CISCO VALIDATED DESIGN

Software-Defined Access Design Guide

December 2018
Solution 1.2

Table of Contents

Cisco Digital Network Architecture and Software-Defined Access introduction

Cisco® Digital Network Architecture (Cisco DNA™) provides a roadmap to digitization and a path to realize immediate benefits of network automation, assurance, and security. Cisco Software-Defined Access (SD-Access) is the Cisco DNA evolution from traditional campus LAN designs to networks that directly implement the intent of an organization. SD-Access is enabled with an application package that runs as part of the Cisco DNA Center™ software for designing, provisioning, applying policy, and facilitating the creation of an intelligent campus wired and wireless network with assurance.

Fabric technology, an integral part of SD-Access, enables wired and wireless campus networks with programmable overlays and easy-to-deploy network virtualization, permitting a physical network to host one or more logical networks as required to meet the design intent. In addition to network virtualization, fabric technology in the campus network enhances control of communications, providing software-defined segmentation and policy enforcement based on user identity and group membership. Software-defined segmentation is seamlessly integrated using Cisco TrustSec® technology, providing micro-segmentation through the use of scalable groups within a virtual network. Using Cisco DNA Center to automate the creation of virtual networks reduces operational expenses, coupled with the advantage of reduced risk with integrated security and improved network performance provided by the assurance and analytics capabilities.

This design guide provides an overview of the requirements driving the evolution of campus network designs, followed by a discussion about the latest technologies and designs that are available for building an SD-Access network to address those requirements. It is a companion to the associated deployment guides for SD-Access, which provide configurations explaining how to deploy the most common implementations of the designs as described in this guide. The intended audience is a technical decision maker who wants to understand Cisco's campus offerings and to learn about the technology options available and the leading practices for designing the best network for the needs of an organization.

If you didn't get this guide from Cisco Design Zone, you can check for the latest version of this guide.

For the associated Software-Defined Access Deployment Guide, related design guides, and white papers, see the following page: https://www.cisco.com/go/designzone

Network requirements for the digital organization

With digitization, software applications are evolving from simply supporting business processes to becoming, in some cases, the primary source of business revenue and competitive differentiation. Organizations are now constantly challenged by the need to scale their network capacity in order to quickly react to application demands and growth. Because the campus LAN is the network through which users and devices within a location access applications, campus wired and wireless LAN capabilities should be enhanced to support those changing needs.

The following are key requirements that are driving the evolution of existing campus networks.

Flexible Ethernet foundation for growth and scale

- **Simplified deployment and automation**—Network device configuration and management through a centralized controller using open APIs allow for very fast, lower-risk deployment of network devices and services.
- **Increased bandwidth needs**—Bandwidth needs are doubling potentially multiple times over the lifetime of a network, resulting in new networks needing to be prepared to aggregate using 10 Gbps Ethernet to 40 Gbps to 100 Gbps capacities over time.
- **Increased capacity of wireless access points**—The bandwidth demands on wireless Access Points (APs) with the latest 802.11ac Wave 2 technology now exceed 1 Gbps, and the IEEE has now ratified the 802.3bz standard that defines 2.5 Gbps and 5 Gbps Ethernet. Cisco Catalyst® Multigigabit technology supports that bandwidth demand without requiring an upgrade of the existing copper Ethernet wiring plant.
- **Additional power requirements from Ethernet devices**—New devices may require higher power to operate, such as lighting, surveillance cameras, virtual desktop terminals, remote access switches, and APs. Your access layer design should have the ability to support power over Ethernet with 60W per port, offered with Cisco Universal Power Over Ethernet, and the access layer should also provide Power over Ethernet (PoE) perpetual power during switch upgrade and reboot events. The Cisco Catalyst 9000 family of access layer switches is perpetual PoE-capable and hardware-ready for 100W per port, as that technology becomes available.

Integrated services and security

- **Consistent wired and wireless security capabilities**—Security capabilities described below should be consistent whether a user is connecting to a wired Ethernet port or connecting over the wireless LAN.
- **Network assurance and analytics**—Proactively predict network-related and security-related risks by using telemetry to improve the performance of the network, devices, and applications, even with encrypted traffic.
- **Identity services**—Identifying users and devices connecting to the network provides the contextual information required to implement security policies for access control, network segmentation by using scalable group membership and mapping of devices into Virtual Networks (VNs).
- **Group-based policies**—Creating access and application policies based on user group information provides a much easier and scalable way to deploy and manage security policies. Traditional Access Control Lists (ACLs) can be difficult to implement, manage, and scale because they rely on network constructs such as IP addresses and subnets.
- **Software-defined segmentation**—Scalable Group Tags (SGTs) assigned from group-based policies can be used to segment a network in order to achieve data plane isolation within physical and virtual networks.
- **Network virtualization**—The capability to share a common infrastructure while supporting multiple VNs with isolated data and control planes enables different sets of users and applications to be isolated securely.

Example use case: Secure segmentation and profiling for healthcare

Our healthcare records are just as valuable to attackers as our credit card numbers and online passwords. In the wake of recent cyberattacks, hospitals are required to have HIPAA-compliant wired and wireless networks that can provide complete and constant visibility into their network traffic in order to protect sensitive medical devices (such as electronic medical records servers, vital monitors or nurse workstations) so that a malicious device cannot compromise the networks.

A patient's mobile device, when compromised by malware, may change network communication behavior to propagate and infect other endpoints. It is considered abnormal behavior when a patient's mobile device communicates with any medical device. SD-Access can address the need for complete isolation between patient devices and medical facility devices by using macro-segmentation and putting devices into different overlay networks, enabling the isolation.

How can we address a similar scenario, but for the case of a compromised medical professional's mobile device requiring connectivity to information systems for some tasks, but not requiring connectivity to other medical devices? SD-Access can take this need for segmentation beyond simple network separation, by profiling devices and users as they come onto the network and applying micro-segmentation within an overlay network. Flexible policy creation allows the ability to have groups of device types and user roles to restricted communication within a group or to enable communication among groups only as needed to implement the intent of the policies of an organization.

Deploying the intended outcomes for the needs of the organization is simplified using the automation capabilities built into Cisco DNA Center, and those simplifications span the wired and wireless domains.

Software-Defined Access architecture

The SD-Access architecture is supported by fabric technology implemented for the campus, which enables the use of virtual networks (overlay networks) running on a physical network (underlay network) in order to create alternative topologies to connect devices. Overlay networks in data center fabrics are commonly used to provide Layer 2 and Layer 3 logical networks with virtual machine mobility (examples: Cisco ACI™, VXLAN/EVPN, and FabricPath). Overlay networks are also used in wide-area networks to provide secure tunneling from remote sites (examples: MPLS, DMVPN, and GRE). This section provides information about the SD-Access architecture elements. SD-Access design recommendations are covered in the Design Considerations section.

Underlay network

The underlay network is defined by the physical switches and routers that are used to deploy the SD-Access network. All network elements of the underlay must establish IP connectivity via the use of a routing protocol. Instead of using arbitrary network topologies and protocols, the underlay implementation for SD-Access uses a well-designed Layer 3 foundation inclusive of the campus edge switches (also known as a ***routed access*** design), to ensure performance, scalability, and high availability of the network.

The Cisco DNA Center LAN Automation feature is an alternative to manual underlay deployments for new networks and uses an IS-IS routed access design. Though there are many alternative routing protocols, the IS-IS selection offers operational advantages such as neighbor establishment without IP protocol dependencies, peering capability using loopback addresses, and agnostic treatment of IPv4, IPv6, and non-IP traffic. In the latest versions of Cisco DNA Center, LAN Automation uses Cisco Network Plug and Play features to deploy both unicast and multicast routing configuration in the underlay, aiding traffic delivery efficiency for services built on top.

In SD-Access, the underlay switches support the end-user physical connectivity. However, end-user subnets are not part of the underlay network—they are part of a programmable Layer 2 or Layer 3 overlay network.

Tech tip

The validated SD-Access 1.2 solution supports IPv4 underlay networks. For IPv6 underlay networks, see the release notes for your software version to verify support.

Overlay network

An overlay network is created on top of the underlay to create a virtualized network. The data plane traffic and control plane signaling is contained within each virtualized network, maintaining isolation among the networks in addition to independence from the underlay network. The SD-Access fabric implements virtualization by encapsulating user traffic in overlay networks using IP packets that are sourced and terminated at the boundaries of the fabric. The fabric boundaries include borders for ingress and egress to a fabric, fabric edge switches for wired clients, and fabric APs for wireless clients. The details of the encapsulation and fabric device roles are covered in later sections. Overlay networks can run across all or a subset of the underlay network devices. Multiple overlay networks can run across the same underlay network to support multitenancy through virtualization. Each overlay network appears as a virtual routing and forwarding (VRF) instance for connectivity to external networks. You preserve the overlay separation when extending the networks outside of the fabric by using VRF-lite, maintaining the network separation within devices connected to the fabric and also on the links between VRF-enabled devices.

In earlier versions of SD-Access, IPv4 multicast forwarding within the overlay operates using headend replication of multicast packets into the fabric for both wired and wireless endpoints. Recent versions of SD-Access use underlay multicast capabilities, configured manually or by using LAN Automation, for more efficient delivery of traffic to interested edge switches versus using headend replication. The multicast is encapsulated to interested fabric edge switches, which de-encapsulate the multicast, replicating the multicast to all the interested receivers on the switch. If the receiver is a wireless client, the multicast (just like unicast) is encapsulated by the fabric edge towards the AP with the multicast receiver. The multicast source can exist either within the overlay or outside of the fabric. For PIM deployments, the multicast clients in the overlay use an RP at the fabric border that is part of the overlay endpoint address space. Cisco DNA Center configures the required multicast protocol support.

Tech tip

The SD-Access 1.2 solution supports both PIM SM and PIM SSM, and the border node must include the IP multicast Rendezvous Point (RP) configuration. For available multicast underlay optimizations, see the release notes.

Multicast support in the underlay network enables Layer 2 flooding capability in overlay networks. SD-Access now supports overlay flooding of ARP frames, broadcast frames, and link-local multicast frames, which addresses some specific connectivity needs for ***silent hosts***, requiring receipt of traffic before communicating, and mDNS services.

Layer 2 overlays

Layer 2 overlays emulate a LAN segment to transport Layer 2 frames, carrying a single subnet over the Layer 3 underlay. Layer 2 overlays are useful in emulating physical topologies and depending on the design can be subject to Layer 2 flooding. SD-Access supports transport of IP frames without Layer 2 flooding of broadcast and unknown multicast traffic. Without broadcasts from the fabric edge, ARP functions by using the fabric control plane for MAC-to-IP address table lookups.

Figure 1. Layer 2 overlay—connectivity logically switched

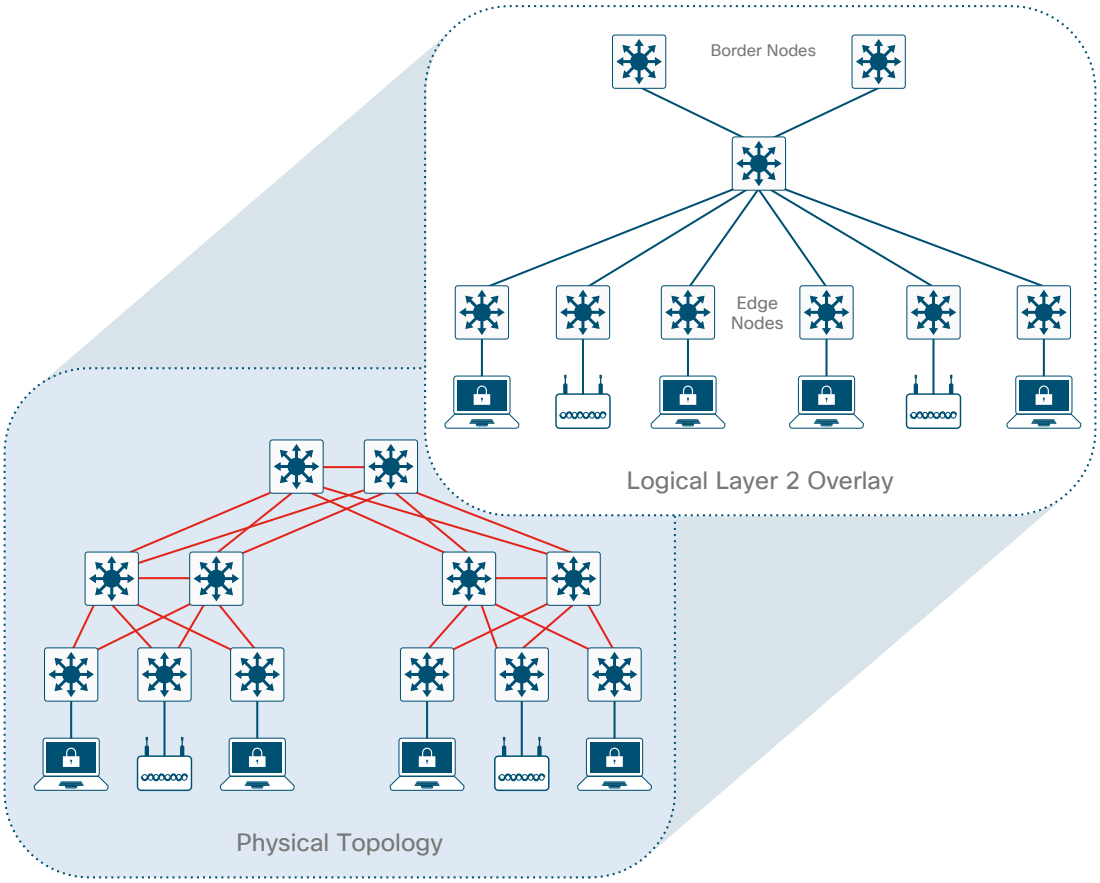

Layer 3 overlays

Layer 3 overlays abstract IP-based connectivity from physical connectivity and allow multiple IP networks as part of each virtual network. Overlapping IP address space across different Layer 3 overlays is outside the scope of validation, and should be approached with the awareness that the network virtualization must be preserved for communications outside of the fabric, while addressing any IP address conflicts.

Tech tip

The SD-Access 1.2 solution supports IPv4 overlays. Overlapping IP addresses are not supported for wireless clients on the same WLC. For IPv6 overlays, see the release notes for your software version to verify support.

Figure 2. Layer 3 overlay—connectivity logically routed

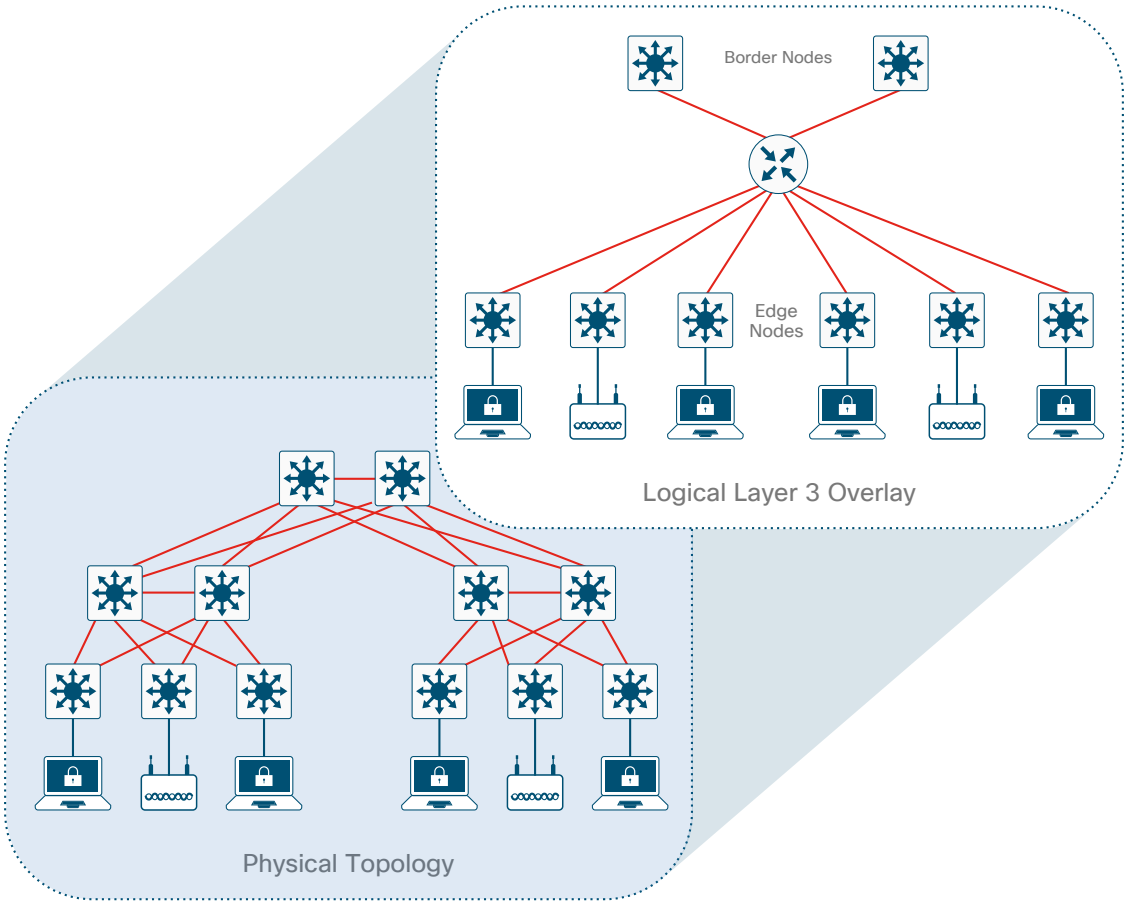

Fabric data plane and control plane

SD-Access configures the overlay network with a fabric data plane by using virtual extensible LAN (VXLAN) technology. VXLAN encapsulates and transports complete Layer 2 frames across the underlay, with each overlay network identified by a VXLAN Network Identifier (VNI). The VXLAN header also carries the SGTs required for micro-segmentation.

The mapping and resolving of endpoints requires a control plane protocol, and SD-Access uses Locator/ID Separation Protocol (LISP) for this task. LISP brings the advantage of routing based not only on the IP address or MAC address as the Endpoint Identifier (EID) for a device but also on an additional IP address that it provides as a Routing Locator (RLOC) to represent the network location of that device. The EID and RLOC combination provides all the necessary information for traffic forwarding, even if an endpoint uses an unchanged IP address when appearing in a different network location. Simultaneously, the decoupling of the endpoint identity from its location allows addresses in the same IP subnetwork to be available behind multiple Layer 3 gateways, versus the one-to-one coupling of IP subnetwork with network gateway in traditional networks.

The following diagram shows an example of two subnets that are part of the overlay network. The subnets stretch across physically separated Layer 3 devices. The RLOC interface is the only routable address that is required to establish connectivity between endpoints of the same or different subnet.

Figure 3. Example topology—subnet stretching

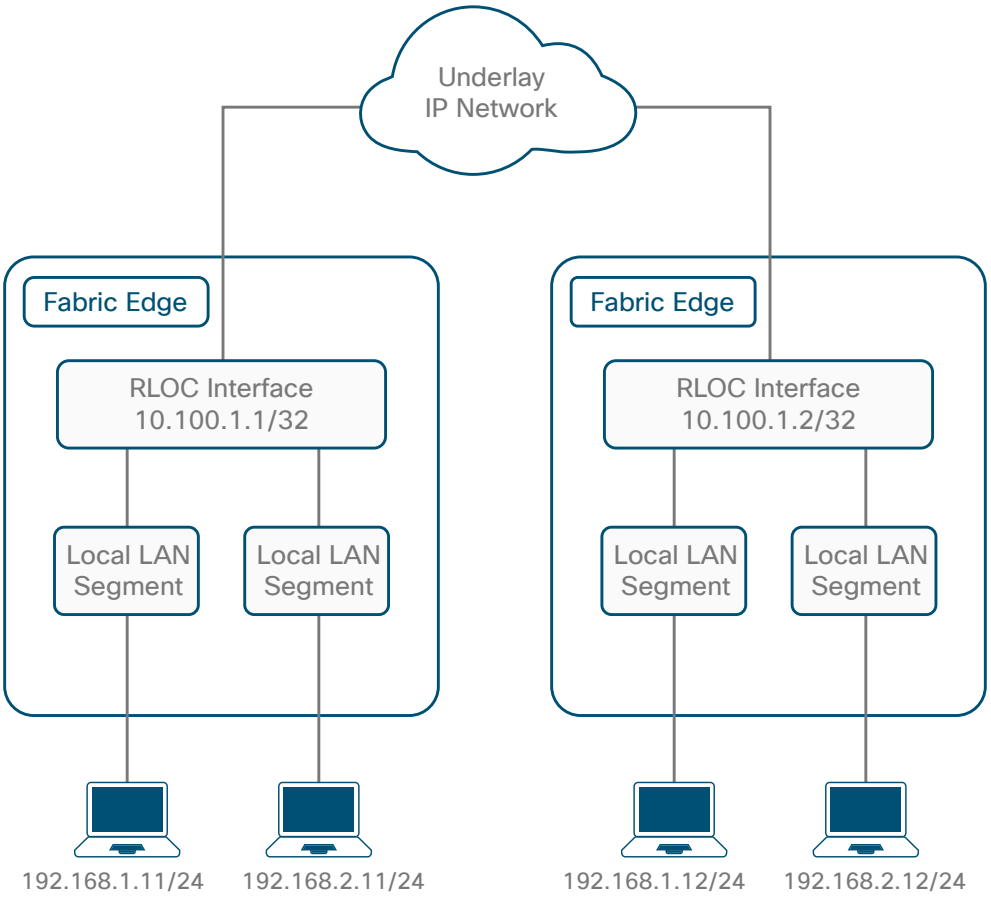

For details about the fabric control plane and data plane constructs, as well as a glossary of terms, see the appendices.

Wireless integration

SD-Access supports two options for integrating wireless access into the network. One option is to use traditional Cisco Unified Wireless Network local-mode configurations "over the top" as a non-native service. In this mode, the SD-Access fabric is simply a transport network for the wireless traffic, which can be useful during migrations. The other option is fully integrated SD-Access Wireless, extending the SD-Access benefits to include wireless users.

You gain advantages by integrating wireless natively into SD-Access using two additional components—fabric wireless controllers and fabric mode APs. Supported Cisco Wireless LAN Controllers (WLCs) are configured as fabric wireless controllers to communicate with the fabric control plane, registering Layer 2 client MAC addresses, SGT, and Layer 2 VNI information. The fabric mode APs are Cisco 802.11ac Wave 2 and Wave 1 APs associated with the fabric wireless controller and configured with fabric-enabled SSIDs. The APs are responsible for communication with wireless endpoints, and in the wired domain, the APs assist the VXLAN data plane by encapsulating and de-encapsulating traffic at the connected edge node.

Tech tip

SD-Access provides optimized features such as Application Visibility and Control (AVC) when deployed with Wave 2 APs. For the differences between Wave 2 and Wave 1 AP feature support, see the release notes for your wireless software version.

Fabric wireless controllers manage and control the fabric mode APs using the same model as the traditional centralized model of local-mode controllers, offering the same operational advantages, such as mobility control and radio resource management. A significant difference is that client traffic carried from wireless endpoints on fabric SSIDs avoids Control And Provisioning of Wireless Access Points (CAPWAP) encapsulation and forwarding from the APs to the central controller. Instead, communication from wireless clients is VXLAN-encapsulated by fabric-attached APs. This difference enables a distributed data plane with integrated SGT capabilities. Traffic forwarding takes the optimum path through the SD-Access fabric to the destination with consistent policy, regardless of wired or wireless endpoint connectivity.

The control plane communication for the APs uses a CAPWAP tunnel to the WLC, similar to the traditional Cisco Unified Wireless Network control plane. However, the WLC integration with the SD-Access control plane supports wireless clients roaming to APs across the fabric. The SD-Access fabric control plane inherently supports the roaming feature by updating its host-tracking database with any changes for a wireless client EID associated with a new RLOC.

Although the fabric mode APs are used for VXLAN traffic encapsulation for wireless traffic while it moves between the wireless and the wired portions of the fabric, the APs are not edge nodes. Instead, APs connect directly to edge node switches using VXLAN encapsulation and rely on those switches to provide fabric services, such as the Layer 3 anycast gateway.

Integrating the wireless LAN into the fabric enables the fabric advantages for the wireless clients, including addressing simplification, mobility with stretched subnets, and end-to-end segmentation with policy consistency across the wired and wireless domains. Wireless integration also enables the WLC to shed data plane forwarding duties while continuing to function as the centralized services and control plane for the wireless domain.

Guest wireless services

If you are not doing Cisco Unified Wireless Network wireless over the top and require fabric wireless guest access services to the Internet, separate the wireless guests from other network services by creating a dedicated virtual network supporting the guest SSID. Extend the separation of the guest traffic between the fabric border and DMZ, using VRF Lite or similar techniques.

If your wireless deployment requires guest traffic to be delivered to the DMZ using a control plane and data plane dedicated to that purpose, deploy a set of fabric border and control plane nodes (described in the Solution Components section) for guest services within the DMZ segment. For this case, the traffic separation is maintained using encapsulation from edge switches connecting the APs all the way to the DMZ border, with the advantage that the guest control plane has independent scale and performance and avoids techniques such as VRF Lite. Considerations for deploying this guest wireless design include configuration of in-path firewalls to permit the fabric traffic and accommodating fabric MTU requirements end-to-end.

In SD-Access 1.2, Cisco DNA Center automates and manages the workflow for implementing the wireless guest solution for fabric devices only, and wired guest services are not included in the solution.

Solution management

A full understanding of LISP and VXLAN is not required to deploy the fabric in SD-Access. Nor is there a requirement to know the details of how to configure each individual network component and feature to create the consistent end-to-end behavior offered by SD-Access. Instead, you use Cisco DNA Center—an intuitive centralized management system—to design, provision, and apply policy across the wired and wireless SD-Access network.

In addition to automation for SD-Access, Cisco DNA Center offers traditional applications to improve an organization's efficiency, such as software image management, along with new capabilities, such as device health dashboards and 360-degree views, as listed in the Solutions Components section.

Cisco DNA Center is integral to SD-Access, enabling automation of device deployments into the network to provide the speed and consistency required for operational efficiency. Organizations then benefit from lower costs and reduced risk when deploying and maintaining their networks.

Policy management with identity services integrates into the SD-Access network using an external repository hosted by the Cisco Identity Services Engine (ISE). ISE couples with Cisco DNA Center for dynamic mapping of users and devices to scalable groups, simplifying end-to-end security policy management and enforcement at a greater scale than traditional network policy implementations relying on IP access lists.

Solution components

The SD-Access solution combines the Cisco DNA Center software, identity services, and wired and wireless fabric functionality. Within the SD-Access solution, a fabric site is composed of an independent set of fabric control plane nodes, edge nodes, intermediate (transport only) nodes, and border nodes. Wireless integration adds fabric WLC and fabric mode AP components to the fabric site. Fabric sites can be interconnected using an SD-Access transit network to create a larger fabric domain. This section describes the functionality for each role, how the roles map to the physical campus topology, and the components required for solution management, wireless integration, and policy application.

Figure 4. SD-Access solution and fabric components

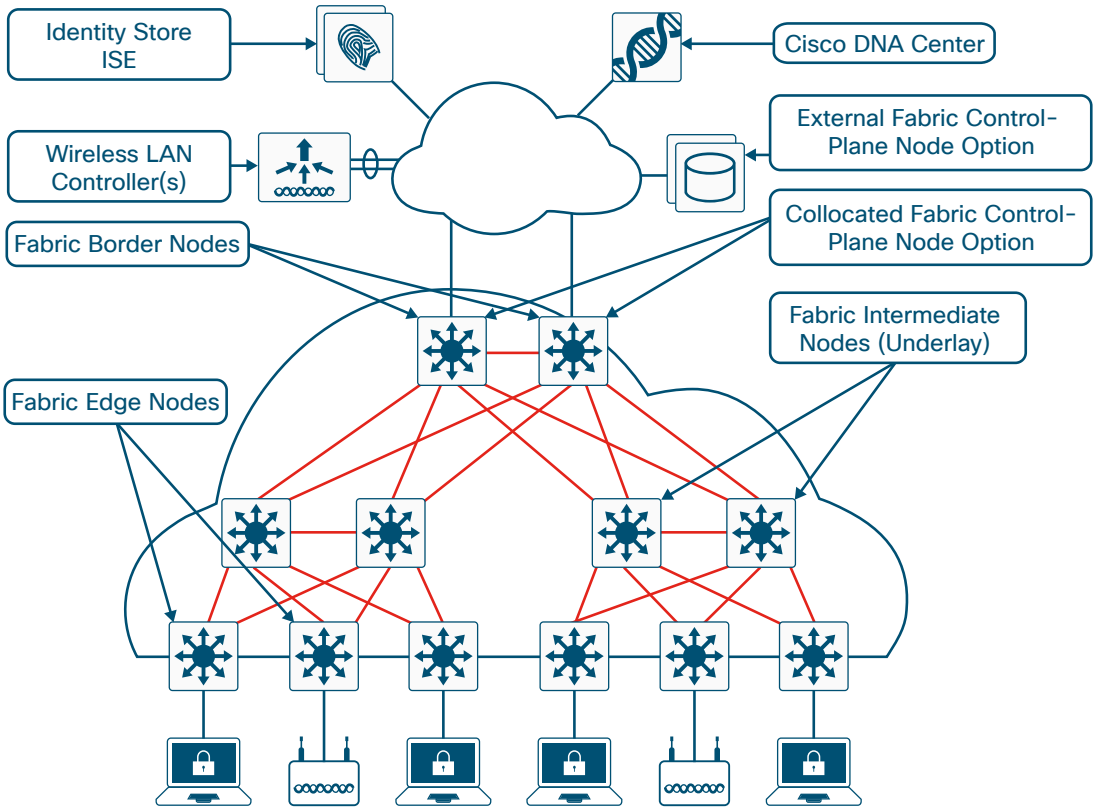

Control plane node

The SD-Access fabric control plane node is based on the LISP Map-Server (MS) and Map-Resolver (MR) functionality combined on the same node. The control plane 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 (closest router) in the network. The control plane node functionality can be collocated with a border node or can use dedicated nodes for scale and between two and six nodes are used for resiliency. Border and edge nodes register with and use all control plane nodes, so resilient nodes chosen should be of the same type for consistent performance.

The control plane node enables the following functions:

- **Host tracking database**—The Host Tracking Database (HTDB) is a central repository of EID-to-fabric-edge node bindings.
- **Map server**—The LISP MS is used to populate the HTDB from registration messages from fabric edge devices.
- **Map resolver**—The LISP MR is used to respond to map queries from fabric edge devices requesting RLOC mapping information for destination EIDs.

Edge node

The SD-Access fabric edge nodes are the equivalent of an access layer switch in a traditional campus LAN design. The edge nodes implement a Layer 3 access design with the addition of the following fabric functions:

- **Endpoint registration**—After an endpoint is detected by the fabric edge, it is added to a local host tracking database called the EID-table. The edge device also issues a LISP map-register message in order to inform the control plane node of the endpoint detected so that it can populate the HTDB.
- **Mapping of user to virtual network**—Endpoints are placed into virtual networks by assigning the endpoint to a VLAN associated with a LISP instance. The mapping of endpoints into VLANs can be done statically or dynamically using 802.1X. An SGT is also assigned, and an SGT can be used to provide segmentation and policy enforcement at the fabric edge.

> **Tech tip**
>
> Cisco IOS® Software enhances 802.1X device capabilities with Cisco Identity Based Networking Services (IBNS) 2.0. For example, concurrent authentication methods and interface templates have been added. Likewise, Cisco DNA Center has been enhanced to aid with the transition from IBNS 1.0 to 2.0 configurations, which use Cisco Common Classification Policy Language (commonly called C3PL). See the release notes and updated deployment guides for additional configuration capabilities. For more information about IBNS, see: https://cisco.com/go/ibns.

- **Anycast Layer 3 gateway**—A common gateway (IP and MAC addresses) can be used at every node that shares a common EID subnet providing optimal forwarding and mobility across different RLOCs.
- **LISP forwarding**—Instead of a typical routing-based decision, the fabric edge nodes query the map server to determine the RLOC associated with the destination EID and use that information as the traffic destination. In case of a failure to resolve the destination RLOC, the traffic is sent to the default fabric border in which the global routing table is used for forwarding. The response received from the map server is stored in the LISP map-cache, which is merged to the Cisco Express Forwarding table and installed in hardware. If traffic is received at the fabric edge for an endpoint not locally connected, a LISP solicit-map-request is sent to the sending fabric edge in order to trigger a new map request; this addresses the case where the endpoint may be present on a different fabric edge switch.
- **VXLAN encapsulation/de-encapsulation**—The fabric edge nodes use the RLOC associated with the destination IP address to encapsulate the traffic with VXLAN headers. Similarly, VXLAN traffic received at a destination RLOC is de-encapsulated. The encapsulation and de-encapsulation of traffic enables the location of an endpoint to change and be encapsulated with a different edge node and RLOC in the network, without the endpoint having to change its address within the encapsulation.

Intermediate node

The fabric intermediate nodes are part of the Layer 3 network used for interconnections among the edge nodes to the border nodes. In case of a three-tier campus design using a core, distribution, and access, the intermediate nodes are the equivalent of distribution switches, though the number of intermediate nodes is not limited to a single layer of devices. Intermediate nodes route and transport IP traffic inside the fabric. No VXLAN encapsulation/de-encapsulation or LISP control plane messages are required from an intermediate node, which has only the additional fabric MTU requirement to accommodate the larger-size IP packets encapsulated with VXLAN information.

Border node

The fabric border nodes serve as the gateway between the SD-Access fabric site and the networks external to the fabric. The fabric border node is responsible for network virtualization interworking and SGT propagation from the fabric to the rest of the network. The fabric border nodes can be configured as an ***internal border***, operating as the gateway for specific network addresses such as a shared services or data center network, or as an ***external border***, useful as a common exit point from a fabric, such as for the rest of an enterprise network along with the Internet. Border nodes can also have a combined role as an ***anywhere border*** (both internal and external border).

Border nodes implement the following functions:

- **Advertisement of EID subnets–**SD-Access configures Border Gateway Protocol (BGP) as the preferred routing protocol used to advertise the EID prefixes outside of the fabric and traffic destined to EID subnets from outside the fabric goes through the border nodes. These EID prefixes appear only on the routing tables at the border–throughout the rest of the fabric, the EID information is accessed using the fabric control plane.
- **Fabric domain exit point–**The external fabric border is the gateway of last resort for the fabric edge nodes. This is implemented using LISP Proxy Tunnel Router functionality. Also possible are internal fabric borders connected to networks with a well-defined set of IP subnets, adding the requirement to advertise those subnets into the fabric.
- **Mapping of LISP instance to VRF–**The fabric border can extend network virtualization from inside the fabric to outside the fabric by using external VRF instances in order to preserve the virtualization.
- **Policy mapping–**The fabric border node also maps SGT information from within the fabric to be appropriately maintained when exiting that fabric. SGT information is propagated from the fabric border node to the network external to the fabric either by transporting the tags to Cisco TrustSec-aware devices using SGT Exchange Protocol (SXP) or by directly mapping SGTs into the Cisco metadata field in a packet, using inline tagging capabilities implemented for connections to the border node.

Tech tip

The roles of borders are expanding to scale to larger distributed campus deployments with local site services interconnected with a transit control plane and managed by Cisco DNA Center. Check your release notes for general availability of these new roles and features.

Fabric in a box

For sites where a single switch or switch stack (examples: Catalyst 9400, Catalyst 9300) is supporting all the Ethernet connectivity at that site, SD-Access is available without having to deploy separate devices for each fabric role. Create a ***fabric in a box*** by assigning control plane node, edge node, and border node functionality to a single switch device. Because no additional fabric devices are required or permitted for the fabric in a box deployment, solution resiliency depends on the redundant switches in a stack or redundant supervisor modules in a chassis.

Fabric wireless LAN controller

The fabric WLC integrates with the control plane for wireless and the fabric control plane. Both fabric WLCs and non-fabric WLCs provide AP image and configuration management, client session management, and mobility services. Fabric WLCs provide additional services for fabric integration by registering MAC addresses of wireless clients into the host tracking database of the fabric control plane during wireless client join events and by supplying fabric edge RLOC location updates during client roam events.

A key difference with non-fabric WLC behavior is that fabric WLCs are not active participants in the data plane traffic-forwarding role for the SSIDs that are fabric enabled—fabric mode APs directly forward traffic through the fabric for those SSIDs.

Typically, the fabric WLC devices connect to a shared services distribution or data center outside of the fabric and fabric border, which means that their management IP address exists in the global routing table. For the wireless APs to establish a CAPWAP tunnel for WLC management, the APs must be in a VN that has access to the external device. In the SD-Access solution, Cisco DNA Center configures wireless APs to reside within the VRF named INFRA_VRF, which maps to the global routing table, avoiding the need for route leaking or ***fusion router*** (multi-VRF router selectively sharing routing information) services to establish connectivity.

Fabric mode access points

The fabric mode APs are Cisco 802.11AC Wave 2 and Wave 1 APs associated with the fabric WLC that have been configured with one or more fabric-enabled SSIDs. Fabric mode APs continue to support the same 802.11ac wireless media services that traditional APs support; apply AVC, quality of service (QoS), and other wireless policies, and establish the CAPWAP control plane to the fabric WLC. Fabric APs join as local-mode APs and must be directly connected to the fabric edge node switch to enable fabric registration events, including RLOC assignment via the fabric WLC. The APs are recognized by the fabric edge nodes as special wired hosts and assigned to a unique overlay network with in a common EID space across a fabric. The assignment allows management simplification by using a single subnet to cover the AP infrastructure at a fabric site.

When wireless clients connect to a fabric mode AP and authenticate into the fabric-enabled wireless LAN, the WLC updates the fabric mode AP with the client Layer 2 VNI and an SGT supplied by ISE. Then the WLC registers the wireless client Layer 2 EID into the control plane, acting as a proxy for the egress fabric edge node switch. After the initial connectivity is established, the AP uses the Layer 2 VNI information to VXLAN-encapsulate wireless client communication on the Ethernet connection to the directly-connected fabric edge switch. The fabric edge switch maps the client traffic into the appropriate VLAN interface associated with the VNI for forwarding across the fabric and registers the wireless client IP addresses with the control plane database.

Identity Services Engine

Cisco ISE is a secure network access platform enabling increased management awareness, control, and consistency for users and devices accessing an organization's network. ISE is an integral part of SD-Access for policy implementation, enabling dynamic mapping of users and devices to scalable groups and simplifying end-to-end security policy enforcement. Within ISE, users and devices are shown in a simple and flexible interface. ISE integrates with Cisco DNA Center by using Cisco Platform Exchange Grid (pxGrid) and REST APIs for exchange of client information and automation of fabric-related configurations on ISE. The SD-Access solution integrates Cisco TrustSec by supporting group-based policy end-to-end, including SGT information in the VXLAN headers for data plane traffic, while supporting multiple VNs using unique VNI assignments. Groups, policy, Authentication, Authorization, and Accounting (AAA) services, and endpoint profiling are driven by ISE and orchestrated by Cisco DNA Center's policy authoring workflows.

Scalable groups are identified by the SGT, a 16-bit value that is transmitted in the VXLAN header. SGTs are centrally defined, managed, and administered by Cisco ISE. ISE and Cisco DNA Center are tightly integrated through REST APIs, with management of the policies driven by Cisco DNA Center.

ISE supports standalone and distributed deployment models. Additionally, multiple distributed nodes can be deployed together supporting failover resiliency. The range of options allows support for hundreds of thousands of endpoint devices, with a subset of the devices used for SD-Access to the limits described later in the guide. Minimally, a basic two-node ISE deployment is recommended for SD-Access deployments, with each node running all services for redundancy.

SD-Access fabric edge node switches send authentication requests to the Policy Services Node (PSN) persona running on ISE. In the case of a standalone deployment, with or without node redundancy, that PSN persona is referenced by a single IP address. An ISE distributed model uses multiple active PSN personas, each with a unique address. All of the PSN addresses are learned by Cisco DNA Center, and the Cisco DNA Center user maps fabric edge node switches to the PSN that supports each edge node.

Transit

A transit network can be either SD-Access transit to interconnect sites with the native SD-Access encapsulation and functionality or can be an IP-based transit. IP transits offer IP connectivity without native SD-Access encapsulation and functionality, potentially requiring additional VRF and SGT mapping for stitching together the macro and micro segmentation needs between sites.

Use SD-Access transits to create larger fabric domains sharing a common policy and provisioning by interconnecting a fabric site with one or more additional sites. Each fabric site has an independent set of control plane nodes, edge nodes, border nodes, wireless LAN controllers, and ISE nodes. By virtue of each site having all components required to make it independently survivable, the SD-Access transit is useful for building a ***distributed campus*** — multiple independent fabrics representing different buildings at a location or as part of a metropolitan network. The key consideration for the distributed campus design using SD-Access transit is that the network between fabric sites and to Cisco DNA Center should be created with campus-like connectivity. The connections should be high-bandwidth (Ethernet full port speed with no sub-rate services), low latency, and should accommodate the MTU setting used for SD-Access in the campus network (typically 9100 bytes).

You create an SD-Access transit by associating it with two things: a network that has connectivity to the fabric sites that are to be included as part of the larger fabric domain, and a control plane node dedicated to the transit functionality. For resiliency, you should use redundant network connections and dedicate a second control plane node to the transit.

Cisco DNA Center

At the heart of automation of the SD-Access solution is Cisco DNA Center. SD-Access is enabled with an application package that runs as part of the Cisco DNA Center software for designing, provisioning, applying policy, and facilitating the creation of an intelligent campus wired and wireless network with assurance.

Figure 5. Cisco DNA Center Dashboard

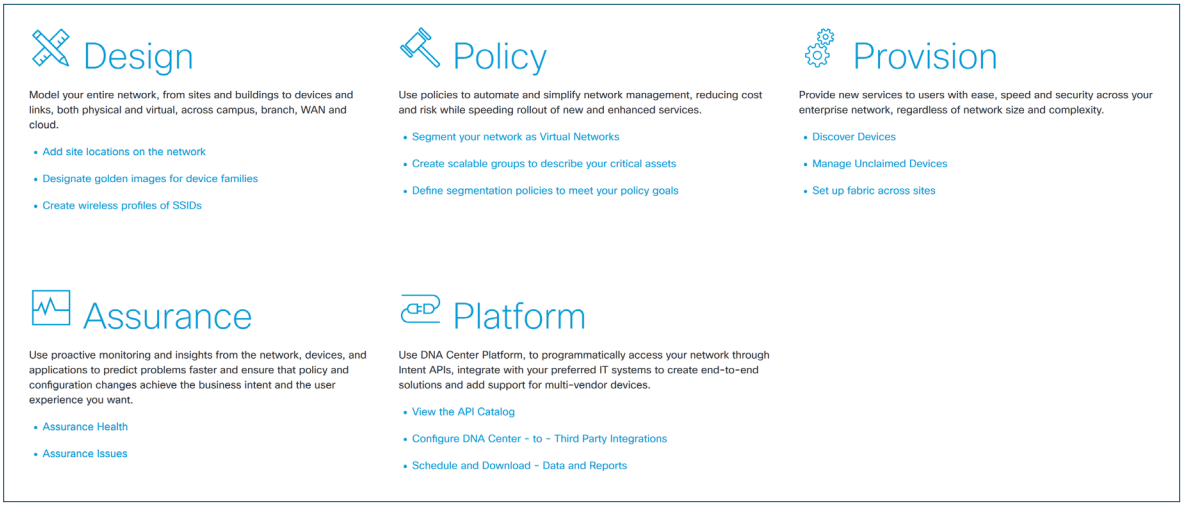

Cisco DNA Center centrally manages major configuration and operations workflow areas.

- **Design**—Configures device global settings, network site profiles for physical device inventory, DNS, DHCP, IP addressing, software image management, plug-and-play, and user access.
- **Policy**—Defines business intent for provisioning into the network, including creation of virtual networks, assignment of endpoints to virtual networks, and policy contract definition for groups.
- **Provision**—Provisions devices for management and creates fabric domains, control plane nodes, border nodes, edge nodes, fabric wireless, Cisco Unified Wireless Network wireless, transit and external connectivity.
- **Assurance**—Enables proactive monitoring and insights to confirm user experience meets configured intent, using network, client, and application health dashboards, issue management, and sensor-driven testing.
- **Platform**—Allows programmatic access to the network and system integration with third-party systems using APIs, using feature set bundles, configurations, a runtime dashboard, and a developer toolkit.

Cisco DNA Center supports integration using APIs. For example, Infoblox and Bluecat IP address management and policy enforcement integration with ISE are available through Cisco DNA Center. A comprehensive set of northbound REST APIs enables automation, integration, and innovation.

- All controller functionality is exposed through northbound REST APIs.
- Organizations and ecosystem partners can easily build new applications.
- All northbound REST API requests are governed by the controller RBAC mechanism.

Cisco DNA Center is key to enabling automation of device deployments into the network providing the speed and consistency required for operational efficiency. Organizations using Cisco DNA Center benefit from lower cost and reduced risk when deploying and maintaining their networks.

Cisco DNA Center Appliance

The Cisco DNA Center software with the SD-Access application package is designed to run on the Cisco DNA Center Appliance. The appliance is available in form factors sized to support not only the SD-Access application but also network assurance and new capabilities as they are available.

When you deploy only a single Cisco DNA Center Appliance, and then that node becomes unavailable, an SD-Access network still functions, but automated provisioning capabilities are lost until the single node availability is restored. For high-availability purposes, configure three Cisco DNA Center appliances to form a three-node cluster. The Cisco DNA Center cluster is accessed using a single GUI interface hosted on a virtual IP, which is serviced by the resilient nodes within the cluster.

Within a three-node cluster, you enable service distribution to automatically provide distributed processing, database replication, security replication, and file synchronization. Software upgrades are also automatically replicated across the nodes in a cluster. A cluster will survive the loss of a single host and requires two hosts to remain operational. Some maintenance operations, such as software upgrades and file restoration from backup, are restricted until the full three-node cluster is restored.

First generation M4-based appliances (DN1-HW-APL) are clustered with the same appliance type and with the second generation M5-based appliances only if they have a compatible CPU core configuration (DN2-HW-APL).

For additional information about the Cisco DNA Center Appliance and compatible versions of hardware and software with clustering, visit Cisco.com and search for "Cisco Digital Network Architecture Center Appliance."

SD-Access design considerations

Designing for an SD-Access fabric is extremely flexible to fit many environments, which means it is not a one-design-fits-all proposition. The scale of a fabric can be as small as an access-distribution block or as big as a three-tier campus deployment. In a single network, multiple fabrics can be deployed as long as the fabric elements are assigned to a single fabric only.

Platform roles and recommendations

Choose your SD-Access network platform based on capacity and capabilities required by the network, considering the recommended functional roles. Roles tested during the development of this Cisco Validated Design (CVD) guide are noted in the tables. Sizing and scale values for the roles are listed in subsequent tables in the section.

Tech tip

To achieve the functionality shown, you must meet minimum software release requirements. Some of the platforms may include functionality not specifically tested as part of CVD verification. For more information, see the software release notes for your platform and refer to CVD deployment guides for as-tested code versions.

Tech tip

The SD-Access solution has specific software version requirements for WLAN control-plane support when using the Cisco Catalyst 6500 and 6800 Series. For the latest information, see the software release notes for your platform.

Table 1. SD-Access 1.2 switching platforms and deployment capabilities

Platform	Supported supervisor	Supported fabric-facing interfaces	Edge node	Border node	Control plane node
Cisco Catalyst 9500 Series	-	Onboard ports and network module ports	No	Yes—CVD verified	Yes—CVD verified
Cisco Catalyst 9400 Series	Supervisor Engine-1	Supervisor and line card ports	Yes—CVD verified	No	No
Cisco Catalyst 9400 Series	Supervisor Engine-1 XL	Supervisor and line card ports	Yes	Yes	Yes
Cisco Catalyst 9300 Series	—	Onboard ports and network module ports	Yes	Yes	Yes
Cisco Catalyst 9200 Series	—	Onboard ports and network module ports	Yes	No	No
Cisco Catalyst 3850 Series	—	Onboard ports and 10G/40G network module ports	Yes—CVD verified	Yes—3850 XS 10-Gbps fiber versions CVD verified (small-scale deployments)	Yes—3850 XS 10-Gbps fiber versions CVD verified (small-scale deployments)
Cisco Catalyst 3650 Series	—	Onboard ports and uplink ports	Yes—CVD verified	No	No
Cisco Catalyst 4500-E Series	Supervisor 8-E	Supervisor uplink ports	Yes—CVD verified	No	No
Cisco Catalyst 4500-E Series	Supervisor 9-E	Supervisor uplink ports	Yes	No	No

Platform	Supported supervisor	Supported fabric-facing interfaces	Edge node	Border node	Control plane node
Cisco Catalyst 6807-XL Switch and 6500-E Series	Supervisor 6T and Supervisor 2T	Supervisor uplink ports (Supervisor 6T only) C6800 Series WS-X6900 Series	No	Yes–CVD verified	Yes–CVD verified
Cisco Catalyst 6880-X and 6840-X Series	–	Onboard ports and port card ports	No	Yes–CVD verified	Yes–CVD verified
Cisco Nexus® 7700 Series	Supervisor 2E	M3 Series	No	Yes--External Border Only (For high-density 40G/100G deployments)	No (requires adding and configuring a dedicated external control plane node)

Table 2. SD-Access 1.2 routing and wireless platforms and deployment capabilities

Platform	Supported fabric-facing interfaces	Edge node	Border node	Control plane node
Cloud Services Router 1000V Series	-	-	-	Yes—CVD verified
Cisco 4400 and 4300 Series Integrated Services Routers	Onboard LAN ports and routed LAN Network Interface Module and enhanced service Module Ports	No	Yes—CVD verified	Yes—CVD verified
Cisco ASR 1000-X and 1000-HX Series Aggregation Services Routers	Onboard LAN ports, Ethernet line cards, and Ethernet shared port adapters	No	Yes	Yes—CVD verified (large-scale deployments)
Cisco 8540, 5520, and 3504 Series Wireless LAN Controllers	Via associated 802.11ac Wave 2 and Wave 1 fabric mode AP network ports	No	No	Proxy to control plane for wireless clients—CVD verified

Physical topologies

SD-Access topologies should follow the same design principles and best practices associated with a hierarchical design by splitting the network into modular groups, as described in the Campus LAN and Wireless LAN Design Guide. You create design elements that can be replicated throughout the network by using modular designs. The following example shows the physical topology of a three-tier campus design in which all nodes are dual homed with equal-cost links that will provide for load-balancing, redundancy, and fast convergence. Though the topology depicts the border at a campus core, the border can be configured separate from the core at another aggregation point. A cross link at each aggregation layer is used for optimal routing in case of an uplink failure.

Figure 6. Three-tier SD-Access fabric topology

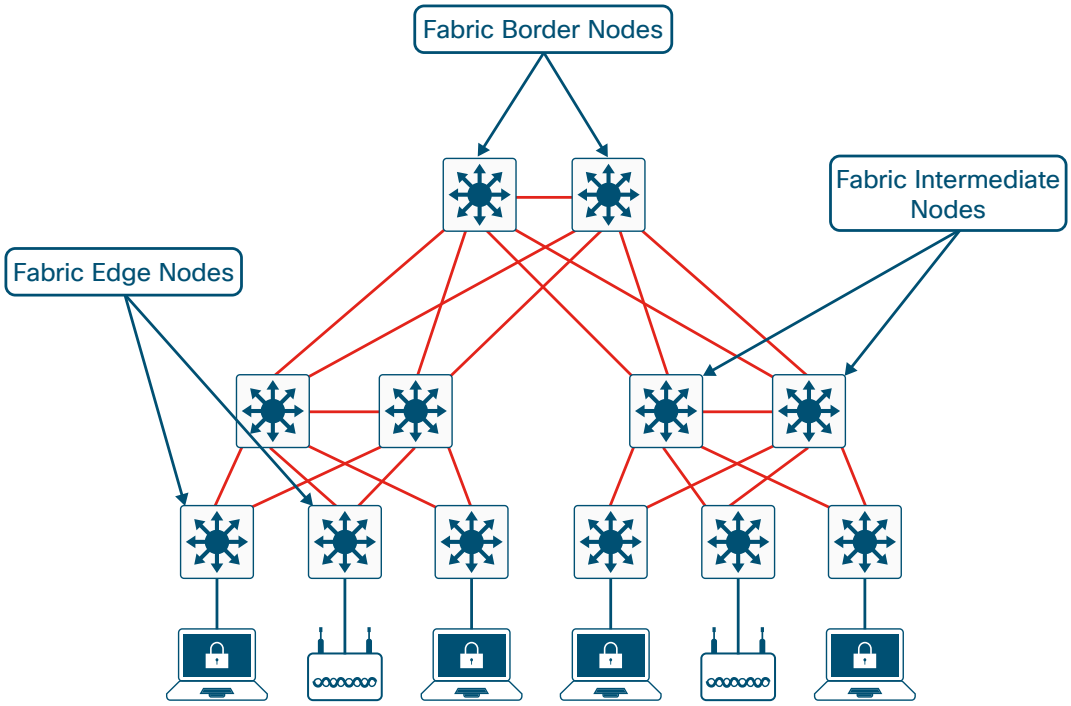

For smaller deployments, an SD-Access fabric can be implemented using a two-tier design. The same design principles should be applied but without the need for an aggregation layer implemented by intermediate nodes.

Figure 7. Two-tier SD-Access fabric topology

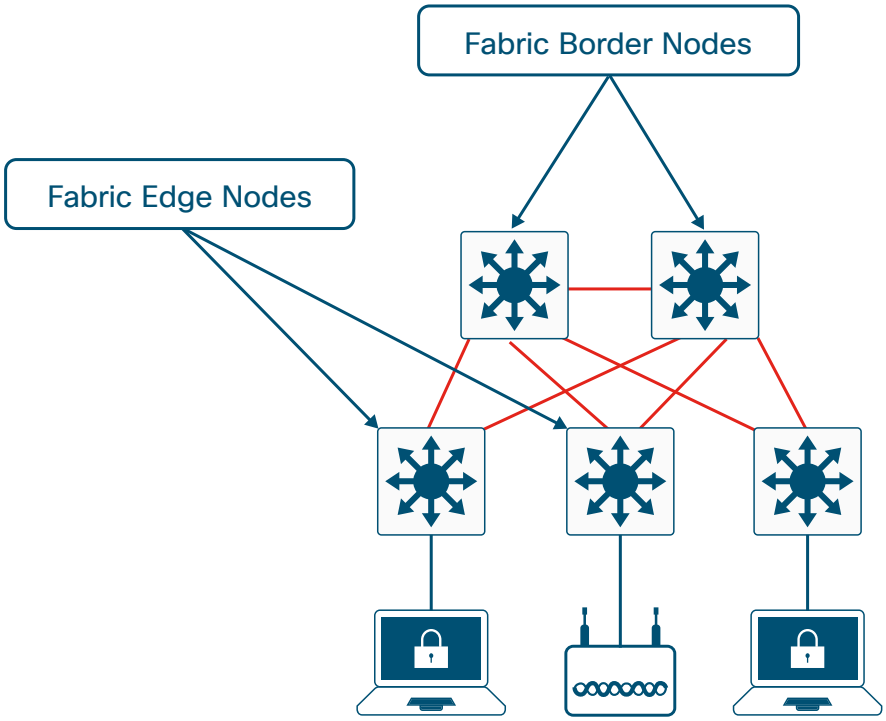

In general, SD-Access topologies should be deployed as spoke networks with the fabric border node at the exit point hub for the spokes, although other physical topologies can be used. Topologies in which the fabric is a transit network should be planned carefully in order to ensure optimal forwarding. If the border node is implemented at a node that is not the aggregation point for exiting traffic, sub-optimal routing results when traffic exits the fabric at the border and then doubles back to the actual aggregation point.

Underlay network design

Having a well-designed underlay network will ensure the stability, performance, and efficient utilization of the SD-Access network. Automation for deploying the underlay is available using Cisco DNA Center.

Underlay networks for the fabric have the following design requirements:

- **Layer 3 to the access design–**The use of a Layer 3 routed network for the fabric provides the highest level of availability without the need to use loop avoidance protocols or interface bundling techniques.
- **Increase default MTU–**The VXLAN header adds 50 and optionally 54 bytes of encapsulation overhead. Some Ethernet switches support a Maximum Transmission Unit (MTU) of 9216 while others may have an MTU of 9196 or smaller. Given that server MTUs typically go up to 9,000 bytes, enabling a network wide MTU of 9100 ensures that Ethernet jumbo frames can be transported without any fragmentation inside and outside of the fabric.

- **Use point-to-point links—**Point-to-point links provide the quickest convergence times because they eliminate the need to wait for the upper layer protocol timeouts typical of more complex topologies. Combining point-to-point links with the recommended physical topology design provides fast convergence after a link failure. The fast convergence is a benefit of quick link failure detection triggering immediate use of alternate topology entries preexisting in the routing and forwarding table. Implement the point-to-point links using optical technology and not copper, because optical interfaces offer the fastest failure detection times to improve convergence. Bidirectional Forwarding Detection should be used to enhance fault detection and convergence characteristics.
- **Dedicated IGP process for the fabric—**The underlay network of the fabric only requires IP reachability from the fabric edge to the border node. In a fabric deployment, a single area IGP design can be implemented with a dedicated IGP process implemented at the SD-Access fabric. Address space used for links inside the fabric does not need to be advertised outside of the fabric and can be reused across multiple fabrics.
- **Loopback propagation—**The loopback addresses assigned to the underlay devices need to propagate outside of the fabric in order to establish connectivity to infrastructure services such as fabric control plane nodes, DNS, DHCP, and AAA. As a best practice, use /32 host masks. Apply tags to the host routes as they are introduced into the network. Reference the tags to redistribute and propagate only the tagged loopback routes. This is an easy way to selectively propagate routes outside of the fabric and avoid maintaining prefix lists.

LAN Automation

You can fully automate the configuration of the underlay by using LAN Automation services in Cisco DNA Center. In non-greenfield deployment cases, you manually create the underlay. Manual underlays allow variations from the automated underlay deployment (for example, a different IGP could be chosen), but the previously listed underlay design principles still apply.

To automate the deployment of the underlay, Cisco DNA Center uses IP to access a Cisco Network Plug and Play seed device directly connected to the new underlay devices. The remaining devices are accessed using hop-by-hop CDP discovery and provisioning.

Overlay fabric design

In the SD-Access fabric, the overlay networks are used for transporting user traffic within the fabric. The fabric encapsulation also carries scalable group information that can be used for traffic segmentation inside the overlay. The following design considerations should be taken into account when deploying virtual networks:

- **Virtualize as needed for network requirements—**Segmentation using SGTs allows for simple-to-manage group-based policies and enables granular data plane isolation between groups of endpoints within a virtualized network, accommodating many network policy requirements. Using SGTs also enables scalable deployment of policy, without having to do cumbersome updates for policies based on IP addresses, which can be prone to breakage. VNs support the transport of SGTs for group segmentation. Use virtual networks when requirements dictate isolation at both the data plane and control plane. For those cases, if communication is required between different virtual networks, you use an external firewall or other device to enable inter-VN communication. You can choose either or both options to match your requirements.

- **Reduce subnets and simplify DHCP management—**In the overlay, IP subnets can be stretched across the fabric without flooding issues that can happen on large Layer 2 networks. Use fewer subnets and DHCP scopes for simpler IP addressing and DHCP scope management. Subnets are sized according to the services that they support versus being constrained by the location of a gateway. Enabling optional broadcast flooding features can limit the subnet size based on the additional bandwidth and endpoint processing requirements for the traffic mix within a specific deployment. Subnets are sized according to the services that they support versus being constrained by the location of a gateway. Enabling optional broadcast flooding features can limit the subnet size based on the additional bandwidth and endpoint processing requirements for the traffic mix within a specific deployment.
- **Avoid overlapping IP subnets—**Different overlay networks can support overlapping address space, but be aware that most deployments require shared services across all VNs and other inter-VN communication. Avoid overlapping address space so that the additional operational complexity of adding a network address translation device is not required for shared services and inter-VN communication.

Fabric control plane design

The fabric control plane contains the database used to identify endpoint location for the fabric elements. This is a central and critical function for the fabric to operate. A control plane that is overloaded and slow to respond results in application traffic loss on initial packets. If the fabric control plane is down, endpoints inside the fabric fail to establish communication to remote endpoints that do not already exist in the local database.

Cisco DNA Center automates the configuration of the control plane functionality. For redundancy, you should deploy two control plane nodes to ensure high availability of the fabric, as a result of each node containing a duplicate copy of control plane information. The devices supporting the control plane should be chosen to support the HTDB, CPU, and memory needs for an organization based on fabric endpoints.

If the chosen border nodes support the anticipated endpoint scale requirements for a fabric, it is logical to collocate the fabric control plane functionality with the border nodes. However, if the collocated option is not possible (example: Nexus 7700 borders lacking the control plane node function or endpoint scale requirements exceeding the platform capabilities), then you can add devices dedicated to this functionality, such as physical routers or virtual routers at a fabric site.

Fabric border design

The fabric border design is dependent on how the fabric is connected to the outside network. VNs inside the fabric should map to VRF-Lite instances outside the fabric. Depending on where shared services are placed in the network the border design will have to be adapted. For more information, see “End-to-End Virtualization Considerations,” later in this guide.

Larger distributed campus deployments with local site services are possible when interconnected with a transit control plane. You can search for guidance for this topic after these new roles are a generally available feature.

Infrastructure services

SD-Access does not require any changes to existing infrastructure services, though fabric border devices have implemented capabilities to handle the DHCP relay functionality differences assisting fabric deployments. In a typical DHCP relay design, the unique gateway IP address determines the subnet address assignment for an endpoint, in addition to the location where the DHCP server should direct the offered address. In a fabric overlay network, that gateway is not unique—the same anycast IP address exists across all fabric edge devices within an overlay. Without special handling either at the border or by the DHCP server itself, the DHCP offer returning from the DHCP server through the border may not be relayed to the correct fabric edge switch where the DHCP request originated.

To identify the specific DHCP relay source, Cisco DNA Center automates the configuration of the relay agent at the fabric edge with DHCP option 82 including the information option for circuit ID insertion. Adding the information provides additional sub-options to identify the specific source relay agent. DHCP relay information embedded in the circuit ID is used as the destination for DHCP offer replies to the requestor—either by a fabric border with advanced DHCP border relay capabilities or, alternatively, by the DHCP server itself.

Using a border with the advanced DHCP border relay capability allows DHCP server scope configuration to remain unchanged for scopes covering fabric endpoints versus standard non-fabric scope creation. When you are using border nodes with this additional DHCP capability, the borders inspect the DHCP offers returning from the DHCP server. The border node receiving the DHCP offer references the embedded circuit ID and directs the DHCP offers to the correct relay destination.

Fabric wireless integration

As described earlier, when you integrate a fabric WLC and fabric mode APs into the SD-Access architecture, fabric WLCs are not active participants in the data plane traffic-forwarding role, and fabric mode APs are responsible for delivering wireless client traffic into and out of the wired fabric. The WLC control plane still keeps many of the characteristics of a local-mode controller, including the requirement to have a low-latency connection between the WLC and the APs. The colocation requirement precludes a fabric WLC from being the controller for fabric mode APs at a remote site across a typical WAN. As a result, a remote site desiring SD-Access with integrated wireless needs to have a local controller at that site.

When integrating wireless into SD-Access, another consideration is fabric WLC placement and connectivity. In larger scale deployments, WLCs typically connect to a shared services distribution block that is part of the underlay. The preferred distribution block has chassis redundancy and also the capability to support L2 multichassis EtherChannel connections for link and platform redundancy to the WLCs. Often Virtual Switching System or switch-stacking is used to accomplish these goals.

APs connect into a pre-defined VRF, named INFRA_VRF. The VRF has connectivity into the global routing table, allowing the WLC connection into the network to remain unchanged, while still being able to manage APs at the edge of a fabric domain.

Figure 8. Wireless components integrated into SD-Access

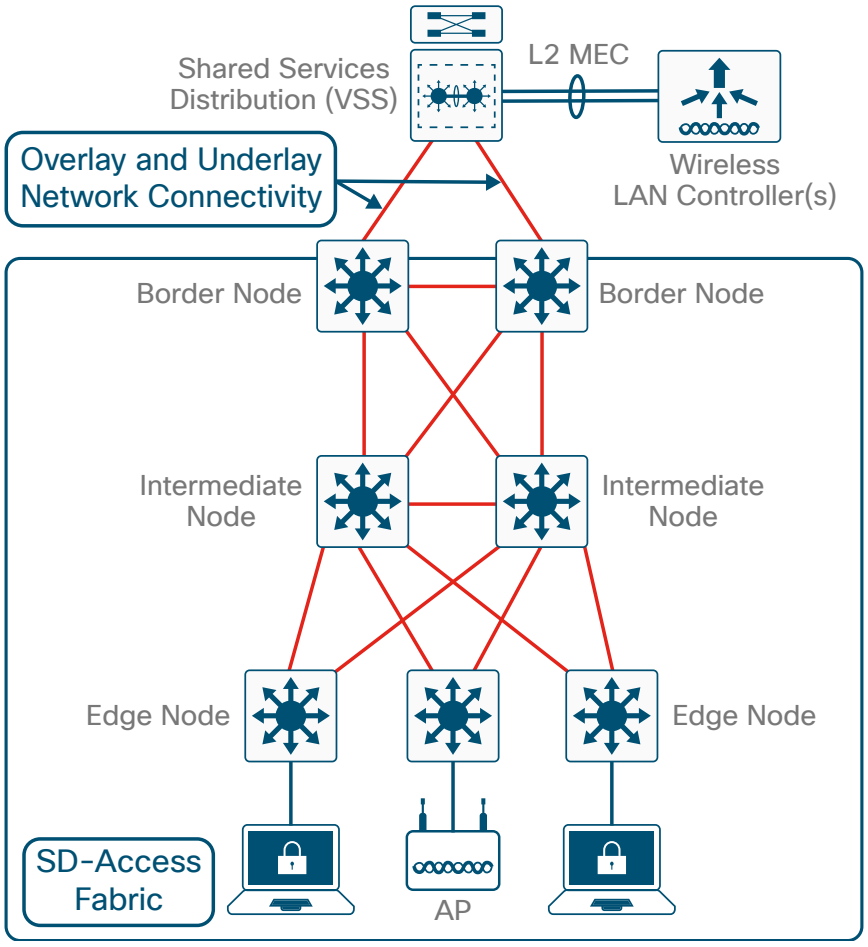

Non-fabric centralized wireless option

In cases where you cannot dedicate WLCs and APs in a seamless roaming area to participate in fabric, a traditional Cisco Unified Wireless Network design model, also known as a ***local-mode model***, is an option. SD-Access is compatible with Cisco Unified Wireless Network "over the top" as a non-native service option, without the benefits of fabric integration.

An over-the-top centralized design still provides IP address management, simplified configuration and troubleshooting, and roaming at scale. In a centralized model, the WLAN controller and APs are both located within the same site. You can connect the WLAN controller to a data center services block or a dedicated block adjacent to the campus core. Wireless traffic between WLAN clients and the LAN is tunneled by using the Control And Provisioning of Wireless Access Points (CAPWAP) protocol between the controller and the AP. APs can reside inside or outside the fabric without any change to the recommended centralized WLAN design, keeping in mind that the benefits of fabric and SD-Access are not extended to and integrated with the wireless when the fabric is used only as an over-the-top transport.

For additional information about campus wireless design, see the Campus LAN and Wireless LAN Design Guide.

Tech tip

Converged Access and Cisco FlexConnect® are not supported inside the SD-Access fabric.

Mixed SD-Access Wireless and centralized wireless option

Many organizations may deploy SD-Access with centralized wireless over-the-top as a first transition step before integrating SD-Access Wireless into the fabric. For this case, an organization should dedicate a WLC for enabling SD-Access Wireless. A WLC dedicated to SD-Access allows use of the same SSID in both the fabric and non-fabric domains, without modifying the existing centralized deployment with changes such as new software versions and AP group configurations.

For organizations wishing to deploy both centralized and SD-Access Wireless services as a migration stage without the requirement to keep the existing centralized wireless software and configuration unchanged, Cisco DNA Center can automate a new installation supporting both on the same WLC. In this case, the new installation from Cisco DNA Center does not take into consideration existing running configurations, but automates the creation of a fresh configuration.

Security policy design

Security policies vary by organization—it is not possible to define one-size-fits-all security design. Security designs are driven by information security policies and legal compliance. The planning phase for a security design is key to ensuring the right balance of security and user experience. You should consider the following aspects designing your security policy for the SD-Access network:

- **Openness of the network—**Some organizations allow only organization-issued devices in the network, and some support a "Bring Your Own Device" approach. Alternatively, you can balance user choice and allow easier-to-manage endpoint security by deploying a "Choose Your Own Device" model in which a list of IT-approved endpoints is offered to the users for business use. And an identity-based approach is also possible in which the network security policies can be deployed depending of the device ownership. For example, organization-issued devices may get group-based access, while personal devices may get Internet access only.
- **Identity management—**In the simplest form, identity management can be a username and password used for authenticating users. Adding embedded security functions and application visibility in the network devices provides telemetry for advanced policy definitions that can include additional context such as physical location, device used, type of access network, application used, and time of day.
- **Authentication, authorization, and accounting policies—*Authentication*** is the process of establishing and confirming the identity of a client requesting access to the network. ***Authorization*** is the process of authorizing the endpoint to some set of network resources. Segmentation policies do not necessarily have to be enforced at the access layer, and can be deployed in multiple locations. Policies are enforced with the use of SGACLs for segmentation within VNs, and dynamic VLAN assignment for mapping endpoints into VNs at the fabric edge node. Event logs, ACL hit counters, and similar standard accounting tools are available to enhance visibility.
- **Endpoint security—**Endpoints can be infected with malware, compromising data and create network disruptions. Malware detection, endpoint management, and data exports from the network devices provide insight into endpoint behavior. Tight integration of the network with security appliances and analytics platforms enable the network with the necessary intelligence to quarantine and help remediate compromised devices.
- **Data integrity and confidentiality—**Network segmentation using VNs can control access to applications, such as separating employee transactions from IoT traffic; encryption of the data path in the switching environment using IEEE 802.1AE MACsec is used to provide encryption at Layer 2 to prevent eavesdropping and to ensure that the data cannot be modified.

- **Network device security–**Hardening the security of the network devices is essential because they are common targets for security attacks. The use of the most secure device management options, such as enabling device authentication using TACACS+ and disabling unnecessary services, are best practices to ensure the network devices are secured.

Enabling group-based segmentation within each virtual network allows for simplified hierarchical network policies. Network-level policy scopes of isolated control and data planes are possible using virtual networks, and group-level policy scopes are possible using SGTs within VNs, enabling common policy application across the wired and wireless fabric.

SGTs provide the capability to tag endpoint traffic based on a role or function within the network and subject to role-based policies or SGACLs centrally defined at ISE. In many deployments, Active Directory is used as the identity store for user accounts, credentials, and group membership information. Upon successful authorization, endpoints can be classified based on that information and assigned the appropriate scalable group assignments. These scalable groups can then be used to create segmentation policies and virtual network assignment rules.

SGT information is carried across the network in several forms:

- **Inside the SD-Access network–**The SD-Access fabric header transports SGT information. Fabric edge nodes and border nodes can enforce SGACLs to enforce the security policy.
- **Outside of the fabric on a device with Cisco TrustSec capability–**Inline devices with Cisco TrustSec capability carry the SGT information in a CMD header on the Layer 2 frame. This is the recommended mode of transport outside of the SD-Access network.
- **Outside of the fabric over devices without Cisco TrustSec capability–**SXP allows the transport of SGTs over a TCP connection. This can be used to bypass network devices that do not support SGT inline.

For additional information about Cisco TrustSec, see cisco.com/go/trustsec.

SD-Access design sizing considerations

In addition to the platform role recommendations listed in Tables 1 and 2, when designing a network consider the Cisco Validated Design (CVD) tested values in the following tables along with hardware and software release notes for a deployment. If you need to manage more endpoints or fabrics, additional Cisco DNA Center deployments may be necessary.

Tech tip

The numbers listed show as-tested values during CVD validation of the SD-Access 1.2 solution. The actual limits are often higher when considered without other constraints. Check the single-platform scale limitations in the next section and hardware and software release notes for additional limits of specific components.

Table 3. Cisco Validated Design tested values for Cisco DNA Center management of SD-Access

SD-Access component	CVD tested value for single-node cluster using DN1-HW-APL (maximum may be higher)
Clients—across all fabric domains (wired clients and wireless clients)	25,000 validated
Edge, border, and control plane fabric nodes—across all fabric domains (switches/switch stacks, routers, WLCs)	500 validated
All fabric node types across all fabric domains, including fabric intermediate nodes and edge, border and control plane nodes	1,000 validated
Access points—across all fabric domains (each AP counts as an endpoint)	1,500 validated
IP pools—in a single fabric domain or split across all fabric domains	200 pools across 500 edges validated
Sites—each item in the site hierarchy is additive (site, building, floor)	200 validated
Fabric domains	10 validated
Scalable group tags—across all fabric domains	4,000 validated
Access control policies—across all fabric domains	300 validated
Contracts—across all fabric domains	500 validated

Table 4. Cisco Validated Design tested values for SD-Access virtual networks by platform and role

	Fabric edge platform validation	Fabric border validation
Cisco Catalyst 9200 Series	4 validated	
Cisco Catalyst 3850 and 3650 Series	40 validated	-
Cisco Catalyst 3850 XS (10 Gbps fiber)	-	40 validated
Cisco Catalyst 9300 Series	40 validated	40 validated
Cisco Catalyst 9400 Series with Supervisor Engine-1	40 validated	-
Cisco Catalyst 4500 Series with Supervisor 8-E	40 validated	-
Cisco Catalyst 9500 Series	40 validated	40 validated
Cisco Catalyst 6807-XL with Supervisor 6T	-	40 validated
Cisco Nexus 7700 Supervisor 2E	-	40 validated
ASR 1000 Series and 4000 Series ISRs	-	40 validated

SD-Access single-platform scale considerations

Beyond the end-to-end network validation of the CVD, consider the scaling constraints of single platforms in isolation before deploying an organization's network. Use this data when selecting platforms to use during planning for current and future growth of your network.

Cisco DNA Center numbers are per instance, which can be a single-node cluster or a three-node cluster. The maximum numbers are either the platform absolute limits or the recommended limit based on the most current testing of a single platform.

Table 5. Maximum scale recommendations—Cisco DNA Center (DN1-HW-APL)

SD-Access element	Maximum
Fabric domains	10
Fabric sites in one fabric domain or split across multiple domains	200
APs connected to fabric edge	4,000
Wired endpoints connected to fabric edge (includes APs counted as wired endpoints)	25,000
Fabric nodes including border, edge (switch or switch stack), and WLC	500
Non-fabric nodes including intermediate, subtended, and routers	1,000
Control plane nodes per fabric site	2
Default border nodes per fabric site	4
IP pools across all fabric domains	200
Sites across all fabric domains	200
Scalable groups	4,000
Policies	1,000
Contracts	500

Table 6. Single-platform maximum scale recommendations–SD-Access edge node

	Virtual networks	Attached wired endpoints	SGT/DGT table	SGACLs – security ACEs
Cisco Catalyst 3650 Series	64	2,000	4,000	1,500
Cisco Catalyst 3850 Series	64	4,000	4,000	1,500
Cisco Catalyst 9300 Series	256	4,000	8,000	5,000
Cisco Catalyst 4500 Series with Supervisor 8-E	64	4,000	2,000	1,350
Cisco Catalyst 9400 Series with Supervisor Engine-1	256	4,000	8,000	18,000
Cisco Catalyst 9500 Series	256	4,000	8,000	18,000

Table 7. Single-platform maximum scale recommendations–SD-Access border node

	Virtual networks	SGT/ DGT table	SGA-CLs–security ACEs	Fabric control plane entries – border collocated with control plane	IPv4 fabric routes	IPv4 fabric host entries
Cisco Catalyst 3850 XS (10 Gbps fiber)	64	4,000	1,500	3,000	8,000	16,000
Cisco Catalyst 9400 Series with Supervisor Engine-1XL	256	8,000	18,000	80,000	20,000	80,000
Cisco Catalyst 9500 Series	256	8,000	18,000	80,000	48,000	96,000
Cisco Catalyst 6800 Series Supervisor 6T	512–Unicast 100–Multicast	12,000	30,000	25,000	256,000	256,000
Cisco Catalyst 6800 Series Supervisor 6T (XL)	512–Unicast 100–Multicast	30,000	30,000	25,000	1,000,000	1,000,000
Cisco Nexus 7700 Supervisor 2E	256	64,000	64,000	–	32,000 LISP-mapped	32,000 LISP-mapped
ASR 1000 Series (8 GB memory)	4,000	64,000	64,000	200,000	1,000,000	1,000,000
ASR 1000 Series (16 GB memory)	4,000	64,000	64,000	200,000	4,000,000	4,000,000
4000 Series ISRs (8 GB memory)	4,000	64,000s	64,000	100,000	1,000,000	1,000,000
4000 Series ISR (16 GB memory)	4,000	64,000	64,000	100,000	4,000,000	4,000,000
CSR 1000v Series	–	–	–	200,000	–	–

End-to-end design considerations

In a virtualized network, there is full isolation of data and control planes over a shared networking infrastructure. In the case of the SD-Access, a user on one VN is completely isolated and will not be able to communicate with a user on a different VN. The fabric border node is responsible for extending network virtualization beyond the SD-Access fabric. Organizations may have business requirements that call for this type of isolation. Some example of vertical specific use cases where network virtualization maybe useful include:

- **Education–**College campus divided into administrative and student residence networks.
- **Retail–**Isolation for point-of-sale machines supporting payment card industry compliance.
- **Manufacturing–**Isolation for machine-to-machine traffic in manufacturing floors.
- **Healthcare–**Dedicated networks for medical equipment, patient wireless guest access and HIPAA compliance.
- **Enterprise–**Integration of networks during mergers, where overlapping address spaces may exist. Separation of building control systems and video surveillance devices.

Designing for end-to-end network virtualization requires detailed planning to ensure the integrity of the virtual networks. In most cases, there is a need to have some form of shared services that can be reused across multiple virtual networks. It is important that those shared services are deployed correctly to preserve the isolation between different virtual networks sharing those services. The use of a fusion router directly attached to the fabric border provides a mechanism for route leaking shared services prefixes across multiple networks, the use of firewalls provides an additional layer of security and monitoring of traffic between virtual networks. Examples of shared services include:

- **Wireless infrastructure–**Radio frequency performance and cost efficiency is improved using common wireless LANs (single SSID), versus previous inefficient strategies of using multiple SSIDs to separate endpoint communication. Traffic isolation is achieved by assigning dedicated VLANs at the WLC and using dynamic VLAN assignment using 802.1X authentication to map wireless endpoints into their corresponding VNs.
- **DHCP, DNS, and IP address management–**The same set of infrastructure services can be reused as long as they have support for virtualized networks. Special capabilities such as advanced DHCP scope selection criteria, multiple domains, and support for overlapping address space are some of the capabilities required to extend the services beyond a single network.
- **Internet access–**The same set of Internet firewalls can be used for multiple virtual networks. If firewall policies need to be unique for each virtual network, the use of a multi-context firewall is recommended.
- **IP voice/video collaboration services–**When IP phones and other unified communications devices are connected in multiple virtual networks, the call control signaling to the communications manager and the IP traffic between those devices needs to be able to traverse multiple VNs in the infrastructure.

Network virtualization technologies

Extending the SD-Access fabric virtualization beyond the fabric border is enabled using multi-VRF configurations. SD-Access VNs can have 1:1 or N:1 mapping to VRFs outside of the SD-Access fabric. Guidance for virtualizing your end-to-end network is beyond the scope of this guide. However, this section provides a brief introduction to the most commonly used technologies that you can investigate when virtualizing your network.

Device-level virtualization

Within the same device physical device, logical separation capabilities at Layer 2 and Layer 3 can be used to extend virtual networks:

Host-pool communication and virtual LANs

The most basic form of device-level virtualization is isolating network traffic using different virtual LANs (VLANs). This form of virtualization applies to Layer 2 devices and can be extended across switched domains. Also, VLANs are used to virtualize point-to-point links between routers and security appliances that require connectivity to multiple host pools via the same physical interface.

VN communication and virtual routing and forwarding

VRF is a device-level virtualization technology for creating multiple Layer 3 routing tables on the same device. VRFs can be tied to existing Layer 2 domains in order to provide Layer 3 edge functionality to multiple VLANs and also between Layer 3 routed interfaces in order to extend a multiple virtualized control plane over the same set of interfaces. The VN-to-VRF mapping is used to extend and enable VN communication beyond the border.

Path isolation

To maintain isolation on the paths of links interconnecting devices, there are many technology options that provide network virtualization among devices. For SD-Access, the recommended path-isolation technologies are VRF-Lite and MPLS VPN. The number of virtualized networks required typically dictates the design. If you forecast a need for more than a few VRFs, deploying MPLS VPNs simplifies configuration and management.

VRF-Lite end-to-end

VRF-Lite is deployed on a hop-by-hop basis in a campus network, making use of 802.1Q trunks between devices in order to isolate data and control plane for each virtual network. For ease of operation, you should use the same set of VLANs across every hop and use BGP with per-VN address families providing attributes that can be leveraged for easy route-leaking for shared services.

MPLS

Although often considered a service-provider technology, MPLS is common on larger enterprises needing a large number of virtualized networks, most commonly in the WAN but also extended to the campus network. While VRF-Lite is common to most routing platforms, MPLS is not supported across all platforms. A combination of VRF-Lite at the edge with MPLS VPN is another design that could be considered.

Tech tip

The SD-Access solution supports VRF-Lite handoff at the fabric border node. For other options, reference the release notes for your software version to verify support.

Migration to SD-Access

You can readily create SD-Access greenfield networks by adding the infrastructure components, interconnecting them, and using Cisco DNA Center with Cisco Plug and Play features to automate provisioning of the network architecture from the ground up. Migrating an existing network requires some additional planning. Here are some example considerations:

- Migration typically implies that a manual underlay is used. Does an organization's underlay network already include the elements described in the "Underlay Network" section? Or do you have to reconfigure your network into a Layer 3 access model?
- Do the SD-Access components in the network support the desired scale for the target topologies, or do the hardware and software platforms need to be augmented with additional platforms?
- Is the organization ready for changes in IP addressing and DHCP scope management?
- If you plan to enable multiple overlays, what is the strategy for integrating those overlays with common services (for example: Internet, DNS/DHCP, data center applications)?
- Are SGTs already implemented, and where are the policy enforcement points? If SGTs and multiple overlays are used to segment and virtualize within the fabric, what requirements exist for extending them beyond the fabric? Is infrastructure in place to support Cisco TrustSec, VRF-Lite, MPLS, fusion routers, or other technologies necessary to extend and support the segmentation and virtualization?
- Can wireless coverage within a roaming domain be upgraded at a single point in time, or do you need to rely on over-the-top strategies?

There are two primary approaches when migrating an existing network to SD-Access. If many of the existing platforms are to be replaced, and if there is sufficient power, space, and cooling, then building an SD-Access network in parallel may be an option allowing for easy user cutovers. Building a parallel network that is integrated with the existing network is effectively a variation of a greenfield build. Another approach is to do incremental migrations of access switches into an SD-Access fabric. This strategy is appropriate for networks that have equipment capable of supporting SD-Access already in place or where there are environmental constraints.

To assist with network migration, SD-Access supports a Layer 2 border construct that can be used temporarily during a transition phase. Create a Layer 2 border handoff using a single border node connected to the existing traditional Layer 2 access network, where existing Layer 2 access VLANs map into the SD-Access overlays. You can create link redundancy between the single Layer 2 border and the existing Layer 2 access network using EtherChannel, and chassis redundancy on the existing Layer 2 access network using switch stacks, Virtual Switching System, or StackWise Virtual configurations. Support for 4,000 hosts on a Layer 2 border is available in initial releases, and the DHCP services for the fabric support both the fabric and existing non-fabric Layer 2 network at connection time.

For detailed coverage of migration topics, see Software-Defined Access Migration on Cisco.com.

Appendix A–SD-Access fabric details

Fabric data plane

RFC 7348 defines the use of virtual extensible LAN (VXLAN) as a way to overlay a Layer 2 network on top of a Layer 3 network. Using VXLAN, you tag the original Layer 2 frame using UDP/IP over the Layer 3 network. Each overlay network is called a ***VXLAN segment*** and is identified using a 24-bit VXLAN network identifier, which supports up to 16 million VXLAN segments.

Figure 9. RFC 7348 VXLAN header

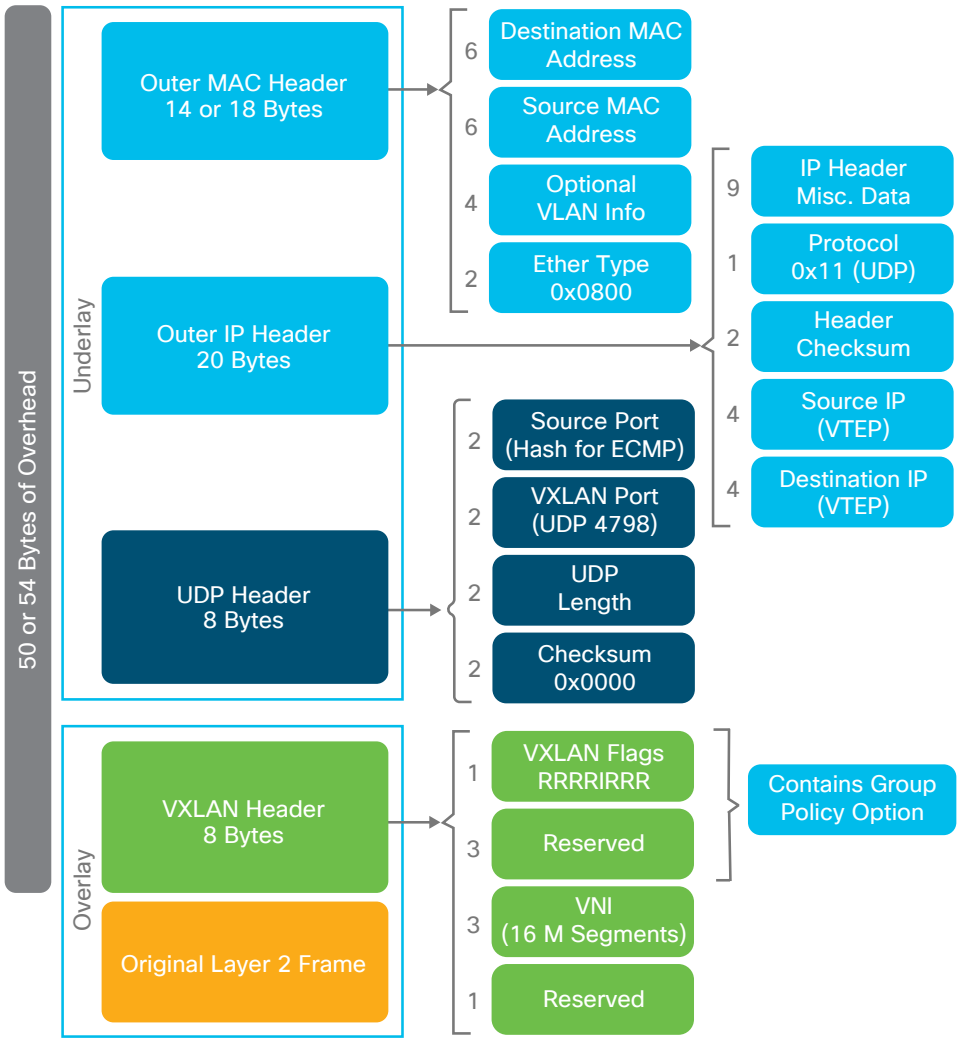

The SD-Access fabric uses the VXLAN data plane in order to provide transport of the full original Layer 2 frame and additionally uses Locator/ID Separation Protocol as the control-plane in order to resolve endpoint-to-location mappings. The SD-Access fabric replaces 16 of the reserved bits in the VXLAN header in order to transport up to 64,000 SGTs, using a modified VXLAN-GPO format described in https://tools.ietf.org/html/draft-smith-vxlan-group-policy-04.

The VNI maps to a virtual routing and forwarding instance for Layer 3 overlays, whereas a Layer 2 VNI maps to a VLAN broadcast domain, both providing the mechanism to isolate data and control plane to each individual virtual network. The SGT carries group membership information of users and provides data-plane segmentation inside the virtualized network.

Fabric control plane

RFC 6830 and other RFCs define LISP as a network architecture and set of protocols that implement a new semantic for IP addressing and forwarding. In traditional IP networks, the IP address is used to identify both an endpoint and its physical location as part of a subnet assignment on a router. In a LISP-enabled network, an IP address or MAC address is used as the endpoint identifier for a device, and an additional IP address is used as an RLOC to represent the physical location of that device (typically a loopback address of the router to which the EID is attached). The EID and RLOC combination provides the necessary information for traffic forwarding. The RLOC address is part of the underlay routing domain, and the EID can be assigned independently of the location.

The LISP architecture requires a mapping system that stores and resolves EIDs to RLOCs. This is analogous to using DNS to resolve IP addresses for host names. EID prefixes (either IPv4 addresses with /32 "host" masks or MAC addresses) are registered into the map server along with their associated RLOCs. When sending traffic to an EID, a source RLOC queries the mapping system in order to identify the destination RLOC for traffic encapsulation. Just like with DNS, a local node probably does not have the information about everything in a network but instead asks for the information only when local hosts need it to communicate (pull model), and the information is then cached for efficiency.

Although a full understanding of LISP and VXLAN is not required to deploy a fabric in SD-Access, it is helpful to understand how these technologies support the deployment goals. Included benefits provided by the LISP architecture are:

- **Network virtualization–**A LISP Instance ID is used to maintain independent VRF topologies. From a data-plane perspective, the LISP Instance ID maps to the VNI.
- **Subnet stretching–**A single subnet can be extended to exist at multiple RLOCs. The separation of EID from RLOC enables the capability to extend subnets across different RLOCs. The RLOC in the LISP architecture is used to encapsulate EID traffic over a Layer 3 network. As a result of the availability of the anycast gateway across multiple RLOCs, the EID client configuration (IP address, subnet, and gateway) can remain unchanged, even as the client moves across the stretched subnet to different physical attachment points.
- **Smaller routing tables–**Only RLOCs need to be reachable in the global routing table. Local EIDs are cached at the local node while remote EIDs are learned through conversational learning. ***Conversational learning*** is the process of populating forwarding tables with only endpoints that are communicating through the node. This allows for efficient use of forwarding tables.

Appendix B–Glossary

AAA authentication, authorization, and accounting

ACL access control list

AP access point

BGP border gateway protocol

CAPWAP control and provisioning of wireless access points protocol

Cisco DNA Cisco Digital Network Architecture

CMD Cisco Meta Data

DMZ firewall demilitarized zone

EID endpoint identifier

HTDB host tracking database

IGP interior gateway protocol

ISE Cisco Identity Services Engine

LISP Locator/ID Separation Protocol

MR Map-Resolver

MS Map-Server

MTU maximum transmission unit

RLOC routing locator

SD-Access Software-Defined Access

SGACL scalable group access control list

SGT scalable group tag

SXP scalable group tag exchange protocol

VLAN virtual local area network

VN virtual network

VNI virtual extensible LAN network identifier

VRF virtual routing and forwarding

VXLAN virtual extensible LAN

Please use the feedback form to send comments and suggestions about this guide.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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