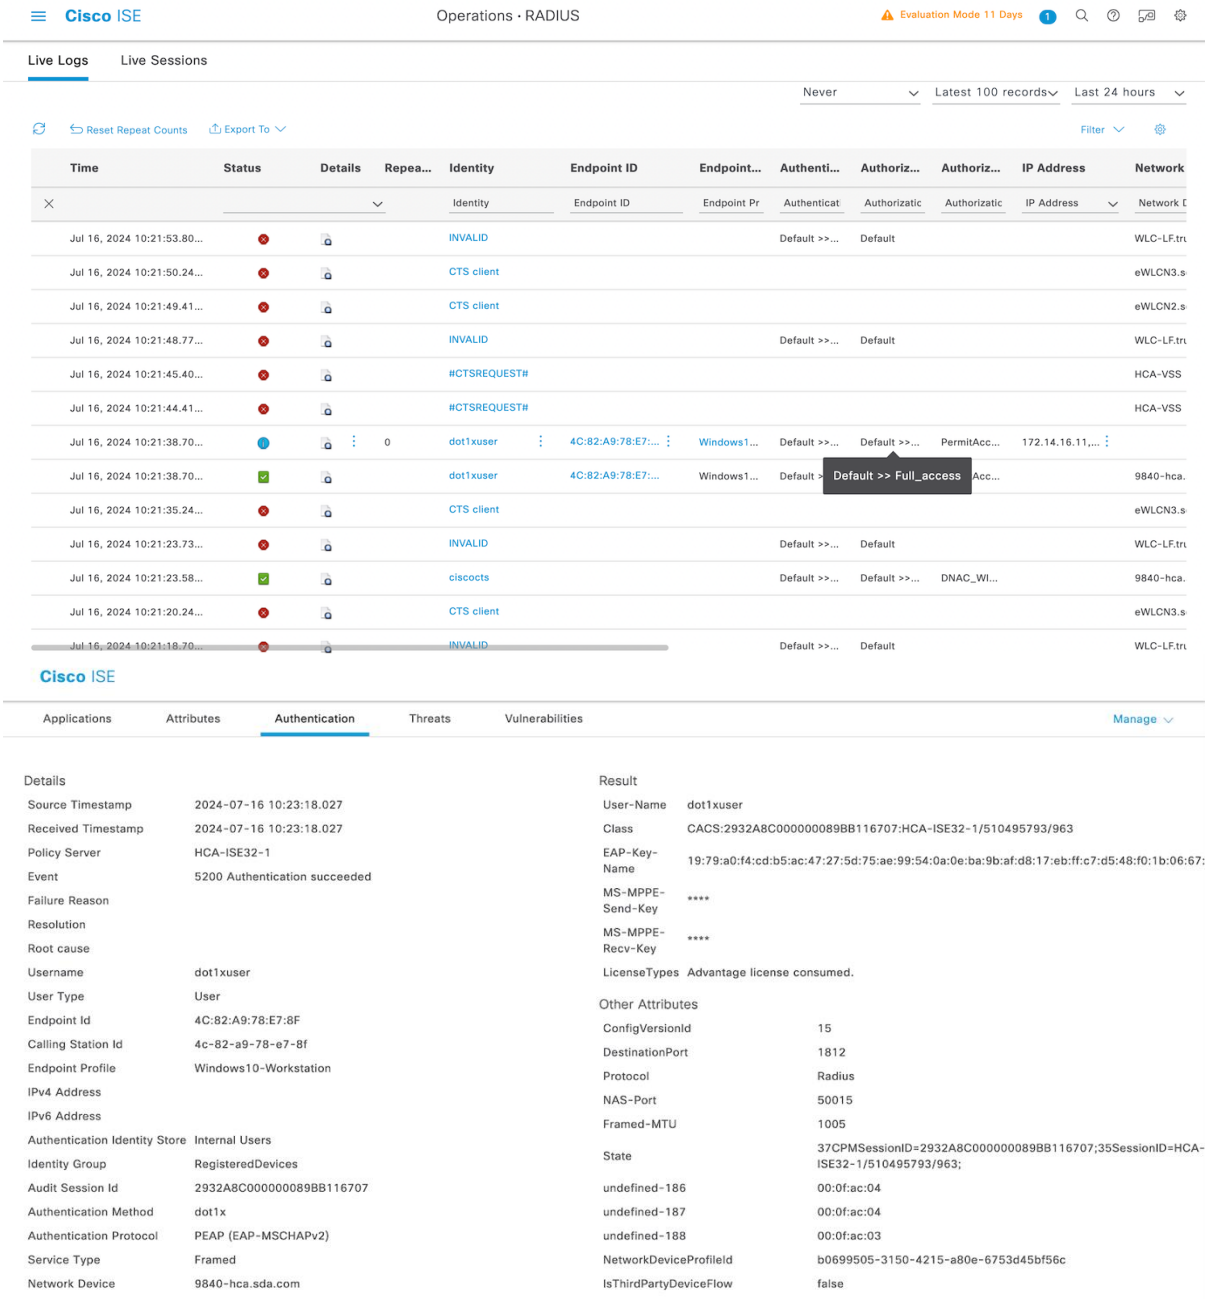

Hardware and software specifications

The university vertical has been tested with the hardware and software specified in the table. For a comprehensive list of hardware supported by the Cisco SD-Access solution, see the Cisco Software-Defined Access Compatibility Matrix.

Hardware or software component	Supported software version	
Catalyst Center Appliance (Part Number: DN2-HW-APL-XL)	2.3.7.7	2.3.7.9
Identity Services Engine (ISE)	3.3 Patch 4	3.3 Patch 4
Control Plane Node: Catalyst 9500 Series Switches	17.9.5, 17.12.4	17.9.6a, 17.12.5, 17.15.3
Fabric Border Node: Catalyst 9200, 9300 & 9400 Series Switches	17.9.5, 17.12.4	17.9.6a, 17.12.5, 17.15.3
Wireless Controller: Catalyst 9800-40 and 9800-CL	17.9.5, 17.12.4	17.12.5, 17.9.6
Cisco SD-Access Extended Node: Catalyst 9200 IE-4000	15.2(7)E10	17.15.3 15.2(8)E5

Multidimensional scale numbers

Category	Value
Device Inventory	2000
Devices per Fabric Site	600
Buildings and Floors	3000
VNs per Fabric Site	64
IP Pool per Fabric Site	500
WLCs per fabric Site	2
Fabric Sites	4
APs in Inventory	8000
Endpoints	100,000 (80,000 wireless, 20,000 wired)
SSIDs	8
SDGs	25
Bonjour Service Instances	16,000

Note: The officially supported scales are outlined in this data sheet, while the data provided here has been validated in the lab.

Links to relevant Cisco documentation

- Cisco SD-Access Solution Design Guide (CVD)
- Catalyst Center User Role Permissions
- Implement Disaster Recovery
- Release Notes for Cisco Catalyst Center
- Cisco Catalyst Center Security Best Practices Guide
- Software Defined Access (SD-Access) Provisioning Best Practice Guide

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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