



Validated Profile: IT-Enabled Services Vertical

[Solution Overview](#) 3

[Scope](#) 3

[Traditional Network Versus Cisco SD-Access](#) 3

[Challenges in Traditional Networks](#) 4

[The Importance of Cisco SD-Access](#) 4

[IT Enabled Services Network Overview](#) 5

[Cisco Catalyst Center](#) 5

[Cisco Identity Service Engine](#) 6

[Cisco SD-Access Fabric](#) 6

[Fabric Architecture Overview](#) 7

[Network Architecture](#) 7

[Fabric Roles](#) 7

[Fabric in a Box](#) 9

[Extended Nodes](#) 10

[Fabric Wireless LAN Controllers and Access Point](#) 10

[SD-Access Embedded Wireless](#) 10

[Transits](#) 11

[Cisco Catalyst 9000 Series Switches](#) 12

[Cisco Catalyst Wireless LAN Controller and Access Point](#) 12

[Compatibility Matrix](#) 12

[ITES Deployment and Design Solutions](#) 12

ITES Business Outcome and Challenges	14
Solutions to ITES Business Outcomes	15
Financial Efficiency	25
Experience Improvement	29
Network Deployment Options	31
Validated Solution Use Cases	42
ITES ODC Deployment Models	44
Supported Hardware and Software	75
Multidimensional Scale Numbers	76
Links to Relevant Cisco Documentation	76

Solution Overview

Information technology-enabled services (ITES) organizations are companies that provide numerous IT services, such as customer support, technical help desks, software development, and data processing. ITES organizations leverage technology to deliver these services, often over digital networks like the internet. ITES companies are an integral part of the global outsourcing industry, offering cost-effective solutions to businesses around the world. They are known for their ability to provide high-quality services, often at a lower cost compared to in-house operations in developed countries. These organizations play a crucial role in driving innovation, improving efficiency, and enabling businesses to focus on their core competencies.

This guide offers guidance on navigating challenges, solutions, deployment options, operational management, and migration within an ITES network deployment. This guide encapsulates validated end-to-end use cases, scalability insights, and hardware and software recommendations to facilitate optimal decision-making during network and organizational deployment processes. This guide directs readers to associated design and deployment guides for enterprise networks, furnishing valuable insights into deploying common implementations of Cisco SD-Access.

Scope

This guide serves as a roadmap for understanding ITES network challenges, common use cases, and how Cisco SD-Access can address them. Although this guide doesn't provide in-depth configuration steps, it equips you with valuable insights for your ITES 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 at 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 policy-based automation.
- It simplifies network design by carrying 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

Organizations implementing traditional network architectures face increasing challenges due to the expanding number of users, devices, and device types. Organizations prioritize managing and monitoring traffic from these entities to safeguard their corporate infrastructure in case of device or network compromise. In traditional networks, the reliance on numerous VLANs and manual Access Control Lists (ACLs) across various devices can lead to potential misconfiguration issues. As businesses grow and incorporate more devices and locations, the complexity and potential for errors also increase. Updating security rules across the enterprise becomes a labor-intensive task.

When the enterprise adds a new branch, the network operations team may need to update the ACLs in both the headquarters and branch locations. Any errors during these updates can lead to inconsistent security policies and potential security breaches. Network administrators must dedicate substantial time and effort to plan and implement network changes, ensuring that each device is securely integrated into the network using the appropriate network segment. Traditional network building methods do not address the evolving nature of networks and the growing concerns regarding security.

The Importance 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 the following technological advantages:

- **Simplified Operations:** The IP subnet stretching simplifies network design. The configuration is centralized at the fabric edge and border nodes, eliminating hop-by-hop configuration.
- **Automation:** Automates routine network operations such as configuration, provisioning, and management. This reduces the risk of human error and increases efficiency. Cisco Catalyst Center streamlines the deployment, eliminating the need for command-line interfaces (CLI) or CLI configuration templates.
- **Agility:** Network operations become more agile and align with business requirements by minimizing manual configuration steps.
- **Security:** Provides enhanced security and segmentation through Virtual Network (VN) and Security Group Tags (SGT). SD-Access provides a strong framework for securing and managing complex enterprise networks through macro-segmentation with Virtual Routing and Forwarding (VRFs) and VNs, and micro-segmentation with 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.

IT Enabled Services 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 an ITES network, proceed to the following sections to delve deeper into the SD-Access fabric components. Learn about the benefits that Cisco SD-Access solutions offer in addressing the requirements and challenges specific to the ITES sector.

You can manage traditional networks using Cisco Prime Infrastructure or Catalyst Center. Catalyst Center provides automation, monitoring, and telemetry capabilities for both traditional networks and SD-Access environments. If you are managing a network with Cisco Prime Infrastructure and plan to migrate to Catalyst Center, see [Cisco Prime Infrastructure to Cisco Catalyst Center Migration](#).

To migrate existing Cisco Catalyst legacy networks to a Cisco SD-Access fabric, see [Migration to Cisco SD-Access, on page 40](#), which explains options to migrate existing networks with 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.

Using Catalyst Center, network administrators can perform the following tasks:

- **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 the following 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.

Following are the key features of Cisco ISE:

- **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.
- **BYOD (Bring Your Own Device) Support:** Enable secure Bring Your Own Device (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. For SD-Access single-site deployments, it is recommended to have a basic two-node ISE deployment, with each ISE node running all services (or functions/roles) for redundancy.

For details on ISE deployment models, see Cisco Identity Services Engine Administrator Guide [here](#).

For additional information on ISE deployment and scalability, see Performance and Scalability Guide for Cisco Identity Services Engine [here](#).

Cisco SD-Access Fabric

Cisco SD-Access (Software-Defined 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 components and features of Cisco SD-Access Fabric include the following:

- **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.
- **Automation:** Automates network provisioning, configuration, and management tasks, reducing manual errors and increasing operational efficiency.
- **Identity Services Engine (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 (Software-Defined 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 virtual networks (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

In Cisco's Software-Defined Access (SD-Access) architecture, various roles and components work together to provide a secure, automated, and scalable network solution. Following are the primary roles within the SD-Access fabric:

- Control Plane Node
- Border Node

- Edge Node
- Intermediate Node

Control Plane Node

The SD-Access fabric control plane node operates using the Locator/ID Separation Protocol (LISP), Map-Server(MS) and Map-Resolver(MR) functionality combined on a single node. The control plane node's database tracks all endpoints in the fabric site and associates the endpoints to fabric nodes, decoupling the endpoint IP address or MAC address from the location of the closest router in the network.

Control Plane node enables the following functions:

- Host Tracking Database
- Endpoint identifiers (EID)
- Map-Server (MS)
- Map-Resolver (MR)

Border Node

The fabric border nodes serve as the gateway between the SD-Access fabric site and the networks external to the fabric. The border node is responsible for network virtualization interworking and SGT propagation from the fabric to the rest of the network.

Border nodes can implement the following functions:

- Advertisement of Anycast Layer 3 Gateway subnets
- Fabric site exit point
- Network virtualization extension to the external world
- Policy mapping
- VXLAN encapsulation/de-encapsulation

Edge Node

The SD-Access fabric edge nodes are the equivalent of an access layer switch in a traditional campus LAN design. The edge node functionality is based on the Ingress and Egress Tunnel Routers (xTR) in LISP. The edge nodes must be implemented using a Layer 3 routed access design.

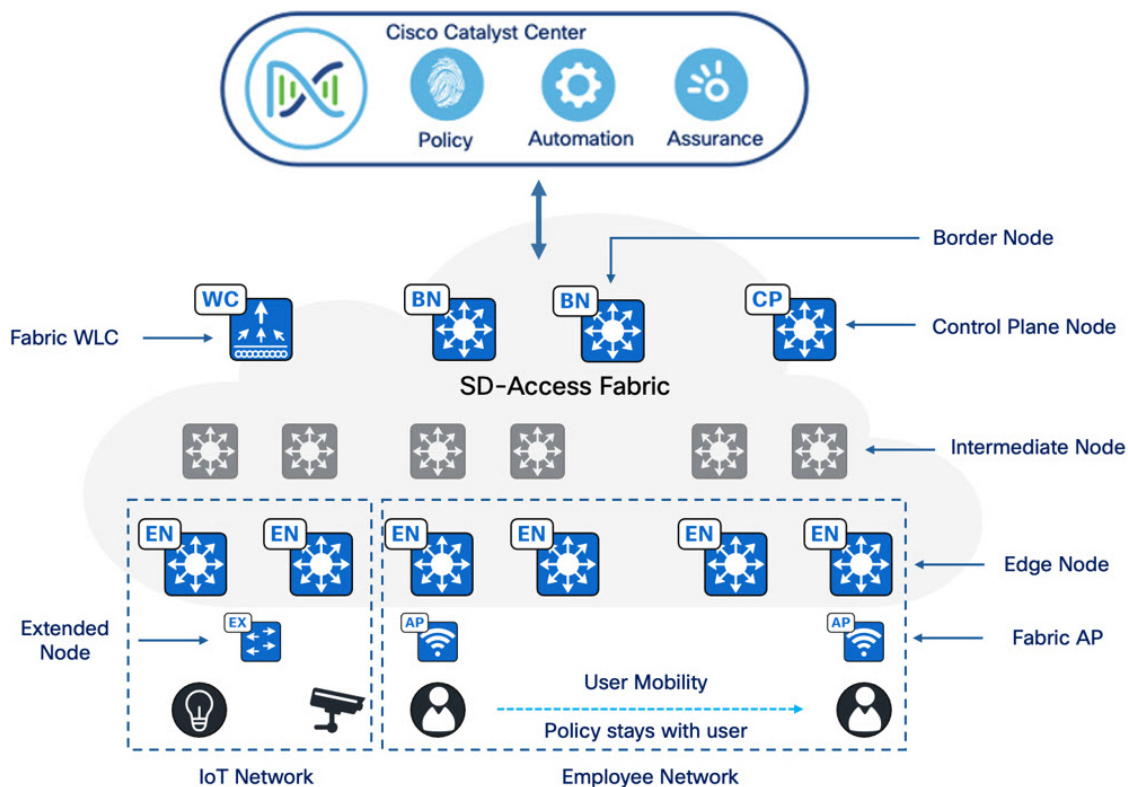
Fabric Edge nodes support the following functions:

- Register the endpoint in the local database and the control plane
- Create Layer 3 Access Gateway
- Mapping of users to virtual networks
- Authentication and authorization of endpoint in concert with ISE
- VXLAN encapsulation and de-encapsulation

Intermediate Node

Intermediate nodes are part of the Layer 3 network that interconnects devices operating in fabric roles, such as connections between border nodes and edge nodes. These nodes provide IP reachability, physical connectivity, and support the increased MTU requirements needed to accommodate larger IP packets encapsulated with fabric VXLAN information. However, it is important to note that Intermediate Nodes do not have overlay capabilities or participate in overlay functions.

The diagram below highlights the key components involved in an SD-Access fabric deployment and shows their respective 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.

Fabric In a Box (FIAB) provides the following benefits:

- Simplicity
- Cost-Effectiveness
- Faster Deployment
- Ideal for branches and small-sized deployments

For details on StackWise Virtual deployment white paper, see [here](#).

Extended Nodes

SD-Access Extended Nodes enable the extension of the enterprise network to non-carpeted areas, often referred to as the Extended Enterprise. These 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 a detailed guide on Extended Node design, see [here](#).

Fabric Wireless LAN Controllers and Access Point

Fabric Wireless LAN Controllers (WLCs) and traditional WLCs manage Access Point (AP) images and configurations, handle client sessions, and offer mobility services. Fabric WLCs additionally support fabric integration by registering MAC addresses of wireless clients into the host tracking database of fabric control plane nodes.

Fabric-mode APs (Access Points) that are Cisco Wi-Fi 6 (802.11ax) and 802.11ac Wave 2 APs associated with the fabric WLC, 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 on wireless operations and interactions with SD-Access Wireless, Fabric WLCs, and Fabric APs, see SD-Access Wireless Design and Deployment Guide [here](#).

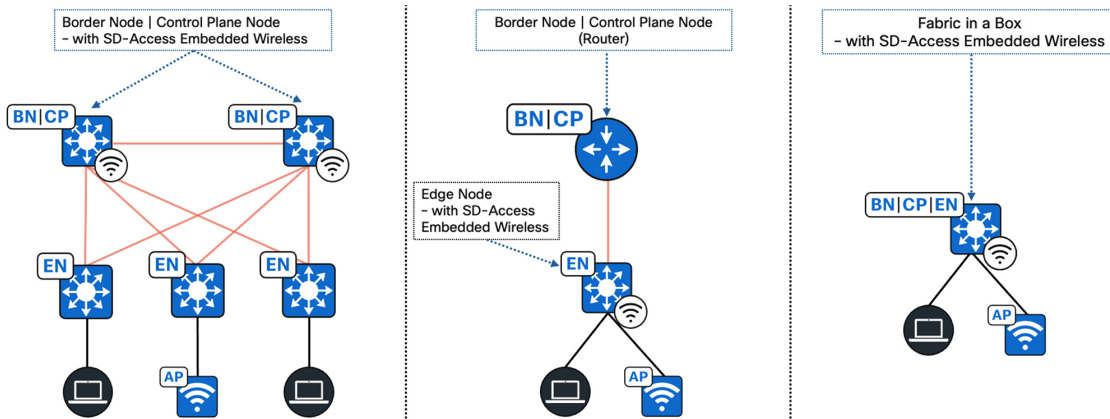
SD-Access Embedded Wireless

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

The Catalyst 9800 Embedded Wireless Controller is supported for 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.

Following are the 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 will automate the Border Node configuration for connections between Fabric Sites or between a Fabric Site and an external domain. Following are the two types of Transits:

- **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-Aware 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 WAN/MAN network, much like SD-WAN and DMVPN.

Following is a concise comparison between IP-Based Transit and SD-Access Transit:

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.
- **Suitable for 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

- **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.
- **Centralized control:** Uses a domain-wide Control Plane node for simplified management.
- **Requires dedicated infrastructure:** Requires additional infrastructure for the SD-Access transit control plane.

Ensure the following key considerations when using an SD-Access Transit:

- Connections should accommodate the recommended MTU settings used for Cisco SD-Access in the Campus Network.
- IP reachability must exist between Fabric Sites. Specifically, there must be a known underlay route between all Fabric Nodes.



Note The default route cannot be used for this purpose.

- Support for Underlay SSM is necessary if multicast traffic will traverse SD-Access Transit.

For detailed information on Cisco SD-Access components and architecture, see [Cisco SD-Access](#).

Cisco Catalyst 9000 Series Switches

Cisco Catalyst 9000 series switches offers 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 [data sheet](#).

Cisco Catalyst Wireless LAN Controller and Access Point

Cisco Catalyst 9800 series Wireless LAN Controllers (WLC) and Access Points (AP) provide seamless network management and deployment in both on-premises and cloud for wireless clients.

To refer the data sheet for Catalyst 9800 and Catalyst 9100 devices, see the following links:

- [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. Refer to the compatibility matrix for a complete list of supported Cisco products:

- [Cisco Catalyst Center Compatibility Matrix](#)
- [Cisco SD-Access Compatibility Matrix](#)

ITES Deployment and Design Solutions

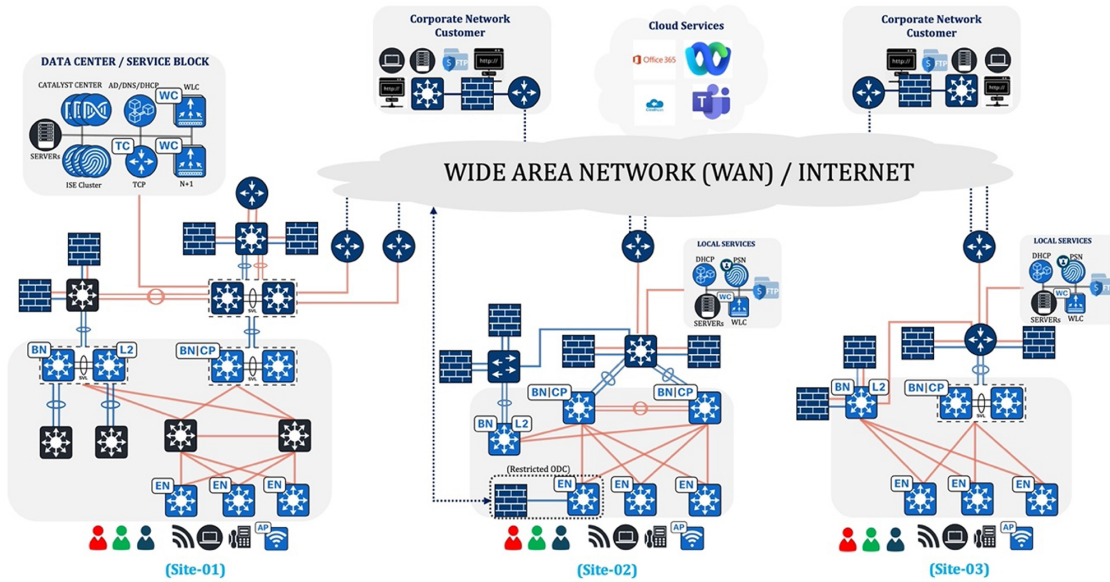
Profile Deployment

This section provides design guidance for the ITES sector, emphasizing its requirements and the utilization of Cisco SD-Access to create a network that is simple, secure, and flexible.

This section discusses the topologies, use cases, and solutions that focus on meeting the standard deployment options for ITES while addressing their themes and requirements.

Solution Topology

Following is an overview of the Device and Firewall Placement in the Topology:



Site-01: Large Site

- Cat9600 SVL switch serves as both the Border Node and the Control Plane Node.
- Cat9500 SVL switch serves as a dedicated Layer 2 Border Node.
- Cat9300 and 9400 switches serve as Edge Nodes.
- C2S/S2S firewalls functioning as gateways are connected to an aggregation switch positioned beyond the Layer 2 Border.
- C2S/S2S firewalls that do not function as gateways are connected to an aggregation switch located beyond the Fusion Node.

Site-02: Medium Site

- Cat9500 switch serves as both the Border Node and the Control Plane Node.
- Cat9500 switch serves as a dedicated Layer 2 Border Node.
- Cat9200 and 9400 switches serve as Edge Nodes.
- C2S/S2S firewalls functioning as gateways are connected to an aggregation switch positioned beyond the Layer 2 Border.
- C2S/S2S firewalls that do not function as gateways are connected directly to the Fusion Node.

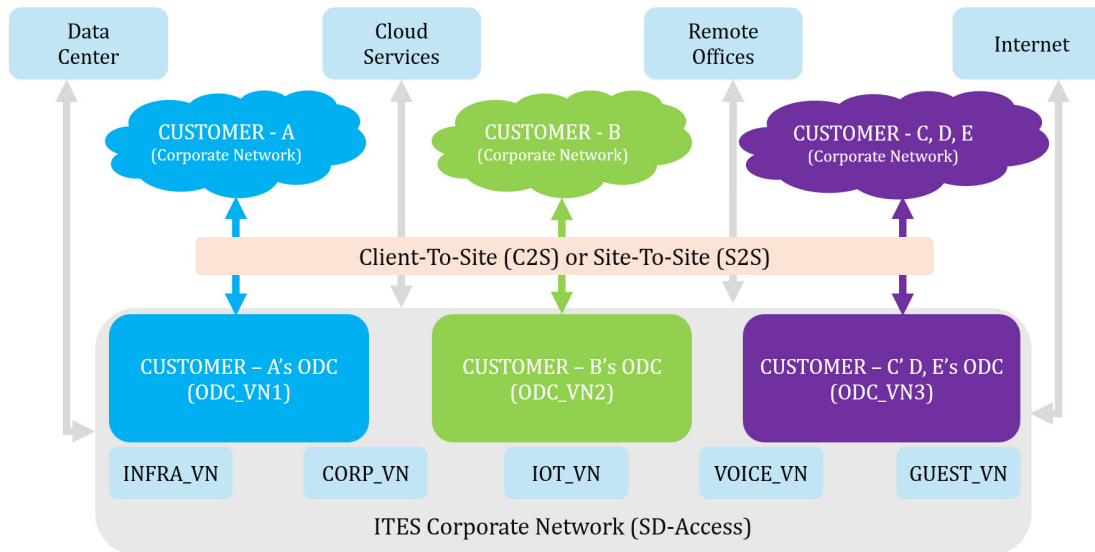
Site-03: Small Site

- Cat9500 SVL switch serves as both the Border Node and the Control Plane Node.2
- Cat9500 switch serves as a dedicated Layer 2 Border Node.
- Cat9300 switches serve as Edge Nodes.

- C2S/S2S firewalls functioning as gateways are connected directly to the Layer 2 Border.
- C2S/S2S firewalls that do not function as gateways are connected directly to the Fusion Node.

Logical Diagram

The logical diagram for ITES outlines the network architecture of an Information Technology Enabled Services (ITES) environment, depicting the connections between customer networks and ITES corporate networks. It highlights the secure and efficient data flow across various segments within an SD-Access infrastructure.



ITES Business Outcome and Challenges

ITES (Information Technology Enabled Services) refers to the outsourcing of various processes and services, enabled by technology. Businesses across industries are increasingly leveraging ITES to improve efficiency, reduce costs, and enhance customer experience. However, following are some of the ITES challenges and potential outcomes:

- Security
- Compliance
- Operational
- Financial
- Experience

Security

For an ITES company, enhancing security measures, mitigating risks, and ensuring compliance with regulatory standards can be achieved by implementing robust security protocols, conducting regular risk assessments, and adhering to industry-specific regulations and standards. Cybersecurity threats pose the greatest concern for an ITES company's Chief Information Security Officer (CISO). The rapid shift to hybrid work and the evolution of digital business services for customers have significantly increased the attack

surface vectors available to cybercriminals. Unchecked, these malicious actors can exploit vulnerabilities, leading to substantial losses both financially and in terms of reputation. The CISO group regularly reviews fundamental security practices and processes.

Compliance

For an ITES company, compliance with regulatory standards is paramount to maintain trust, security, and legality in their operations. These companies are often entrusted with handling sensitive data and providing critical services to clients across various industries. Thus, adherence to compliance regulations is essential to ensure the confidentiality, integrity, and availability of data. Non-compliance not only exposes ITES companies to legal repercussions but also risks damaging their reputation and losing valuable client trust. Therefore, a proactive approach to compliance is crucial for the success and sustainability of ITES companies in today's regulatory landscape.

Operational

For ITES companies, network uptime is paramount to smooth operations and achieving business goals. Since ITES networks are mission-critical, the ultimate goal is to get as close to 100% availability as possible. Five-nines availability (99.999% uptime) represents a significant step towards this objective, allowing only 5 minutes and 16 seconds of downtime annually. Seamless and uninterrupted services are essential for ITES customer productivity and business success. By implementing automation, monitoring, load balancing, and failover mechanisms, ITES firms can achieve or even surpass the five-nines availability target.

Financial

Operational expenses are a major focus for ITES businesses. Streamline expenses and boost earnings through the automation of deployment across thousands of sites while minimizing the need for on-site network operations whenever feasible. Large-scale multisite deployments are common in the ITES sector, often encompassing hundreds of Offshore Development Centers (ODCs) distributed across extensive geographic areas. Managing such networks box-by-box or site-by-site with onsite teams poses significant challenges

To address the complex requirements of ITES, a solution is required to quickly set up any site or Offshore Development Center (ODC) in any location and manage it remotely. This enables ITES organizations to maintain an efficient IT staff. Achieving this entails implementing network automation and monitoring to streamline deployment and troubleshooting procedures.

Experience

Enhance user and application experiences by strategically utilizing modern technologies that support essential business capabilities. Beyond security, compliance, and availability concerns, a network with inconsistent or slow Quality of Service (QoS) can lead to poor customer satisfaction and financial losses. In environments such as time-sensitive operations, where delays are critical, low latency and consistent QoS are crucial to meet organizational requirements.

Solutions to ITES Business Outcomes

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

Security Challenges

Information Technology Enabled Services (ITES) sector faces significant security challenges due to its complex and dynamic environments, such as increased attack surfaces, data breaches, insider threats, regulatory compliance, sophisticated cyber attacks, and the security of remote work. Cisco's Software-Defined Access (SD-Access) framework addresses these challenges through a comprehensive set of tools and capabilities:

- Macro Segmentation

- Micro Segmentation
- TrustSec Models
- Group-Based Policy Analytics (GBPA)
- AI (Artificial Intelligence) Endpoint Analytics
- Endpoint Security with Zero-Trust Solution
- Isolation of Guest Users

Macro Segmentation

For ITES networks, assign different virtual routing and forwarding instances (VRFs) network endpoints such as employees, monitoring devices, and guests to implement a recommended segmentation strategy. SD-Access offers the ability to macro-segment endpoints into different VRFs, which can be configured within the network using Catalyst Center.

Following are few examples demonstrating the implementation of Virtual Networks (VNs):

- INFRA VN: This VN is exclusively for access points, classic and policy extended nodes for connectivity and is mapped to the global routing table.
- Employee VN: Use this VN for regular employee 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.
- ODC VN: Use this VN for employees working on client projects, ensuring secure and segregated connectivity to the client's corporate network.

An ITES company can effectively segment and secure diverse types of traffic, enhance overall network performance and security by implementing VNs in an SD-Access network.

Micro Segmentation

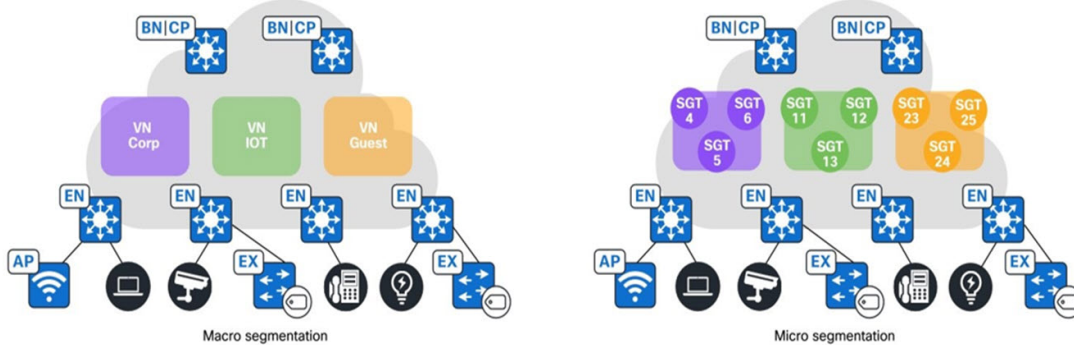
Micro-segmentation simplifies the provisioning and management of network access control by using security groups to classify traffic and enforce policies, allowing for more granular security within SD-Access virtual networks

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 micro-segmentation, allowing the use of the same subnet with a user and endpoint-centric approach. Dynamic authorization assigns different Security Group Tags (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 further information on Macro and Micro Segmentation, see the Software-Defined Access Macro Segmentation Deployment Guide [here](#).

The following image provides an example illustrating both macro segmentation and micro segmentation:



TrustSec Models

Cisco TrustSec is a security solution that simplifies the provisioning and management of network access, enforcing security policies across an organization. It provides comprehensive segmentation and access control based on roles and policies rather than traditional IP-based methods. This approach enhances security and operational efficiency in both wired and wireless environments.

Following are the two main TrustSec models:

- 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.

The default action is to deny 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.

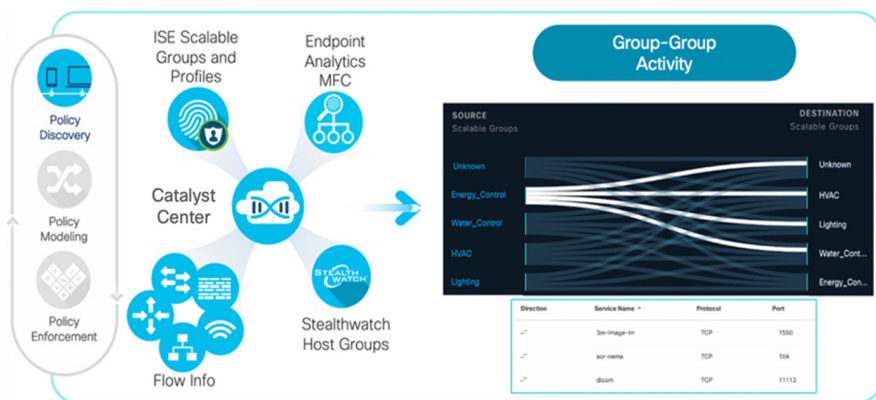
For more details on the TrustSec model, see Cisco ISE TrustSec Allow-List Model (Default Deny IP) with SDA [here](#).

Group-Based Policy Analytics

High-profile cyber attack news is driving ITES organizations 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 empowers network administrators with the following capabilities:

- Discover and visualize group interactions: GBPA analyzes network traffic flows to identify how different network groups such as departments, functions, etc., communicate.
- 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.



As seen in the diagram, GBPA leverages information from the following sources to create a holistic view of your network:

- Cisco Identity Services Engine (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 Artificial Intelligence (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 further details on Cisco AI Endpoint Analytics, see the following resources:

- [Cisco SD-Access AI Endpoint Analytics](#)
- [Cisco Catalyst Center Guide - AI Endpoint Analytics](#)

Endpoint Security with Zero-Trust Solution

Endpoint security with zero-trust solutions in Software-Defined Access (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. We apply zero-trust principles, which 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 the following 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 information on SBEN, see the Configure Supplicant-Based Extended Nodes section in the [Cisco Catalyst Center User Guide](#).


Isolation of Guest Users

Managing guest Wi-Fi across all campuses with a single subnet can be challenging for network administrators. Cisco SD-Access addresses this issue with Multisite Remote Border (MSRB) using Virtual Network Anchors. MSRB enables guest traffic from the same virtual network at different locations to be routed back to a central 'Anchor Site.' This approach allows the network to use one common subnet for all guest devices, eliminating the need for separate subnets at each site.

When a guest endpoint connects to the guest SSID and passes authentication through Cisco ISE, it is associated with the anchored guest Virtual Network. The system tunnels guest traffic to the Anchor Site Border Node and exits to the Internet through a firewall. Cisco ISE recommends the following options for ITES customers who need isolated guest environments:

- Host a Separate Policy Service Node (PSN) in the demilitarized zone (DMZ) for guest users
- Dedicate PSN for guest users

Separating authentication, authorization, and Multisite Remote Border for guest users fully isolates the guest network from the control plane, data plane, and policy plane perspectives. This ensures that the guest network is separated from other ITES users, devices, and resources.



The figure demonstrates the Guest Control Plane and Border Node positioned in the demilitarized zone (DMZ). In this solution, traffic is still encapsulated at the AP in the VXLAN to the fabric edge switch, but the FE is configured to use a different border node. This border node can reside in your DMZ, providing complete traffic isolation similar to the guest anchor solution. Guest users register with a dedicated Control Plane which may or may not be co-located with the border, and receive an IP address in the DMZ.

Catalyst Center configures Central Web Authentication (CWA), External Web Authentication (EWA), and hotspot SSIDs on Cisco AireOS and Polaris-based Wireless LAN Controllers (Cisco Catalyst 9800 and embedded wireless on Cisco Catalyst 9000 platforms) to enable guest access flow in Cisco SD-Access network.

See the configuration guide for Cisco SD-Access Guest Automation [here](#).

Compliance Regulations

Compliance regulations refer to the set of rules and standards that organizations must adhere to operate lawfully within a particular industry or jurisdiction. These technologies aid in ensuring compliance by automating regulatory processes, enhancing data security, and providing real-time monitoring and reporting capabilities 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 (RBAC)
- Audit Logs
- Configuration Compliance
- Configuration Drift

Role-Based Access Control

Role-Based Access Control 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 refer to 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 information and instructions on accessing the audit logs in Catalyst Center, see [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 perform on demand using the following schedule options:

- Automated Compliance Check: Use 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: Lets 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 information and instructions on enabling or managing compliance for network devices using Catalyst Center, see [Compliance User Guide](#).

Configuration Drift

Configuration drift occurs when the actual configuration settings of network devices deviate from their intended or predefined state over time. In ITES organizations, compliance mandates often require maintaining archives of configurations for all network devices. Catalyst Center offers support for configuration drift, allowing users to track and monitor changes in device configurations. This feature enables users to review the current configuration of each device and analyze historical changes from the past month to understand how configurations have evolved over time on a specific device.

For more information, see [Configuration Drift of a Device](#) in the *Cisco Catalyst Center User Guide*.

Operational Efficiency

Operational efficiency is vital for ITES businesses, directly enhancing productivity, cost-effectiveness, and service quality. This efficiency enables them to maximize employee output, streamline digital transformation efforts, and ultimately, enhance their reputation and brand value. SD-Access addresses the following key aspects of operational efficiency.

- High Availability
- 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. Here's an overview of achieving high availability for the following components:

- Disaster Recovery
- Resilient Network Architecture
- Fallback Segments

Disaster Recovery

ITES organizations have a low tolerance for management, control, or data plane failure. Catalyst Center supports both intracluster and intercluster resiliency. The Disaster Recovery implementation in Catalyst Center consists of three components: the main site, the recovery site, and the witness site. At any given time, the main and recovery sites operate in either active or standby roles. The active site manages your network, while the standby site maintains a continuously updated copy of the active site's data and managed services. Whenever an active site goes down, Catalyst Center automatically initiates a failover, completing the tasks necessary to designate the former standby site as the new active site.

For additional details on Disaster Recovery in Catalyst Center, see [Implement Disaster Recovery](#).

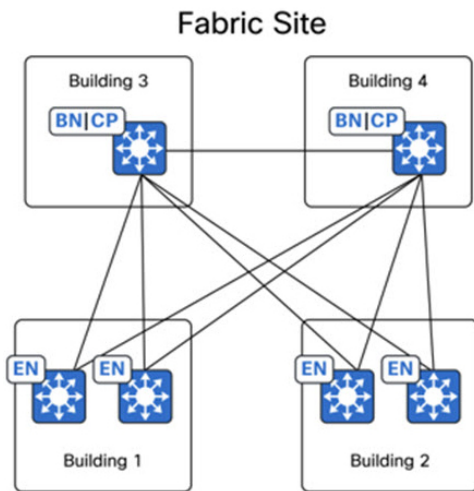
Resilient Network Architecture

Resilient network architecture in SD-Access is designed to provide a highly available and reliable network infrastructure, ensuring that critical services remain operational even in the face of disruptions.

- Similar to 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 level. This can remove the requirement for spanning tree and first hop redundancy protocols, and their related configurations.
- With Layer 3 routed access, the boundary between Layer 2 and Layer 3 shifts from the distribution layer to the access layer. This eliminates the need for the distribution and collapsed core layers to handle Layer 2 adjacency and redundancy.

In ITES networks, alongside traditional resilience methods like stacking and StackWise Virtual, regional hubs and campus headquarters often need protection from building failures to ensure continuous connectivity to data centers for critical applications.

Cisco SD-Access provides a flexible deployment architecture that allows fabric borders to be positioned in different physical sites while integrating them under a single Fabric Site, as demonstrated in the following figure:



As depicted in the figure, Buildings 1-4 belong to the same Fabric Site, with the Co-located 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 a 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 clients that have already been onboarded, if the connection to the ISE Policy Service Node 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 micro-segmentation 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. Here is an overview of how to achieve this for the following components:

- Catalyst Center High Availability
- Identity Services Engine High Availability
- Cisco Wireless LAN Controller Redundancy

Catalyst Center High Availability

Catalyst Center's High Availability (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. High Availability 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 additional details on Catalyst Center High Availability, see [Cisco Catalyst Center High Availability Guide](#).

Identity Services Engine High Availability

Cisco Identity Services Engine (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 additional details on Cisco ISE High Availability, see "Distributed Deployment Scenarios" in the [Cisco Identity Services Engine Installation Guide](#).

Cisco Wireless LAN Controller Redundancy

Cisco Wireless LAN Controller (WLC) redundancy is essential for maintaining continuous wireless network services. In an HA pair setup, two wireless controllers are configured as a pair. One WLC 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 WLC hardware or software failures. Consequently, users experience minimal disruption and maintain connectivity to the wireless network.

For additional details on Cisco Wireless LAN Controller, see [Cisco Catalyst 9800 Series Wireless Controllers High Availability SSO Deployment Guide](#).

Reports

Catalyst Center's 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 includes the following:

- **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 detailed information on Reports, see the [Cisco Catalyst Center Platform User Guide](#).

Efficient Troubleshooting

Efficient troubleshooting is a critical component to support business operations for ITES customers. Catalyst Center offers comprehensive debugging ability features designed to meet these needs effectively. These features empower IT administrators to quickly identify, diagnose, and resolve Catalyst Center issues, ensuring continuous and optimal performance of the network infrastructure. The following tools help with troubleshooting:

- **Validation Tool:** Before Catalyst Center 2.3.5.x, the Audit and Upgrade Readiness Analyzer (AURA) tool assessed the upgrade readiness of a cluster. With the restricted shell fully implemented in 2.3.5.x, most of the AURA upgrade checks are now implemented in Catalyst Center. The Validation Tool tests both Catalyst Center appliance hardware and connected external systems and identifies any issues that need to be addressed before they seriously impact your network.

For details on the Catalyst Center Validation Tool, see the following links:

- [Validate Cisco DNA Center Upgrade Readiness](#)
 - [Use the Validation Tool](#)
- **System Analyzer:** To address troubleshooting needs, the System Analyzer tool provides efficient log file retrieval. The System Analyzer performs comprehensive assessments and diagnostics to ensure the optimal functioning and reliability of Catalyst Center and its connected network components. By leveraging the System Analyzer capabilities for monitoring, diagnostics, and performance optimization, organizations can enhance operational efficiency, ensure compliance with security standards, and deliver reliable ITES services.

For more information, see [Use the System Analyzer Tool](#).

Overall, the Catalyst Center Validation Tool and System Analyzer are invaluable assets for ITES network administrators. These tools promote proactive maintenance, efficient troubleshooting, and enhanced network stability, significantly boosting operational efficiency for ITES delivery.

Financial Efficiency

Reducing operational expenses and increasing earnings are major priorities for ITES businesses. By automating the deployment and monitoring of large-scale, multisite networks, ITES organizations can significantly reduce operational expenses, streamline processes, and maintain efficient IT operations. This approach enables the management of complex networks with minimal manual intervention, enhancing overall productivity and profitability. Here are some of the approaches adopted for achieving financial efficiency in ITES organizations:

- 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 the following components:

- LAN Automation
- Plug and Play (PnP) and Return Material Authorization (RMA)
- Software Image Management (SWIM)
- Intelligent Capture (iCAP)
- 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 the following 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 modeled 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 details on LAN automation, 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 access points. 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 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 details, see the [Network Device Onboarding for Cisco Catalyst Center Deployment Guide](#).

Software Image Management

Catalyst Center's 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 Catalyst Center's 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 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. Intelligent Capture 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, Intelligent Capture allows you to access data from APs that is not available from wireless controllers.

For 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/Client 360:** View device or client connectivity, which includes information on topology, throughput, and latency from various times and different applications.
- **Network Time Travel:** Ability to go back in time and see the cause of a network issue.
- **Application Experience:** Provide unparalleled visibility and performance control on the applications critical to your core business on a per-user basis.
- **Network Analytics:** Provide recommended corrective actions for found issues in the network. These actions can involve guided remediation, where the engine specifies steps for a network administrator to perform.

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 details on the IPAM integration with Catalyst Center, 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 the following 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 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, enterprise administrators can 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)**

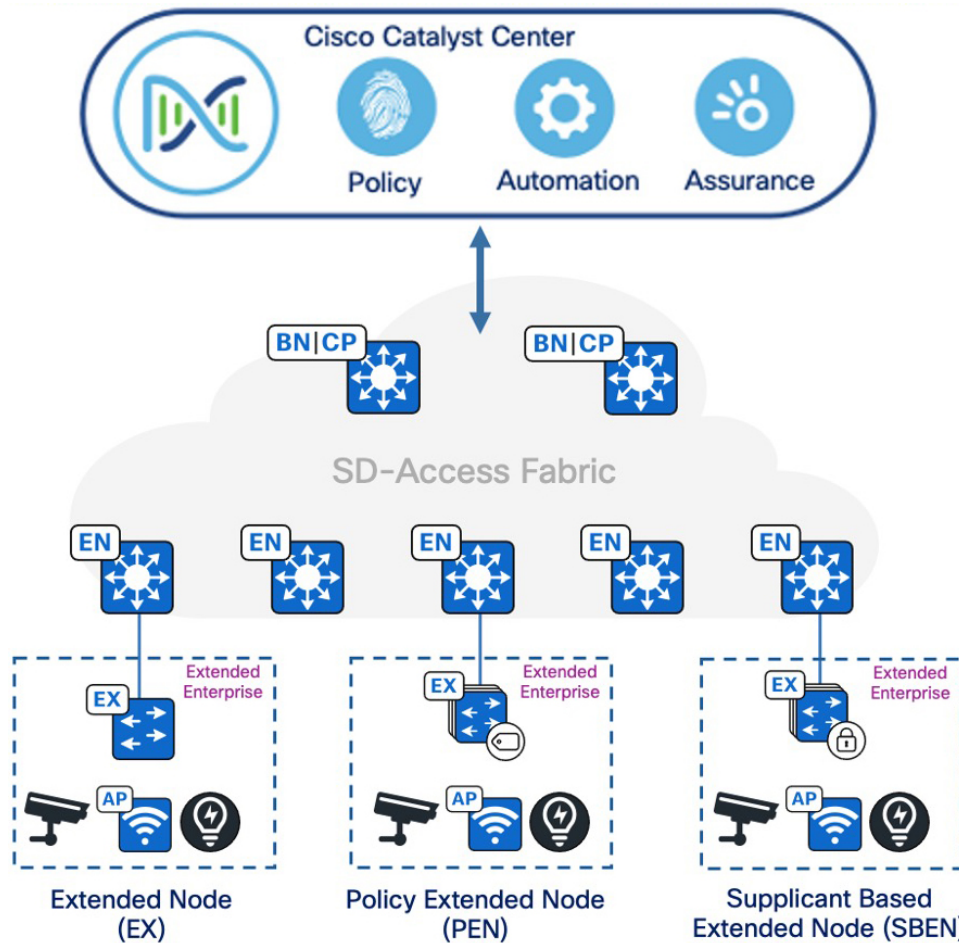
Supplicant-Based Extended Node is an extended node that undergoes a stricter onboarding process. It requires an IEEE 802.1X supplicant configuration and completes a full authentication and authorization process before being allowed into the SD-Access network. This approach enhances security by ensuring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

Key points to remember:

- 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.

The diagram below illustrates the extended enterprise network utilizing Extended Node (EX), Policy Extended Node (PEN), and Supplicant-Based Extended Node (SBEN).



For more details, see the "Extended Node Design" section of the [Cisco SD-Access Solution Design Guide](#).

For further details on Extended Nodes and Policy Extended Nodes, see the [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. Below is an overview of strategies for enhancing these areas:

- Quality of Service (QoS)
- Application Visibility
- Video Streaming across Sites

Quality of Service

Quality of Service (QoS) refers to a network's 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 the following core parameters:

- 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 perform optimally, enhancing overall productivity and user satisfaction.

You can install the following 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

An ITES organization needs to conduct regular employee training sessions across multiple branch offices located in different regions. These training sessions include live video broadcasts of presentations, demonstrations, and interactive Q&A sessions. To efficiently distribute the video content to all branches simultaneously without overloading the network, the organization utilizes multicast technology.

Such multicast data can be streamed from various sources, including regional data centers and corporate data centers. The Cisco SD-Access architecture provides the flexibility for end-to-end seamless multicast data traffic to flow from anywhere within the larger

enterprise network to any location globally. Cisco SD-Access supports both headend replication and native multicast modes, offering the flexibility to assign Multicast RP (Rendezvous Points) nodes either within the SD-Access fabric or externally.

SD-Access supports two different transport methods for forwarding multicast traffic: one uses the Overlay, referred to as Head-End Replication, and the other uses the Underlay, known 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 does not require the ingress fabric node to perform unicast replication. Instead, the entire underlay, including intermediate nodes, is used to handle the replication. To support Native Multicast, the first-hop routers (FHRs), last-hop routers (LHRs), and all network infrastructure components between them must be enabled.

By leveraging multicast technology within the SD-Access framework, the ITES organization can effectively conduct large-scale employee training sessions, enhancing communication and learning across all branch offices.

Network Deployment Options

The following sections describe deployment options for an ITES network.

Fabric Site Reference Models

In deployments with physical locations, you can use different templates for each of the different site types such as a large branch, a regional hub, headquarters, or small remote office. The underlying design challenge is to look at the existing network deployment and wiring, and propose a method to layer SD-Access fabric sites in these areas. This process can be simplified and streamlined by creating templates of the reference models. The templates help in understanding the common site designs by offering reference categories based on the multidimensional design elements along with the endpoint count to provide design guidelines for sites of similar size. The numbers are used as guidelines only and do not necessarily match the maximum specific scale and performance limits for devices within a reference design.

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

Common types of fabric site reference models include the following:

- Very Small Site (Fabric in a Box)
- Small Site
- Medium Site
- Large Site
- Extra Large Site



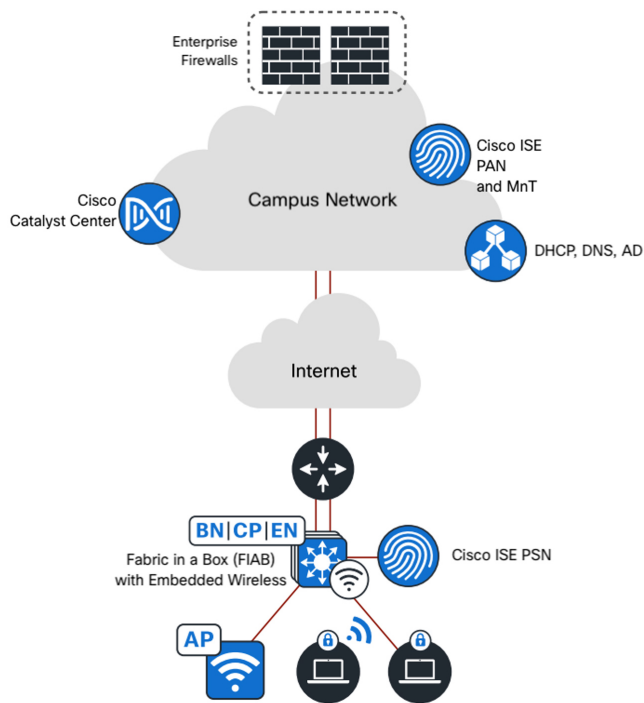
Note

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](#).

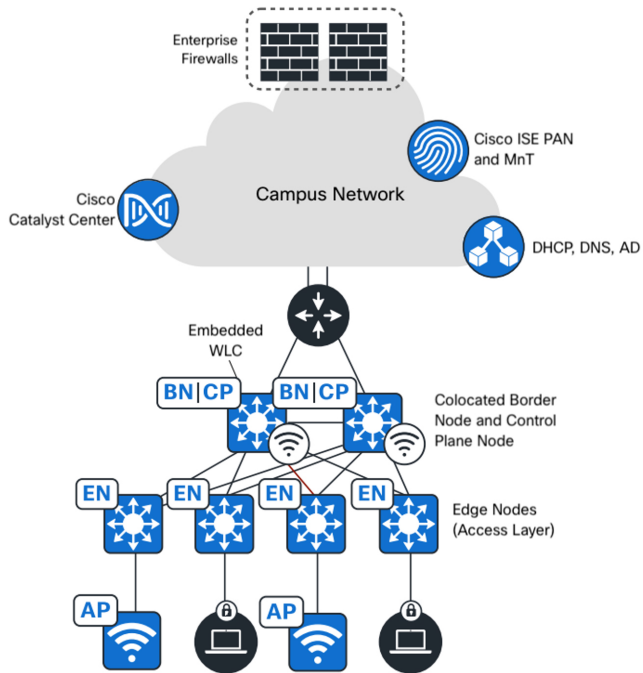
Fabric in a Box Deployment

The Fabric in a Box (FIAB) reference model covers a single fabric site, with resilience supported by switch stacking or StackWise Virtual to support up to 200 endpoints and 40 APs. For FIAB deployments, SD-Access embedded wireless is used to provide site-local wireless controller functionality. The site may have an ISE PSN (Policy Service Node) depending on the WAN/Internet circuit and latency.



Small Site Deployment

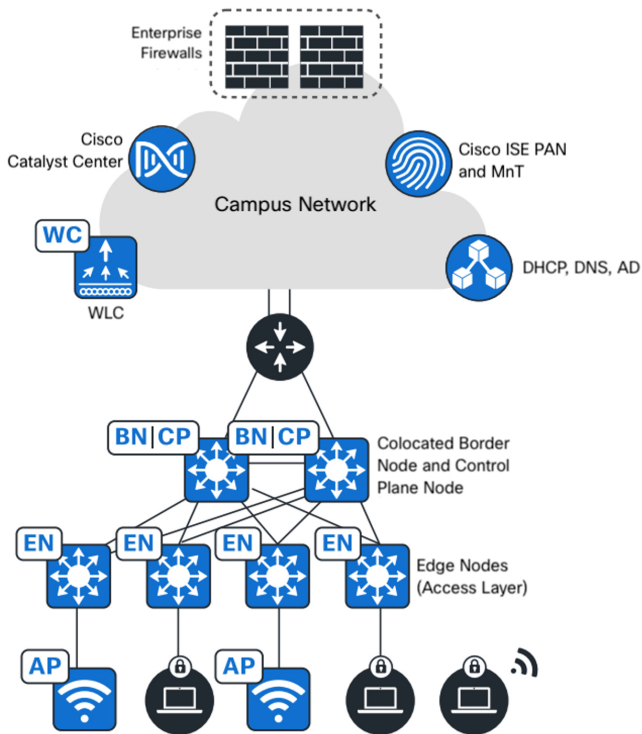
The small site reference model covers a single office or building with single wiring closets, usually up to 4,000 and up to 100 APs. The border node function is collocated with the control plane node function on one or two devices and usually uses embedded wireless with the option of hardware wireless controllers. The physical network is usually a two-tier collapsed core or distribution layer with an access layer servicing several wiring closets. Rather than collocating all roles in one device, the small site reference model provides added resiliency and redundancy along with a larger number of endpoints by separating the edge node roles onto dedicated devices in the access layer. The border and control plane nodes are collocated in the collapsed core layer. For SD-Access wireless, the embedded wireless controller is provisioned on one of the collocated border and control plane nodes. Optionally, a virtual or hardware-based wireless controller is used.



Medium Site Deployment

The medium site reference model usually covers a building with multiple wiring closets with a physical topology consisting of a two-tier collapsed core/distribution with an access layer.

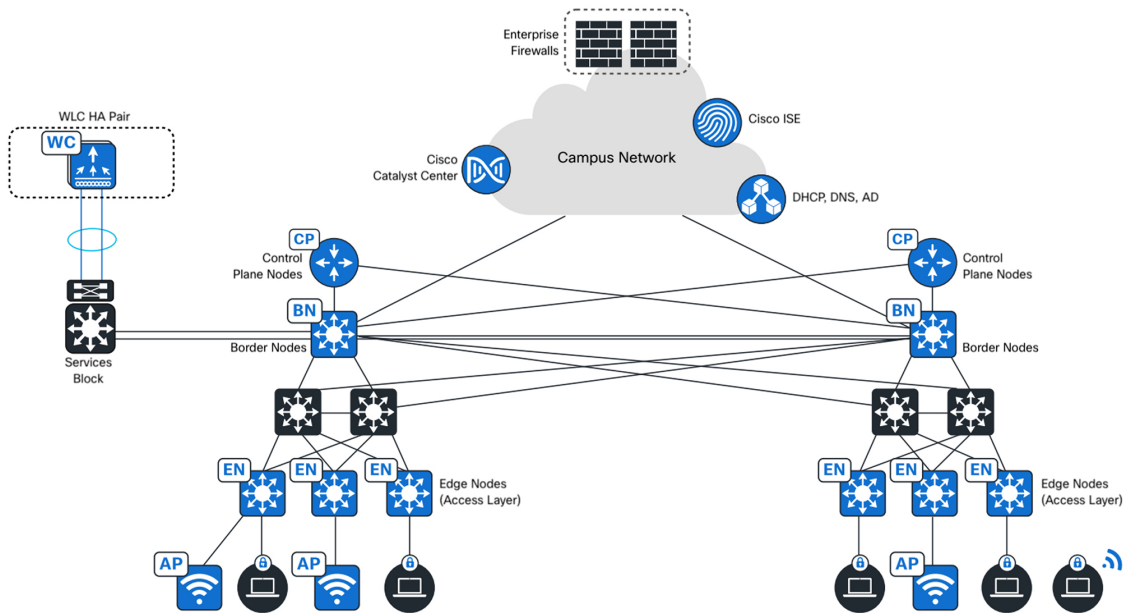
The medium site is designed to support fewer than 25,000 endpoints and fewer than 2000 APs. The border node function is colocated with the control plane node function on one or two devices or a highly resilient single device, and a separate wireless controller is ideally deployed in a HA configuration.



Large Site Deployment

The large site reference model covers multiple buildings or a building with multiple wiring closets. The physical network is usually three-tier with core, distribution, and access layers. It may even have a routed super core that aggregates many buildings and serves as the network egress point to the WAN and Internet. The border and control plane nodes functions are provisioned on separate devices rather than colocating.

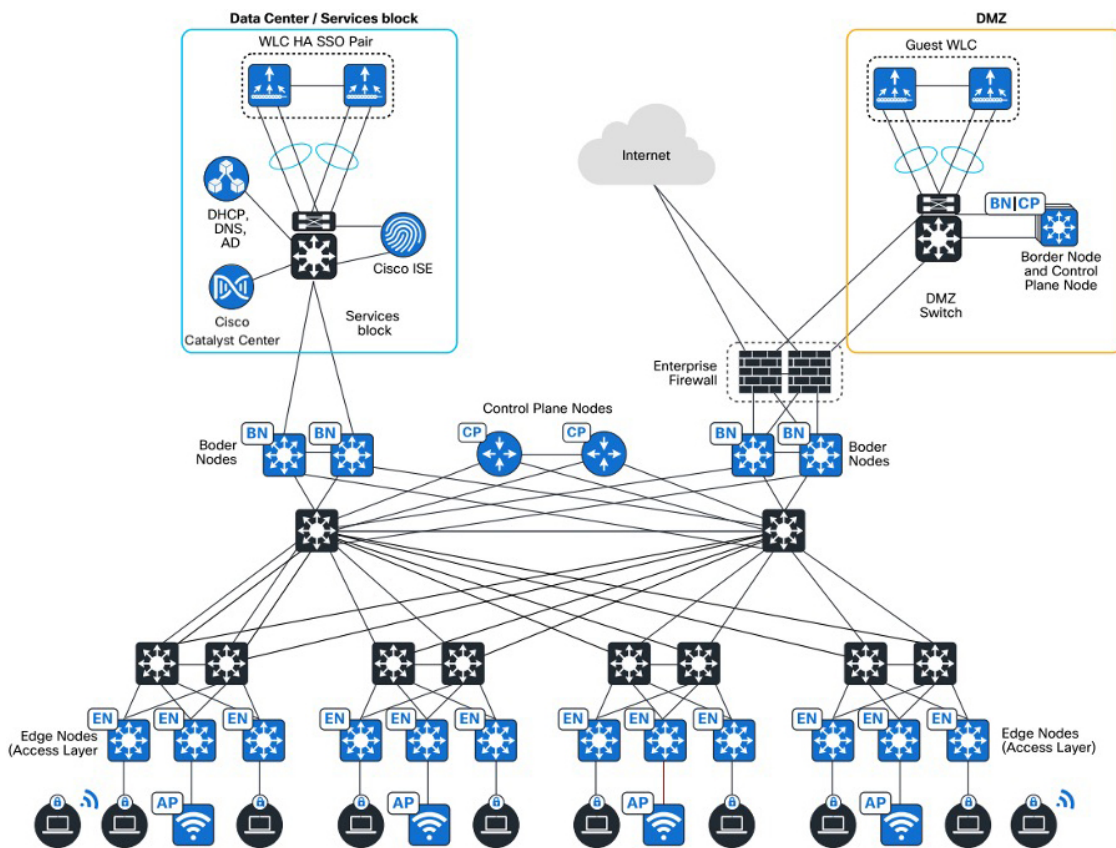
The large site supports up to 100,000 endpoints and 6000 APs. The border is distributed using redundant devices from the control plane function, and a separate wireless controller in an HA configuration.



Extra Large Site Deployment

The extra large site reference model covers a building with multiple wiring closets or multiple facilities spread across a large campus. The physical network is a three-tier network with core, distribution, and access layers, and may sometimes include a super core in a fourth tier. An extra large network requires dedicated service exit points, such as a dedicated data center, shared services block, and internet services.

The extra large site supports up to 200,000 endpoints and 10,000 APs, with multiple border nodes distributed from the control plane node on redundant devices and a separate wireless controller in an HA configuration.



Wireless Design

Designing wireless solutions within Cisco's SD-Access framework requires configuring and integrating multiple components to ensure smooth operation and management. ITES organizations utilizing an SD-Access fabric for their wired network have two choices for incorporating wireless access:

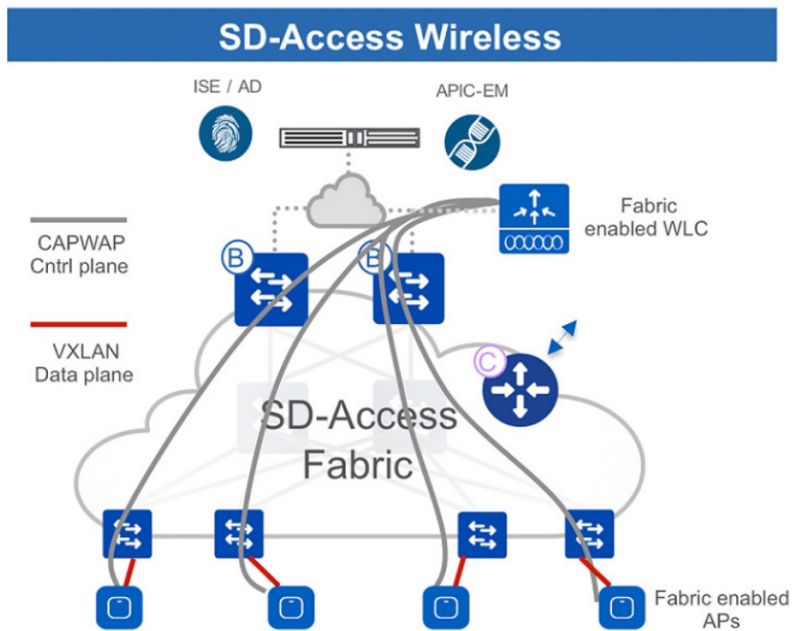
- SD-Access Wireless Architecture
- Cisco Unified Wireless Network Wireless Over-the-Top (OTT)

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: Networks can have a single subnet for both wired and wireless clients.
- Consistency of Policy: Wired policies are extended to wireless traffic, with both enforced at the edge node.
- Improved Performance: Wireless roams are Layer 2 and do not require any form of anchoring.
- Distributed Data Plane: Enables higher overall wireless throughput compared to centrally switched wireless architectures.

The diagram below depicts the 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:

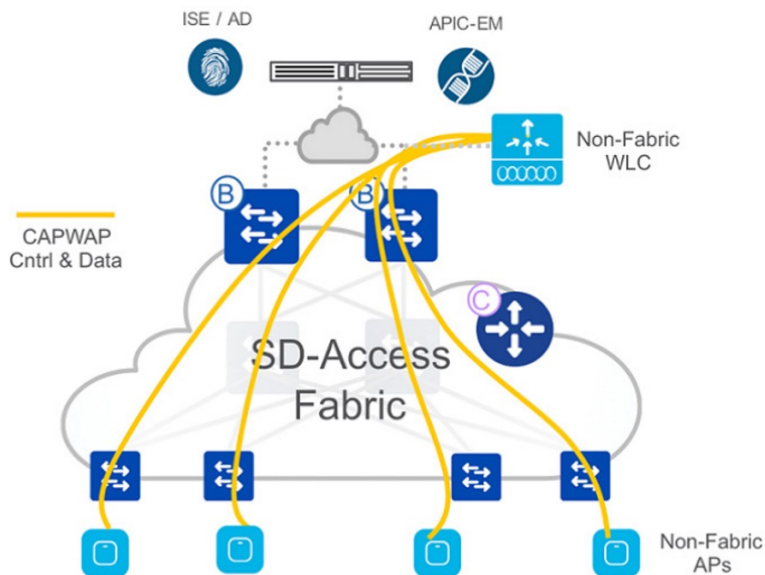
- Existing Cisco wireless controllers and APs that are not SD-Access wireless-capable.
- Presence of third-party wireless devices in the network.
- Asymmetric migration pace between wired and wireless networks.

In wireless OTT deployments, wireless control, management, and data plane traffic travel through the fabric in a Control and Provisioning of Wireless Access Points (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 tunneling protocols, the concept of using the SD-Access fabric as a transport remains the same.

Cisco SD-Access also supports Cisco FlexConnect or distributed wireless deployments. In these deployments, APs locally switch data traffic to the edge nodes to which they are connected.

The diagram below illustrates the flow of control plane and data plane traffic in wireless OTT:

CUWN wireless Over The Top (OTT)



For further information on wireless design and deployment in SD-Access, 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 Virtual Routing and Forwarding (VRF) and Scalable Group Tags (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 inter-site 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 micro-segmentation 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 further information on distributed site design and deployment, see [Cisco SD-Access Distributed Campus Deployment Guide](#) and [Cisco SD-Access Distributed Campus Design](#).

Multisite Remote Border

Multisite remote border (MSRB) centralizes the routing of untrusted traffic within the fabric network to a designated location such as a firewall or DMZ (demilitarized zone). For example, in a scenario where a guest virtual network 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 multisite remote border manages traffic to and from a specific virtual network extended across multiple sites. This configuration enables the deployment of a virtual network 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 virtual network, providing several advantages:

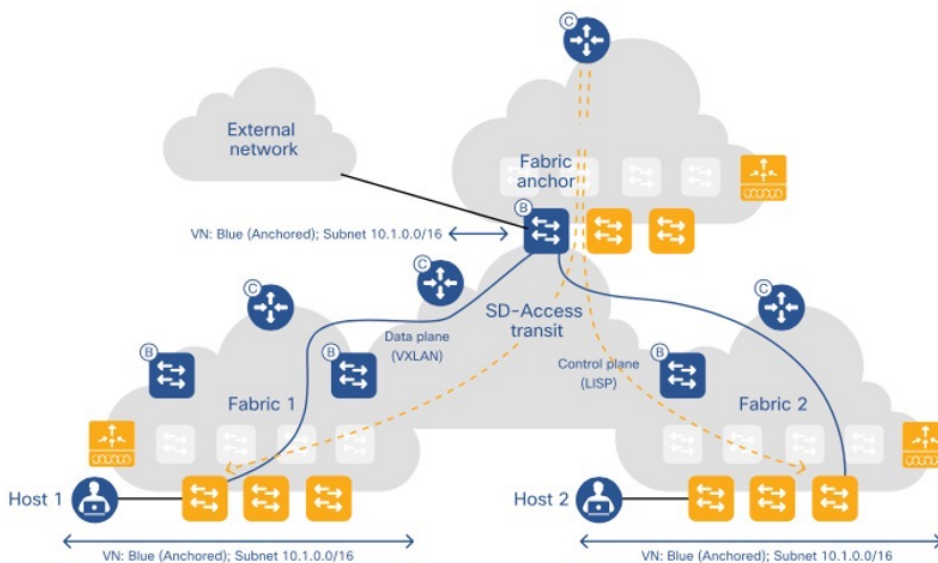
- Centralized control: You can designate a common border switch, called the anchor border, to handle all traffic for a particular virtual network across various sites. This simplifies management and policy enforcement.
- Subnet consistency: Multisite remote border enables you to use the same subnet for the virtual network across all sites. This eliminates the need to manage different subnets at each location, saving IP address space and simplifying configuration.
- Traffic isolation: Multisite remote border 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 a multisite remote border:

- Anchor virtual network: A virtual network 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 virtual network. anchor site handles the ingress and egress traffic for the anchor virtual network.
- Anchoring sites: Fabric sites other than the anchor site where the anchor virtual network is deployed.
- Anchor border node or multisite remote border: The fabric border node at the anchor site that provides the ingress and egress location for traffic to and from the anchor virtual network.
- 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 virtual network.

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

The diagram below illustrates a multisite remote border deployment.



For additional details regarding multisite remote border, see [LISP VXLAN Fabric Configuration Guide](#).



Note It is crucial to consider the maximum transmission unit (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 (PubSub) 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 PubSub model is a crucial component in modern, large-scale network designs.

LISP PubSub design eliminates the need for an additional protocol to register the LISP site registration table to control plane nodes in the fabric. LISP PubSub 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 PubSub 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 PubSub 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.

LISP PubSub design recommendations:

- LISP/BGP fabric sites and LISP PubSub fabric sites cannot co-exist with the same transit control plane nodes.
- Migration from one to another is not supported yet.
- LISP PubSub is recommended only for new network implementation.

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 the following aspects:

- Network: Maximum transmission unit (MTU), network topology, IP addressing for underlay and overlay, and location of shared services.
- Policy: Existing policy definition and enforcement points, virtual networks, and SGTs.
- Hardware platform: Switches, routers, wireless controllers, and access points 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](#).

Virtual Private Network for ODC Deployment

Establishing a secure Offshore Development Center (ODC) within an ITES organization necessitates the implementation of two Virtual Private Network (VPN) configurations:

- Site-to-Site VPN
- Client-to-Site VPN

Site-to-Site VPN

Site-to-Site VPN is a type of VPN setup that creates a secure connection between two or more geographically separate offices or physical locations of the same organization. This connection, often referred to as a VPN tunnel, allows for the secure transmission of data and resources across a public network, such as the internet, as if the sites were connected by a private network.

The Site-to-Site VPN is typically established between network gateways, such as routers or firewalls, that are configured to encrypt and decrypt traffic entering and exiting the sites. This setup is commonly used by businesses to ensure that communication between their offices is secure and cannot be intercepted by unauthorized parties.

Client-to-Site VPN

Client-to-Site VPN, also known as a remote access VPN, is a type of VPN that allows individual users to connect to a remote network, such as their workplace's network, from any location. This is achieved through VPN client software installed on the user's device, which establishes a secure connection to the VPN server located at the remote site.

Once connected, the user's device becomes a part of the remote network, allowing access to network resources as if the user were physically present within the network's local environment. This type of VPN is commonly used by remote workers, travellers, or anyone needing secure access to their organization's network from an external location.

Both Site-to-Site and Client-to-Site VPNs use encryption and other security mechanisms to ensure that data transmitted over the public internet remains confidential and protected from unauthorized access.

Validated Solution Use Cases

The following sections describe some of the important use cases validated for ITES that serve as trusted templates, empowering organizations to build their IT infrastructure, ensuring that these designs have undergone thorough testing and are customized to meet their specific business needs.

Day-0 and Day-1 Network Bring Up Use Cases

- Bring up network infrastructure and integrate all features for greenfield campus.
- 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 access points using Catalyst Center.
- Onboard devices via plug and play for network devices and access points.
- 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 new sites using Fabric in a Box (FiaB) with an embedded wireless controller.
- Add small new fabric sites with another Flex Over-The-Top (OTT) deployment.
- Perform day-*n* credential changes such as, device password changes and network device updates.
- Allow VLAN on the Layer 2 and Layer 3 handoff link using a template.



Note

The template must be reprovisioned after any operation that triggers the "switchport mode trunk" and "switchport trunk allowed vlan all" configuration for the uplink port.

Segmentation and Policy Use Cases

- Implement virtual networks across the organization to achieve consistent macrosegmentation.
- Use group-based access policy for microsegmentation within a virtual network using SGTs.
- Ensure secure onboarding of wired and wireless clients using authentication.

- Enforce policy at different entry and exit traffic points in the fabric.
- Operational scenarios that include the addition of new segments and group-based access policies for devices and users.

Security Use Cases

- Apply and use trusted CA FQDN-based certificates with Catalyst Center.
- Create granular role-based users with external AAA authentication and use audit logging to check Catalyst Center activities.
- Monitor audit policy changes, deployment of policy changes, and status of these deployments.

SD-Access Wireless Use Cases

- Onboard wireless clients using enterprise SSID for branch employees and users.
- Provide guest wireless access for users at branches using a Central Web Authentication (CWA) portal.
- Deploy wireless infrastructure with Catalyst 9800 Series Wireless Controller in HA mode and an embedded wireless LAN controller.
- Enforce guest SSID traffic policies on the firewall using multisite remote border.

Robustness Use Cases

- Recovery from device or link failure automatically with minimal impact on existing applications, traffic, and users.
- Catalyst Center in three-node HA mode. In case of services or node failure in Catalyst Center, the system should recover without user intervention.
- Cisco ISE distributed nodes failover with PSN, pxGrid node failover.
- Cisco Catalyst wireless LAN controller and access points failover.
- Failover scenarios with link and network device failure within the fabric.
- Back up Catalyst Center controller configuration and data either, one time or on schedule.
- Restore backup on a new Cisco Catalyst cluster and verify the ability to manage the devices.
- Longevity with churn in the network, policy, and device connectivity.

Monitoring and Troubleshooting with Assurance and Analytics Use Cases

- Monitor the state of the network, wired users, and wireless users from a unified interface.
- Monitor severe, critical, and other ongoing issues with the network and devices and follow the suggested actions in Assurance to resolve the issues.
- Obtain a comprehensive view of individual devices, wired user, or wireless user and retrieve detailed information.
- Track detailed application data usage by users using application visibility.

Performance and Scale Use Cases

- Multi-dimensional scale configuration with all solutions integrated and check for stability of the network.

ITES ODC Deployment Models

Site-to-Site ODC with Dedicated Firewall and Gateway Outside the Fabric

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure that all traffic transits through the firewall and is logged for auditing and compliance.
- Ensure secure data transfer between networks.
- Restrict ODC users' access to systems within their respective ODC.
- Implement a dedicated firewall for the ODC.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Implement Layer 2 virtual networks with a gateway outside of the fabric.
- Configure a site-to-site VPN for secure data transfer.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to manage traffic from the dedicated ODC.

Procedure

Step 1 Using an appropriate VLAN ID, configure a Layer 2 VN and then assign it to a site.

Configuration Attributes

Provide a name for each Layer 2 Virtual Network and define its attributes.

LAYER 2 VIRTUAL NETWORK

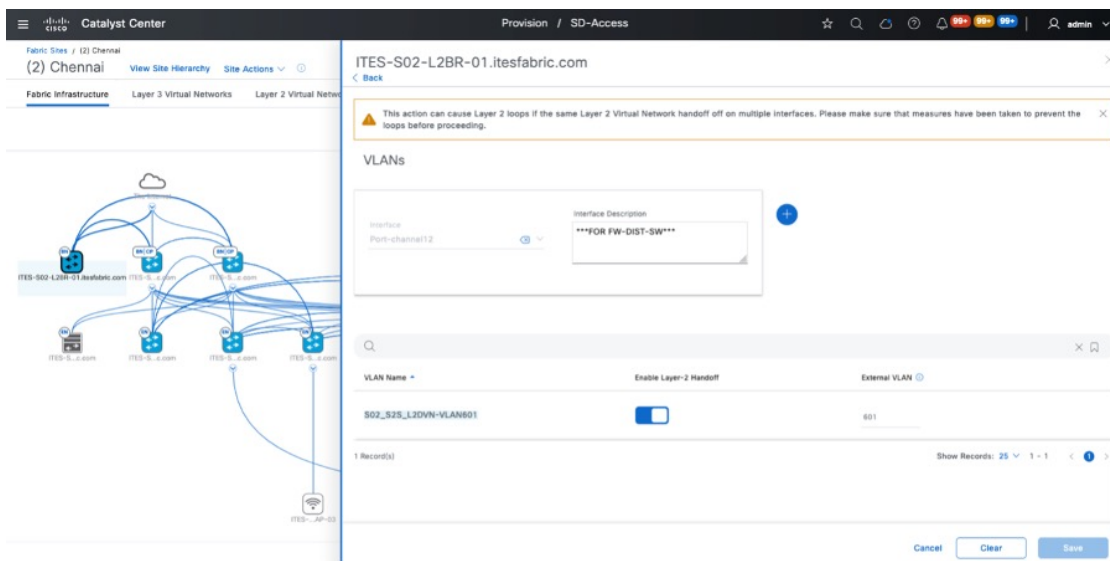
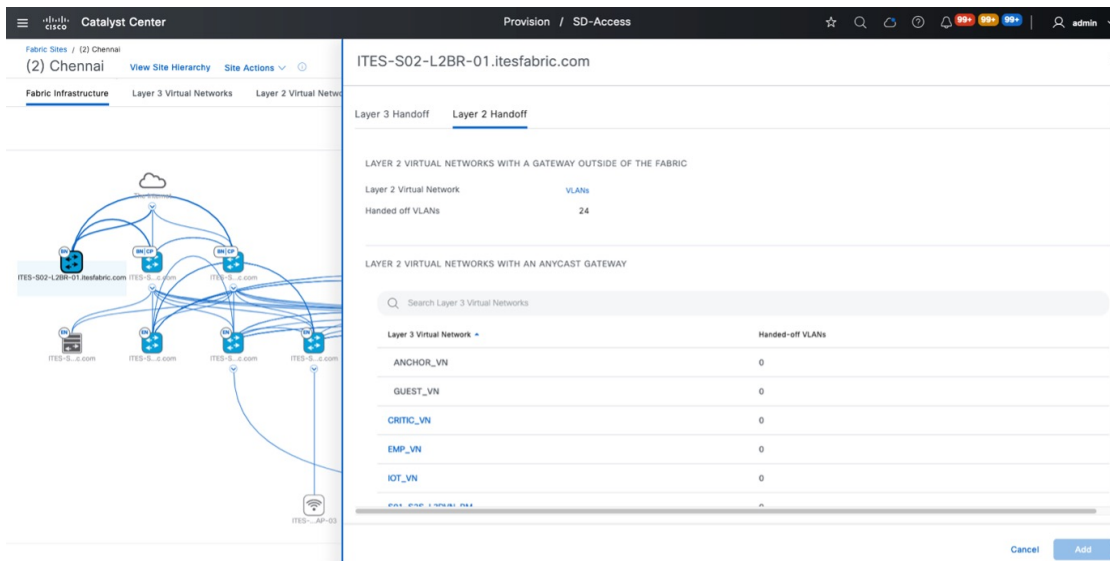
VLAN Name: S03_S2S_L2DVN-VLAN601 VLAN ID: 601 Traffic Type: Data Voice

Fabric-Enabled Wireless Layer 2 Flooding

Advanced Attributes

Step 2 Configure Layer 2 Handoff on the dedicated border node.

Enable the VLANs that you can access from the **Layer 2 Virtual Networks with a Gateway Outside of the Fabric** area.



Note On the Layer 2 border node facing the firewall, the uplink port is configured as a trunk port with either a standalone or port-channel interface.

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability
- Site-to-Site VPN

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

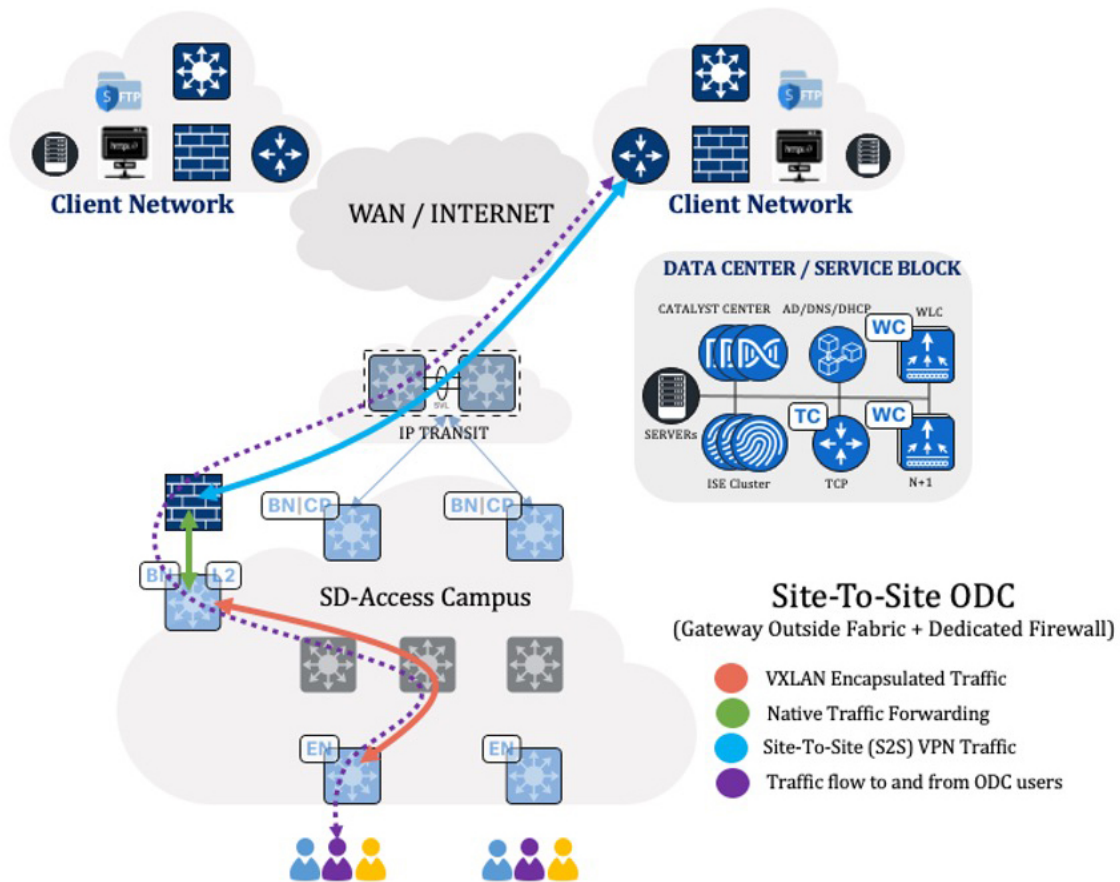
Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish an SXP session between Cisco ISE and the Layer 2 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Site-to-Site ODC with Shared Firewall and Gateway Outside the Fabric

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure that all traffic transits through the firewall and is logged for auditing and compliance.

- Ensure secure data transfer between networks.
- Restrict ODC users' access to systems within their respective ODC.
- Use the same firewall for multiple ODCs.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Implement Layer 2 virtual networks with a gateway outside of the fabric.
- Configure a site-to-site VPN for secure data transfer.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to handle traffic from each ODC separately.

Procedure

Step 1 Using an appropriate VLAN ID, configure a Layer 2 VN and then assign it to a site.

The screenshot shows the 'Create Layer 2 Virtual Networks' page in Cisco Catalyst Center. The page title is 'Configuration Attributes'. Below the title, there is a instruction: 'Provide a name for each Layer 2 Virtual Network and define its attributes.' The main content area contains three configuration cards for Layer 2 Virtual Networks.

The first card is titled 'LAYER 2 VIRTUAL NETWORK' and has a dropdown arrow on the right. It contains the following fields and options:

- VLAN Name: S01_S2S_L3SVN-VLAN501
- VLAN ID: 501
- Traffic Type: Data Voice
- Fabric-Enabled Wireless
- Layer 2 Flooding ⓘ
- Advanced Attributes ⓘ

The second card is identical to the first, with the following fields and options:

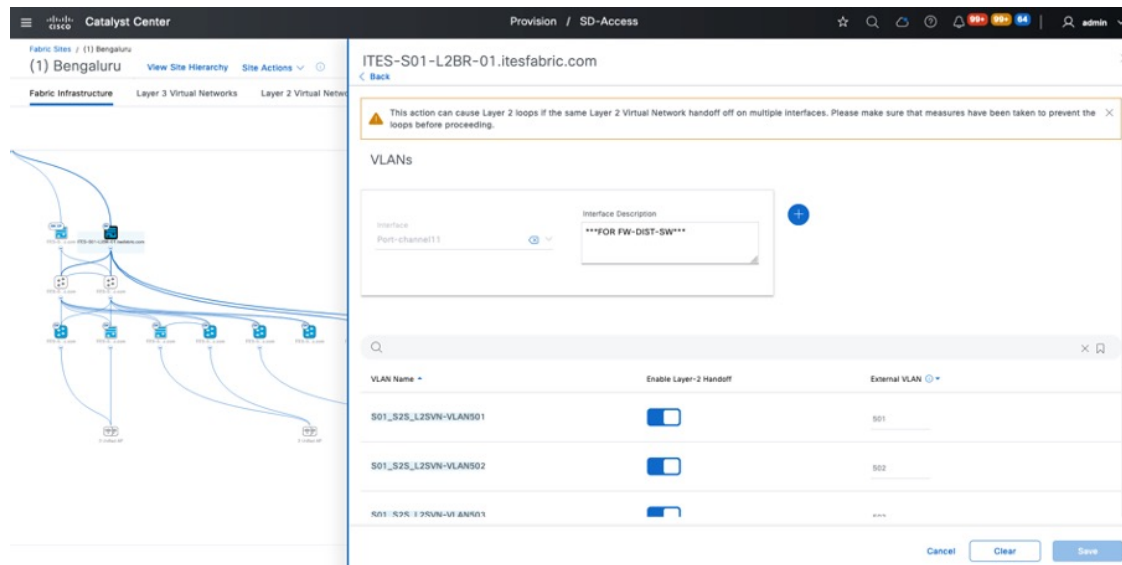
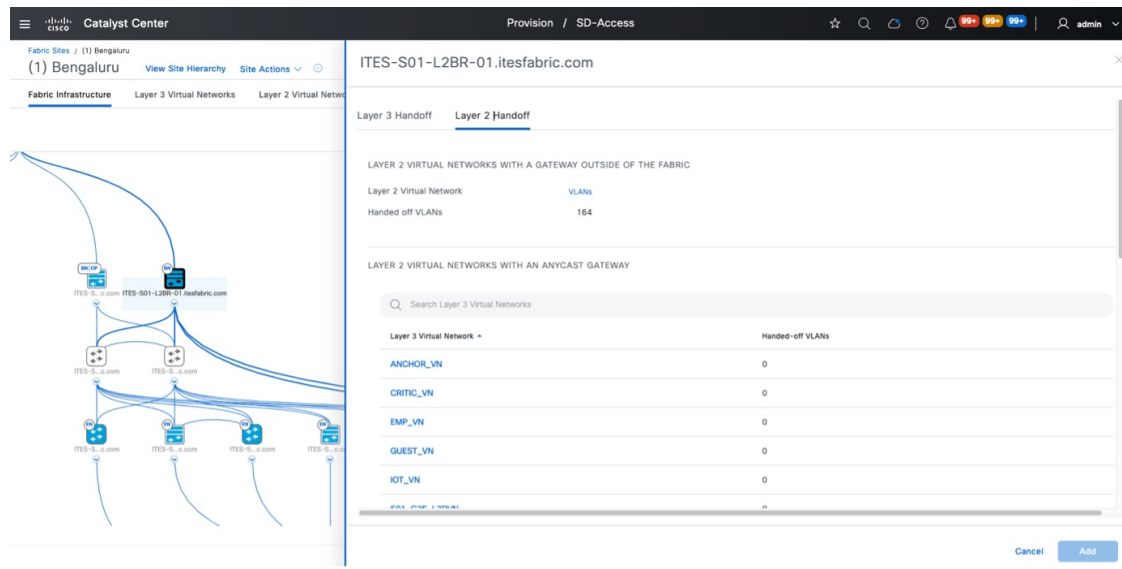
- VLAN Name: S01_S2S_L3SVN-VLAN502
- VLAN ID: 502
- Traffic Type: Data Voice
- Fabric-Enabled Wireless
- Layer 2 Flooding ⓘ
- Advanced Attributes ⓘ

The third card is partially visible and shows the following fields and options:

- VLAN Name: _____
- Traffic Type: _____

Step 2 Configure Layer 2 Handoff on the dedicated border node.

Enable the VLANs that you can access from the **Layer 2 Virtual Networks with a Gateway Outside of the Fabric** area.



Note On the Layer 2 border node facing the firewall, the uplink port is configured as a trunk port with either a standalone or port-channel interface.

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability
- Site-to-Site VPN

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

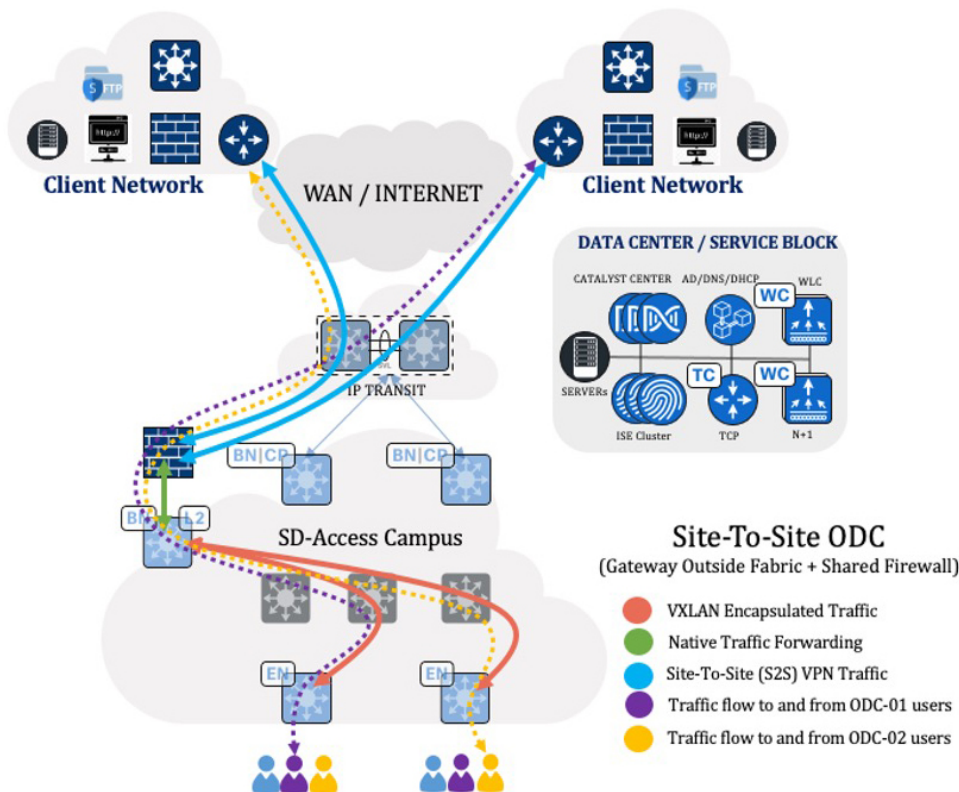
Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish an SXP session between Cisco ISE and the Layer 2 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Client-to-Site ODC With Dedicated Firewall and Gateway Outside Fabric

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Set up a firewall to act as the client's gateway and log all traffic for auditing and compliance.
- Use a VPN client to securely connect to the client network.

- Restrict ODC users' access to systems within their respective ODC.
- Implement a dedicated firewall for the ODC.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Implement Layer 2 VNs with a gateway outside of the fabric.
- Install and configure a VPN client on ODC users' laptops and desktops.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall to manage traffic from the dedicated ODC.

Procedure

Step 1 Using an appropriate VLAN ID, configure a Layer 2 VN and then assign it to a site.

The screenshot shows the Cisco Catalyst Center interface for configuring a Layer 2 Virtual Network. The page title is "Configuration Attributes" and the subtitle is "Provide a name for each Layer 2 Virtual Network and define its attributes." The main configuration area is titled "LAYER 2 VIRTUAL NETWORK" and contains the following fields and options:

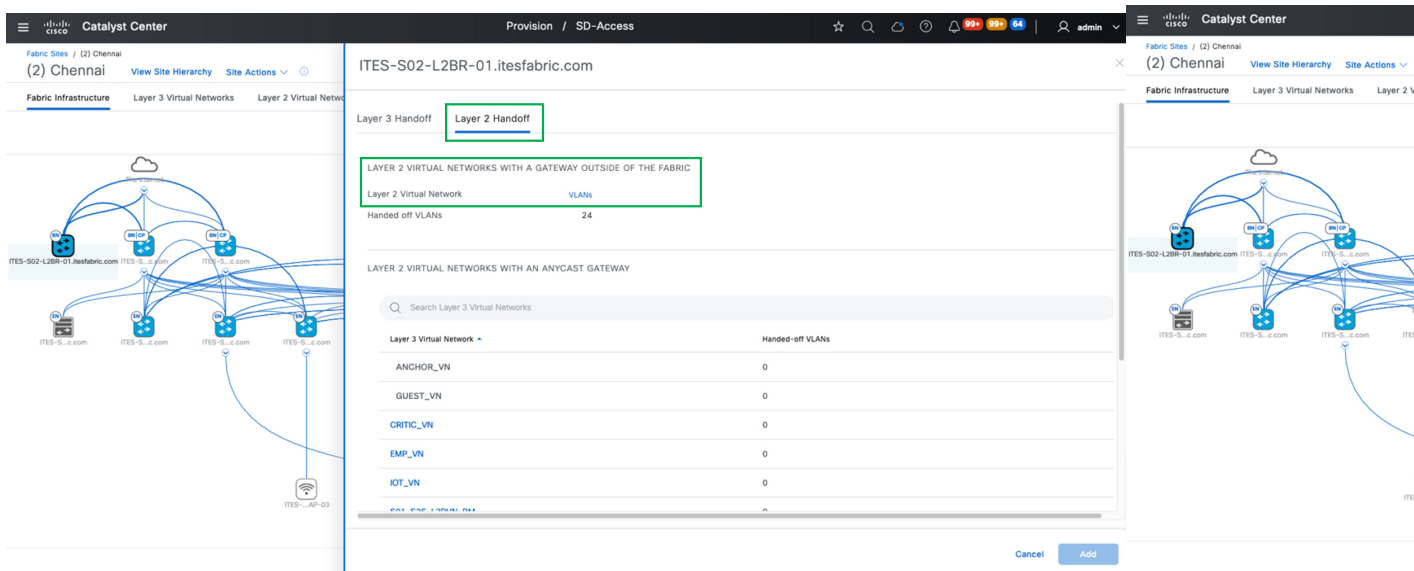
VLAN Name	VLAN ID	Traffic Type
S02_C2S_L2DVN-VLAN401	401	<input checked="" type="radio"/> Data <input type="radio"/> Voice

Below the table, there are several checkboxes and a toggle:

- Fabric-Enabled Wireless
- Layer 2 Flooding ⓘ
- Advanced Attributes ⓘ

Step 2 Configure Layer 2 Handoff on the dedicated border node.

Enable the VLANs that you can access from the **Layer 2 Virtual Networks with a Gateway Outside of the Fabric** area.



Note On the Layer 2 border node facing the firewall, the uplink port is configured as a trunk port with either a standalone or port-channel interface.

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

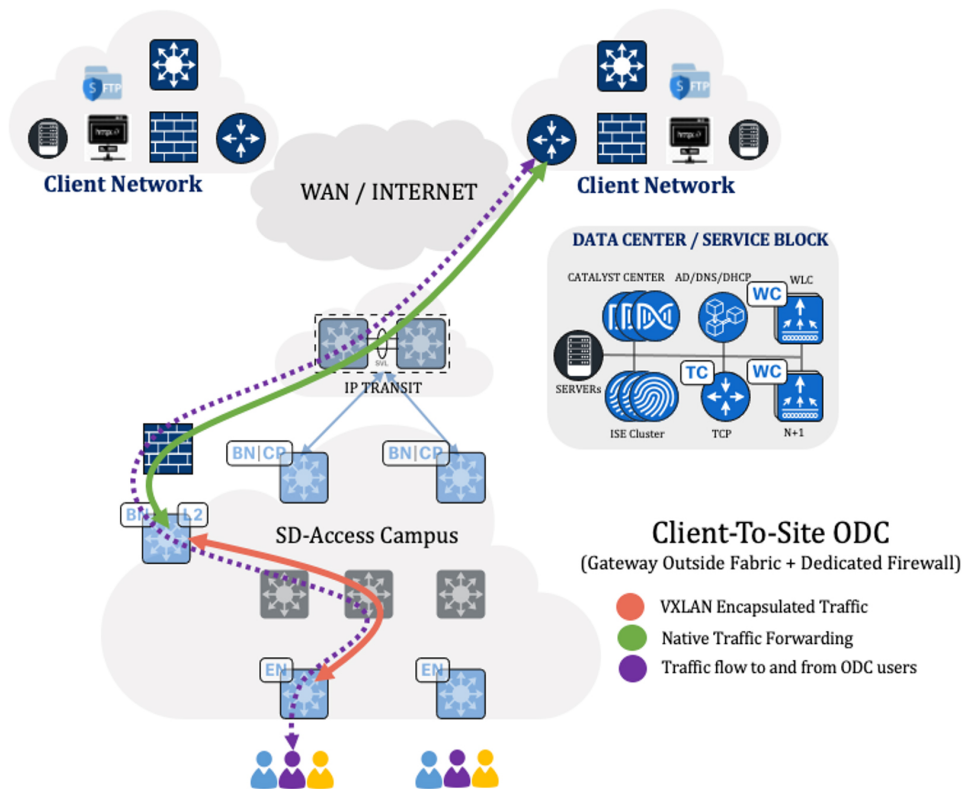
Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish an SXP session between Cisco ISE and the Layer 2 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Client-to-Site ODC With Shared Firewall and Gateway Outside Fabric

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Set up a firewall to act as the client's gateway and log all traffic for auditing and compliance.
- Use a VPN client to securely connect to the client network.
- Restrict ODC users' access to systems within their respective ODC.
- Use the same firewall for multiple ODCs.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Implement Layer 2 VNs with a gateway outside of the fabric.
- Install and configure a VPN client on ODC users' laptops and desktops.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to handle traffic from each ODC separately.

Procedure

Step 1 Using an appropriate VLAN ID, configure a Layer 2 VN and then assign it to a site.

The screenshot shows the Catalyst Center interface for configuring Layer 2 Virtual Networks. The page title is "Configuration Attributes" and the subtitle is "Provide a name for each Layer 2 Virtual Network and define its attributes." Two configuration panels are visible, each for a different Layer 2 Virtual Network.

Layer 2 Virtual Network 1:

- VLAN Name: S01_C2S_L2SVN-VLAN301
- VLAN ID: 301
- Traffic Type: Data Voice
- Fabric-Enabled Wireless
- Layer 2 Flooding
- Advanced Attributes

Layer 2 Virtual Network 2:

- VLAN Name: S01_C2S_L2SVN-VLAN302
- VLAN ID: 302
- Traffic Type: Data Voice
- Fabric-Enabled Wireless
- Layer 2 Flooding
- Advanced Attributes

Step 2 Configure Layer 2 Handoff on the dedicated border node.

Enable the VLANs that you can access from the **Layer 2 Virtual Networks with a Gateway Outside of the Fabric** area.

The screenshot shows the Catalyst Center interface for configuring Layer 2 Handoff on a dedicated border node. The page title is "Provision / SD-Access" and the subtitle is "ITES-S01-L2BR-01.itesfabric.com". The "Layer 2 Handoff" tab is selected, and the "Layer 2 Handoff" configuration area is highlighted.

Layer 2 Handoff Configuration:

- Layer 2 Handoff:
- Layer 2 Virtual Networks with a Gateway Outside of the Fabric:
- Layer 2 Virtual Network: S01_C2S_L2SVN-VLAN301
- Handed off VLANs: 164
- Layer 2 Virtual Networks with an Anycast Gateway:

The "Layer 2 Virtual Networks with an Anycast Gateway" section is also visible, showing a table of Layer 3 Virtual Networks and their associated Handed-off VLANs.

Layer 3 Virtual Network	Handed-off VLANs
ANCHOR_VN	0
CRITIC_VN	0
EMP_VN	0
GUEST_VN	0
IOT_VN	0
RES_C2S_L2SVN	0

Note On the Layer 2 border node facing the firewall, the uplink port is configured as a trunk port with either a standalone or port-channel interface.

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

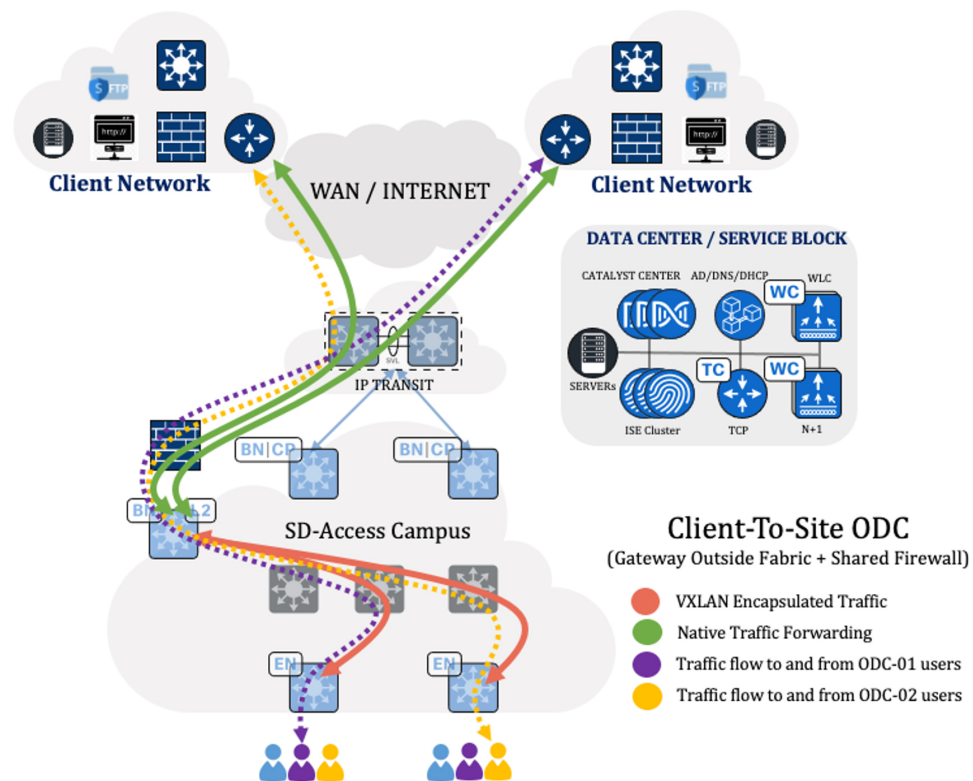
Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish an SXP session between Cisco ISE and the Layer 2 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Site-to-Site ODC With Dedicated Firewall and Anycast Gateway

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure that all traffic transits through the firewall and is logged for auditing and compliance.
- Ensure the secure transfer of data between networks.
- Restrict ODC users' access to systems within their respective ODC.
- Implement a dedicated firewall for the ODC.

Technical requirements:

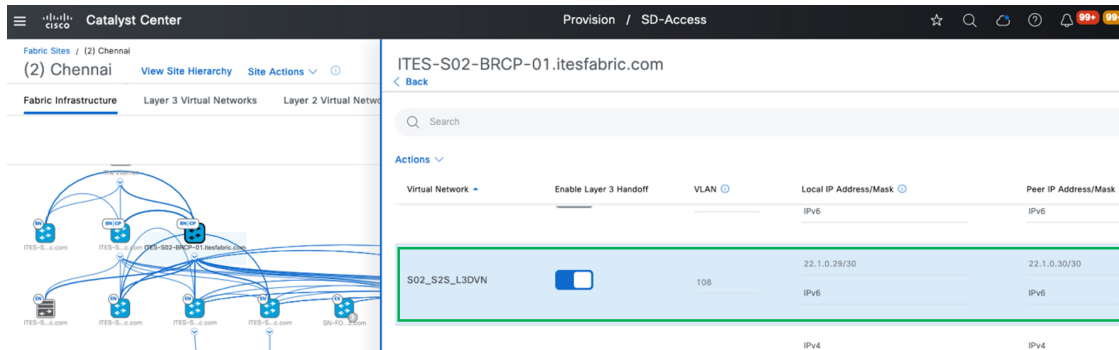
- Provision the Cisco SD-Access network using Catalyst Center.
- Redirect all traffic from the fusion device to the firewall and then forward it to the destination.
- Configure a site-to-site VPN for secure data transfer.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to manage traffic from the dedicated ODC.

Procedure

Step 1 Configure a Layer 3 VN at the site and then provision the required anycast gateway.

The screenshot shows the Catalyst Center interface for configuring a Layer 3 Virtual Network. The main heading is "Fabric Sites and Fabric Zones (Optional)". Below this, a text block explains that a Layer 3 Virtual Network can be assigned to multiple Fabric Sites and Fabric Zones, or created without assigning it to any. A diagram illustrates the configuration flow: a Layer 3 Virtual Network named "S02_S2S_L3DVN" is linked to a Fabric Site named ".../Tamil Nadu(2) Chennai" (with a "Select Fabric Sites" link below it), which in turn is linked to Fabric Zones, currently showing "No Fabric Zones Created". On the right side, a "Configuration Attributes" panel is visible, containing a search bar and a list of "LAYER 3 VIRTUAL NETWORKS" with the entry ".../Site-...u/(2) Chennai" and "S02_S2S_L3DVN" highlighted with a green box and a green checkmark.

Step 2 Configure Layer 3 handoff on the border nodes for the newly added VN, then ensure that the eBGP sessions are up.



Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability
- Site-to-site VPN

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

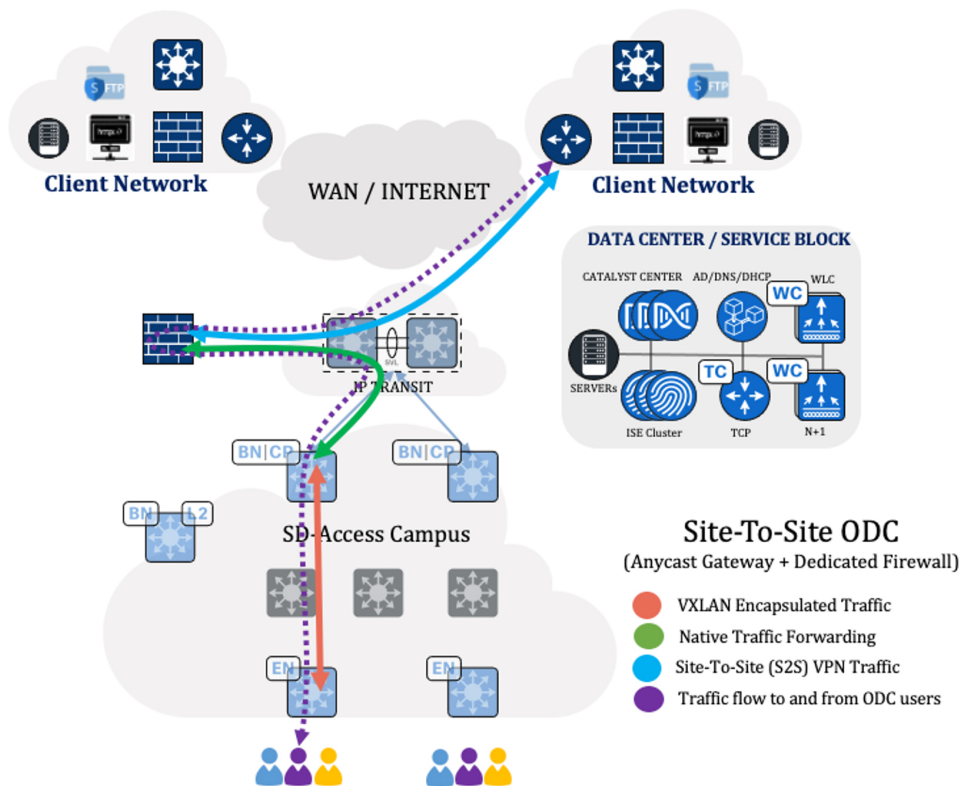
- Establish a VRF-aware SXP session between Cisco ISE and the VN's Layer 3 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Note Configure the VN traffic that traverses the firewall to allow TCP bypass specifically for SXP communications between the Cisco ISE server and Layer 3 border nodes.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Site-to-Site ODC With Shared Firewall and Anycast Gateway

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure that all traffic transits through the firewall and is logged for auditing and compliance.
- Ensure the secure transfer of data between networks.
- Restrict ODC users' access to systems within their respective ODC.
- Use the same firewall for multiple ODCs.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Redirect all traffic from the fusion device to the firewall and then forward it to the destination.
- Configure a site-to-site VPN for secure data transfer.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to handle traffic from each ODC individually.

Procedure

Step 1 Configure a Layer 3 VN at the site and then provision the required anycast gateways.

The screenshot shows the Cisco Catalyst Center interface for creating Layer 3 Virtual Networks. The main content area displays the title "Fabric Sites and Fabric Zones (Optional)" and a descriptive paragraph: "A Layer 3 Virtual Network can be assigned to multiple Fabric Sites and Fabric Zones. They can also be assigned to parent Fabric Sites without being assigned to a Fabric Zone within the Site. A Layer 3 Virtual Network can also be created without assigning it to a Fabric Site or Fabric Zone." Below this text is a diagram showing the relationship between a Layer 3 Virtual Network (S01_S2S_L3SVN), Fabric Sites (Karnataka/(1) Bengaluru), and Fabric Zones (No Fabric Zones Created). A right-hand sidebar titled "Configuration Attributes" shows a list of Layer 3 Virtual Networks, with S01_S2S_L3SVN highlighted.

Step 2 Configure Layer 3 handoff on the border nodes for the newly added VN, then ensure that the eBGP sessions are up.

The screenshot shows the Cisco Catalyst Center interface for provisioning SD-Access. The main content area displays the configuration for the virtual network S02_S2S_L3DVN. The "Enable Layer 3 Handoff" checkbox is checked. The "Local IP Address/Mask" is 22.1.0.29/30 and the "Peer IP Address/Mask" is 22.1.0.30/30. A network diagram on the left shows the physical infrastructure with various nodes and connections.

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability
- Site-to-site VPN

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish a VRF-aware SXP session between Cisco ISE and the VN's Layer 3 border node to obtain the destination addresses' security tags.

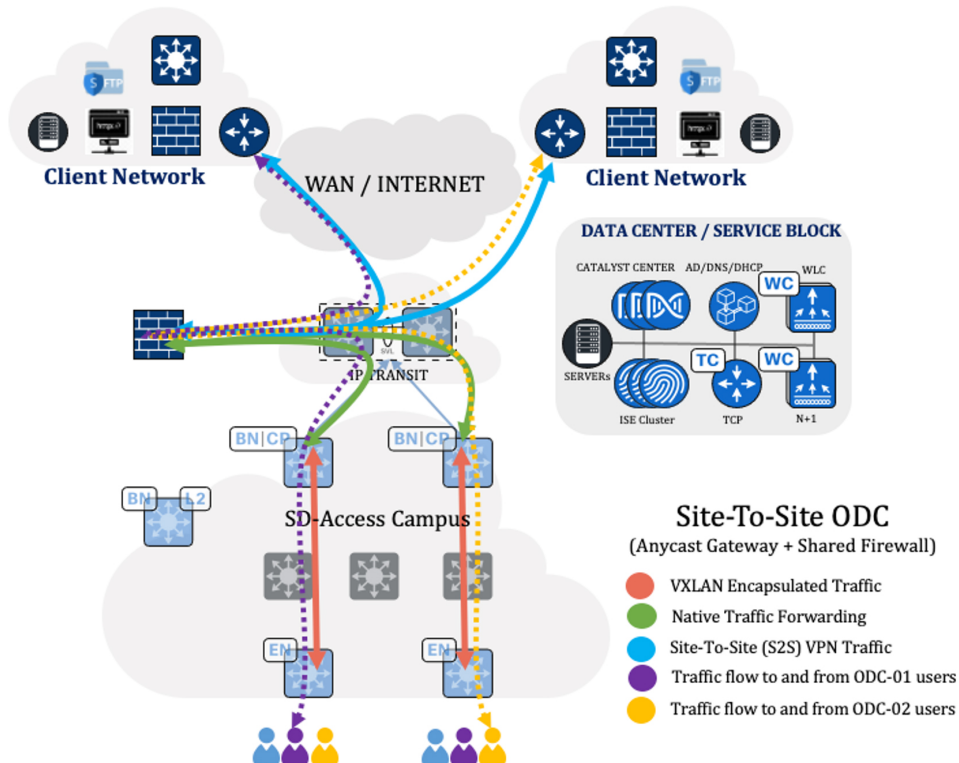
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Note Configure the VN traffic that traverses the firewall to allow TCP bypass specifically for SXP communications between the Cisco ISE server and Layer 3 border nodes.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Client-to-Site ODC With Dedicated Firewall and Anycast Gateway

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure that all traffic transits through the firewall and is logged for auditing and compliance.
- Use a VPN client to securely connect to the client network.
- Restrict ODC users' access to systems within their respective ODC.
- Implement a dedicated firewall for the ODC.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Redirect all traffic from the fusion device to the firewall and then forward it to the destination.
- Install and configure a VPN client on ODC users' laptops and desktops.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to manage traffic from the dedicated ODC.

Procedure

Step 1 Configure a Layer 3 VN at the site and then provision the required anycast gateway.

The screenshot shows the Catalyst Center interface for creating a Layer 3 Virtual Network. The main heading is "Fabric Sites and Fabric Zones (Optional)". Below it, a text block explains that a Layer 3 Virtual Network can be assigned to multiple Fabric Sites and Fabric Zones, or created without assigning it to a Fabric Site or Fabric Zone. A diagram shows the relationship: "Layer 3 Virtual Network S02_C2S_L3DVN" is linked to "Fabric Sites .../Tamil Nadu(2) Chennai" (with a "Select Fabric Sites" link), which is then linked to "Fabric Zones No Fabric Zones Created". On the right, a "Configuration Attributes" panel shows a search bar and a list of Layer 3 Virtual Networks, with "S02_C2S_L3DVN" highlighted and marked with a green checkmark.

Step 2 Configure the Layer 3 handoff on the border nodes for the newly added VN, then ensure that the eBGP sessions are up.

The screenshot shows the Catalyst Center interface for provisioning SD-Access. The main heading is "Provision / SD-Access". The left sidebar shows "Fabric Sites / (2) Chennai" and "(2) Chennai" with "View Site Hierarchy" and "Site Actions" options. The main content area shows "Fabric Infrastructure" and "Layer 3 Virtual Networks". A network diagram is visible on the left. The right panel shows the configuration for "ITES-S02-BRCP-01.itesfabric.com". A table lists the configuration for the virtual network "S02_C2S_L3DVN":

Virtual Network	Enable Layer 3 Handoff	VLAN	Local IP Address/Mask	Peer IP Address/Mask
S02_C2S_L3DVN	<input checked="" type="checkbox"/>	106	22.1.0.21/30 IPv6	22.1.0.22/30 IPv6

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

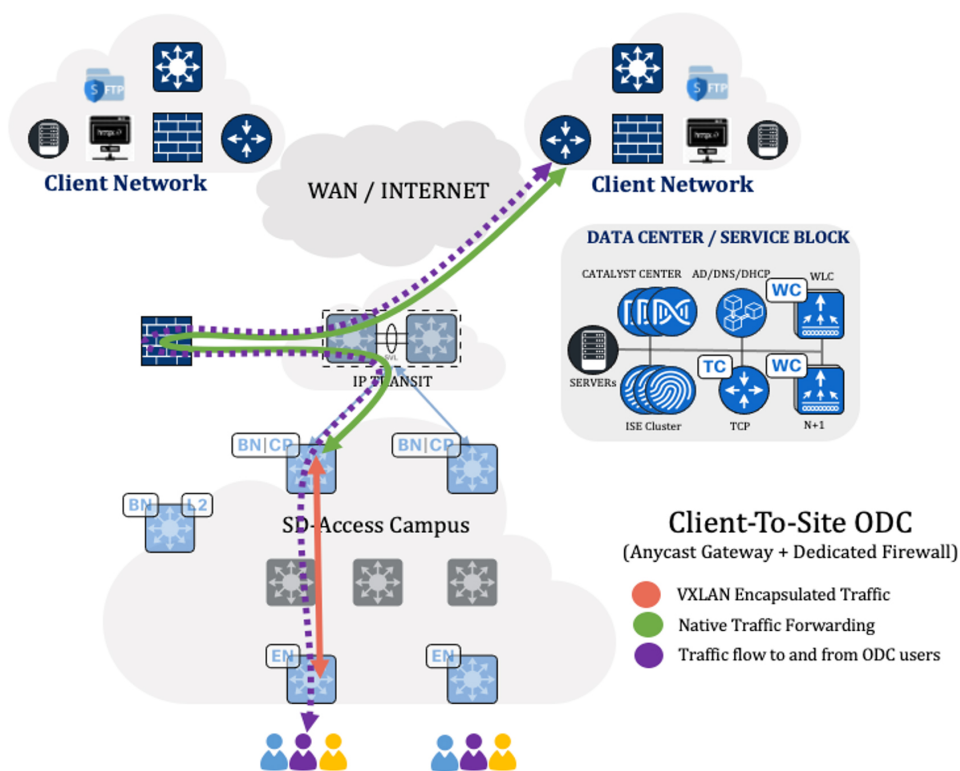
- Establish a VRF-aware SXP session between Cisco ISE and the VN's Layer 3 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Note Configure the VN traffic that traverses the firewall to allow TCP bypass specifically for SXP communications between the Cisco ISE server and Layer 3 border nodes.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Client-to-Site ODC With Shared Firewall and Anycast Gateway

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure that all traffic transits through the firewall and is logged for auditing and compliance.
- Use a VPN client to securely connect to the client network.

- Restrict ODC users' access to systems within their respective ODC.
- Use the same firewall for multiple ODCs.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Redirect all traffic from the fusion device to the firewall and then forward it to the destination.
- Install and configure a VPN client on ODC users' laptops and desktops.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to handle traffic from each ODC separately.

Procedure

Step 1 Configure a Layer 3 VN at the site and then provision the required anycast gateway.

The screenshot shows the Catalyst Center interface for creating a Layer 3 Virtual Network. The main content area is titled "Fabric Sites and Fabric Zones (Optional)" and explains that a Layer 3 Virtual Network can be assigned to multiple Fabric Sites and Fabric Zones. Below this text is a flow diagram showing the relationship between the virtual network, fabric sites, and fabric zones.

```

    graph LR
      VN[S01_C2S_L3SVN] --> FS[.../Karnataka(1) Bengaluru  
Select Fabric Sites]
      FS --> FZ[No Fabric Zones Created]
  
```

On the right side, the "Configuration Attributes" panel shows a search for "LAYER 3 VIRTUAL NETWORKS" and lists the created virtual network "S01_C2S_L3SVN" with a green checkmark indicating it is successfully provisioned.

Step 2 Configure Layer 3 handoff on the border nodes for the newly added VN, then ensure that the eBGP sessions are up.

The screenshot shows the Catalyst Center interface for configuring Layer 3 handoff on the border nodes. The main content area is titled "Provision / SD-Access" and shows the configuration for the virtual network "S01_C2S_L3SVN". The "Enable Layer 3 Handoff" checkbox is checked, and the "Local IP Address/Mask" is set to "21.1.0.25/30" and the "Peer IP Address/Mask" is set to "21.1.0.26/30".

Virtual Network	Enable Layer 3 Handoff	VLAN	Local IP Address/Mask	Peer IP Address/Mask
S01_C2S_L3SVN	<input checked="" type="checkbox"/>	107	21.1.0.25/30 IPv6	21.1.0.26/30 IPv6

Step 3 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)

- Outside interface
- Firewall policies
- Routing for reachability

Step 4 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

Step 5 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

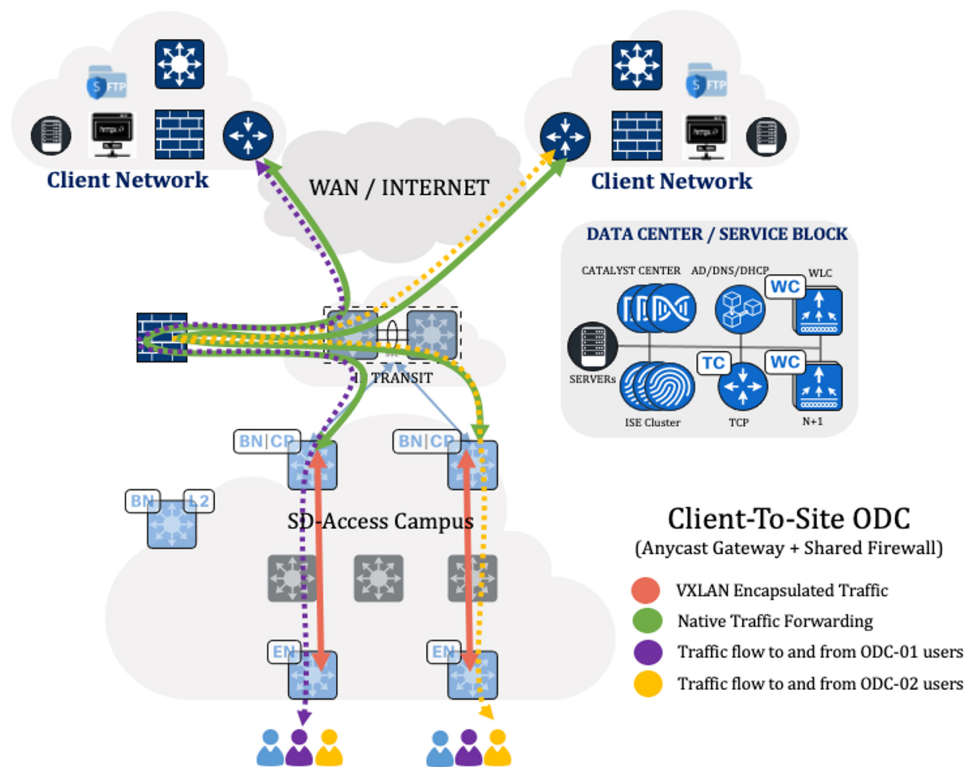
- Establish a VRF-aware SXP session between Cisco ISE and the VN's Layer 3 border node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Note Configure the VN traffic that traverses the firewall to allow TCP bypass specifically for SXP communications between the Cisco ISE server and Layer 3 border nodes.

Step 6 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Site-to-Site ODC With Dedicated Firewall Connected to the Edge and Acting as Gateway

Business requirements:

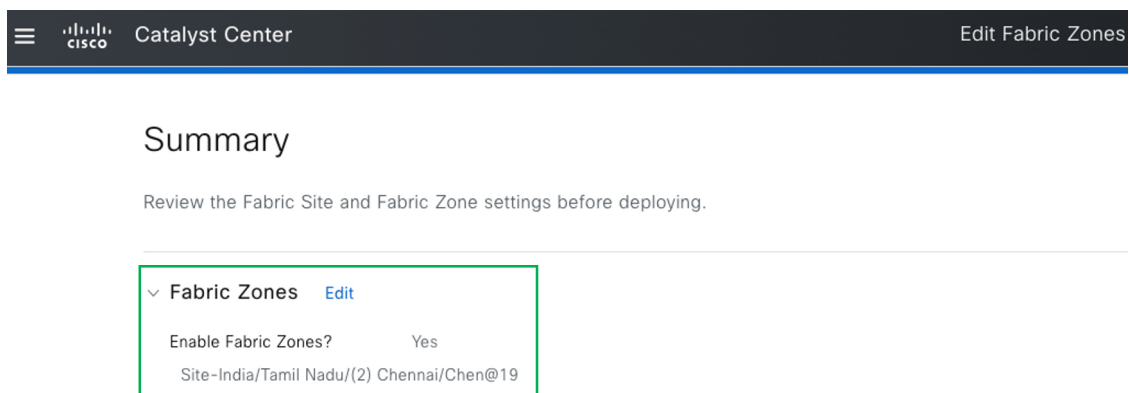
- Establish an ODC using Cisco SD-Access.
- Set up a firewall to act as the client's gateway and log all traffic for auditing and compliance.
- Confirm that the dedicated firewall is directly connected to the fabric edge node.
- Ensure the secure transfer of data between networks.
- Restrict ODC users' access to systems within their respective ODC.
- Implement a dedicated firewall for the ODC.

Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Provision a Layer 2 VN and add it to the fabric zone.
- Ensure that all ODC users and systems, as well as the firewall, are connected on the edge node within the fabric zone. The port connected to the firewall must be configured as a trunk port.
- Configure a site-to-site VPN for secure data transfer.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall interface to manage traffic from the dedicated ODC.

Procedure


Step 1 Enable the fabric zone for the site, where the ODC users, systems, and firewall are connected to the edge node.



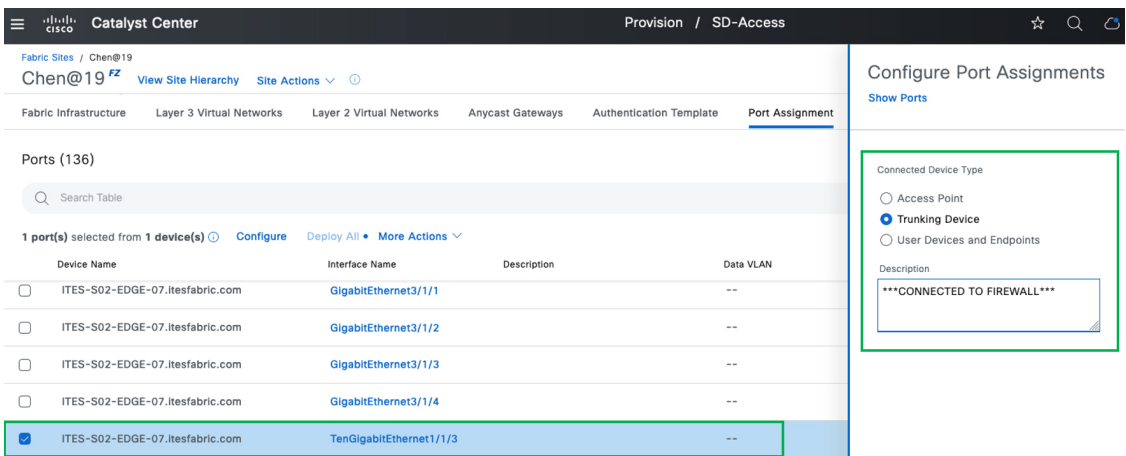
Step 2 Create a Layer 2 VN with an appropriate VLAN ID. Then assign the VN to both the fabric site and fabric zone.

Fabric Sites and Fabric Zones

A Layer 2 Virtual Network must be assigned a Fabric Site and can optionally be assigned to one or more Fabric Zones within the Site.



Step 3 Configure the port connected to the firewall on the edge node as a trunk port.



Step 4 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability
- Site-to-site VPN

Step 5 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

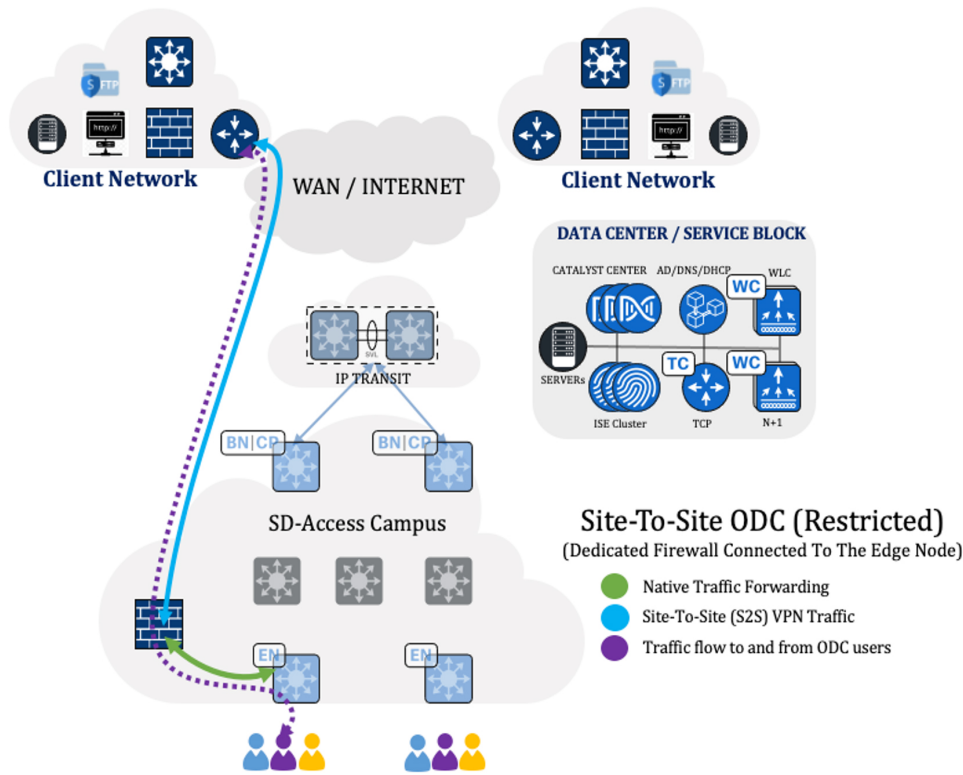
Step 6 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish an SXP session between Cisco ISE and the fabric edge node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Note Configure the VN traffic that traverses the firewall to allow TCP bypass specifically for SXP communications between the Cisco ISE server and Layer 3 border nodes.

Step 7 After authentication completes successfully, the end user should have access to all authorized resources. After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Client-to-Site ODC With Dedicated Firewall Connected to the Edge and Acting as Gateway

Business requirements:

- Establish an ODC using Cisco SD-Access.
- Ensure the dedicated firewall is directly connected to the fabric edge node.
- Set up a firewall to act as the client's gateway and log all traffic for auditing and compliance.
- Use a VPN client to securely connect to the client network.
- Restrict ODC users' access to systems within their respective ODC.
- Implement a dedicated firewall for the ODC.

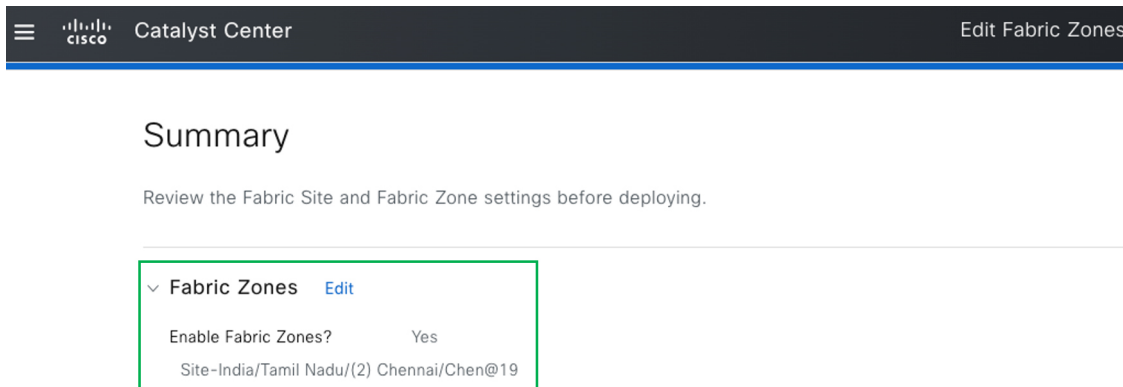
Technical requirements:

- Provision the Cisco SD-Access network using Catalyst Center.
- Provision a Layer 2 VN and add it to the fabric zone.

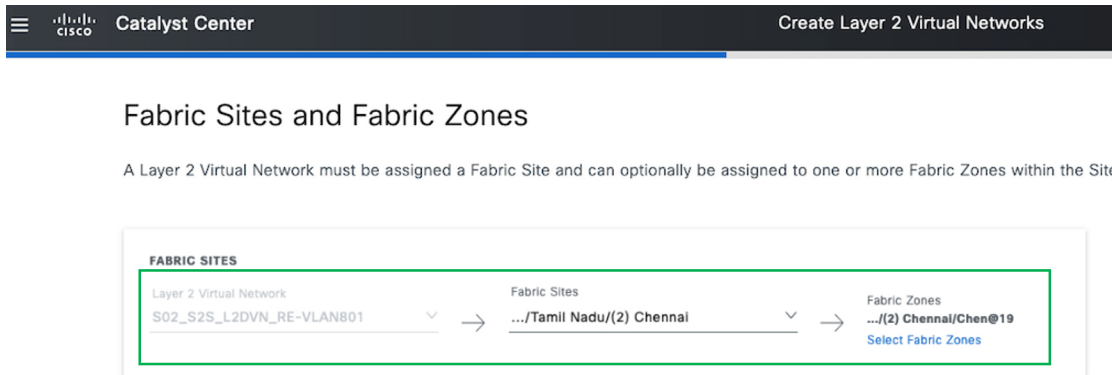
- Ensure that all ODC users, systems, and the firewall are connected to the edge node within the fabric zone. The port connected to the firewall must be configured as a trunk port.
- Install and configure a VPN client on ODC users' laptops and desktops.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Configure the firewall to manage traffic from the dedicated ODC.

Procedure

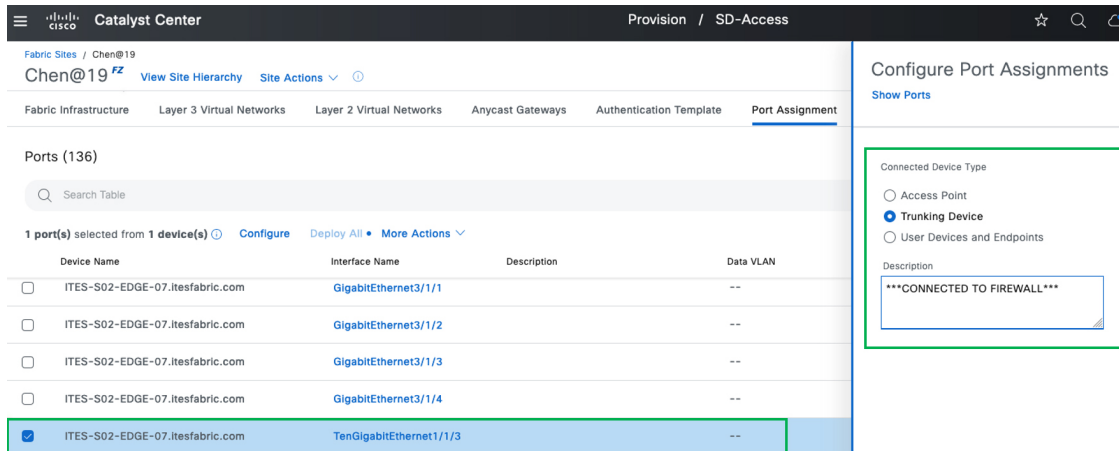
Step 1 Enable the fabric zone for the site, where ODC users, systems, and the firewall are connected to the edge node.



Step 2 Create a Layer 2 VN with an appropriate VLAN ID. Then assign the VN to both the fabric site and fabric zone.



Step 3 Configure the port connected to the firewall on the edge node as a trunk port.



Step 4 Enable end-to-end communication on the firewall.

At a minimum, you'll need to configure these items:

- Inside interface (specifying a physical interface or subinterface as the ODC's client gateway)
- Outside interface
- Firewall policies
- Routing for reachability

Step 5 Configure Cisco ISE to ensure proper authentication and authorization for both ODC and corporate users.

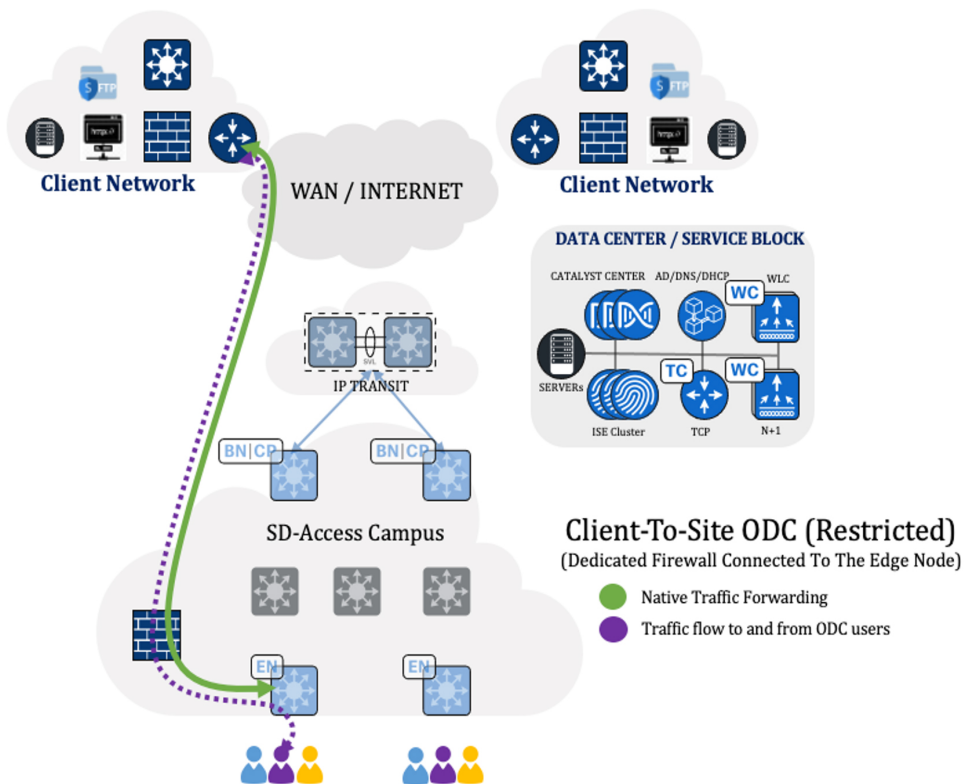
Step 6 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish an SXP session between Cisco ISE and the fabric edge node to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Step 7 After authentication completes successfully, the end user should have access to all authorized resources.

After a successful deployment, the data flow between the ODC and the client's network can be depicted as shown below:



Note Refer to the legend and arrows, which indicate the flow of traffic.

Location Agnostic Access for Site-to-Site ODC With Dedicated Firewall and Anycast Gateway

Business requirements:

- Return To Office (RTO) requirements for ITES customers:
 - Remove ODC client locations to enable work from any ITES site.
 - Classify and tunnel client traffic to the home location for exit.
- Confirm that existing site-to-site ODC users are able to work from any ITES site.
- Restrict ODC users' access to systems within their respective ODC, regardless of their location.
- Use the existing dedicated firewall for the site-to-site ODC at the home location.

Use case requirements:

- Provision a Layer 3 VN for site-agnostic clients:



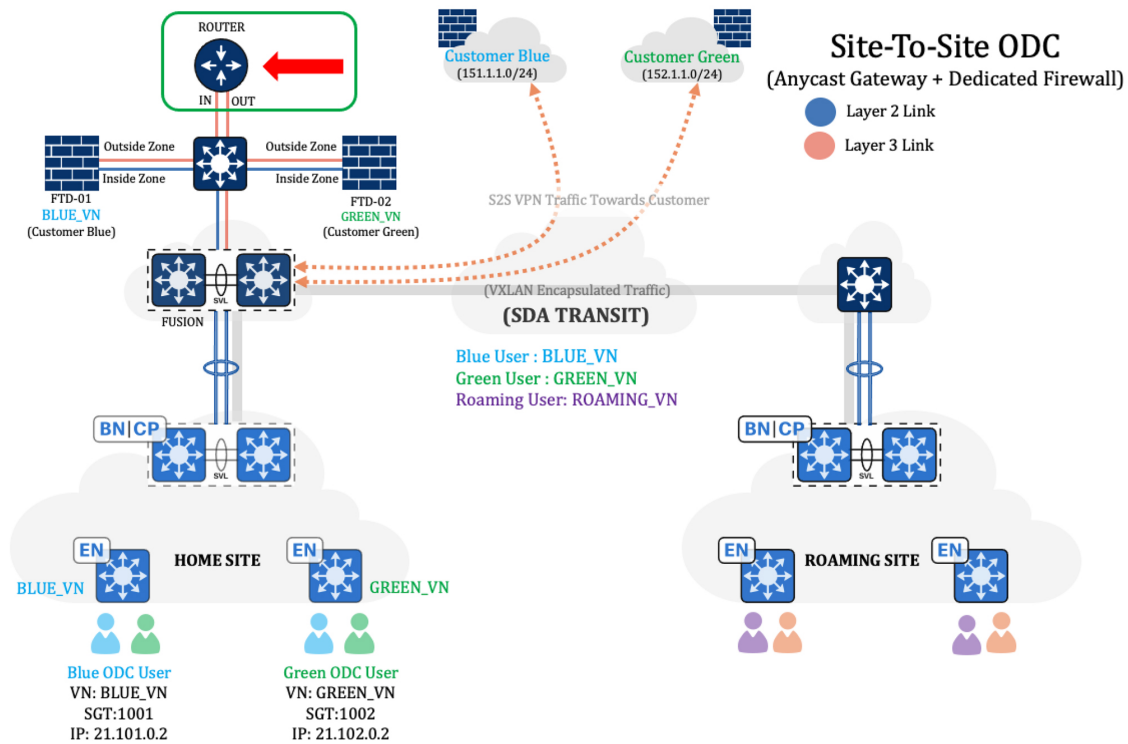
Note This type of VN will be referred to as a roaming VN in the remainder of this document.

A roaming VN is typically a virtual network that facilitates seamless connectivity and mobility across different geographical locations. It enables users to maintain consistent network access and services as they move between sites. The roaming VN is crucial for ensuring that a user's experience remains uninterrupted and secure, regardless of their physical location within the ITES infrastructure.

- Deploy an anycast gateway under the roaming VN at all necessary remote sites (excluding the home site).
- Install a router to implement Security Group Tag (SGT)-based policy-based routing (PBR) for roaming VN traffic.
- Configure Cisco ISE to ensure that roaming users are properly authenticated and authorized.
- Use macro (VN) or micro (SGT) segmentation to isolate ODC users and systems.
- Update the existing firewall policies and site-to-site VPN configuration to accommodate the newly added subnet for roaming users.

Procedure

- Step 1** Provision a roaming VN and deploy it to all necessary remote sites connected via Cisco SD-Access transit (LISP Pub/Sub), including the home site.
- Step 2** Deploy an anycast gateway under the roaming VN at all necessary remote sites, excluding the home site.
- Step 3** Configure Layer 3 handoff on the home site for exiting traffic.
- Step 4** At the home location, install a router to implement SGT-based PBR. Refer to this topology for guidance on router placement.



Step 5

To ensure the proper authentication and authorization of roaming users when they're onboarding from the roaming site, complete these tasks for Cisco ISE:

- On each roaming site's Cisco ISE server, create a location in the **Network Device Groups** window.
- Assign existing network access devices (NADs) to the appropriate location.
- Configure the authorization policies and profiles for onboarding roaming users.

Here are sample authorization policies created for home and roaming users:

S01-S2S-GIF-D1-CUST-01 (ODC-ROAMING-SITE)	OR	ITES-FABRIC-ExternalGroups EQUALS Itesfabric.com/Users/S01-S2S-GIF-D1-CUST-01	S01-S2S-GIF-D-ROAM (W...)	CTS1001
		InternalUser-IdentityGroup EQUALS User Identity Groups:S01-S2S-GIF-D1-CUST-01 (IXIA)		
	AND	Radius-NAS-Port-Type EQUALS Ethernet		
	OR	Normalised Radius-RadiusFlowType EQUALS Wired802_1x		
	OR	Normalised Radius-RadiusFlowType EQUALS WiredMAB		
	OR	DEVICE-Location EQUALS All Locations#Chennai		
	OR	DEVICE-Location EQUALS All Locations#Kolkata		
S01-S2S-GIF-D1-CUST-01 (ODC-HOME-SITE)	OR	ITES-FABRIC-ExternalGroups EQUALS Itesfabric.com/Users/S01-S2S-GIF-D1-CUST-01	S01-S2S-GIF-D1-CUST-0...	CTS1001
		InternalUser-IdentityGroup EQUALS User Identity Groups:S01-S2S-GIF-D1-CUST-01 (IXIA)		
	AND	Radius-NAS-Port-Type EQUALS Ethernet		
	OR	Normalised Radius-RadiusFlowType EQUALS Wired802_1x		
	OR	Normalised Radius-RadiusFlowType EQUALS WiredMAB		

Note To ensure that the appropriate attributes are matched before authenticating end users, the policy for roaming users is positioned above the policy for home users.

Step 6

Update the existing site-to-site firewall policy to permit traffic from the roaming VN's subnet.

Also set the roaming VN's subnet as a protected network for the site-to-site VPN and update the NAT policies accordingly.

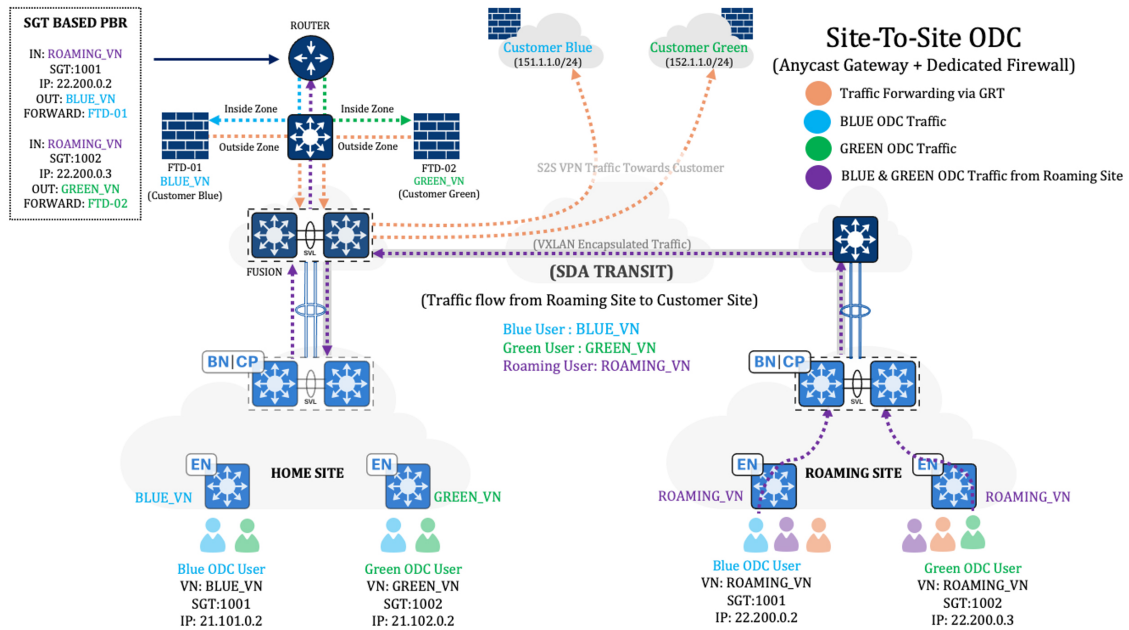
Step 7

Configure the routing of traffic destined for the customer site:

- Forward traffic from the site-to-site ODC user at the home site to the dedicated firewall of that ODC.
- Direct traffic from the site-to-site ODC user at the roaming site (roaming VN traffic) to the router.

SGT-based PBR ensures that this traffic is routed to the appropriate firewall.

Note Confirm that the roaming client's tag is preserved all the way to the steering router.



Note Refer to the legend and arrows, which indicate the flow of traffic.

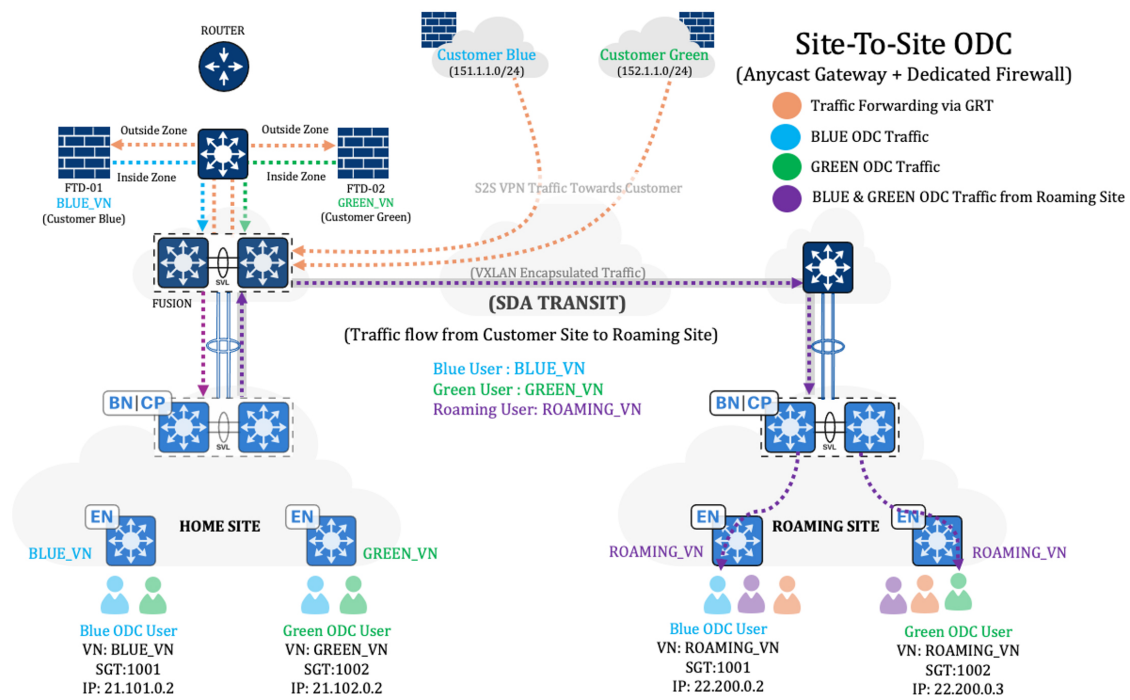
Traffic from the roaming site flows to the customer site displayed in this topology in the following manner:

- Roaming client traffic from the roaming site (the purple flow) is forwarded to the router via the home site's border node and fusion node.
- The router then performs SGT-based PBR, forwarding ingress traffic with SGT 1001 to the inside zone of the Customer Blue's firewall (FTD-01). Traffic with SGT 1002 is forwarded to the inside zone of the Customer Green firewall (FTD-02).

Step 8

Configure the routing for traffic destined for the roaming site:

- Configure the site-to-site ODC firewall to forward return traffic, or traffic initiated from the customer site and intended for the roaming site, to the fusion node.
- Configure VRF route leaking on the fusion node to enable the site-to-site ODC VN to learn about the roaming subnets.



Note Refer to the legend and arrows, which indicate the flow of traffic.

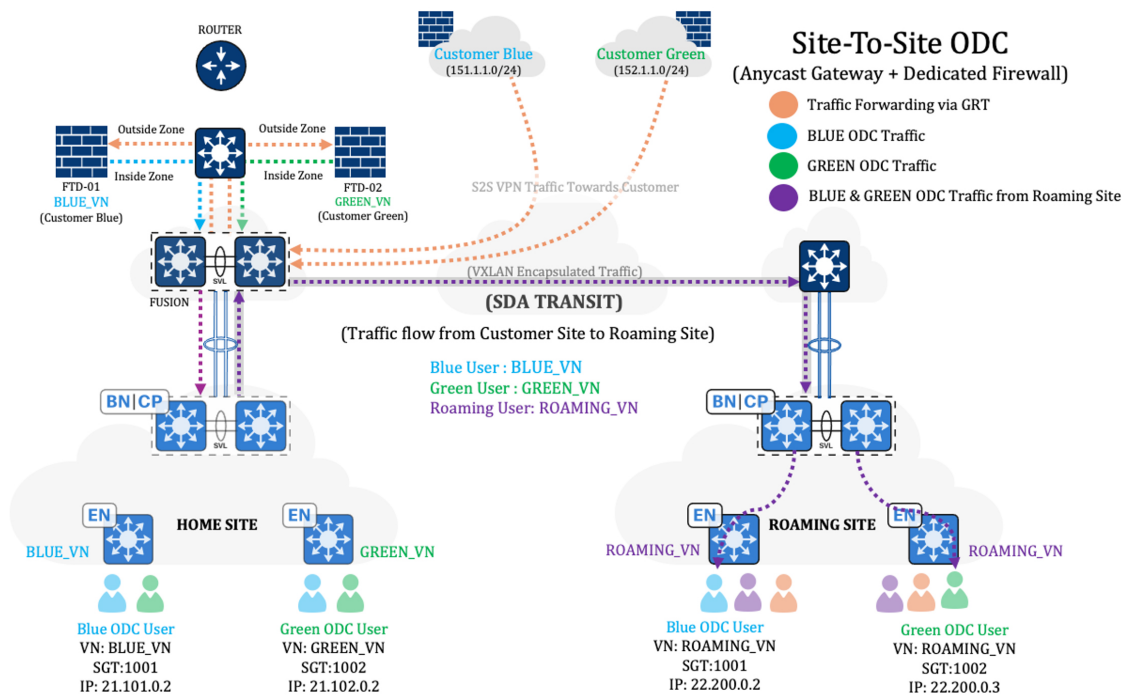
Traffic from the customer site flows to the roaming site displayed in this topology in the following manner:

- Traffic destined for Customer Blue and Customer Green ODC clients at the roaming site is first received by the customer's dedicated firewall. The firewall then forwards this traffic to the fusion node.
- The fusion node passes the traffic to the roaming VN at the home site's border node, which subsequently directs it to the roaming site.

Step 9

Communication between site-to-site ODC users onboarded from the home site and the roaming site should function properly as part of the route exchange, provided that the home site offers internet access to other sites through Cisco SD-Access. If this is not the case, you'll need to perform these tasks:

- Make the roaming VN on the home site's border node aware of the site-to-site ODC subnet.
- Import its routes into the LISP instance.



Note Refer to the legend and arrows, which indicate the flow of traffic.

Step 10 If the allow-list (Default Deny IP) TrustSec model is enabled, do one of the following:

- Establish a VRF-aware SXP session between Cisco ISE and the home site border node for the roaming VN to obtain the destination addresses' security tags.
- Manually configure the subnet-to-SGT mapping.

Also ensure that the correct policy matrix is configured for the destination and source security tags.

Supported Hardware and Software

The ITES solution is tested with the hardware and software listed in this table. For the complete list of the hardware that the Cisco SD-Access solution supports, see the [Cisco Software-Defined Access Compatibility Matrix](#).

Hardware or Software Component	Supported Software Version
Catalyst Center appliance (Cisco part number: DN2-HW-APL-XL)	Catalyst Center 2.3.7.6
Cisco Identity Services Engine (ISE)	Cisco ISE 3.3 Patch 3
Control plane node: Cisco Catalyst 8000V Edge Software	Cisco IOS-XE 17.9.5a, 17.12.3a

Hardware or Software Component	Supported Software Version
Fabric border node: Cisco Catalyst 9500 and 9600 Series switches	Cisco IOS-XE 17.9.5, 17.12.4
Fabric edge node: Cisco Catalyst 9200, 9300, and 9400 Series switches	Cisco IOS-XE 17.9.5, 17.12.4
Wireless controller: Cisco Catalyst 9800-40 and 9800-CL	Cisco IOS-XE 17.9.5, 17.12.3
Cisco Secure Firewall Management Center (FMC): Management Center Virtual	7.2.8
Cisco Secure Firewall Threat Defense (FTD): Firewall Threat Defense Virtual and Cisco Firepower 1150	7.2.8

Multidimensional Scale Numbers

Category	Value
Devices in inventory (includes routers, switches, and wireless controllers)	5000
Number of fabric sites	100
Number of buildings and floors	4000
Number of IP pools per site	1000
Number of VNs per site	128
Number of wireless controllers per site	2 with HA
Number of APs (fabric and nonfabric)	8000
Number of SGTs	4000
Number of group-based policies	12,000
Number of endpoints	100,000 (30% wired and 70% wireless)
Route-map entries (match & set)—for SGT-based policy-based routing (PBR)	125

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*
- Cisco SD-Access Design Tool
- *Software Defined Access (SDA) Provisioning Best Practice Guide*

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2024 Cisco Systems, Inc. All rights reserved.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

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