Cisco Software-Defined Access
Enabling intent-based networking
2nd edition

Craig Hill
Darrin Miller
Dave Zacks
Josh Suhr
Karthik Kumar Thatikonda
Kedar Karmarkar
Sanjay Hooda
Satish Kondalam
Saurav Prasad
Shawn Wargo
Simone Arena
Vaibhav Katkade
Vikram Pendharkar

———————————

Bill Rubino
Imran Bashir
Jeffrey Meek
Jeevak Bhatia
Kanu Gupta
Meghna Muralinath
Shane DeLong
Tarunesh Ahuja

CISCO™

Cisco Software-Defined Access

Enabling intent-based networking

2nd edition

Preface

Authors

This book represents a collaborative effort between Technical Marketing, Product Management, Advanced Services, Engineering, and Sales Engineers during a week-long intensive session at Cisco® Headquarters in San Jose, CA.

- Craig Hill — Systems Engineering
- Darrin Miller — Technical Marketing
- Dave Zacks — Technical Marketing
- Josh Suhr — Customer Experience
- Karthik Kumar Thatikonda — Technical Marketing
- Kedar Karmarkar — Technical Marketing
- Sanjay Hooda — Engineering
- Satish Kondalam — Technical Marketing
- Saurav Prasad — Technical Marketing
- Shawn Wargo — Technical Marketing
- Simone Arena — Technical Marketing
- Vaibhav Katkade — Product Management
- Vikram Pendharkar — Product Management

Another group of authors contributed to the most recent update to this book, which was completed in April 2019.

- Bill Rubino — Marketing

- Dave Zacks — Technical Marketing

- Imran Bashir — Technical Marketing

- Jeffrey Meek — Marketing

- Jeevak Bhatia — Product Management

- Kanu Gupta — Technical Marketing

- Meghna Muralinath — Technical Marketing

- Sanjay Hooda — Engineering

- Shane DeLong — Customer Experience

- Shawn Wargo — Technical Marketing

- Tarunesh Ahuja — Technical Marketing

Acknowledgments

Organization of this book

This book is intended to be read sequentially. While each chapter could be reviewed in any order, the concepts discussed in the book build sequentially, so the reader is recommended to review the topics in the order presented.

The topics that the book covers include an introduction to SD-Access, discussing the business drivers for enterprises, and the challenges to enterprise IT to enable those business outcomes today. Next, we examine an overview of SD-Access, providing a synopsis of the SD-Access solution and discussing how SD-Access overcomes some of the challenges that cannot be solved with existing network tools and techniques. Following this, we dive deep into the technology behind SD-Access, including components, operation, and deployment options. We then examine the automation and orchestration framework that drives the SD-Access workflows — Cisco DNA Center®. Areas including policy, automation, and assurance are discussed in detail. Finally, we explore the integration with partner ecosystems and wrap up with a summary and a review of recommended next steps.

Intended audience

Network and IT professionals are always looking for ways to improve the design and operations of their networks. This book focuses on the SD-Access based network architecture for wired-wireless integration with WAN, DC, and services. While the network architects and administrators will reap maximum benefits from this book, the IT professionals will be able to utilize the book to understand the new networking technologies and concepts in their network environments and how these simplify things for them. Security architects can also utilize the book to understand how the SD-Access architecture can help them extend the security perimeter in their enterprise networks to the edge.

The elements in this book cover fabric technologies, policy, automation, assurance and how these technologies can be utilized by network professionals to help the security and simplification of IT deployments.

Book writing methodology

Great design is complexity presented via simplicity − M. Cobanli

Simplicity has been an overriding theme in designing SD-Access. The idea of the book is to present readers with the current challenges in enterprise networking, where the existing networking technologies are lacking, and the underpinnings of SD-Access that solve these challenges for NetOps and SecOps teams without losing sight of simplicity.

A group of Cisco Engineers from diverse backgrounds accepted the challenge of writing a book that changes the paradigm of enterprise networking. At the end of day one, the task seemed even more daunting, given the breadth of areas that SD-Access covers in an enterprise network. However the team persisted, and after hundreds of hours of diligent penmanship, this book was born! The Book Sprints (www.booksprints.net) methodology captured each of our unique strengths, fostered a team-oriented environment, and accelerated the overall time to completion. And now, with the second edition of this book that you hold in your hands, the experiences and expertise of even more Cisco engineers has been brought to bear to capture some of the latest and greatest capabilities that Cisco SD-Access has to offer!

#HardtoTalkAboutSimplicity

What is new in this edition of the book?

This book has been updated to reflect several of the latest and most advanced features that are available with Cisco Software-Defined Access. These new capabilities include the following:

LAN automation: A part of Cisco DNA Center, LAN automation will deploy a standards-based IGP routing protocol to automatically bring up the underlay network. Traditionally, this has been IS-IS, but SD-Access now offers support for OSPF as the automated underlay routing protocol.

Multicast (native): Multicast is used to distribute copies of data to mutiple different network destinations. Cisco SD-Access has offered overlay multicast (head-end replication) since its inception. Now, SD-Access also offers native multicast support, which provides replication in the underlay network as an option. This improves the efficiency of multicast deployment for fabric networks by distributing the load of multicast replication to multiple network elements.

Layer 2 flooding: SD-Access offers support for Layer 2 flooding by forwarding broadcasts for certain traffic and application types which may require leveraging of Layer 2 connectivity, such as silent hosts, card readers, door locks, etc. Such devices and applications may require flooding of traffic in a Layer 2 domain, which SD-Access deployments can now accommodate.

Layer 2 border: Designed as a migration solution, the Layer 2 border feature provides functionality which enables hosts to communicate from the VXLAN-based SD-Access fabric to a traditional VLAN switchport connected to the enterprise network (outside of the fabric). This capability simplifies migrations by enabling the same IP subnet to be used both inside and outside of the SD-Access fabric.

Fabric-in-a-box: The fabric-in-a-box feature enables a single SD-Access device to be used for all three fabric roles (border, control plane and fabric edge node). This feature is especially valuable for support of smaller sites and/or remote branch deployments.

Embedded wireless LAN controller: This new and exciting capability provides the ability to support an embedded wireless LAN controller capability on a Catalyst® 9000 family switch, for use with SD-Access deployments. This simplifies wireless deployments with SD-Access fabric, especially for smaller sites and branch locations.

IoT extension for SD-Access: This capability is used to attach downstream non-fabric Layer 2 network devices to the SD-Access fabric edge node (thus, extending the fabric). This is done by using a device (such as a smaller Layer 2-only switch) designated as an extended node, connected to and leveraging the upstream fabric edge switch for fabric connectivity and policy enforcement. This is especially useful in industrial deployments or deployments outside of the traditional "carpeted" space.

Extranet: SD-Access extranet provides a flexible, and scalable method for achieving inter-VN communications, simplifying the SD-Access fabric deployment and providing a more efficient and policy-based method of communication between devices and services located in separate virtual networks (VNs).

VN anchoring: VN anchoring allows for the traffic from a given VN at multiple dispersed sites to be aggregated back to a central location, using a single common subnet, rather than having to define and use subnets per-site for that VN as would otherwise be the case. This simplifies the overall deployment for several key use cases, including centralized guest access and similar deployments.

IPv6 Support: SD-Access now supports IPv6-based access for a client attached to the fabric network overlay, a critical requirement as more and more hosts support the next generation of the Internet Protocol.

Access control application (ACA): The ACA application, residing on Cisco DNA Center, is designed to provide a greater level of interoperability with non-Cisco identity and cloud solutions, as well as simplifying the design, deployment, and use of policies within a fabric.

This revised edition of the SD-Access book adresses all of the above areas and capabilities. Read on to know more!

Introduction

Executive summary

Digital transformation is creating new opportunities in every industry. In healthcare, doctors are now able to monitor patients remotely and to leverage medical analytics to predict health issues. In education, technology is enabling a connected campus and a more personalized, equal access to learning resources. In retail, shops are able to provide an omnichannel experience by being able to engage online and in-store, with location awareness. In the world of finance, technology now enables users to securely bank anywhere, anytime, on the device of their choice. In today's world, digital transformation is absolutely necessary for businesses to stay relevant!

For any organization to successfully transition to a digital world, investment in their network is critical. It is the network that connects all things and it is the cornerstone where digital success is realized or lost. It is the pathway for productivity and collaboration and an enabler of improved end-user experience. It is also the first line of defense to securing enterprise assets and intellectual property.

Software-Defined Access is the industry's first intent-based networking solution for the enterprise. An intent-based network treats the network as a single system that provides the translation and validation of the business intent (or goals) into the network and returns actionable insights.

DIAGRAM   Intent-based network for a digital business

Digital business

Business goals

Insights

Network learning

Intent   Context

Mobile   Security   IoT   MultiCloud

Security

SD-Access provides automated end-to-end services (such as segmentation, quality of service, and analytics) for user, device, and application traffic. SD-Access automates user policy so organizations can ensure the appropriate access control and application experience are set for any user or device to any application across the network. This is accomplished with a single network fabric across LAN and WLAN which creates a consistent user experience, anywhere, without compromising on security.

SD-Access benefits

- **Automation**: Plug-and-play for simplified deployment of new network devices, along with consistent management of wired and wireless network configuration provisioning

- **Policy**: Automated network segmentation and group-based policy

- **Assurance**: Contextual insights for fast issue resolution and capacity planning

- **Integration**: Open and programmable interfaces for integration with third-party solutions

Network evolution – the challenges

Today's networks support a very different IT environment compared to just a few years ago. There has been a significant rise in the use of mobile clients, increased adoption of cloud-based applications, and the increasing deployment and use of Internet of Things (IoT) in the network environment.

As businesses have digitized, the scale of networks and networking needs have grown over the years, without a proportionate increase in IT resources. End-user expectations have also been rising, with businesses continually expecting IT to keep up with their evolving technology and growth needs.

Over the years, the networking technologies that have been the foundation of interconnectivity between clients, devices, and applications have remained fairly static. While today's IT teams have a number of technology choices to design and operate their networks, there hasn't been a comprehensive, turnkey solution to address today's rapidly-evolving enterprise needs around mobility, IoT, cloud, and security.

> ⌐ the bottom line
>   Network requirements have evolved, but technologies and operations have not.

In this section, we'll explore several broad areas of modern networking challenges in the context of a number of common use cases:

Network deployment

- Implementation complexity

- Wireless considerations

Service deployment

- Network segmentation

- Access control policies

- User and device onboarding

Network operations

- Slow issue resolution

Network deployment

Implementation complexity

Over time, network operators have had to accommodate new network services by implementing new features and design approaches. However, they have had to do so on top of a traditional, inflexible network infrastructure. In addition, the network has had to continually be optimized for high availability, new applications, etc., resulting in "snowflake" networks — no two are alike. While this may meet functional goals, this also makes the networks complex to understand, troubleshoot, predict, and upgrade.

A slow-to-deploy network impedes the ability of many organizations to innovate rapidly and adopt new technologies such as video, collaboration, and connected workspaces. The ability of a company to adopt any of these is impeded if the network is slow to change and adapt. It is difficult to automate the many, many potential variations of "snowflake" network designs and this limits the ability to adopt automation in today's networks to drive greater operational efficiencies for an organization.

> **the bottom line**
> Too many network variations and combinations (snowflakes) make it challenging to adopt new capabilities and services.

Wireless considerations

In addition, one of the major challenges with wireless deployment today is that it does not easily utilize network segmentation. While wireless can leverage multiple SSIDs for traffic separation over-the-air, these are limited in the number that can be deployed and are which are ultimately mapped back into VLANs at the WLC. The WLC itself has no concept of VRF / Layer 3 segmentation, making deployment of a true wired and wireless network virtualization solution very challenging.

> └─ the bottom line
> Traditional wireless networks are managed separately and are difficult to segment.

Service deployment

Network segmentation

Let's look at some of the options available today and their challenges for creating network segmentation.

Segmentation using VLANs

The simplest form of network segmentation relies on VLANs. You may not be used to thinking about a VLAN as network segmentation technology, but that is what a VLAN is: a segmented Layer 2 domain. By placing users and devices in different VLANs, traffic controls can be enforced between them at the Layer 3 boundary. For wireless, different SSIDs might be used for separation into the air, but then these are mapped into VLANs on the wired side.

The challenge with VLANs as a segmentation method is two-fold: their span and the topology-related issues they bring along with them. In terms of span, most organizations choose to constrain VLANs to a relatively small area (often limited to one wiring closet). Because of this, many organizations end up managing hundreds or even thousands of VLANs in a typical deployment, making IP address planning much more complex to deploy and manage.

Key challenges with using VLANs include the following:

- A widely-spanned VLAN is vulnerable to Layer 2 loops in redundant network designs, and are thus at risk of meltdowns if an uncontrolled Layer 2 loop should occur for any reason

- Large Layer 2 designs are very inefficient (50% of ports blocking typically)

- Traffic filtering options for intra-VLAN traffic are also typically much more limited than those available at a Layer 3 boundary

└ **the bottom line**
VLANs are simple but, in this case, simple is not best — a flat Layer 2 design exposes the organization to too many potential events that could take down the network, and in addition, managing hundreds of VLANs is daunting for most organizations.

Segmentation using VRF-Lite

Another approach — one that leverages Layer 3 — is to segment the network by using VRFs (Virtual Routing and Forwarding instances — essentially, separate versions of the IP routing table). This has the benefit that segmentation can be provided without the need to build large, complex ACLs to control traffic flows — since traffic between different VRFs can only flow as the network manager dictates via the network topology (typically, via route leaking or through a firewall).

Challenges with the VRF-Lite approach to segmentation are as follows:

- VRF-Lite using 802.1q trunks between devices is relatively simple to implement on a few devices but becomes very cumbersome very quickly when implemented more widely

- VRF-Lite requires separate routing protocol processes per VRF, resulting in increased CPU load and complexity

- The typical rule-of-thumb is that VRF-Lite deployments should not be scaled beyond 8-10 VRFs, as they become far too unwieldy to manage end-to-end in an enterprise deployment at a larger scale

Segmentation using MPLS VPNs

An alternative technology available for network segmentation, MPLS VPNs, has a steep learning curve since they require the network manager to become familiar with many new MPLS-specific capabilities and network protocols, including LDP for label distribution and Multi-Protocol BGP as a control plane. Moreover, they need to understand how to troubleshoot MPLS-enabled networks when issues arise.

Challenges with MPLS VPNs are as follows:

- MPLS VPNs will scale much better than VRF-Lite; however, they are often too complex for many network managers to tackle, especially across an end-to-end network deployment.

- MPLS VPN support is not available pervasively across all network platforms.

> **the bottom line**
> Despite having VRF capabilities for more than ten years, only a small percentage of organizations have deployed VRF segmentation in any form. Why is this? In a word – complexity.

Network policies

Policy is one of those abstract words that can mean many different things to many different people. However, in the context of networking, every organization has multiple policies that they implement. Use of security ACLs on a switch, or security rule-sets on a firewall, is security policy. Using QoS to sort traffic into different classes, and using queues on network devices to prioritize one application versus another, is QoS policy. Placing devices into separate VLANs based on their role, is device-level access control policy.

Today's network manager typically uses a few sets of common policy tools every day: VLANs, subnets, and ACLs. For example:

- Adding voice to a network? This implies carving a new set of voice VLANs and associated subnets.

- Adding IoT devices? Door locks, badge readers, et cetera? More VLANs and sub-nets.

- Adding IP cameras and streaming video endpoints? More VLANs and subnets again.

This is why an enterprise network today ends up with hundreds, or even thousands, of VLANs and subnets. The level of complexity in designing and maintaining this is obvious in and of itself — and yet it also requires the further maintenance of many DHCP scopes, IPAM tools, and the complexity associated with managing a large IP address space across all of these various VLANs and functions.

Today's network, faced with many internal and external threats, also needs to be se-cure. This makes it necessary to create and implement — and maintain on an ongoing basis — large Access Control Lists, implemented on network devices including switches, routers, and firewalls, most often at Layer 3 boundaries in the network deployment.

> └ the bottom line
> The traditional methods used today for policy administration (large and complex ACLs on devices and firewalls) are very difficult to implement and maintain.

User and device onboarding

No matter which solution is chosen today — a Layer 2 or Layer 3 network design, a seg-mented or non-segmented network approach — there is always the issue of the optimal approach to onboard users and devices into the network.

This could be as simple as hard-coding a VLAN/subnet to a wired port or wireless SSID, but there are some common challenges with this approach:

- While functional, this offers little real security, since anyone connecting into that port or SSID is associated with that "role" in the network.

- Either on the first-hop switch, or on a firewall ten hops away, that user's IP ad-dress will be examined and the appropriate security policy will be controlled and

enforced. Essentially, the IP address ends up being used as a proxy for identity. However, this is hard to scale, and to manage.

Otherwise, a VLAN/subnet could be assigned dynamically using 802.1x or another authentication method, but there are some common challenges with this as well:

- While the use of 802.1x is common in wireless deployments, it is less common in wired network use.

- Many issues exist as deployment blockers, such as 802.1x supplicant settings on the device, 802.1x support on the device, switching VLANs/subnets dynamically based on roles, and 802.1x support and features on the network equipment.

Finally, once that user/device identity is established, how can it be carried end-to-end within the network today? There is no place within an IP packet header to carry this user/device mapping. So, once again, IP addresses are used as a proxy for this. However, this leads to a proliferation of user/device subnets, and all of the attendant complexity that goes along with this.

> └ the bottom line
>   Most organizations want to establish user/device identity and use it end-to-end
>   for policy. However, many find this to be a daunting task.

Network operations

Slow issue resolution

Many networks today provide very limited visibility into network operation and use. The wide variety of available network monitoring methods — SNMP, Netflow, screen scraping, and the like — and the mixture of availability of these tools across various platforms, makes it very difficult to provide comprehensive, real-time, end-to-end insights derived from ongoing monitoring in today's network deployments.

Without insight into ongoing operational status, organizations often find themselves reacting to network problems, rather than addressing them proactively — whether

these problems are caused by issues or outages, or simply brought on by growth or changes in user / application patterns.

Many organizations would place significant value on being able to be more knowledgeable about how their network is being used, and more proactive in terms of network visibility and monitoring. A more comprehensive, end-to-end approach is needed — one that allows insights to be drawn from the mass of data that potentially can be reported from the underlying infrastructure.

> ⌐ **the bottom line**
> Most organizations lack comprehensive visibility into network operation and use — limiting their ability to proactively respond to changes.

Tying it all together

So, what does it take to roll out networks and the associated policies end-to-end today?

First, the network manager has to settle on a given network design: multi-layer access, routed access, access virtualized using VRF-Lite or MPLS VPNs, etc. There are many considerations and tradeoffs associated with this today, so this is not necessarily a simple choice.

Based on the diagram below, the following steps represent a typical service deployment:

1 Map to user groups in active directory (AD) or a similar database for user authentication.

2 Link these AD identities to the AAA server (such as Cisco Identity Services Engine, ISE) if using dynamic authentication. This provides each identity with an appropriate corresponding VLAN/subnet.

3 Define and carve out new VLANs and associated subnets for the new services to be offered. Then, implement these VLANs and subnets on all necessary devices (switches, routers, and WLCs).

4   Secure those subnets with the appropriate device or firewall ACLs, or network segmentation. If using a segmented, virtualized network approach, extend these VRFs end-to-end using VRF-Lite or MPLS VPNs.

5   To do all of this, it is necessary to work across multiple user interfaces — the AD GUI, the AAA GUI, the WLC GUI for wireless; the switch or router CLI for wired — and stitch together all of the necessary constructs manually.

And when it becomes necessary to implement another group of users or devices or to alter a policy associated with them, all of these steps must be repeated all over again.

└ the bottom line
      No wonder it takes days or weeks to roll out new network services today!

DIAGRAM   Service deployment overview
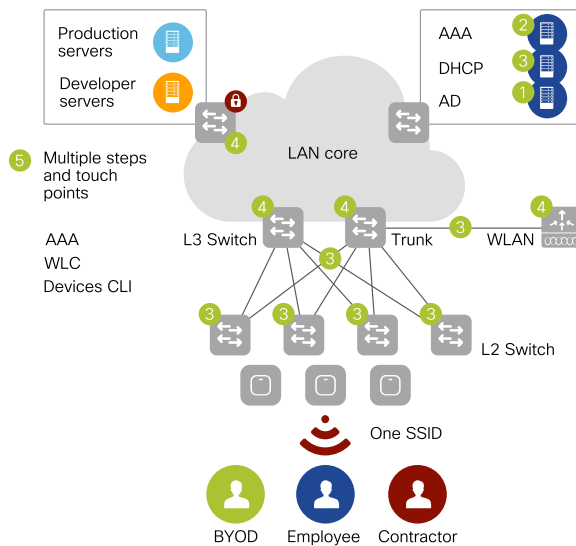
SD-Access overview

Software-Defined Access overview

Cisco's Software-Defined Access (SD-Access) solution is a programmable network architecture that provides software-based policy and segmentation from the edge of the network to the applications. SD-Access is implemented via Cisco Digital Network Architecture Center (Cisco DNA Center) which provides design settings, policy definition and automated provisioning of the network elements, as well as assurance analytics for an intelligent wired and wireless network.

In an enterprise architecture, the network may span multiple domains, locations or sites such as main campuses and remote branches, each with multiple devices, services, and policies. The Cisco SD-Access solution offers an end-to-end architecture that ensures consistency in terms of connectivity, segmentation, and policy across different locations (sites).

These can be described as two main layers:

- **SD-Access fabric:** physical and logical network forwarding infrastructure
- **Cisco DNA Center**: automation, policy, assurance and integration infrastructure

This chapter will provide an overview of each of the major SD-Access solution components, with additional details for these areas provided in following chapters.

SD-Access fabric

As described earlier, part of the complexity in today's network comes from the fact that policies are tied to network constructs such as IP addresses, VLANs, ACLs, etc.

What if the enterprise network could be divided into two different layers, each for different objectives? One layer would be dedicated to the physical devices and forwarding of traffic (known as an underlay), and another entirely virtual layer (known as an overlay) would be where wired and wireless users and devices are logically connected together, and services and policies are applied.

This provides a clear separation of responsibilities and maximizes the capabilities of each sublayer while dramatically simplifying deployment and operations since a change of policy would only affect the overlay, and the underlay would not be touched.

DIAGRAM   Software-Defined Access overview
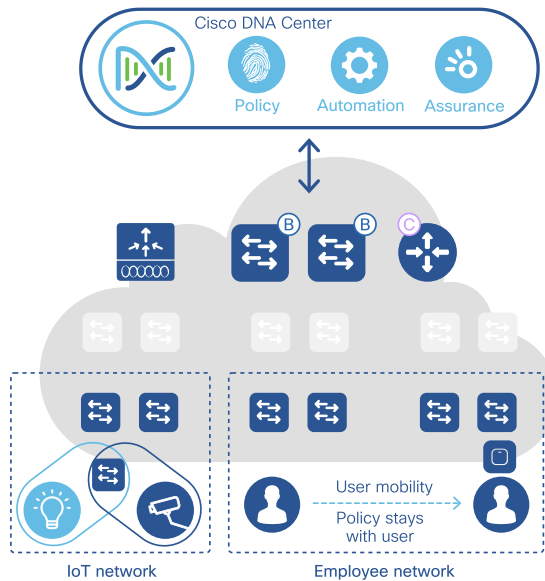
The combination of an underlay and an overlay is called a "network fabric".

The concepts of overlay and fabric are not new in the networking industry. Existing technologies such as MPLS, GRE, LISP, and OTV are all examples of network tunneling technologies which implement an overlay. Another common example is Cisco Unified Wireless Network (CUWN), which uses CAPWAP to create an overlay network for wireless traffic.

DIAGRAM   Underlay and overlay networks

To understand what is unique about the SD-Access fabric, we will first define key SD-Access components.

SD-Access network underlay

An SD-Access network underlay (or simply: underlay) is comprised of the physical network devices, such as routers, switches, and wireless LAN controllers (WLCs) plus a traditional Layer 3 routing protocol. This provides a simple, scalable and resilient foundation for communication between the network devices. The network underlay is not used for client traffic (client traffic uses the fabric overlay).

All network elements of the underlay must establish IPv4 connectivity between each other. This means an existing IPv4 network can be leveraged as the network underlay. Although any topology and routing protocol could be used in the underlay, the implementation of a well-designed Layer 3 access topology (i.e. a routed access topology) is highly recommended to ensure consistent performance, scalability, and high availability.

Using a routed access topology (i.e. leveraging routing all of the way down to the access layer) eliminates the need for STP, VTP, HSRP, VRRP, and other similar protocols in the network underlay, simplifying the network dramatically and at the same time increasing resiliency and improving fault tolerance. In addition, running a logical fabric topology

on top of a prescriptive network underlay provides built-in functionality for multi-path and optimized convergence, and simplifies the deployment, troubleshooting, and management of the network.

Cisco DNA Center provides a prescriptive LAN automation service to automatically discover, provision, and deploy network devices according to Cisco-validated design best practices. Once discovered, the automated underlay provisioning leverages Plug and Play (PnP) to apply the required IP address and routing protocol configurations.

Cisco DNA Center LAN Automation uses a best practice intermediate system to intermediate system (IS-IS) or open shortest path first (OSPF) routed access design. The main reasons for using IS-IS or OSPF are:

- Both IS-IS and OSPF are standards-based shortest-path first (SPF) link state routing protocols.

  - IS-IS is common in large SP and DC networks, and the *de facto* underlay protocol for fabric environments.

  - OSPF is common in large enterprise and WAN networks, for most traditional campus routing environments.

- Link state routing protocols don't advertise the entire routing table. Instead, they advertise information about the network topology (directly connected links, neighboring routers, etc.), so that all routers within an area have the same topology database.

- Link state routing protocols provide a multi-level hierarchy (called levels or areas), so routing updates within a defined area are hidden from routers outside that area.

- Link state routing protocols support classless routing, send updates using multicast addresses, and use triggered routing updates.

- Link state routing protocols use an algorithm to determine the shortest path to every other node in the topology.

- Link state routing protocols converge much faster than distance vector routing protocols

SD-Access fabric overlay

The SD-Access fabric overlay (or simply: overlay) is the logical, virtualized topology built on top of the physical underlay. As described earlier, this requires several additional technologies to operate.

SD-Access fabric overlay has 3 main building blocks:

- **Fabric data plane**: the logical overlay is created by packet encapsulation using virtual extensible LAN (VXLAN), with group policy option (GPO)

- **Fabric control plane**: the logical mapping and resolving of users and devices (associated with VXLAN tunnel endpoints) is performed by locator/ID separation protocol (LISP)

- **Fabric policy plane**: where the business intent is translated into a network policy, using address-agnostic scalable group tags (SGT) and group-based policies

VXLAN-GPO provides several advantages for SD-Access, such as support for both Layer 2 and Layer 3 virtual topologies (overlays), and the ability to operate over any IP-based network with built-in network segmentation (using VRFs / VNs) and group-based policy (using scalable groups tags / SGTs).

LISP dramatically simplifies traditional routing environments by removing the need for each router to process every possible IP destination address and route. It does this by moving remote destination information to a centralized map database (the LISP map server / map resolver, also known as the fabric control plane) that allows each router to manage only its local routes and query the map system to locate destination endpoints.

SD-Access policy

A fundamental benefit of SD-Access is the ability to instantiate logical network policy, based on services offered by the fabric. Some examples of services that the solution offers are the following:

- Security segmentation services

- Quality of service (QoS)

- Capture/copy services

- Application visibility services

These services are offered across the entire fabric independently of device-specific address or location.

SD-Access segmentation

Segmentation is a method or technology used to separate specific groups of users or devices from other groups for the purpose of security, overlapping IP subnets, etc. In SD-Access fabric, VXLAN data plane encapsulation provides network segmentation by using the VNI (virtual network identifier) and scalable group tag (SGT) fields in the VXLAN-GPO header.

By leveraging these constructs, SD-Access fabric provides a simple way to implement hierarchical network segmentation: macro segmentation and micro segmentation.

Macro segmentation: logically separating a network topology into smaller virtual networks, using a unique network identifier and separate forwarding tables. This is instantiated as a virtual routing and forwarding (VRF) instance and referred to as a virtual network (VN).

Micro segmentation: logically separating user or device groups within a VN, by enforcing source-to-destination access control permissions. This is commonly instantiated using scalable access group access control lists (SGACLs), also known as an access control policy.

A **virtual network** (VN) is a logical network instance within the SD-Access fabric, providing Layer 2 or Layer 3 services and defining a Layer 3 routing domain. The VXLAN VNI is used to provide both the Layer 2 (L2 VNI) and Layer 3 (L3 VNI) segmentation.

A **scalable group** is a logical object ID assigned to a "group" of users and/or devices in the SD-Access fabric, and used as the source and destination classifier in SGACLs. The SGT is used to provide address-agnostic group-based policies.

SD-Access fabric for wireless

Administrators may notice that a traditional Cisco Unified Wireless Network (CUWN) design provides some of the same advantages of SD-Access. For example, some of the common attributes include:

- Tunneled overlay network (via CAPWAP enscapsulation and separate control-plane)

- Some levels of infrastructure automation (e.g. AP management, configuration management, etc.)

- Simple wireless user or device mobility (also known as client roaming)

- Centralized management via a wireless controller (WLC)

However, the CUWN approach also comes with some trade offs:

- Only wireless users can benefit from the CAPWAP overlay — does not apply to wired users.

- In the most commonly deployed wireless mode (centralized or local mode), wireless traffic must be tunneled to a centralized anchor point (the WLC), which may not be optimal for many applications.

In addition, there are several advantages unique to wired users, which users of a traditional wireless design using CUWN would lack:

- Wired users can benefit from the performance and scalability that a distributed switching data plane provides

- Wired users also benefit from advanced QoS and innovative services such as Encrypted Traffic Analytics (ETA), available in the switching infrastructure

As can be seen, each domain (wired and wireless) has different advantages. One of the unique and powerful aspects of SD-Access wireless is that it allows network operators and users to get the "best of both worlds", leveraging an SD-Access fabric deployment.

- SD-Access fabric provides the best of the distributed wired and centralized wireless architectures by providing a common overlay and extending the benefits to both wired and wireless users.

- With SD-Access fabric, customers can have a common policy and one unified experience for all their users independently of the access media.

SD-Access management with Cisco DNA Center

Cisco DNA Center provides a central management plane for building and operating a SD-Access fabric. The management plane is responsible for forwarding configuration and policy distribution, as well as device management and analytics.

There are two main functions of Cisco DNA Center: automation and assurance.

Automation and orchestration

Cisco DNA Center automation provides the definition and management of SD-Access group-based policies, along with the automation of all policy-related configurations. Cisco DNA Center integrates directly with Cisco Identity Services Engine (ISE) to provide host onboarding and policy enforcement capabilities.

Automation can be generally defined as a technology or system that performs an action or task without human assistance. A single task may require multiple actions but will have a single expected outcome. Orchestration is automating the execution of an overall workflow or process and may require multiple related tasks and involve multiple systems.

This is the basis on which the industry terminology "Software-Defined" can be applied to the automation and orchestration of enterprise campus access network environ-

ments, as well as translate the user "intent" into meaningful configuration and verification tasks.

With SD-Access, Cisco DNA Center uses controller-based automation as the primary configuration and orchestration model, to design, deploy, verify, and optimize wired and wireless network components for both non-fabric and fabric-based deployments.

With Cisco DNA Center completely managing the infrastructure, IT teams can now operate on an abstracted intent-based layer and not have to worry about the implementation details. This results in simplifying operations for the IT teams by minimizing the chances of making a human error and more easily standardizing the overall network design.

Network assurance

Network assurance quantifies availability and risk from an IT network perspective, based on a comprehensive set of network analytics. Beyond general network management, network assurance measures the impact of network change on security, availability, and compliance.

Cisco DNA Center Assurance has been developed as a full management and operations solution to address the most common customer challenges. Cisco DNA Center provides multiple forms and levels of assurance and analytics, for both non-fabric and fabric-based components.

The key enabler to Cisco DNA Assurance is the analytics piece: the ability to continually collect data from the network and transform it into actionable insights. To achieve this, Cisco DNA Center collects a variety of network telemetry, in traditional forms (e.g. SNMP, Netflow, syslogs, etc) and also emerging forms (NETCONF, YANG, streaming telemetry, etc). Cisco DNA Assurance then performs advanced processing to evaluate and correlate events to continually monitor how devices, users, and applications are performing.

Correlation of data is key since it allows for troubleshooting issues and analyzing network performance across both the overlay and underlay portions of the SD-Access fabric. Other solutions often lack this level of correlation and thus lose visibility into un-

derlying traffic issues that may affect the performance of the overlay network. By providing correlated visibility into both underlay and overlay traffic patterns and usage via fabric-aware enhancements to Netflow, SD-Access ensures that network visibility is not compromised when a fabric deployment is used.

For further details on Cisco DNA Center Assurance, please refer to the Cisco e-book on Assurance, located at:

https://cs.co/assurancebook

SD-Access benefits

The transformational capabilities of SD-Access make it possible to enable many important capabilities and use cases outlined in the following text.

Automated deployments at scale

SD-Access leverages controller-based automation to build out large enterprise networks without requiring the network operator to have an intricate understanding of the underlying network forwarding constructs. SD-Access provides a single set of networking constructs that cater to many connectivity scenarios.

Most importantly, SD-Access provides this flexible, automated connectivity across large enterprise domains in a resilient manner that minimizes instabilities and reduces risks of downtime. SD-Access is based on a distributed framework to support scalability and does not require re-architecting as the scale of the deployment grows.

SD-Access also works and scales seamlessly across a large number of sites with automated inter-site connectivity, and extends to environments outside of traditional wiring closets, such as connected workspaces, Operational Technology (OT) environments, and manufacturing floors.

As a result, SD-Access enables fast IT / lean IT initiatives critical for business agility by easily accommodating any new connectivity requirements with a common, consistent, and fully automated network.

DIAGRAM   Simplification through automation
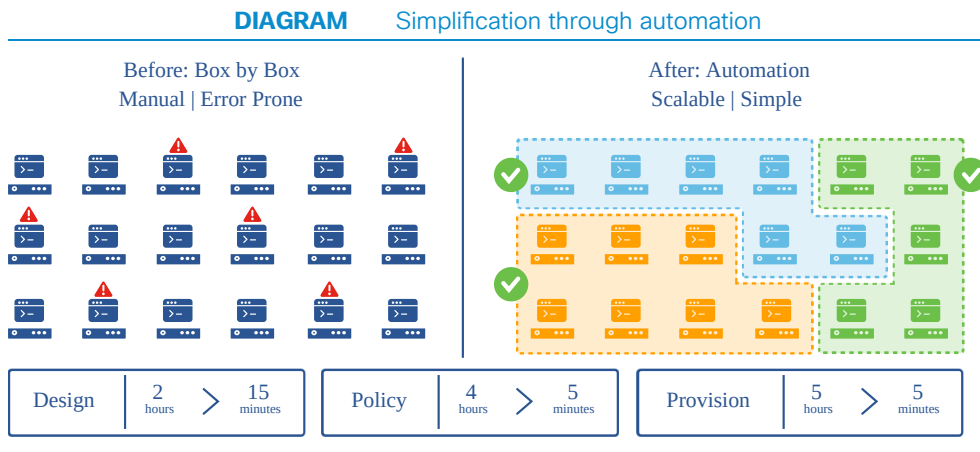
Some of the applications of this use case are listed below:

1   Healthcare: remote collaboration and consultations

2   Education: access to teaching and learning resources across remote campuses

3   Manufacturing: easily extend and manage plant floor networks

└ **the bottom line**
Too many network variations and combinations (snowflakes) make it challenging
to adopt new capabilities and services.

└ **the benefit**
SD-Access reduces the complexity by automating deployment activities without
need for detailed knowledge of the underlying network
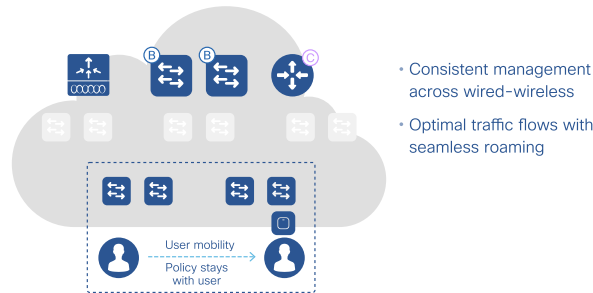
Integrated wired and wireless infrastructure

The wireless data plane in an SD-Access deployment is distributed (i.e. is not centralized at the wireless LAN controller) and shares the same transport and encapsulation as wired traffic. This enables the ability to leverage capabilities that might otherwise pertain only to the wired infrastructure for the wireless traffic as well. For example, the use of optimized multicast or First Hop Security and Segmentation results in a better overall user experience than that of wireless users.

SD-Access wireless provides:

- **Distributed data plane**: The wireless data plane is distributed at the edge switches for optimal performance and scalability, without bringing the hassles usually associated with distributing traffic forwarding, including spanning VLANs, subnetting, etc.

- **Centralized wireless control plane**: The same innovative RF features that Cisco has today in CUWN deployments are leveraged in SD-Access wireless as well. Wireless operations are similar in terms of RRM, client onboarding, and client mobility. This allows simplified IT operations as the single control plane is maintained for wired and wireless, as well as seamless roaming.

- **Simplified guest and mobility tunneling**: A simplified guest access capability is available with a fabric deployment, as an anchor WLC controller is no longer necessarily required, and the guest traffic can directly go to the DMZ without needing to leverage a foreign wireless controller.

- **Policy simplification**: SD-Access breaks the dependencies between policy and network constructs (IP address and VLANs), simplifying the definition and implementation of policies for both wired and wireless clients.

- **Segmentation made easy**: Segmentation is carried end-to-end in the fabric and is hierarchical, based on VNs and SGTs. The same segmentation policy is applied to both wired and wireless users.

SD-Access wireless provides an IoT-ready infrastructure allowing the IoT devices to be segmented across the enterprise in their own segments without interfering with the corporate network.

DIAGRAM   Best of wired and wireless

- Consistent management across wired–wireless
- Optimal traffic flows with seamless roaming

User mobility
Policy stays with user

Some of the applications of this use case are listed below:

1   Healthcare: Segment patient/guest from physician and clinical networks

2   Education: Improve learning experience in classrooms

3   Retail: Enhanced guest experience over store WiFi

└ the bottom line
Traditional wireless networks are managed separately from wired, making wireless difficult to segment and implement changes.

└ the benefit
SD-Access provides a single control plane that is maintained for wired and wireless. This allows for consistent management, policy enforcement and connectivity and provides a unified experience independent of the access media.

Provide secure access to users and devices

SD-Access provides topology-agnostic, identity-based methods to define access control and network segmentation policies. This simplifies policy definition, updates, and compliance reporting (refer to the following diagram).

The automation framework translates the high-level business intent into low-level configuration of devices in the network infrastructure to enable rapid, consistent, and validated roll-out of policies throughout the network.

DIAGRAM   Simple and secure access
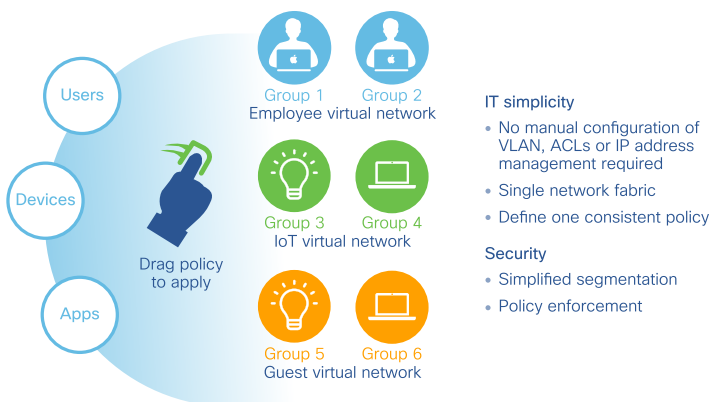
Some of the applications of this use case are:

1   Healthcare: Keep patient, devices, and data secure

2   Education: Create a secure campus

3   Manufacturing: Converge siloed IT and OT networks

└ **the bottom line**
Most organizations want to establish user/device identity and use it end-to-end for policy. However, many find this to be a daunting task.

└ **the benefit**
SD-Access simplifies policy definition segmentation and updates while enabling rapid consistent policy implementation end-to-end.

Correlated insights and analytics

Issue resolution today is reactive, slow and inefficient. The reasons could be multiple tools (each with limited visibility), network complexity, user mobility, or even lack of consistent policies.

Cisco DNA Assurance provides deep visibility into networks by collecting and correlating fine-grained telemetry from a rich variety of sources such as syslog, SNMP, NetFlow, AAA, DHCP, and DNS. This provides IT with rich actionable insights to optimize network infrastructure and support better business decisions.

DIAGRAM   Actionable insights
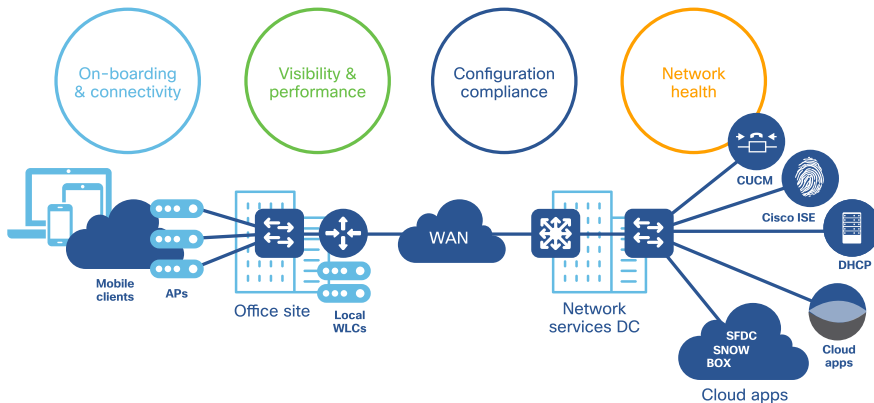
Some of the applications of this use case are listed below:

1   Healthcare: improved clinical workflows and operations

2   Education: ensure network uptime and performance during classroom changes

3   Manufacturing: get new business insights from plant floors to make better IT decisions

└ the bottom line
Most organizations lack comprehensive visibility into network operation and use — limiting their ability to proactively respond to change.

└ the benefit
Cisco DNA Assurance provides insights from every device, application, service, and client on your network and then evaluates the context before recommending a course of action for remediation.

SD-Access fabric

Fabric components

As we begin our examination of the Cisco SD-Access fabric, we need to review the various components that make up fabric deployment, examine the capabilities that they provide, and outline how the various components interact with each other to provide the entire SD-Access solution.

The following diagram outlines the main components that are part of SD-Access fabric deployment, and indicate the various positions that they occupy within a SD-Access fabric system.

DIAGRAM   SD-Access components
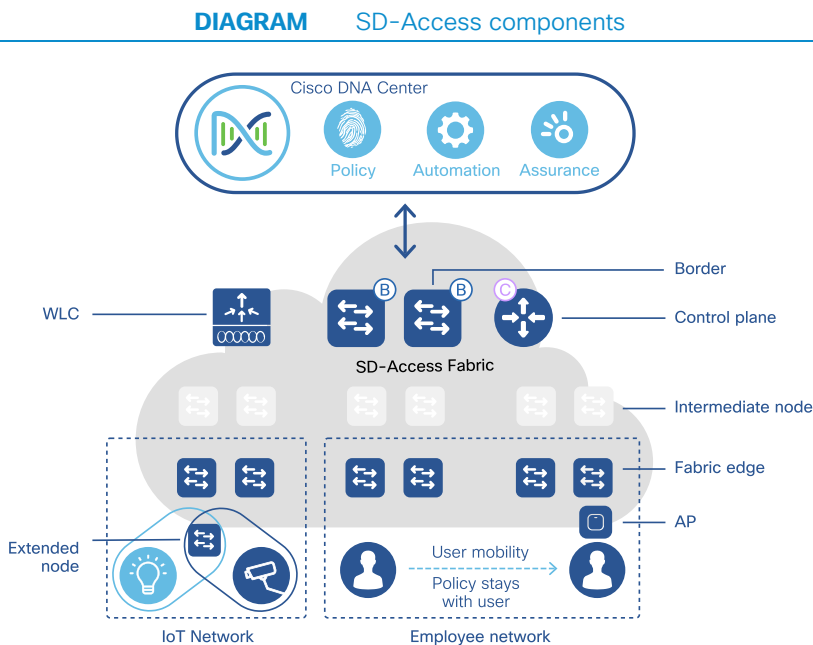
Fabric components and terminology

In this section, we will review the various SD-Access fabric components, and explain the role that each component plays to make the solution as a whole.

Fabric control plane node: The fabric control plane node serves as a central database, tracking all users and devices as they attach to the fabric network, and as they roam around. The fabric control plane allows network components (switches, routers, WLCs, etc) to query this database to determine the location of any user or device attached to the fabric, instead of using a flood and learn mechanism. In this way, the fabric control plane serves as a "single source of truth" about where every endpoint attached to the fabric is located at any point in time. In addition to tracking specific endpoints (/32 address for IPv4, /128 address for IPv6), the fabric control plane can also track larger summarized routes (IP / mask). This flexibility helps in summarization across fabric sites and improves overall scalability. An SD-Access control plane node is shown in diagrams as an icon denoted with a (C), used to indicate that the selected device is operating in a fabric control plane role.

Fabric border node: Fabric border nodes connect the SD-Access fabric to traditional Layer 3 or Layer 2 networks, or to different fabric sites. Fabric border nodes are responsible for the translation of context (user/device mapping and identity) from one fabric site to another fabric site, or to a traditional network. When the encapsulation is the same across different fabric sites, the translation of fabric context is generally 1:1 mapped. The fabric border is also the device where the fabric control planes of different fabric sites exchange reachability and policy information. An SD-Access border node is shown in diagrams as an icon denoted with a (B), used to indicate that the selected device is operating in a fabric border role.

For the migration use case, Layer 2 border functionality (enabled on the border) allows communication between hosts residing in traditional networks with hosts residing in the SD-Access fabric. For an existing deployment involving IoT clients it may prove to be expensive to re-IP the network in order to move them to SD-Access. This Layer 2 border feature allows the flexibility to share the same subnet space inside and outside the fabric. The Layer 2 border, when enabled for a subnet, is configured with an anycast gateway IP address. Non-fabric clients will have their gateway pointing to the anycast

gateway configured on the Layer 2 border. The Layer 2 border will register as the location for these non-fabric hosts within the fabric. Catalyst 9000 family border switches can be configured as a Layer 2 border if desired.

DIAGRAM Layer 2 border
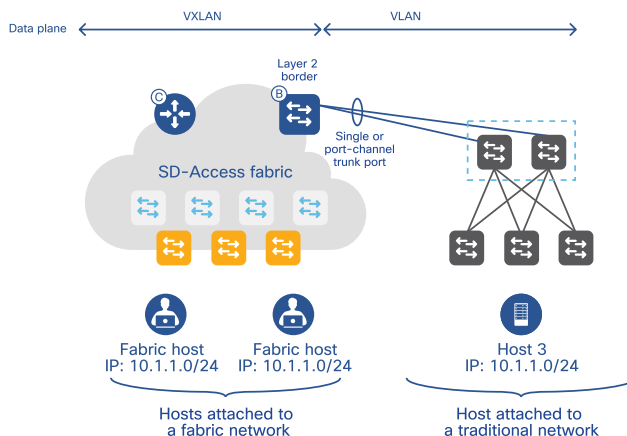
There are two notable Layer 3 fabric border functions: one for internal (defined) networks and one for external (default) networks. Borders that are not explicitly defined as default borders advertise to the fabric control plane a set of known subnets, such as those leading to a group of branch sites, or to a data center. Default borders, on the other hand, are used to reach unknown destinations, typically the internet (similar to the function of a default route). In an SD-Access fabric, an arbitrary number of non-default borders can exist. The total number of default borders supported per SD-Access fabric is smaller (2 or 4, depending on the border node platforms chosen).

Fabric edge node: Fabric edge nodes are responsible for connecting endpoints to the fabric, and encapsulating/decapsulating and forwarding traffic from these endpoints to and from the fabric. Fabric edge nodes operate at the perimeter of the fabric and are the first points for attachment of users and the implementation of policy. It is to be noted that the endpoints need not be directly attached to the fabric edge node (refer to SD-Access Extension).

An important point to note about fabric edge nodes is how they handle the subnets used for endpoint attachment. All subnets hosted in a SD-Access fabric are, by default, provisioned across every edge node in that fabric. For example, if the subnet 10.10.10.0/24 is provisioned in a given fabric, this subnet will be defined across all of the edge nodes in that fabric, and hosts located in that subnet can be placed on any edge node within that fabric. This essentially "stretches" these subnets across all of the edge nodes in that fabric, thus simplifying IP address assignment, allowing fewer but larger IP subnets to be deployed.

Fabric intermediate node: Fabric intermediate nodes are the simplest devices in the SD-Access fabric architecture. Intermediate nodes act as pure Layer 3 forwarders that connect the fabric edge, border, and control plane nodes and provide the Layer 3 underlay for fabric overlay traffic.

SD-Access extension for IoT: This capability is used to attach downstream non-fabric Layer 2 network devices to the SD-Access fabric (thus, extending the fabric). This is done by using a device designated as an extended node — typically, a small switch (compact switch, industrial Ethernet switch, or building automation switch) which connects to the fabric edge node via Layer 2. Devices connected to the SD-Access extended node use the fabric edge node for communication to outside subnets. It is worth noting that policy enforcement (for example, using SGACLs for traffic filtering) is performed by the fabric edge node, not the extended node. As such, traffic requiring policy enforcement must flow via the fabric edge node in order for such policies to be enforced. An SD-Access extension can be created using a single Layer 2 switch attached to a fabric edge node, or (in future) can also support a ring of Layer 2 switches attached below the fabric edge node, with this ring formed using Resilient Ethernet Protocol (REP) or Spanning Tree Protocol (STP).

Fabric wireless LAN controller (WLC): The fabric WLC supports fabric-enabled APs attached to fabric edge switches, handling not only the traditional tasks associated with a WLC but also handling interaction with the fabric control plane for wireless client registration and roaming. It should be noted that a fabric-enabled wireless deployment moves the data-plane termination from a centralized location (with previous overlay CAPWAP deployments), to the AP/fabric edge node using VXLAN. This enables distributed forwarding and distributed policy application for wireless traffic, while retaining the benefits of centralized provisioning and administration.

Fabric-enabled access point (AP): Fabric-enabled APs are attached to fabric edge nodes and connect wireless clients into the fabric network. Fabric APs allow for distributed wireless forwarding in the SD-Access architecture by encapsulating wireless user traffic into the VXLAN-based overlay to their adjacent fabric edge node where it is decapsulated, any necessary policies are applied, and then re-encapsulated to its ultimate destination within the fabric.

Endpoint: The devices that connect to the fabric edge node are called endpoints. Endpoints may be wired clients that directly connect to the fabric edge node, wireless clients attached to a fabric AP, or connecting through a Layer 2 network via an SD-Access extended node. There are of course many types of endpoints. These can include traditional endpoints such as laptop computers, desktop computers, tablets, and smartphones, as well as IoT devices such as cameras, sensors, door locks, badge readers, and many other new and diverse device types. From the point of view of the fabric, these are all endpoints, and are identified by EIDs (endpoint identifiers — typically, the client's IP address (IPv4 / IPv6) and or MAC address).

Cisco Identity Services Engine (ISE): Cisco ISE is a shipping product focused on general profiling, identity, and security policy compliance. ISE is capable of identifying and profiling the network devices and endpoints in a variety of forms, including AAA/RADIUS, 802.1X, MAC Authentication Bypass (MAB), Web Authentication, EasyConnect, and others. It then places the profiled endpoints into the correct security group and host pool. To read more about Cisco ISE, see the ISE product link: *www.cisco.com/go/ise*

Cisco DNA Center: Cisco DNA Center is the command and control system for the SD-Access solution, and houses the automated workflows required to deploy and manage SD-Access fabrics. Cisco DNA Center provides capabilities for both automation and assurance in a fabric deployment (described in more detail in the following chapters). Cisco DNA Center serves both brownfield and greenfield deployments.

Fabric operation

- Control plane operation

- Data plane operation

- Wireless in SD-Access fabric

- SD-Access Extension for IoT

Control plane operation

In SD-Access fabric, the fabric control plane node operates as the database tracking all endpoint connectivity to the fabric, and is responsible for the following functions:

- Registers all endpoints connected to the edge nodes, and tracks their location in the fabric

  - i.e. which edge node the endpoints are located behind.

- Responds to queries from network elements about the location of endpoints in the fabric

- Ensures that when endpoints move from one location to another, traffic is re-directed to the current location

Refer to the following diagram showing control plane operation.

DIAGRAM   Fabric control plane operation
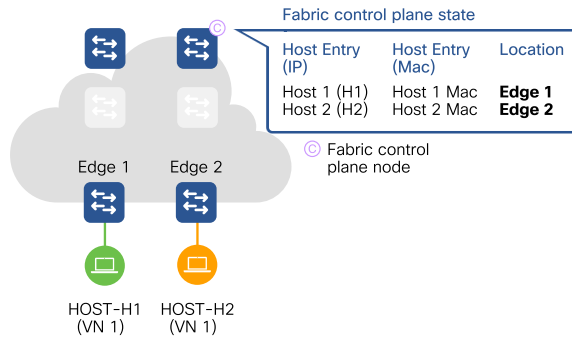
1   Endpoint 1 on edge 1 will be registered to the fabric control plane node. The registration includes Endpoint 1's IP address, MAC address, and location [which is fabric switch edge 1].

2   Endpoint 2 on edge 2 will also be registered to the fabric control plane node. The registration includes Endpoint 2's IP address, MAC address, and location [which is fabric switch edge 2].

3   When Endpoint 1 wants to communicate to Endpoint 2, edge 1 will query the fabric control plane node for the location of Endpoint 2.

4   Upon getting the reply (Endpoint 2 location is behind edge 2) it will encapsulate the traffic from Endpoint 1 using VXLAN, and send it to Endpoint 2 (via edge 2).

5   Once this traffic arrives at edge 2, it will be decapsulated and forwarded along to Endpoint 2.

6   The reverse applies when Endpoint 2 wants to communicate back to Endpoint 1.

Data plane operation

VXLAN requires an underlying transport network (the underlay). Underlay data plane forwarding is required to provide communication between endpoints connected to the

fabric. The underlay network can also be automated by Cisco DNA Center. The follow-
ing diagram illustrates data plane forwarding in a VXLAN-encapsulated network.

DIAGRAM   Fabric data plane operation

SD-Access fabric doesn't change the semantics of Layer 2 or Layer 3 forwarding and al-
lows the fabric edge nodes to perform overlay routing or bridging functions. As such,
the edge nodes offer a set of different gateway functions as outlined below:

- **Layer 2 virtual network interface** (L2 VNI): In this mode, frames from an L2 VNI
  are bridged to another L2 interface. The bridging will be done within the context
  of a bridge domain. A common implementation of a L2 gateway would use a
  VLAN as the bridge domain and will make the L2 VNI a member port of the
  VLAN. The edge nodes will bridge traffic between the L2 VNI and the destination
  L2 port in the VLAN.

- **Layer 3 virtual network interface** (L3 VNI): In this mode, frames from an L3 VNI
  are routed to another L3 interface. The routing will be done within the context
  of a routing instance. A common implementation of an L3 gateway would use a
  VRF as the routing instance and will make the L3 VNI a member port of the VRF.
  The edge nodes will route traffic between the L3 VNI and the destination L3 in-
  terface in the VRF.

In order to provide client mobility and "stretching" of subnets, SD-Access leverages a
capability known as **distributed anycast default gateway**. This provisions the Layer 3

interface (default gateway) onto every edge node in the fabric. For example, if the 10.10.0.0/16 subnet was defined in the fabric, and the default gateway defined for that subnet was 10.10.0.1, then this virtual IP address (with a corresponding virtual MAC address) would be programmed identically on every edge node.

The benefit is that it significantly simplifies the endpoint deployment and facilitates roaming within the fabric infrastructure, since the default gateway is identical on every fabric edge node. This also optimizes traffic forwarding, as traffic from an endpoint going to an off-subnet destination is always L3-forwarded at the first hop. It is never hairpinned to a remote location for L3 traffic forwarding, as some older (non-fabric) stretched subnet solutions require.

In addition, the underlay network can be used to deliver multi-destination traffic to endpoints connected to a common Layer 2 broadcast domain in the overlay network. This includes broadcast and multicast traffic in the fabric. Broadcast traffic in SD-Access fabric is mapped to an underlay multicast group and sent to all the edge nodes within that common broadcast domain. This is described in more detail in the following chapters.

Wireless in SD-Access fabric

There are two primary options for integration of wireless with SD-Access:

- **SD-Access wireless**: provides full integration of wireless traffic into the fabric

- **Traditional wireless** or **over-the-top (OTT)**: legacy wireless traffic is carried over the fabric

In the next sections, we describe the technical implementation of these two modes.

SD-Access wireless

In SD-Access fabric, wired and wireless are part of a single integrated infrastructure and behave the same way in terms of connectivity, mobility, and policy enforcement. This brings a unified experience for users independently of the access media.

In terms of control plane integration, the fabric wireless LAN controller notifies the fabric control plane node of all wireless client joins, roams, and disconnects. In this way, the control plane node always has all the information about both the wired and wireless clients in the fabric, and always serves as the "single source of truth" for client location.

In terms of data plane integration, the fabric WLC instructs the fabric access points to form a VXLAN overlay tunnel to their adjacent fabric edge nodes. This AP VXLAN tunnel carries the segmentation and policy information to/from the edge node, allowing connectivity and functionality identical to that of a wired host.

The fabric WLC is physically located outside the fabric, external to a fabric border node. This can be in the same LAN underlay as SD-Access, but external to the fabric overlay. The fabric WLC may connect either directly to the border node, or be multiple IP hops away (e.g. a local data center). The IP subnet prefix of the fabric WLC must then be advertised into the underlay routing domain, for AP onboarding and management (via traditional CAPWAP control plane). The fabric WLC can also be hosted physically on the same device hosting border and control plane (currently embedded WLC functionality is supported only on Catalyst 9000 family switches).

The fabric APs are connected directly to the fabric edge nodes in the fabric overlay. Alternatively, the APs may be connected to SD-Access extended nodes. The APs leverage the stretched subnet capability and anycast gateway functionality on the fabric edge nodes. This allows all the fabric APs throughout the fabric site to be on the same subnet.

Note   The fabric WLC must be on a network that provides 20ms or less AP-to-WLC latency, since fabric APs operate in local mode.

DIAGRAM   SD-Access wireless deployment
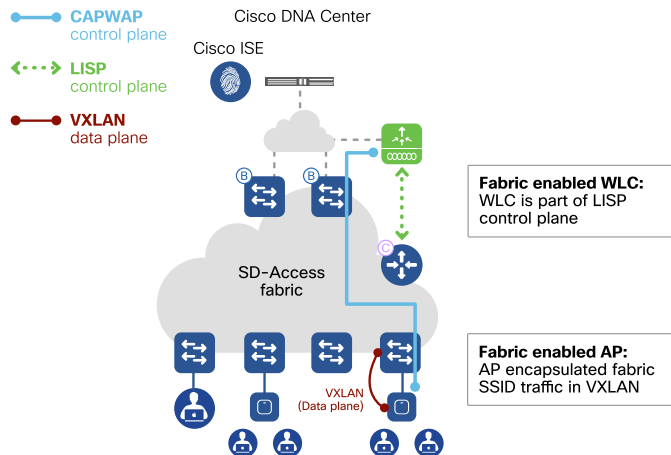
Once fabric capability has been enabled on the WLC, the AP join process works as follows:

- The AP initializes and joins the WLC via CAPWAP — the same way it does today.

    - All management and control traffic such as AP image management, licensing, radio resource management (RRM), client authentication, and other functions, leverage this CAPWAP connection.

- After the AP joins the WLC, the WLC checks whether the AP is capable of supporting fabric.

    - If it does, fabric capability is automatically enabled on the AP.

- Once the appropriate signalling is complete, the AP forms a VXLAN tunnel with the fabric edge node.

Fabric capability is enabled on a per-WLAN basis. The client subnet and L3 gateway are located on the fabric edge nodes in the overlay (in contrast with the current CUWN model, where they exist on the WLC).

When a client joins a fabric-enabled wireless network, the process works as follows:

- The client authenticates with the WLC on an SSID enabled for fabric.

- The WLC notifies the AP to use VXLAN encapsulation to the fabric edge node and to populate the appropriate VN/SGT for that client in the VXLAN packet.

- The WLC registers the client MAC address in the fabric control plane node database.

- Once the client receives an IP address for itself via DHCP, the existing control plane entry is updated by the fabric edge node — and the MAC and IP address are mapped and correlated.

- The client is now free to start communications in the network.

Guest access in SD-Access wireless
Guest access can be enabled in SD-Access wireless in one of the two ways:

- WLC with separate VN for guest

- Dedicated WLC as guest anchor

Separate VN for guest
In this mode, the guest network is just another virtual network in the SD-Access fabric. This leverages end-to-end fabric segmentation by using a separate VN (and SGTs for different guest roles, if needed) to separate the guest traffic from the other enterprise traffic.

There are two ways to enable this model:

1. Use a shared (same) control plane and border as enterprise traffic

2. Use a separate control plane and border which is dedicated to the guest traffic

    - The separate control plane and border can be at the same site with wireless guest clients or at a different site

DIAGRAM   Shared control plane and border

In this option, the guest VN is extended from the shared border and control plane to the firewall via VRF-Lite in the DMZ.

DIAGRAM   Separate (guest) control plane and border

In this option, separation is achieved at all levels for guests to isolate them from the enterprise users. Guest users will be registered in the dedicated "guest" control plane node. The fabric edge node queries the separate guest control plane node and forwards the traffic to the guest border encapsulated in VXLAN. In this mode the guest VN is extended from the guest border to the firewall via VRF-Lite in the DMZ.

WLC controller as guest anchor


DIAGRAM   Legacy guest anchor
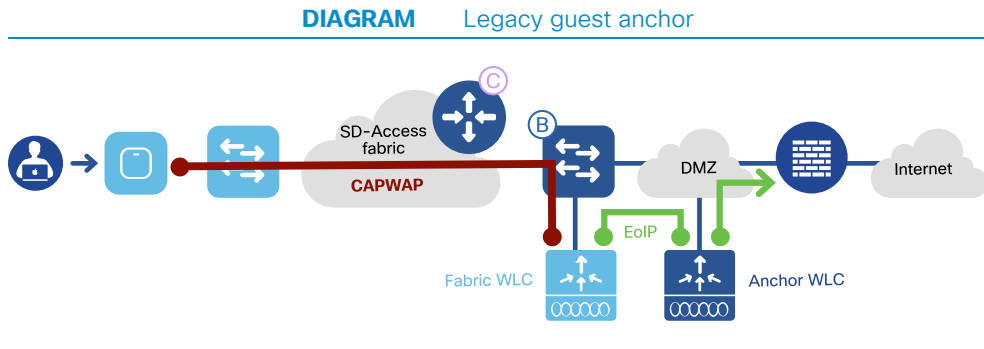
The existing WLC guest anchor solution continues to work as it does today. This mode can be used as a migration step when guest anchor controllers already exist in the DMZ. In this case, the mobility tunnel is formed between the fabric WLC and the guest anchor controller and the guest SSID is anchored at the anchor controller. All guests are tunnelled to the guest anchor by the controller.

Traditional wireless or over-the-top

For backward compatibility, the current Cisco Unified Wireless Network (CUWN) or other traditional centralized wireless architectures are fully supported with SD-Access. In this solution, the wireless infrastructure rides on top of the SD-Access fabric architecture, but the wireless infrastructure is not aware of or integrated with the SD-Access wired network (hence the term "over-the-top") from a control plane, data plane, or policy plane perspective.

This deployment method provides a graceful migration step toward fully-integrated SD-Access wireless. This can also be used to deploy non-Cisco wireless solutions with SD-Access. However, such integration must be tested and validated prior to deployment and use.

The WLC connects to the SD-Access network either directly to (or multiple hops away from) the border.

DIAGRAM   Traditional CUWN " over the top"

CAPWAP
control plane

CAPWAP
data plane

Cisco DNA Center

Cisco ISE
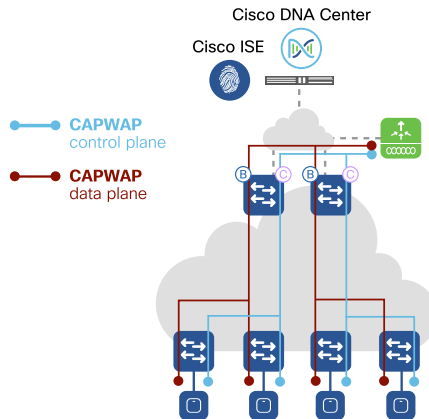
SD-Access extension for IoT

Extended nodes are switches that run in pure Layer 2 mode, and do not natively support fabric technology. These Layer 2 switches will connect to a fabric edge node via traditional Layer 2 methods over single or multiple links (using a portchannel). The VLANs / IP subnets configured on the extended node switches will obtain similar policy segmentation and automation benefits that the fabric natively provides.

The fabric extends the appropriate subnets of the fabric to SD-Access extended nodes using 802.1q Layer 2 trunking. This allows the extended node to perform normal local switching. When traffic leaves the extended node, to the connected fabric edge node, the traffic benefits from the centralized policy and scalability of the fabric.

Hosts connected to the extended nodes can be wired endpoints and/or APs to aggregate wireless endpoints.

DIAGRAM   Fabric extended nodes

Fabric considerations

There are various network underlay considerations to be aware of when implementing SD-Access as these may impact the operation of the fabric overlay. The most important of these are summarized below.

Maximum transmission unit (MTU)

It is recommended to avoid fragmentation and reassembly of traffic between network devices. It is therefore required to increase the maximum transmission unit (MTU) in the underlay network by at least 50 bytes in order to accommodate the VXLAN header (or 54 bytes if an 802.1q header is also required), on all network devices connecting to the fabric underlay.

Jumbo frame support in the underlay network is strongly recommended if the overlay uses frame sizes larger than 1500 bytes. The recommended global or per-interface MTU setting for an SD-Access fabric deployment is 9100 bytes.

In the event that jumbo frame support is not available, the TCP-adjust-MSS option can be used instead of using jumbo frames in the underlay. It should be noted that tcp-adjust-mss will only work for TCP traffic (not UDP).

Underlay interface addressing

The recommended network interface and address design is Layer 3 routed point-to-point interfaces which can be addressed with a /30 or /31 subnet mask.

Underlay routing protocol

Cisco DNA Center LAN automation will deploy a standards-based IGP routing protocol (IS-IS or OSPF) to automatically bring up the underlay.

Similarly, for manual underlay configurations, it is recommended to deploy the IS-IS routing protocol. Other routing protocols (e.g. OSPF) are also supported but may require additional configuration.

IS-IS deployment

As mentioned, Cisco DNA Center will deploy IS-IS as the best practice underlay routing protocol. This link state routing protocol is gaining popularity for large-scale fabric environments, although has primarily been deployed in service provider (SP) environments.

IS-IS uses connection-less network protocol (CLNP) for communication between peers and doesn't depend on IP for this purpose. With IS-IS, there is no SPF calculation on link change. SPF calculation only occurs when there is a topology change which helps with faster convergence and stability in the underlay. No significant tuning is required for IS-IS to achieve an efficient, fast-converging underlay network.

OSPF deployment

In many cases, OSPF (open shortest path first) is a common choice in enterprise deployments. Like IS-IS, OSPF is a link-state routing protocol. The OSPF default interface type used for Ethernet interfaces is "Broadcast," which inherently results in a designated router (DR) and/or backup designated router (BDR) election, thus reducing routing update traffic.

This is unnecessary in a point-to-point network. In a point-to-point network, the "broadcast" interface type of OSPF adds a DR/BDR election process and an additional type 2 link state advertisement (LSA). This results in unnecessary additional overhead, which can be avoided by changing the interface type to "point-to-point".

Fabric deployment models

Multiple deployment options exist for SD-Access fabric. In this chapter, we will explore several of the options available.

- **Fabric site**: a single fabric contained within a single site

    - **Fabric-in-a-box**: All fabric functionality (control plane, border, edge) contained within a single platform, for very small sites

- **SD-Access for distributed campus**: multiple fabrics connected using one or mote transit networks

Fabric site

A fabric site is a portion of the fabric which has its own set of control plane nodes, border nodes, and edge nodes.

Key characteristics of a single fabric site are:

- A given IP subnet is part of a single fabric site (except when VN anchoring is in use)

- L2 extension is only within a fabric

- L2 / L3 mobility is only within a fabric

- No context translation is necessary within a fabric

A fabric site is, in principle, autonomous from other fabric sites from the connectivity perspective.

The diagram below depicts a fabric site.

DIAGRAM   Fabric site

Many SD-Access deployments will be satisfied within the scale associated with a single-site fabric. However, some fabric deployments may need to scale to larger — or to smaller — sizes.

A given fabric site can have different scale characteristics:

- Large-scale fabric site: multiple horizontally-scaled devices, per fabric site

- Fabric-in-a-box: all fabric functions are on a single device (site)

Fabric-in-a-box

Fabric-in-a-box allows the border node, edge node and control plane node functions to operate on the same fabric device. In addition if wireless is used, the Catalyst 9000 family of switches can be configured to host an embedded wireless LAN controller function. Thus, a small site can gain the advantage of fabric benefits, while still maintaining local resiliency and failover mechanisms.

SD-Access across multiple sites

An SD-Access fabric may be composed of multiple sites. Each site may require different aspects of scale, resiliency, and survivability. The overall aggregation of sites (i.e. the fabric) must also be able to accommodate a very large number of endpoints, scale horizontally by aggregating sites, and having the local state be contained within each site.

Multiple fabric sites corresponding to a single fabric will be interconnected by a transit network area. The transit network area may be defined as a portion of the fabric which interconnects the borders of individual fabrics, and which has its own control plane nodes — but does not have edge nodes. Furthermore, the transit network area shares at least one border node from each fabric site that it interconnects. The following diagram depicts a multi-site fabric.

DIAGRAM Multi-site fabric

Role of the transit network area

In general terms, a transit network area exists to connect to the external world. There are several approaches to external connectivity, such as:

- IP transit

- Software-Defined WAN transit (SD-WAN)

- SD-Access transit (native)

In the fabric multi-site model, all external connectivity (including internet access) is modeled as a transit network. This creates a general construct that allows connectivity to any other sites and/or services.

The traffic across fabric sites, and to any other type of site, uses the control plane and data plane of the transit network to provide connectivity between these networks. A local border node is the handoff point from the fabric site, and the traffic is delivered across the transit network to other sites. The transit network may use additional features. For example, if the transit network is a WAN, then features such as performance routing may also be used.

To provide end-to-end policy and segmentation, the transit network should be capable of carrying the endpoint context information (VRF and SGT) across this network. Otherwise, a re-classification of the traffic will be needed at the destination site border.

Fabric control plane state distribution

The local control plane in a fabric site will only hold state relevant to endpoints that are connected to edge nodes within the local fabric site. The local endpoints will be registered to the local site control plane by the local edge devices, as with a single fabric site. Any endpoint that isn't explicitly registered with the local control plane will be assumed to be reachable via the border nodes connected to the transit area.

At no point should the local control plane for a fabric site hold state for endpoints attached in other fabric sites (i.e. the border nodes do not register information from the transit area). This allows the local control plane to be independent of other fabric sites, thus enhancing the overall scalability of the solution.

The control plane in the transit area will hold summary state for all fabric sites that it interconnects. This information will be registered to the transit area control plane by the border nodes from the different fabric sites. The border nodes register EID (endpoints IDs) information from their local fabric site into the transit network control plane for summary EIDs only, further improving overall scalability.

Note   It is important to note that endpoint roaming is only within a local fabric site, and not across sites.

How to create a fabric across multiple sites?

Normally, when there are many small branches, the biggest challenge is how to manage these sites. Creating a separate fabric domain for each site is a cumbersome task. The creation of a multi-site fabric is a three-step process:

1. Create transit network(s):

Transit network creation depends on type of transit network.

    1   SD-Access transit networks are created by enabling fabric control plane node(s) for the transit network

    2   SD-WAN transit networks are created using an SD-WAN control plane protocol. As of this writing, in the case of Cisco SD-WAN, the control plane protocol is overlay management protocol (OMP)

    3   IP transit networks are created using any traditional routing protocol (e.g. BGP)

2. Create fabric site(s):

Creation of a fabric site involves the following steps (as described above):

    1   Add control plane node(s) for this site

    2   Add border node(s) to this site

    3   Add edge nodes to this site

3. Connect local fabric site(s) to the transit network(s):

Fabric site(s) are connected to the transit network using local border node(s).

Wireless in fabric across multiple sites

Fabric wireless is an integral part of the fabric multi-site design. Each local fabric site has its own dedicated WLC and is responsible for managing the wireless infrastructure at that local site. Since there is no Layer 2 mobility across sites in a multi-site design, the WLCs that are dedicated to each site are not configured in a mobility group. Hence, seamless roaming is not supported across sites.

VN anchoring

In some deployments, it may be advantageous to anchor traffic from a VN (one which is deployed across multiple sites) back to a common, central location. This type of deployment, referred to as VN anchoring, allows for the traffic from a given VN at those dispersed sites to be aggregated back to a central location, using a single common subnet, rather than having to define and use subnets per-site for that VN as would otherwise be the case.

An example of this is for guest traffic for wired users. Without VN anchoring, if 20 sites were in use, and guest wired traffic was desired, it would be necessary to define a guest-wired-VN across all of the sites, and then assign a subnet per-site mapped into this VN (20 subnets in total) — since subnets are unique to a site and cannot be stretched between sites. While such a deployment would be functional, it obviously adds to the complexity of designing the solution, as well as deploying and managing it. With VN anchoring, all traffic for this designated guest-wired-VN at each site will be tunneled back to a central location over VXLAN, allowing a single subnet to be deployed at that central site for all of the clients attached to this dispersed VN. Note that not all VNs at a site need to be anchored in this way — most of the VNs for any given site would be deployed in a standard (non-anchored) fashion, with anchoring used only for designated use cases and the VNs associated with them.

Since it allows for a more centralized and simplified subnet structure, the use of VN anchoring simplifies the deployment involved, and is ideal for use cases (such as guest access) where it would be desirable to drop off all of the traffic for this VN at a central location in any case (for example, at a firewall for guest traffic). Other use cases where a less-centralized approach is desirable will continue to use the standard SD-Access VN-separation methods as previously described, without anchoring.
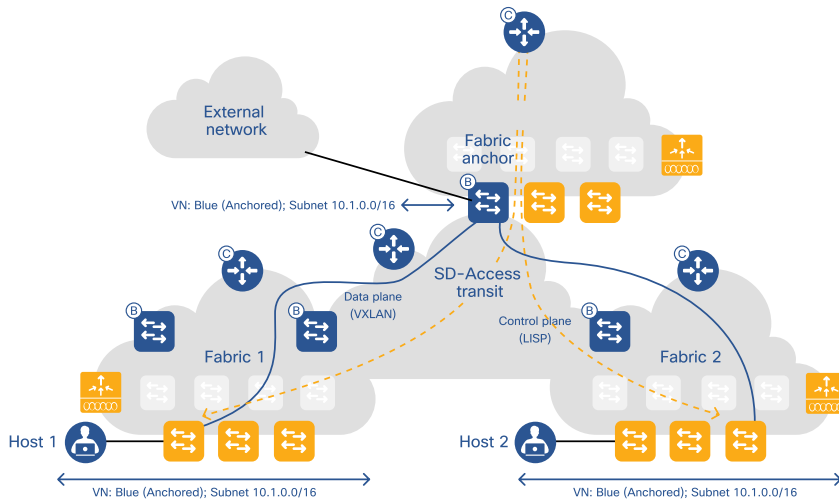
IoT deployment models

As described earlier, SD-Access supports specific Layer 2 connected switches, known as SD-Access extended nodes, allowing the use of indirectly connected endpoints (i.e. endpoints that are not directly attached to a fabric edge node).

SD-Access extended nodes may be connected to a fabric edge node in one of the following ways:

- Connected over point-to-point links to an edge node

- Connected over an STP or REP ring to an edge node

The STP or REP ring–based deployment is important for use with IoT, since the ring infrastructure is used to provide additional resiliency for connection of the various IoT devices. The extended node does all of the endpoint onboarding for devices connected

to its ports, but policy is only applied by the edge nodes. This means that the traffic between endpoints directly connected to the same extended node, or between extended nodes connected to a ring, are not subject to policy enforcement.

Thus, the support of fabric edge node allows the IoT devices to utilize the fabric benefits, including Layer 2 extension and segmentation.

External connectivity

Obviously, the SD-Access fabric will need to connect to external locations. Many external connection options exist, depending on the specific environment. Some common examples include:

- Other campuses

- Branch offices (over a WAN)

- Data centers

- Cloud networks

For each of these external connection options, we will describe solutions for extending SD-Access policy elements (VNs and SGTs) to remote sites. The first consideration is VN, which is translated into a standard VRF format. The second consideration is SGT, which may be transported natively, or re-classified at the remote site.

Let's begin with discussing WAN integration. There are three main WAN hand-off options for an SD-Access fabric:

1. SD-Access transit

2. SD-WAN transit

3. IP-based transit

As described earlier, three architecture aspects need consideration:

- Control plane (routing/signaling protocols)

- Data plane (encapsulation)

- Policy plane (endpoint context)

SD-Access transit

In an SD-Access transit deployment model, the same control, data, and management plane is used end-to-end, and the concept is nothing more than extending SD-Access across multiple fabrics. The SD-Access transit model is typically used across metro Ethernet, targeting close geographical proximity of locations.

The SD-Access transit model natively extends both the VN and SGT policy, by encoding these within the VXLAN encapsulation and by leveraging the LISP control plane, over any transport network (generally in a metropolitan area, for high bandwidth and low latency) that supports IP routing. In this design model, the SD-Access border node serves as a border connecting each local fabric to the SD-Acccess transit fabric.

Note   SD-WAN transport is the long-term preferred design model for WAN environments, including for extending SD-Access domains. The SD-Access transit model offers network designers with additional options for extending SD-Access domains to one or more close geographical locations (particularly for Metro Ethernet environments).

SD-WAN transit
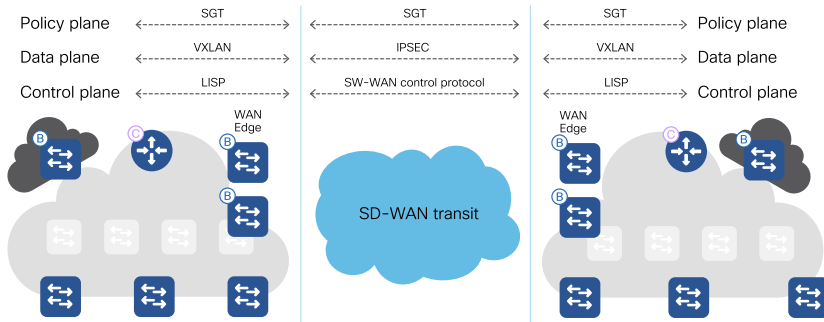
In an SD-WAN deployment model, the SD-Access border node and SD-WAN edge router are combined on the same device. The VN is carried from an edge node to this SD-Access border / SD-WAN edge router and then the context (VN and SGT) is copied from the VXLAN header into the SD-WAN header.

The context thus carried across the SD-WAN deployment allows the SD-Access system to maintain the virtual network (VRF) and group (SGT) classification end-to-end across the network, which allows the policies to be extended across the network.

IP-based transit

While there are many possibilities of "IP-based" transit areas, the SD-Access IP-based transit model covers two common architectures

- IP transit — MPLS VPN transport

- IP transit — DMVPN transport

In both IP-based transit models, VRF-Lite with BGP is used to provide routing information (for each VN) to the external neighbor.

IP transit — MPLS VPN transport

DIAGRAM   MPLS VPN transport example
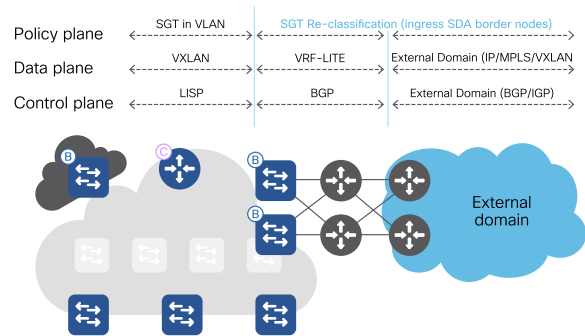
In the MPLS-VPN deployment model, VRF-Lite is used to extend the VN from the SD-Access border node (CE in an MPLS-VPN model) to the MPLS PE node.

MPLS encapsulation is unable to natively transport the SGT tag. To overcome this, the remote SD-Access border node can leverage the SGT re-classification method (an SGT can be defined for an incoming packet on either an interface/sub-interface or subnet), or can alternatively leverage SXP (SGT exchange protocol, a TCP-based connection which extends IP-to-SGT mappings) between border nodes to provision the SGT tags between SD-Access domains. However, it should be noted that SXP requires additional configuration, and has additional complexity and scale considerations.

IP transit — VRF-Lite over DMVPN transport

DIAGRAM   VRF-Lite over DMVPN transport example
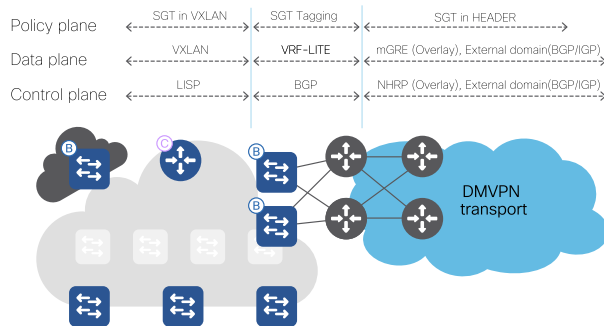
In the dynamic multipoint VPN (DMVPN) deployment model, VRF-Lite is used to extend the VN from the SD-Access border node to the DMVPN edge router.

DMVPN mGRE encapsulation is capable of carrying the SGT tag natively over the WAN. This offers a much simpler solution as there is no need to leverage SGT re-classification or SXP as is the case with the IP / MPLS transit option. However, it should be noted that SD-WAN transit is generally preferred to DMVPN transit for new implementations.

Data center connectivity

For data center interconnection, we will primarily focus on Application Centric Infrastructure (ACI) interoperability.

Note   This section describes the connection if ACI is used in the data center. If the datacenter is a non-ACI (e.g. a traditional IP-based), then a border connected via standard IP-based transit works as-is (i.e. no context transfer between the SD-Access and DC domains).

Cisco ACI is a modern data center architecture which provides centralized automation and policy-driven application profiles, similar to Cisco SD-Access. ACI uses endpoint

groups (EPGs) as the main policy construct for defining policy within the ACI fabric, and carries these in VXLAN, similar to SD-Access. EPGs identify application workloads in the ACI fabric.

Cisco offers the ability to federate endpoint identifiers (SGT and EPG) between the SD-Access and ACI fabric domains, beginning in Cisco ISE 2.1. Cisco ISE will share user-selected SGTs to ACI, as well as read EPGs from ACI. The sharing of the SGT takes the form of writing SGT names into the ACI fabric as EPG names, allowing the ACI solution to then build specific policies based on these constructs.

ACI can then provide all of its policy services to relationships between the SD-Access SGTs and the ACI EPGs. Similar to what is described above, Cisco ISE reads EPGs and shares them to pxGrid ecosystem partners, which includes the Cisco DNA Center managing the SD-Access fabric. This allows SD-Access to build and enforce both user/device and application policies in the SD-Access fabric.

DIAGRAM   Policy group identity federation
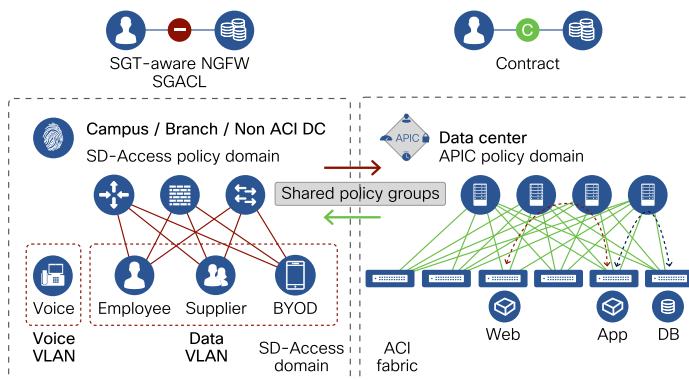
Note   Some consideration in the ACI fabric is necessary since ACI supports multiple tenants (each tenant can have multiple VRFs). The recommended way to connect SD-Access to the ACI fabric is to have a single "shared services" VRF applied across all

of the tenants, to provide a common Layer 3 external connectivity for them. Cisco ISE then writes all the SGTs into this shared VRF and the SGTs can be applied to all tenant policies.

DIAGRAM   Shared VRF between SD-Access and ACI
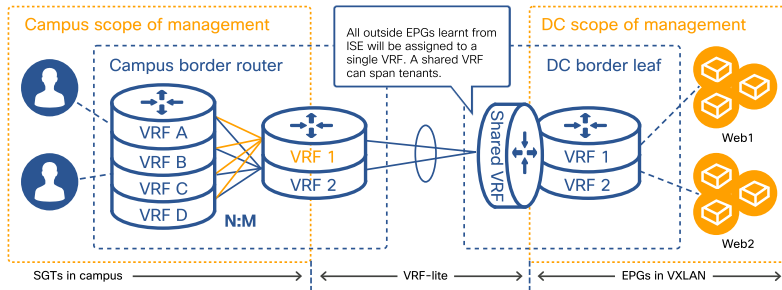
DIAGRAM   Group identity sharing via pxGrid
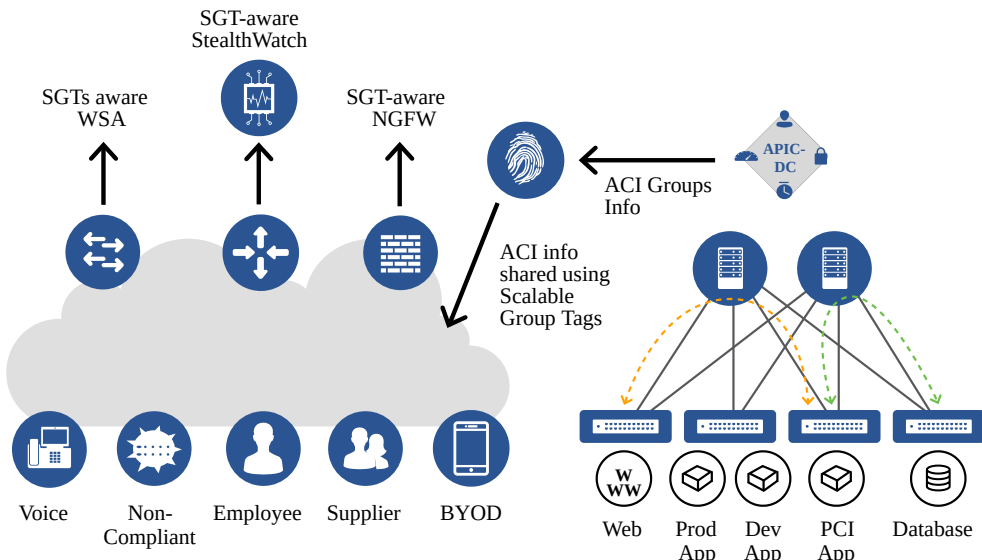
An additional benefit of the SD-Access and ACI integration is that it allows for the sharing of IP/SGTs derived from ACI with other IT security partners (via the pxGrid ecosystem). This allows a normalized identification for users/devices and applications to be represented in telemetry, and used by partner systems such as Cisco Stealthwatch, and/or to be used in a security policy for functions such as the Cisco Web Security Appliance (WSA) and Firepower Next Generation Firewall.

Shared services (DHCP, DNS, IPAM etc.)

In all network deployments there is a common set of resources needed by every endpoint. The following are some common examples:

- Identity services (e.g. AAA/RADIUS)

- Domain name services (DNS)

- Dynamic host configuration protocol (DHCP)

- IP address management (IPAM)

- Monitoring tools (e.g. SNMP)

- Data collectors (e.g. Netflow, syslog)

- Other infrastructure elements

These common resources are often called "shared services". These shared services will generally reside outside of the SD-Access fabric. In the majority of cases, such services reside in the global routing table (GRT) of the existing network (and are not in a separate VRF).

Note   As previously described, SD-Access fabric clients operate in overlay virtual networks. Thus, if the shared services are part of the global routing space, some method of inter-VRF routing is required.

There are two ways to do inter-VRF routing

- Use of SD-Access VN extranet

- Use of a fusion router and/or firewall

Note   The SD-Access VN extranet feature is the preferred method to provide inter-VN communications, as it is both simpler and more scalable than the alternative methods.

SD-Access VN extranet

SD-Access VN extranet provides a flexible and scalable method to achieve inter-VN communications.

The advantages of using VN extranet include:

- **Single touchpoint:** Cisco DNA Center automates the inter-VN lookup policy, thus providing a simplified single point of management.

- **Avoids route duplication:** inter-VN lookup is done in the fabric control plane software, which avoids duplicating hardware routing entries.

- **Maintains SGT context:** SGTs are also maintained, thus providing a simplified and consistent policy and enforcement.

- **No hairpinning:** inter-VN forwarding occurs at the fabric edge (after lookup), and thus traffic does not need to hairpin at the border node, improving performance and scalability.

Another point worth noting is that, depending on the requirements involved, a separate VN can be used for each of the common resources that are needed (i.e. for shared services VN, internet VN, and data center VN, etc.).

DIAGRAM   Multiple VN design using SD-Access extranet
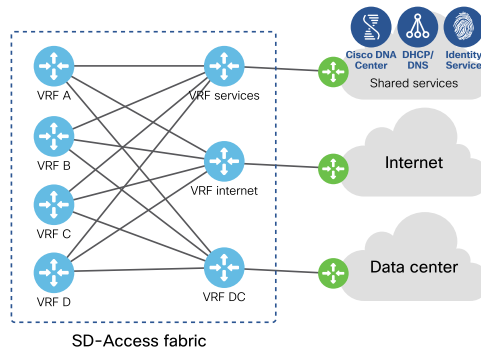
An important concept in SD-Access extranet is the separation of provider VNs and subscriber VNs. A common example of a provider VN is shared services (which are needed by multiple other VNs). A subscriber VN is where the clients (that need the services from the provider VN) reside.

VN extranet example

To understand how provider and subscriber VNs are used with SD-Access VN extranet, consider the scenario below:

VRF RED is a provider VN and VRFs BLUE and GREEN are subscriber VNs.

1   Host C1 connects to edge 1 on VRF BLUE

2   Host C2 connects to edge 3 on VRF GREEN

3   The shared service is an endpoint that connects to edge 2 in VRF RED

4   The control plane node registers the endpoint entries in its database

5   The inter-VN policy dictates —

   -   VRFs BLUE and GREEN should access resources in VRF RED

- VRFs BLUE and GREEN cannot access each other's resources

DIAGRAM   SD-Access VN extranet operation

1 Host C1 now wants to access the shared service

2 Edge 1 receives the packet and sends a query to the control plane node

- The control plane node checks its database for the shared service host in VRF BLUE (but does not find an entry)

- The control plane node then refers to the extranet policy, which instructs it to check entries in VRF RED

3 The control plane node does find the entry for shared service host in VRF RED

4 The control plane node then instructs edge 1 to encapsulate the packet in VRF RED and forward to edge 2

- The result is cached by edge 1, for any future traffic destined for the shared service host

A similar sequence occurs for the host C2 in VRF GREEN to access resources in VRF RED.

Fusion router or firewall inter-VRF routing

Another option for inter-VRF routing is to leverage a fusion router. A fusion router is simply an external router or firewall (located outside of the SD-Access fabric), which performs basic inter-VRF route leaking (import/export of routes in different VRF tables) in order to "fuse" the VRFs together.

DIAGRAM   Shared services with an external fusion router.

Several fusion router design considerations apply, depending on whether the shared services are in located in the global routing table (GRT) or located in another VRF.

Shared services in the GRT

- The fabric border node forms an external BGP routing adjacency with the fusion router, using the global routing table

- On the border node, the same routing adjacency is formed in each VRF context (BGP address-family)

- On the fusion router, routes between the SD-Access VNs are then fused (imported) with the GRT of the external network

Note   Multi-protocol BGP is the routing protocol of choice for this route exchange since it provides an inherent way of preventing routing loops (using AS_PATH attribute). Other routing protocols can be used, but require complex distribute-lists and prefix-lists to prevent loops.

Shared services in separate VRF:

- A separate routing adjacency is formed for each BGP address family, between the border node and fusion router

There are four main challenges using the fusion router method to achieve inter-VN communication:

- **Multiple touchpoints:** manual configuration must be done at multiple points (wherever the route-leaking is implemented)

- **Route duplication:** routes leaked from one VRF to another are also programmed in the hardware tables for both VRFs, resulting in greater TCAM utilization

- **Loss of SGT context:** SGT group tags are not maintained across VRFs and must be re-classified once the traffic enters the other VRF

- **Traffic hairpinning:** Inter-VN traffic needs to be routed to the fusion router, and then back to the fabric border node

Fabric packet walks

Now that we have reviewed the various SD-Access fabric components, and examined their operation and interactions, let's examine fabric packet forwarding and packet flows.

In this chapter the following packet walks are discussed in an SD-Access context:

- Client DHCP operation

- ARP operation

- Unicast wired-to-wireless

- Wireless mobility

- Unicast from fabric to external client with sender and receiver in the same subnet (between SD-Access and external networks)

- Unicast from fabric to external client with sender and receiver in different subnets (between SD-Access and external networks)

- Fabric multicast (in the overlay)

- Native multicast (in the underlay)

- Broadcast support

We will examine why each packet walk is relevant in the context of the SD-Access fabric, and any notable aspects of that packet walk to overall fabric operation and use.

Client DHCP operation in SD-Access

Once the fabric is configured, you can start onboarding clients into the network.

The connected clients will start sending DHCP requests to obtain an IP address. The DHCP flow in fabric is different compared to traditional networks, so let's go through the process step-by-step.

DIAGRAM   Client DHCP operation
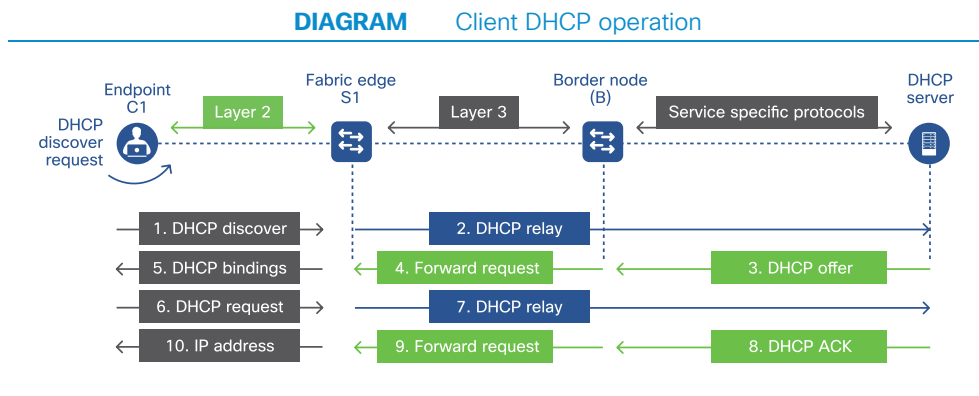
The switchport connecting to the client is configured in a VLAN. The fabric edge has an anycast gateway for the client IP pool.

1 **DHCP discover**: The client C1 sends a DHCP discover packet (broadcast) to the S1 node.

2 **DHCP relay:** S1 adds its routing locator (RLOC) IP address in Option 82 to the DHCP discover. DHCP relay on S1 coverts the broadcast DHCP discover to unicast and forwards it out to the fabric border.

3 **DHCP offer**: The DHCP server (located outisde of the fabric) replies with a DHCP offer destined to the anycast gateway IP address. The Option 82 information encoded by S1 is copied back unaltered into the DHCP offer by the DHCP server.

4 **Forward request**: The DHCP offer arrives back to the border node. The border node is configured with a loopback which is the same IP address as the anycast

gateway. The border will inspect the packet for Option 82 and based on the RLOC information encoded there, will forward the DHCP request back to S1. Note that this is necessary since the anycast gateway address which the DHCP server sent the offer back to is not unique within the fabric (i.e. it is shared by all fabric edge nodes).

5   **DHCP bindings**: Once S1 gets the DHCP offer from the border, it updates its DHCP binding database and forwards the DHCP offer packet to C1.

6   C1 will now send a DHCP request packet (unicast).

7   The request is forwarded to the border, which in turn forwards this to the DHCP server.

8   The DHCP server replies to the border with a DHCP offer to the anycast gateway.

9   The border will inspect the packet for Option 82 to find the RLOC to forward the ACK to the appropriate fabric edge S1.

10 C1 now has an IP address.

ARP operation in SD-Access

The SD-Access fabric provides many optimizations to improve unicast traffic flow, and to reduce the unnecessary flooding of data. One of the first optimizations is ARP suppression. While traditional ARP flooding is workable with the smaller-scale subnets deployed in a traditional enterprise networks, this would be very wasteful of bandwidth in larger networks. In an SD-Access deployment this is fully addressed.

Within SD-Access, anycast gateways are populated on all fabric edge nodes. The endpoints in these subnets could reside anywhere in the fabric (or even roam between edge nodes). ARP operation in SD-Access has been designed to improve the efficiency of this process in a fabric deployment.

In the figure below, a wired client C1 is connected to a switchport of the edge node S1. It has already obtained an IP address from DHCP. The client wants to communicate to another client C2 connected in the same subnet, on a different edge node S2.

DIAGRAM   ARP operation in SD-Access

1   To start communicating to client C2, client C1 sends out an ARP request to discover the MAC address corresponding with the IP address of client C2.

2   Edge node S1 processes this ARP request

- S1 registers the client C1, to the control plane.

- S1 floods the ARP request out the local ports in the same VLAN as the client that sent the ARP request. S1 also sends a query to the control plane node to check if it has the MAC address corresponding to the IP address of C2.

- The control plane has the entry for C2's MAC address and IP address, and returns the MAC address information to S1. Along with the MAC address the control plane also returns Client C2's location (i.e S2)

- S1 caches this information to suppress subsequent queries for this client.

- S1 then replaces the broadcast address in the ARP request, with the MAC address of client C2.

3   S1 encapsulates the ARP request in unicast VXLAN (making the broadcast ARP packet a directed unicast) with a destination of S2.

- S1 applies the appropriate policy context (VN and SGT) and transmits the VXLAN frame to S2.

4   S2 decapsulates the VXLAN header from the incoming packet and knows that the packet is destined for C2

- S2 then forwards the ARP request to C2.

- Client C2 examines the incoming ARP packet then responds with its MAC address in an ARP reply packet.

- S2 accepts the incoming ARP reply packet, encapsulates in VXLAN towards S1, and forwards it to S1.

- S1 decapsulates the VXLAN header and forwards the ARP reply packet to C1 (completing the ARP discovery process).

Thus the SD-Access fabric architecture optimizes ARP discovery and avoids unnecessary ARP broadcast flooding in the fabric.

Unicast wired-to-wireless

As described earlier, wired and wireless endpoints reside in the same common SD-Access fabric. It is important to understand how traffic flow between these endpoints works within an SD-Access deployment.

In this scenario, a wired client C1, is connected to a switchport of fabric edge node S1. Wireless client C2, in a different subnet, is connected to an access point (AP) that is connected to fabric edge node S2. The MAC and IP address of both clients are already registered with the fabric control plane node. Refer to the following diagram.

DIAGRAM   Wired to wireless unicast

1   When C1 wants to communicate with C2, it will send an IP packet with the default gateway's MAC address as the destination in the packet.

2   S1 processes this packet and queries the fabric control plane to resolve the location of client C2.

3   The control plane checks in its host database and returns the IP address of S2 (the switch which the AP that C2 is associated with is located behind).

   -   S1 then caches this information (to suppress subsequent queries for this client).

- S1 applies the appropriate policy context (VN and SGT) and transmits the VXLAN frame to S2.

4   S2 receives the packet and decapsulates the VXLAN header.

- S2 examines the underlying packet that C1 sent to C2, to locate C2 (the client connected to the AP). If any policies need to be enforced on the packet based on the policies defined, S2 enforces them.

- S2 re-encapsulates the packet in VXLAN, with the policy context (L2 VNI, SGT) and forwards it to the AP.

5   The AP decapsulates the VXLAN header, and converts the packet to 802.11 format.

- The AP then forwards the packet (via RF) to the wireless client.

Wireless mobility in SD-Access

Wireless mobility is handled within the context of the SD-Access fabric itself, with wireless endpoint mobility handled between the WLC and the fabric control plane node. The following outlines how this mobility within the SD-Access fabric takes place.

We will examine two mobility use cases — one for intra-switch roaming, and one for inter-switch roaming.

Intra-switch roaming

To begin with, consider two wireless clients (C1 and C2) are connected to different APs (AP1 and AP2 respectively) on the same fabric edge node S1. Assume that there are already communications occurring between the two clients.

DIAGRAM   Wireless mobility — intra-switch
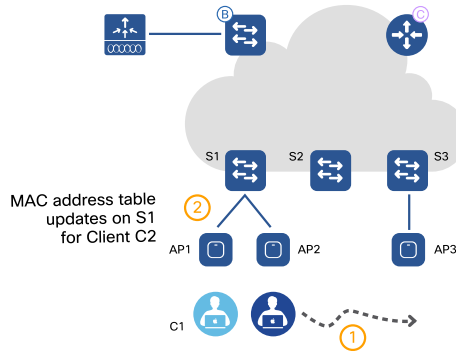
Refer to the above diagram.

1   C2 roams from AP2 to AP1, on S1.

   -   Since this is a roam within the same edge node, no additional signalling
       needs to occur in the fabric.

2   The MAC address table on S1 gets updated using the L2 frames sent by AP2.

Inter-switch roaming

Now, consider two wireless clients (C1 and C2) connected to different APs (AP1 and AP2
respectively) on the same fabric edge node (S1), with client C2 about to roam over to a
different AP (AP3) located on a different fabric edge node (S3)

DIAGRAM   Wireless mobility – inter-switch
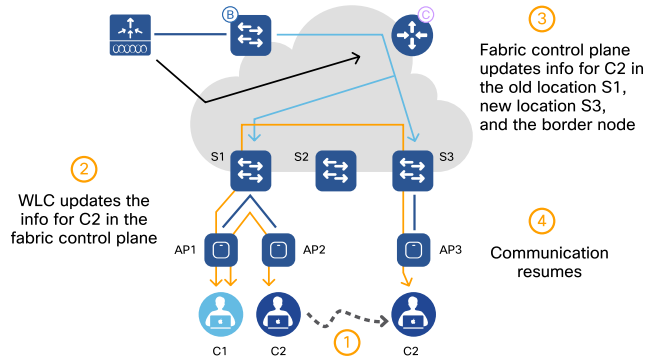
Refer to the above diagram.

1   Now client C2 roams from AP2 on S1 to AP3 on S3.

  -    The fabric APs and WLC process the client roam first.

2   The WLC notifies the fabric control plane node of the new location, AP3.

  -    The fabric control plane then updates its database with the new loca-tion, S3 (to which AP3 is attached).

3   The control plane updates the new roamed-to edge node, the old roamed-from edge node, and the border nodes with the new location of C2. This is necessary to ensure continued connectivity with any other clients (other than C1) to which C2 may be communicating.

4   Communication continues between C1 and C2.

Unicast to external networks in SD-Access

Naturally, there will always be traffic which needs to flow between the SD-Access fabric endpoints and endpoints located outside the fabric, in external networks. These traffic flows will take place via the fabric border nodes. The following outlines how this forwarding takes place.

SD-Access to external network with C1 and C2 in the same subnet

In this scenario, C1 is connected to S1. C1 wants to communicate with a host C2 which is not in the SD-Access fabric (the host is external to the fabric), C1 and C2 are in the same subnet. In this instance, an L2 Border node is used between the fabric and non-fabric portions of the network. Refer to the following diagram.

DIAGRAM   SD-Access to external network with C1 and C2 in the same subnet

1   A separate device in the fabric is configured as an L2 border. The L2 border is the last fabric hop connecting the fabric to a non-fabric network.

2   The L2 border is configured with the anycast gateway for the non-fabric network. The L2 border registers the clients in the non-fabric VLAN associated with the L2 VNI within the fabric with the fabric control plane.

3   The fabric client C1 sends a broadcast ARP request for non-fabric client C2.

4   S1 receives the ARP request and broadcasts the ARP requests on its local ports that belong to the same VLAN as client C1

5   ARP resolution for Client C2

- S1 sends a query to the control plane node. Based on this query the control plane checks if it has MAC address corresponding to the IP address of C2.

- The control plane finds the entry for C2's MAC address, and C2's IP address. The control plane returns the client MAC address and location (RLOC) information to S1. The RLOC in this case is the the L2 border, since client C2 is located outside of the fabric (in the same subnet).

- S1 caches this information in its local cache (to suppress subsequent queries for this client).

- S1 then replaces the broadcast address in the ARP request, with the MAC address of client C2.

6   S1 encapsulates the ARP request in unicast VXLAN (changing the broadcast packet to unicast) with a destination of the L2 border.

- S1 applies the appropriate policy context (VN and SGT) and transmits the VXLAN frame to the L2 border.

7   The L2 border decapsulates the VXLAN header from the incoming packet and knows that the packet is destined for C2.

- The L2 border then forwards the ARP request to C2.

- Client C2 looks at the incoming ARP packet then responds with its MAC address in an ARP reply packet.

- The L2 border accepts the incoming ARP reply packet, encapsulates in VXLAN towards S1, and forwards it to S1.

- S1 decapsulates the VXLAN header and forwards the ARP reply packet to C1. This completes the ARP discovery process. Next step is to start forwarding the unicast packets between the clients.

8 The unicast packet from client C1, destined to client C2, is forwarded to S1.

9 S1 looks up its cache to find the RLOC for client C2 — the L2 border.

10 The unicast packet is encapsulated in VXLAN and forwarded to the L2 border.

11 The L2 border decapsulates the packet and forwards the packet to C2.

SD-Access to external network with C1 and C2 in different subnets

In this scenario, C1 is connected to S1. C1 wants to communicate with a host C2 which is not in the SD-Access fabric (the host is external to the fabric), C1 and C2 are are in different subnets. In this instance, an L3 Border node is used between the fabric and non-fabric portions of the network. Refer to the following diagram.

1 C1 sends the packet to S1, which is the default gateway for C1, destined to an external IP address C2.

2 S1 queries the control plane node for the destination IP address.

- The control plane replies in the negative (no match), since it cannot find a matching entry in its database.

3 S1 then encapsulates the original packet in VXLAN, with the policy context (VN and SGT), and forwards it to the external border.

4 The external border decapsulates the VXLAN header and does an IP lookup for the destination (using route-leaking or extranet)

- If the destination route is in the global routing table, it forwards the packet to the next-hop router.

- If the destination route is another VRF, it adds the appropriate VRF and forwards the packet to the next-hop router.

DIAGRAM   SD-Access to external network with C1 and external host in different subnets
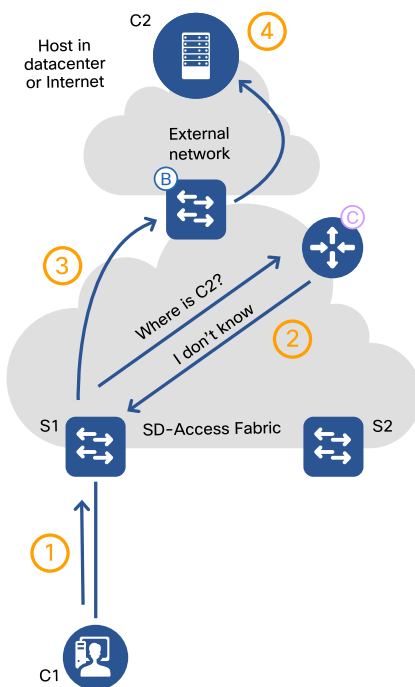
Fabric multicast

Multicast traffic forwarding is used by many applications in enterprise networks today, to simultaneously distribute copies of data to multiple network destinations. Within an SD-Access fabric deployment, multicast traffic flows can be handled in one of two ways (overlay or underlay), depending on whether the underlay network supports multicast replication or not.

Overlay multicast replication in SD-Access

Let's first examine when the underlay network does not support multicast replication. This case, where the multicast is carried in the overlay without enabling multicast in the the underlay, is also known as head-end replication.

In case of head-end replication, the first SD-Access fabric node that receives the multicast traffic must replicate multiple unicast-VXLAN-encapsulated copies of the original multicast traffic, to be sent to each of the remote fabric edge nodes where the multicast receivers are located.

DIAGRAM   Multicast head-end replication
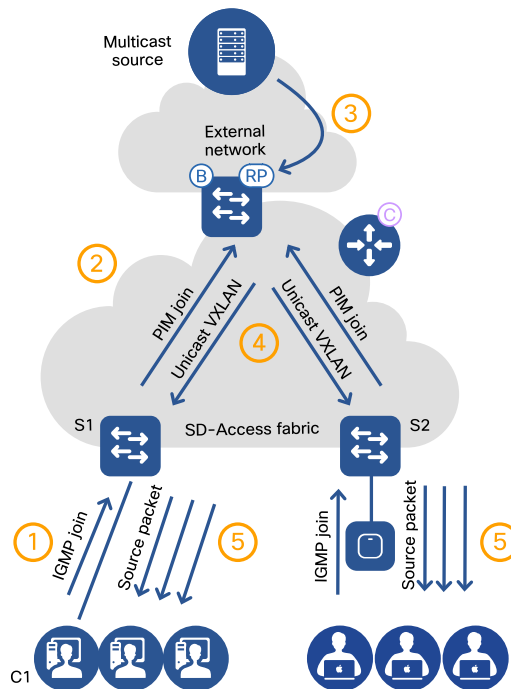
In the figure above, consider client C1 connected to edge node S1. There is a multicast source MS1 in the external network (outside of the border node). PIM sparse-mode op-

eration is used, and the fabric border is configured as the fabric rendezvous point (RP). In this example, there is no multicast configured in the underlay network.

1   S1 processes the IGMP join from C1, for multicast group 239.5.5.5

- S1 sends a corresponding (*,G) PIM join in the overlay (VXLAN encapsulated) towards the fabric RP.

2   The fabric RP creates a (*, 239.5.5.5) state in its multicast FIB (MFIB), with an outgoing interface list entry of S1.

3   The multicast source MS1 starts transmitting data destined to 239.5.5.5, which is registered with the fabric RP.

- The fabric RP creates a specific (MS1, 239.5.5.5) state in its MFIB, with S1 in the outgoing interface list.

4   The fabric RP (in this case, the border) is the first fabric device receiving the multicast traffic flow. The RP replicates the multicast packets and encapsulates them in unicast VXLAN, sending these to all of the fabric edge switches which have clients that have joined the corresponding multicast group via IGMP.

- In this case, the destination IP address is S1 (unicast).

- Traffic is then forwarded directly to S1.

5   S1 receives and decapsulates the unicast VXLAN packet, and forwards the original multicast packet to C1 (based on the IGMP join in step 1).

- If there are 10 receivers for this traffic on S1, it is still a single stream replicated by the head-end (border) node.

¬ S1 performs local replication for all local receivers (e.g. 10, above) interested in this multicast group.

- If there are multiple receivers connected to multiple remote edge nodes, the head-end node will make a separate (unicast) copy for every edge node that has local receivers (based on a PIM join, as in Step 1).

- ¬   Each remote edge node performs local replication for its local
      receivers.

The head-end multicast replication approach provides a method of multicast distribution for networks that do not support multicast in the underlay. The main disadvantage is the potential workload of replication required of the head-end node.

Underlay (native) multicast replication in SD-Access

Native multicast forwards multicast traffic in the underlay. This method of multicast transfer is more efficient compared to head-end replication since the load of replication is now spread across multiple devices in the forwarding path.

Consider the same scenario as above. Client C1 is connected to edge node S1. There is a multicast source (MS1) in the external network (outside of the border node). PIM sparse-mode operation is used (also known as any source multicast or ASM), and the fabric rendezvous point (RP) in this case is on the border node. In this case, though, there is also multicast operating in the underlay network.

Refer to the following diagram.

1   S1 processes the IGMP join from C1, for multicast group 239.1.2.3.

2   S1 sends two PIM joins.

- -   S1 first sends an ASM join, in the overlay, to the fabric RP for the group
      (239.1.2.3).

- -   All the switches hash the multicast group in the overlay to a group in
      the underlay. In this example 239.1.2.3 hashes to 232.1.1.1 in the underlay.

- -   S1 then sends an SSM join, in the underlay, to join the mapped group
      (232.1.1.1).

- -   The RP creates an ASM group (239.1.2.3) in the overlay and maps to the
      corresponding SSM group (e.g. 232.1.1.1) in the underlay.

- All remote edge nodes interested in 239.1.2.3 will join this underlay group (232.1.1.1).

3   The multicast source MS1 starts transmitting data destined to 239.1.2.3, which is registered with the fabric RP.

- The fabric RP creates a specific (MS1, 239.1.2.3) state, which is mapped to 232.1.1.1.

4   The fabric RP encapsulates the original multicast packet (239.1.2.3) in a multicast VXLAN encapsulation (232.1.1.1).

- Traffic is then forwarded in the underlay to all nodes joined to this group.

- The underlay is responsible for performing the necessary multicast replication — thus spreading out the multicast replication load across multiple devices in the underlay.

5   S1 receives and decapsulates the multicast VXLAN packet, maps the outgoing interface list for 239.1.2.3, and forwards the original packet to C1 (based on the IGMP join in Step 1).

- S1 performs local replication for all local receivers interested in this multicast group.

- If there are multiple receivers connected to multiple remote edge nodes, the head-end node replicates once to the underlay multicast group. This multicast packet is replicated natively, by the underlay and forwarded to all the fabric nodes.

DIAGRAM   Native multicast replication

Broadcast in SD-Access

For some traffic and application types, it may be desirable to enable broadcast forwarding within the SD-Access fabric.

By default, this is disabled in the SD-Access architecture. If broadcast propagation is required, it must be specifically enabled on a per-subnet basis. Once broadcast is enabled in a subnet, an underlay multicast group is associated within the VN, and all fabric nodes will join this multicast group.

When a broadcast frame is received by a fabric edge node, it is then encapsulated in VXLAN and forwarded to all remote edge nodes, via the underlay multicast group. The remote edge nodes will then decapsulate the original broadcast frame, and forward it to all local switchports in the appropriate subnet.

Note   This function requires support for multicast in the underlay network as detailed above.

Cisco DNA Center

Cisco DNA Center overview

Cisco DNA Center is a centralized operations platform for end-to-end automation and assurance of enterprise LAN, WLAN, and WAN environments, as well as orchestration with external solutions and domains. It allows the network administrator to use a single dashboard to manage and automate the network.

Cisco DNA Center provides an IT operator with intuitive automation and assurance workflows that make it easy to design network settings and policies, provision and provide assurance for the network and policies along with end-to-end visibility, proactive monitoring and insights — all in order to provide a consistent and high-quality user experience.

DIAGRAM   Cisco DNA Center workflows

Design
- Global settings
- Site profiles
- DDI, SWIM, PNP
- User access

Policy
- Virtual networks
- ISE, AAA, radius
- Endpoint groups
- Group policies

Provision
- Fabric domains
- CP, border, edge
- FEW / OTT WLAN
- External connect

Assurance
- Network health
- Fabric health
- 360° views
- Path trace, sensor

Architecture tenets

Cisco DNA Center has been designed to scale to the needs of medium-to-large enterprise network deployments. It consists of a network controller with automation capabilities and a data analytics functional stack for monitoring, thereby providing a unified platform for deployment, troubleshooting, and management of the network. Cisco DNA Center has been built using a micro-services architecture that is scalable and allows for continuous delivery and deployment.

Some of the key highlights of Cisco DNA Center include:

- High availability — for both hardware component and software packages

- Backup and restore mechanism — to support disaster recovery scenarios

- Role-based access control mechanism, for differentiated access to users based on roles and scope

- Programmable interfaces to enable ecosystem partners and developers to integrate with Cisco DNA Center

Cisco DNA Center is cloud-tethered to enable the seamless upgrade of existing functions and additions of new packages and applications.

DIAGRAM   Cisco DNA Center architecture
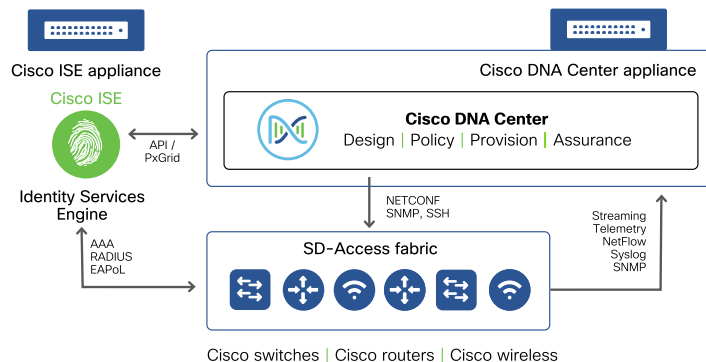
Cisco DNA Center automation

The primary goal of Cisco DNA Center automation workflows is to transform the network administrator's business intent into device-specific network configurations. The Cisco DNA controller consists, at a high level, of the network information database, policy and automation engines, and the network programmer.

The controller has the ability to discover the network infrastructure and periodically scan the network to create a single source of truth that includes the network device details, software images running on the system, network settings, site definitions, device-to-site mapping information, and so on. This also includes the topology information that maps the network devices to the physical topology along with the detailed device-level data. All this information is stored in the controller network information database.

The policy engine provisions various policies across the enterprise network for access control policies and quality of service / application experience. The automation engine provides an abstraction layer for the entire enterprise network using the service and policy framework and leveraging device-specific data models. Finally, the network programmer service does the provisioning on the network devices as required.

Cisco DNA Center assurance

Cisco DNA Center uses advanced machine learning and analytics to provide end-to-end visibility by learning from the network infrastructure, the clients connected to the network, and other contextual sources of information. Cisco DNA Center has an in-built data collector framework that is able to ingest data coming from a variety of sources.

The network infrastructure data may be obtained via streaming telemetry mechanisms designed to optimize network load and reduce the delay in receiving data from the network layer. In addition to this, the data collectors are all built to gather data from various contextual systems such as Cisco ISE, IP Address Management (IPAM) and IT Services Management (ITSM) systems.

The data is processed and correlated in real-time using time series analysis, complex event processing, and machine learning algorithms. It is then stored and visualized for the network operator within the Cisco DNA Center user interface (UI) to provide relevant, useful and timely troubleshooting, insights and trending information, delivered via the assurance workflows.

For more information see *https://cs.co/assurancebook*

SD-Access policy

Policy and services overview

Definition of policy in the context of SD-Access

In any discussion involving policy and services, the starting point is always the business drivers that create the requirement. In the past, the only business requirement of an enterprise network was to provide fast and highly available connectivity (also known as access). With trends in computing and networking over the past several years, both services and policies have evolved. Now the enterprise network policy must meet new requirements to support greater agility, flexibility and increased security.

In the following section, we will focus on security, as an example of how SD-Access policy can address some of the current challenges in enterprise networks. SD-Access security policies come in two types: VN policy and group-based policy. It is important to note that a similar set of requirements and challenges exist for other network services and policies, such as quality of service (QoS), packet capture, traffic engineering, and so on, and SD-Access addresses those in a similar way.

Business drivers

One of the normal business requirements that drives an SD-Access deployment is regulatory compliance for either industry reasons (e.g. Payment Card Industry (PCI), Health Insurance Portability and Accountability Act (HIPAA), etc.) or corporate compliance reasons (e.g. risk mitigation). Multiple requirements can exist within an organization. For example, a healthcare company that not only must be compliant to national regulations (e.g. HIPAA in the U.S.) but also to PCI compliance requirements, while also wanting to mitigate risk to patient care by isolating their medical devices.

Policy scenarios

Below is a set of sample requirements for a common healthcare network — these will be used as the framework for illustrating how SD-Access can solve the requirements while providing business agility, flexibility and lower operational expense.

As described above, the starting point is to evaluate the business drivers and require-ments. For example:

- **Secure patient care**: Only allow approved medical users and medical devices ac-cess to the medical network.

- **Secure business-critical applications**: Identify all user endpoints as they access the enterprise network -and- only allow approved users and devices access to the general enterprise network.

- **Regulatory compliance requirements**: Only allow approved users or devices ac-cess to specific endpoints, servers and applications within the scope of PCI compliance.

- **Provide patient care**: Provide a guest network that is isolated from the enter-prise medical network.

The next step is to evaluate the network to understand where each of the primary re-sources noted above are located.

For example:

- Where are the medical devices located in the network?

- Where are the servers and applications within scope for PCI?

- Where are the medical users located in the network?

If the enterprise was fortunate, they would be able to associate all of the resources in question to a clear set of IP address subnets. This would allow them to build network objects which represent the association between a subnet and a human-readable name. However, over time and due to growth of the network, it is more likely to have discon-tiguous address subnets that are difficult to associate with a specific purpose.

Here are some examples:

- 192.0.2.0/24 = MRI_Devices

- 10.1.1.0/24 = Imaging_Servers

- 198.51.100.0/24 = PCI_Apps

- 10.1.100.0/24 = Staff

- 10.1.200.0/24 = Guests

Policy construction

The network architect would then need to associate these network objects and a set of permissions (i.e. an access control policy) for all of the subnets in the environment. Typically, this is accomplished by using a security management system. Security management systems (for example a firewall or an ACL management system) provide a human-readable abstraction that maps the relationship between the IP prefixes and a "network object".

The network architect would then build access control rules between the network objects, for a specific protocol (IP, TCP, UDP, etc.) and port (http, https, etc.), and then a resulting permission (permit or deny) between the objects.

DIAGRAM   Policy construction example

Prefix 192.0.2.0/24   Prefix 10.1.1.0/24

```
MRI_Devices                              Imaging_Servers
```

Source network object   Destination network object

Classifier

```
tcp dst port 443
```

```
Permit  Deny all
```

Action
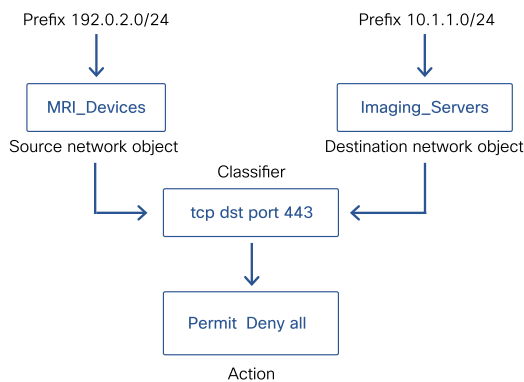
In order to guarantee policy enforcement, the network administrator needs to design the network in such a way as to direct the relevant traffic (for each subnet) to a corresponding policy enforcement point (e.g. distribution switch using ACLs, campus firewall, etc.). The management system would then use the network objects to program the IP addresses back into the policy enforcement point.

In most cases, all of the telemetry that results from an access control security policy is then represented as logs, flow data, hit counters, etc, entirely as IP addresses. This means that all information produced by the policy enforcement points is produced in ways that are only relevant to the network constructs, but not to policy constructs. Which also means that any security management or assurance products must again translate the network constructs back into policy constructs, across multiple enforcement points in the enterprise.

This becomes a very complicated exercise since it normally requires handling of multiple formats of telemetry and different aspects of data. This complicated correlation is needed to perform relatively simple tasks, such as, "IP address 1 (which is part of network object A) is currently communicating to IP address 2 (which is part of network object B), in this log, which means that there is a violation of security policy X" (as proven by having completed this sentence).

> **the bottom line**
> Traditionally, all policy is based on VLAN and/or IP subnet, making it very complex to map to a policy object and retain relevance.

Implementing policy

There is an inherent assumption of what is described above: if you connect a device to the subnet, you inherit all the security access of the subnet. Remember that the traditional purpose of the enterprise network is to provide fast and highly available access.

This has several critical policy implications.

For example (using the network objects above)

- SecOps created an object called "MRI_Devices" associated to IP subnet 192.0.2.0/24 and VLAN 100.

- NetOps assigned ports 1-12 to VLAN 100 on all access switches.

- If anyone plugs into the right port in the wall, then they would be classified by the security system as an MRI device.

With the introduction of Wireless LAN, a new device and identity challenge emerged, but did not change the traditional mapping of IP subnets to network objects for security permissions. Wireless also introduced user and device mobility, which became a challenge to this tight coupling of network topology with security policy because a user and/or device could show up in any part of the network.

In addition, all of the work above needs to be repeated for the addition of IPv6 addresses to a network, along with some new challenges.

- There are more network scopes and aggregate subnets in IPv6, and each user and/or device may use multiple IPv6 addresses.

- IPv6 addresses are often felt to be harder to read and recall than IPv4, due to the length and inclusion of hexadecimal (alphanumerics).

- Depending on the security management tool, this may require creating separate IPv6 network objects and/or upgrading the software.

Today, when the policies are created and applied, they are often locked into the management tool where they were created. While software-defined networking has led to the creation of networks and applications via automation, this automation has not easily extended into security policy. In many cases, the automation is either incomplete (workload only, not for users/devices in the branch, etc), or it is vendor-specific

Often the network objects and policy are not extended to multiple kinds of enforcement points (i.e. firewall and switches or routers). Different types of enforcement points are then managed by different management systems resulting in a manual effort to syn-

chronize the policy between the systems. There are some third-party tools focused on multi-platform and multi-vendor management but they are limited to common constructs and require yet another level of operational complexity.

Furthermore, as we move up into the application layer of security technology and operations, there is relevance for the network security objects. For example, assume the network security policy has allowed a device to communicate to the internet (generically). Unless the network security policy management console and the advanced malware management console manually share the relevance of the network object to the broader enterprise, the advanced malware console will not have any awareness of the business relevance of an endpoint in their generated alarms. Hence, if a device that has been deemed critical to the business is compromised by malware, and the malware sensor detects the infiltration, the malware operator will most likely not see a high-level alert about the event.

Another challenge when dealing with policy is that many access control entries (ACEs) in ACLs on firewalls and/or switches and routers will remain unchanged (or not-optimized) and grow continually over time because the network administrators do not actually know what business drivers or requirements the rules are intended to enforce.

> ### the bottom line
> Changing or removing a policy may have unintended consequences, resulting in considerable risk for the business.

Policy considerations

The primary challenge with the current approach is that we work from objects (IP addresses), to business-relevant objects (network objects), and then back to objects (IP addresses), with no way to carry business relevance.

For example:

- Policy-related telemetry lacks human-readable business relevance (logs and flow statistics only use IP addresses and not user context).

- It is difficult to derive actionable intelligence from what the network telemetry is telling you about the policy.

- Reliance on multi-platform or multi-vendor tools to correlate between the constructs.

- Time-consuming, complicated and error-prone processes can lead to gaps in implementation and/or enforcement.

Current enterprise technologies also lack scalable enforcement for lateral spread within a VLAN / subnet. With recent security events, there is a renewed focus on being able to control traffic within, and between, groups within an enterprise. There have been numerous security incidents involving IoT devices and user devices, where the lack of controls within a VLAN / subnet has led to malware and ransomware infecting business-critical assets.

Due to this natural LAN behavior, new traffic steering mechanisms were introduced to direct devices within the same VLAN / subnet to an enforcement point (e.g. private VLANs, etc.). This also means that to control lateral spread within a subnet, you must create a new ACL policy per VLAN / subnet, even if large parts of these IP subnets are really the same network object.

Finally, with the introduction of mobility, network administrators can no longer assume any given static IP subnet / VLAN structure can be relied upon to accurately represent a given set of end users/devices for the purposes of policy. In addition, network administrators today are hard-pressed to try to manually keep up with the rate of addition and change with mobile users/devices.

To address these many challenges, policy — and specifically, scalable and easy-to-understand policy constructs based on groups, not on IP addresses — was developed as one of the key cornerstones of a Cisco SD-Access deployment. Cisco SD-Access policy is deployed using macro segmentation (using VNs) and / or including micro segmentation (using SGT).

This provides multiple levels to control policies (e.g. devices) which normally may or may not talk to each other. For example, employees and HVAC machines could be sepa-

rated using VNs (as they normally do not communicate), whereas devices and users which work with each other (for example, employees and contractors) could be in the same VN but policy could be enforced between these users and devices as needed using SGTs. Policies applied in this way are decoupled from the endpoint IP addresses, simplifying the deployment and use of the associated policies considerably.

By providing two levels of policy grouping for segmentation — macro and micro — and by decoupling policies from IP addresses and subnets, SD-Access provides a far simpler, and yet more robust and functional, method of implementing segmentation policies within a modern enterprise network.

In the following section, we will explore further how SD-Access uniquely solves the challenges described above, and allow today's enterprises to provide network-integrated policy in a simplified and scalable fashion.

Policy architecture

Policy overview and role of the Cisco Identity Services Engine

Policy management with identity services integrates with the SD-Access network using an external policy repository hosted by Cisco Identity Services Engine (ISE). Cisco ISE couples with Cisco DNA Center for two key functions: dynamic access control (authentication and authorization) of users and devices, and scalable group management (group names and membership) for group-based access policy. Group-based policies in SD-Access are defined on the basis of logical groups of users, devices, things or applications, and as a relationship between two groups, and further define the access control rules based on L3 and L4 classifiers.

Cisco ISE continues to fulfill the first function by a user executing well-defined existing workflows on Cisco ISE.

Cisco DNA Center is now integrated with Access Control Application (ACA) to simplify group-based access control policy management directly within Cisco DNA Center. This also provides a greater level of interoperability with non-Cisco identity solutions. The application (ACA) provides the ability to classify a variety of endpoints (users, enterprise devices, IoT devices or even workloads running in private or public clouds) to be mapped into scalable groups in Cisco DNA Center. These scalable groups can then be used to define group-based access control policies in Cisco DNA Center which are deployed to Cisco ISE for distribution to a SD-Access deployment.

Policy enforcement in SD-Access

SD-Access policy, by design, provides multiple levels of segmentation to address customer requirements. The policy could be constructed by simply using macro segmentation (using VNs) or including micro segmentation (using SGTs).

Here is a simple illustration of macro and micro segmentation
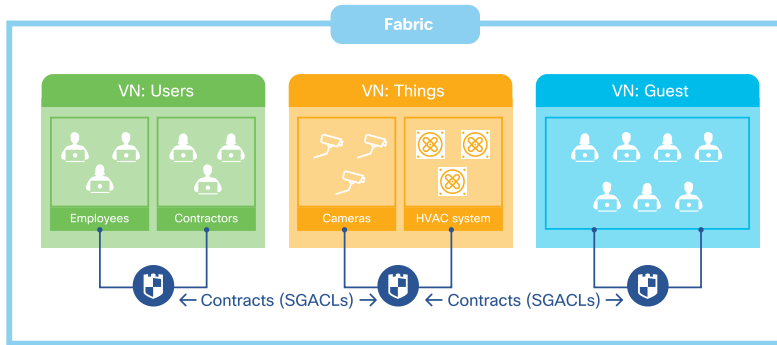
In the illustrations above, the recommendation is to use macro segmentation (segmentation with VNs) for users and devices which typically do not talk to each other. Some of the macro segmentation use cases are

- Virtual network A = Users

- Virtual network B = Things

- Virtual network C = Guest

For example, we do not expect group "Employees" to communicate with the "HVAC system" group or the building security "Cameras" group, but there could be exceptions. It is possible to permit these communications by using a firewall/fusion router or the SD-Access VN Extranet feature.

Furthermore, communications between different groups within the same VN could also be controlled using micro segmentation with SGTs, while being connected to the same fabric edge.

For example, it is possible to create a group-based access control policy defining the "Contractors" group to be denied access to the "Employees" group, or the "Cameras" group be denied to the "HVAC" group.

DIAGRAM   Policy enforcement in SD-Access
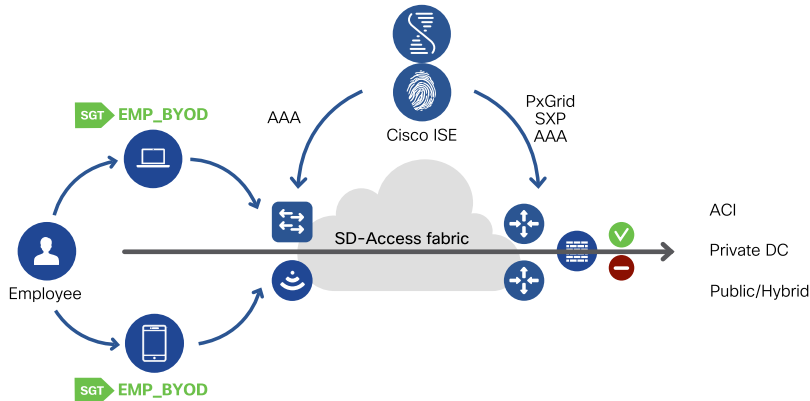
Both wired and wireless policies are centrally defined and managed in SD-Access using Cisco DNA Center. They are enforced at the fabric edge and border nodes, based on the user/device identity, in a topology-agnostic manner. The group classification for the endpoint is embedded into the fabric data plane and carried end-to-end with the SD-Access fabric, so policy for the traffic can be enforced irrespective of its location.

For stateful inspection, group-based policies can also be applied at SGT-aware firewalls or web-proxies.

Endpoint grouping at access

Cisco Identity Services Engine (ISE) helps establish the identity of endpoints connecting to the network through a variety of mechanisms, such as 802.1x, MAC addresses, profiling, Active Directory login, and captive portals.

In general, for any user or device connecting to the fabric edge, the recommendation is to use dynamic group assignment from Cisco ISE which is part of host on-boarding. Dynamic group assigment also ensures that the endpoints are given the same policy when connecting from wired or wireless.

Once the identity of the endpoint is established, Cisco ISE also defines the rules associ-ation of endpoint identity to the group. Attributes from Active Directory groups can be used in defining the group classifications in Cisco ISE, for use in SD-Access. These groups are imported into Cisco DNA Center so that policies can be viewed and adminis-tered from the Cisco DNA Center user interface.

DIAGRAM   User access in SD-Access based on group policy
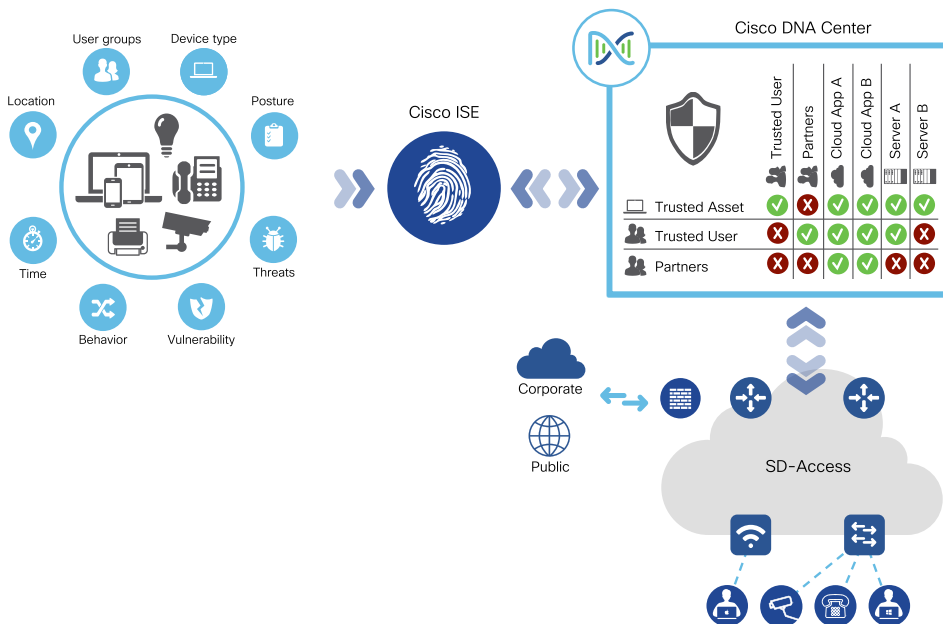
Cisco DNA Center is also capable of gathering endpoint identity information from ex-ternal network access control (NAC) and AAA systems and using the externally derived identity to map the endpoints into groups.

For environments that are not enabled for identity-based access through any of the above mechanisms, Cisco DNA Center allows the network administrator to statically define associations between access ports and groups.

Application groups

Policies can be defined and implemented for users or endpoints to applications in Cisco DNA Center by classifying the external applications into groups based on their IP address or subnets. This is especially relevant in data centers where the applications are grouped for security reasons into pre-defined subnets.

For data centers based on Cisco application centric infrastructure (ACI), the endpoint groups from SD-Access can be imported by the Cisco application policy infrastructure controller (APIC) in ACI. Policies can then be defined based on the same group-policy model end-to-end from user access to the application. This results in highly scalable, automated, simplified policy implementation that caters to user or workload mobility.

Refer to the fabric external connectivity chapter for more details and illustrations.

Cisco's cloud policy platform enables workloads in public cloud environments such as AWS, as well as those in hybrid cloud environments, to be mapped into groups that can be imported into Cisco DNA Center. Cisco DNA Center can thereby define and implement access control policies between user/device endpoints and the applications in public/hybrid cloud using the same group-based policy constructs. These group policies can be instantiated at a policy enforcement point such as the fabric border or a compatible firewall.

Policy benefits

Aside from lowering the complexity and overall costs of operating the network, SD-Access automation and assurance implements a policy-driven model for network operations that reduces the time required to introduce new services and improve overall network security. These benefits are further discussed below.

Decouple policy from infrastructure design

Similar to the way SD-Access abstracts network connectivity through VXLAN overlays, SD-Access also abstracts the notion of policy and decouples it from the underlying network topology. This allows for network design changes without a need for the operator to manually define and update individual policy elements.

Since SD-Access leverages the fabric network infrastructure for policy enforcement, complex traffic engineering mechanisms to forward traffic to firewalls are no longer necessary, and thereby reduce IP-ACL sprawl in firewalls.

The decoupling of the policy from the network topology enables more efficient operations and enables the network to be more effectively used to enforce policies. This leads to a number of associated business benefits around quicker time to enable new business services, enabling seamless network mobility, and overall reduction of effort in the day-to-day network administration.

Simplified policy definition

Management of access control policies on the basis of logical, business relevant, and human-readable groups simplifies ongoing operations and reduces the security risk. It also reduces the time and effort required to demonstrate compliance and simplifies the audit process.

Policy automation

The dynamic association of an endpoint to its group based on its identity reduces the operational overhead required to ensure that the endpoints are on the appropriate net-

work segment. This also enhances overall enterprise security, especially in environments where users and devices are mobile.

The complexity and time needed with the older methods is linear to the number of devices in the enterprise network and execution tasks on each of those devices. In SD-Access, these activities are not only simpler, but also far faster to execute, and much easier to design, deploy, operate, and understand.

Cisco DNA Center also automates AAA configurations (authentication, authorization and accounting) for SD-Access. This includes all the global and port level configurations needed on fabric edge devices in order to start authenticating users and devices on the network. The Cisco SD-Access solution provides AAA templates, providing for easy definition and roll-out of policies into the fabric.

Policy-based enterprise orchestration

The SD-Access policy model provides a platform in which customers can develop a vast number of applications, targeting use cases including segmentation, security, compliance, and responses to real-time security threats, as well as the ability to offer a variety of services within the fabric.

DIAGRAM   Closed-loop automation and programmability process
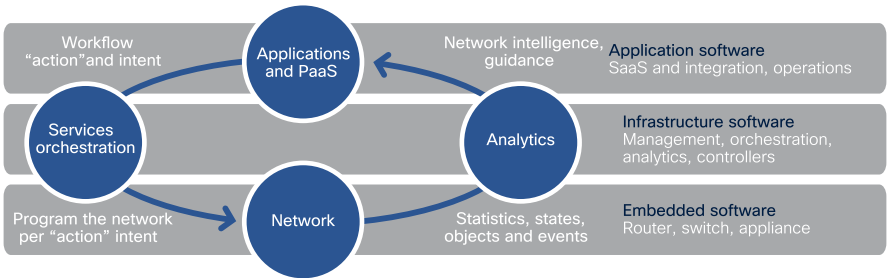
Through the use of APIs in Cisco DNA Center, abundant network telemetry sources, and intelligent machine learning, SD-Access can leverage the "closed-loop" model concept (shown in the diagram above). This closed-loop model can be adopted with SD-Access and will support a multitude of use cases, each leveraging the SD-Access policy model as the method for deploying the "actionable" intent into the SD-Access fabric.

Policy in action

Consider a security operations team's ability to leverage the SD-Access policy model to respond to varying levels of vulnerabilities. For example, assume a new vulnerability has been identified for a popular host OS, and a user of that host logs onto the network. Through a management agent, that host OS is identified as not yet having patched the vulnerability. Based on that criteria, the management agent can identify the vulnerability as a "threat" level, and through APIs into the SD-Access fabric apply a "threat" level policy.

For example, the policy could instantly deny that user access to any business-critical systems throughout the enterprise network, while still allowing them access to non-critical systems and external networks such as the internet.

This example highlights the power of the address-agnostic group-based policy model within SD-Access. Central control from Cisco DNA Center, abstracted in the policy model from network topology, is applied to each network element where policy enforcement actions are required. Attempting the same capabilities with today's network operations tools and methods is simply not possible without hugely time-consuming tasks by human network operators, and could take days.

It is important to understand the implication of policies being applied to the network. Cisco SD-Access provides multiple features including device profiling and also enables pre-built authentication templates for easy roll-out of the security policies as defined by the network manager.

DIAGRAM   Cisco DNA Center and Cisco ISE device profiling

Once the endpoint identity is known, it is easy to assign these endpoints into groups. The endpoints and users may authenticate to the network using various protocols and authentication methods. These include:

- **MAC authentication bypass (MAB):** MAB enables port-based access control using the MAC address of an endpoint.

- **Web authentication:** typically used to onboard guest users for internet access, using a web portal.

- **EasyConnect:** enables enterprises to implement identity-based network access (without the need for 802.1x or supplicants) using alternate credentials (e.g. LDAP).

- **802.1x (monitor, low-impact and closed mode):** the 802.1x standard defines a client/server-based access control and authentication protocol. Monitor and low-impact modes provide for easier integration of 802.1x by allowing the network manager to better gauge and control the impact without the 'all-or-nothing' approach of traditional (closed-mode) 802.1x.

DIAGRAM   Phased policy deployment models

Phase 1 (Visibility only)

Transition

Monitor mode

File servers

Cisco ISE

Low-Impact mode

File servers

Cisco ISE

File servers

Campus network

PREAUTH ACL

permit eap dhcp dns deny any

PERMIT ACL

permit ip to any

Before authentication

After authentication

Campus network

Port Open Unconditionally

Pass / Failed authentication

Transition

Closed mode

Cisco ISE

File servers

Campus network

Only EAP Allowed

Before authentication

After authentication

Phase 2 (Visibility and control)

In general, the recommendation is to use a phased deployment model when implementing policy which will limit the impact on the network, while gradually introducing authentication and authorization on the wired network.

SD-Access gives an administrator the option to assign a policy at a global level (e.g. closed mode for all the ports), but the admininstrator can still configure an exception policy per port or selected ports (e.g. port 1-5 are in open mode).

SD-Access also offers the flexibility for third-party applications to create, instantiate, and push policy into the fabric through an open set of APIs in Cisco DNA Center. Customers could use these APIs in environments that are leveraging applications such as

security information and event and manager (SIEM) systems. While a SIEM system may not configure the network, integrating it with the SD-Access fabric that can drive policy changes in the network can help the security operator accelerate responses to events identified by the SIEM.

As high-risk events are detected within the SIEM, it could then call APIs to Cisco DNA Center requesting a creation or modification of an SD-Access policy to "quarantine" a specific set of users/ports to rapidly contain the threat or to instantiate a traffic copy policy (e.g. encapsulated remote switched port analyzer (ERSPAN) session) for further analysis.

DIAGRAM   3rd-party application calling Cisco DNA Center and Cisco ISE APIs
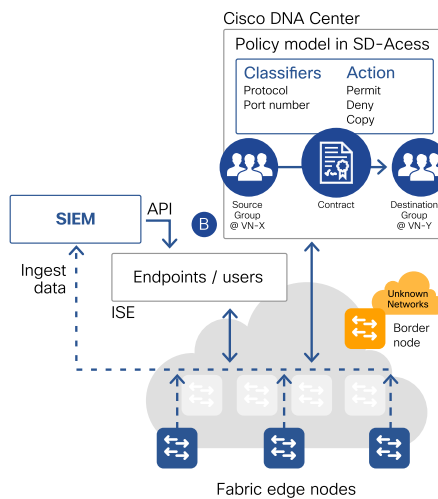
API calls to Cisco DNA Center have triggered the traffic copy policy, and API calls to Cisco ISE have quarantined the user.

SD-Access policy with transits and extended nodes

In a SD-Access distributed deployment, data is transported over a transit network that interconnects sites.

- **SD-Access transit:** uses native SD-Access encapsulation to connect sites.

- **IP-based transit:** relies on traditional IP connectivity (without SD-Access functionality), therefore requiring a mechanism for security constructs to be carried across sites.

A distributed campus design using SD-Access transit connects fabric sites (connected via high bandwidth and low latency) while providing consistent policy and end-to-end segmentation using VRFs and SGTs. Policy defintion and deployment to sites is accomplished from Cisco DNA Center and enforced at the fabric edge nodes in a multi-fabric architecture.

DIAGRAM   SD-Access policy with SD-Access transit
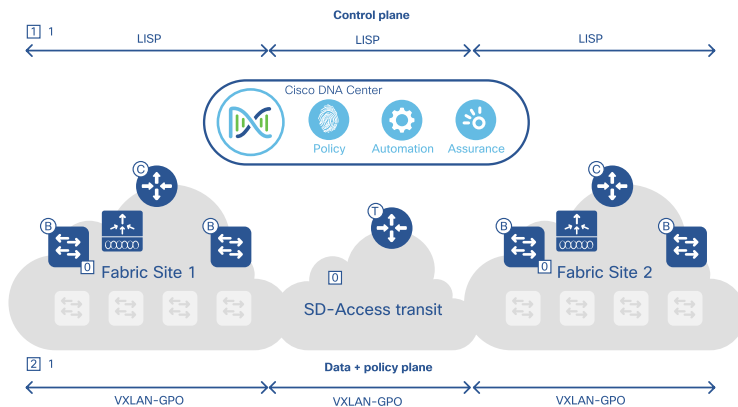
For customers already using an existing WAN and traditional IP/MPLS, an IP-based transit is used to connect the SD-Access fabric to external networks. Fabric constructs (VXLAN, SGT) terminate at the border node and we need to ensure that the SGTs are propagated between sites.

Typically, there two ways of carrying SGT tags outside the fabric:

- inline tagging

- scalable-group-tag exchange protocol (SXP)

Inline tagging allows the SGT to be carried within Cisco metadata (CMD) header; however this requires all devices in the path to support inserting CMD (with SGT) in the packet to be carried to its next hop neighbor. SXP protocol is typically used where network devices do not support inline tagging, and this allows for IP-SGT mappings to be exchanged with its peers over a TCP connection. In SD-Access SXP sessions can be created between border nodes and Cisco ISE, for each VRF, to transport IP-SGT mappings from a site to be exchanged to other sites.

DIAGRAM SD-Access policy with IP-based transit

As we extend the enterprise network beyond conventional users and locations, such as cameras in parking lots, and sensors and lighting in distribution centers, SD-Access extends consistent policy-based automation to industrial ethernet compact switches (known as extended nodes). Cisco DNA Center delivers an automated way to onboard IoT devices (cameras, lighting, etc) to the SD-Access fabric with the ability to set group-based policies for segmentation.

Extended nodes connect to a single fabric edge over a 802.1q trunk which allows fabric VLANs to extend downstream to these switches. Switchports on the extended node are statically mapped to an IP pool via Cisco DNA Center.

For example

- parking lot cameras may be assigned to IP pool 10.10.0.0/16
- distribution center lights may be asigned to IP pool 20.20.0.0/16

IP pool to SGT mapping is done at the fabric edge

- IP pool 10.10.0.0/16 is mapped to SGT cameras (tag of 25)
- IP pool 20.20.0.0/16 is mapped to SGT lighting (tag of 17)

Group-based policy can then be enforced on the egress fabric edge based on the SGTs.

DIAGRAM   Macro and micro segmentation with SD-Access extended nodes
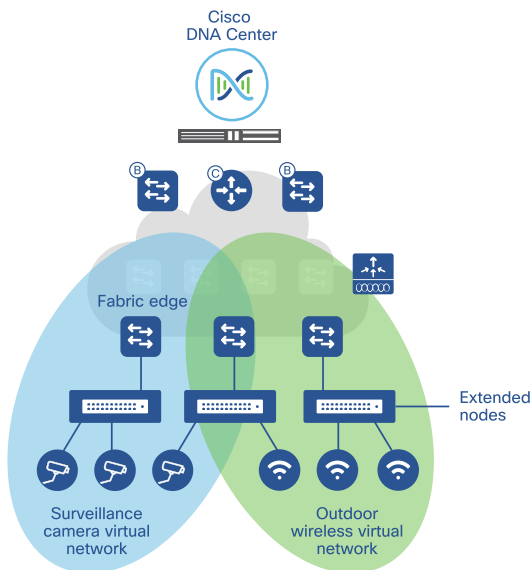
SD-Access automation

Automation and orchestration in Cisco DNA Center

Automation and orchestration, as defined in the Software-Defined Access overview section, brings the "software-defined" concept to the campus "access" network, translating the user's "intent" into meaningful configuration and verification tasks.

Cisco SD-Access uses controller-based (Cisco DNA Center) automation as the primary configuration and orchestration model, to design, deploy, verify, and optimize the wired, wireless and security network components. With Cisco DNA Center, IT teams can now operate at an abstracted level that is aligned with the business objectives and not worry about the implementation details. This results in simplifying operations for the IT teams by minimizing the chances of making a human error and more easily standardizing the overall network design.

Cisco DNA Center provides multiple forms and levels of automation and orchestration, for both non-fabric and fabric-based components. Below is a brief list of the key principles and concepts of Cisco DNA automation and orchestration:

Agility: Reduce the time required to design, deploy and/or optimize a network environment. Cisco DNA Center makes agility real by:

- Centralizing design: Generating, organizing and managing a set of common and/or unique network designs for different operating environments, including specific requirements for global and local network device settings.

- Automating deployment: Rapidly deploying configurations to multiple devices, and providing verification of deployment

- Optimizing design: Ensuring consistency of network state and configurations at scale, to match the desired operator objectives.

Reliability: Consistent deployment of prescriptive "best practice" network configurations. Cisco DNA Center brings reliability with:

- Configuration best-practices: Mature and tested designs and configurations ensure consistent predictable behavior.

- Profile or template-based configuration: Different designs and configurations organized into easy-to-manage sets of profiles or templates.

Simplification: Minimize the complexity of configuring and integrating multiple devices. Cisco DNA Center introduces simplification by:

- Centralizing management: Providing a central location for all of the functions required to design, deploy and manage multiple network components and/or external services, that can serve as the single source of truth, in a single-pane-of-glass.

- Reducing touchpoints: Reducing the traditional box-by-box configuration and management tasks, and providing a single, centralized interface for network operations.

- Programmable interfaces (APIs): Allowing operators to automate network operations in their own customized manner, while leveraging a centralized platform to maintain and drive the changes in network state.

Abstraction: Cisco DNA Center uses easy-to-understand concepts and constructs that abstract out the underlying feature and technology implementation specifics of the network infrastructure. Cisco DNA Center provides this through simple, technology agnostic workflows, supported by views of both physical and logical network topology.

What is automation and orchestration in the context of Cisco SD-Access?

Cisco SD-Access applies the key concepts of automation and orchestration to an enterprise campus network. This includes several major technologies, such as wired access, wireless access, as well as services and policies for security and application optimization. Cisco SD-Access automation and orchestration can be divided into two main categories: network underlay (or non-fabric) and fabric overlay.

Below is a key set of workflows for **network underlay:**

- **Global and site settings**: Hierarchical management of network configurations (e.g. shared services) for different sites.

- **Device discovery (existing networks)**: Automated discovery and inventory of existing network devices.

- **LAN automation (new networks)**: Automated discovery, provisioning, and inventory of new network devices. The Cisco DNA Center creates the configuration of each of the devices based on Cisco best practices.

Below is a key set of workflows for **fabric overlay:**

- **Fabric sites**: An automated configuration of a set of fabric-enabled network devices, with a common fabric control-plane and data-plane.

- **Fabric device roles**: An automated configuration of network devices operating various fabric functions, including control plane, border, edge, WLC, AP or extended nodes.

- **Transits:** A transit enables connectivity between multiple fabric sites and external domains.

    - **IP-Based:** Automated border configuration for traditional IP-based connectivity, which requires manual remapping of VNs and Group context between sites.

    - **SD-Access:** Automated configuration of a domain-wide control plane node for inter-fabric site communication using native SD-Access (LISP, VXLAN, CTS). This includes transport of VN and group context.

- **Virtual networks**: An automated configuration to enable virtual routing and forwarding segmentation in the fabric overlay.

- **VN extranet**: An automated configuration that provides a flexible and scalable method for achieving inter-VN communications.

- **Group-based policies:** An automated configuration to classify and/or enforce group-based policies in the fabric overlay.

- **Host onboarding:** An automated configuration to onboard clients, including VN, IP pool and scalable group assignment, SSID, L2, etc.

- **Multicast services:** An automated configuration to enable IP multicast distribution in the fabric.

- **Pre-verification and post-verification:** Tool to verify capability and support of network devices, prior to deploying fabric overlay automation and to verify proper operation of network devices, following fabric overlay automation.

Automating SD-Access with Cisco DNA Center

This chapter describes the practical application of SD-Access concepts through the Cisco DNA Center platform using the design, policy and provision workflow.

Network design

IT teams in large enterprises often have to manage a large number of distributed sites of varying business functions and natures. For example, a certain enterprise may have retail stores, kiosks, distribution centers, manufacturing sites, and corporate offices. The desire for IT teams in such scenarios often is to simplify their operations through standardizing sites based on their business nature into a network profile. IT teams also need to allow for local teams to manage and customize certain site-specific parameters, such as site-specific logging services, while ensuring other parameters, such as network authentication and policy, are defined consistently across the enterprise.

Cisco DNA Center allows for categorizing network infrastructure in terms of logical sites and also provides granularity to define buildings and floors that closely mirror the physical layout of the organization's network infrastructure. To provide maximum flexibility, Cisco DNA Center also allows for a site's hierarchies to be defined.

Cisco DNA Center allows for automated configuration on either a global- or per-site basis. Cisco DNA Center also supports zero touch provisioning of network infrastructure with Cisco's Plug-n-Play (PnP) solution to automatically onboard new infrastructure components.

To achieve this, the design section of the Cisco DNA Center provides the following:

- Network hierarchy creation
- Site-specific network parameters
- Site-based network profiles

DIAGRAM   Cisco DNA Center network hierarchy

Settings applied at higher levels will cascade to lower levels by default.

Specific configurations may be customized for lower levels where appropriate.

Global

Site

Building

Floor

Apply network settings

Settings defined in the design workflow of Cisco DNA Center provide the primary building blocks that will (1) equip the controller to validate network configurations prior to deployment, and (2) minimize further manual entry of these elements during other phases of the automation process. These settings are applied to a network hierarchy as noted above and will serve multiple purposes in subsequent workflow elements.

Settings defined in this workflow include the following:

- **Shared services configurations** including AAA, DHCP, NTP, and DNS servers.

- **Credentials** with which Cisco DNA Center will access network devices.

- **IP address pools** both for client devices and network devices, including LAN (Underlay) automation and fabric border external connectivity automation.

- **Wireless configurations (**discussed in further detail below).

Plan and build wireless network configurations

Wireless network configurations are automated in Cisco DNA Center and are simply another step in the design/policy/provision workflow. Shared elements such as IP address pools, Virtual Networks, and Scalable Group Tags, are integrated into this workflow and do not need to be defined separately.

Features specific to wireless network deployments, including enterprise and guest SSID configurations, Radio Frequency (RF) optimization parameters, and other key features such as Quality of Service (QoS), Fast Lane, and Adaptive 802.11r configurations are defined and deployed within Cisco DNA Center.

Manage software images

Cisco DNA Center includes Software Image Management (SWIM) — the capability to automate the management of software images for a variety of routers, switches, and wireless LAN controllers. This includes downloading software images from Cisco.com, the distribution of software images and validation checks to ensure devices are adequately prepared for an upgrade or downgrade.

Define policies

Cisco DNA Center enables organizations to create logical network segments and granular group or context-based service policies, which are then automated into prescriptive configurations which get pushed onto the network infrastructure.

There are three primary types of policies that can be automated in the SD-Access fabric as follows:

- **Security:** Access Control policy which dictates who can access what. It has set of rules for cross-group access. For example, deny access for IP from group A to group B.

- **QoS:** Application policy which invokes the QoS service to provision differentiated access to users on the network, from an application experience perspective.

- **Copy:** Traffic copy policy which invokes the Traffic Copy service to configure ERSPAN for monitoring specific traffic flows

Provision your network

Once your network design has been defined, the following provisioning activities enable automated deployment:

- **Add devices to sites:** This step involves assigning network devices (from the Inventory) to the sites created as part of the design workflow. This makes the device ready to accept the site-specific design parameters.

- **Provision devices:** This step involves the provision of the configurations based on the design workflow. When the provision step is executed, all the parameters which were set in the design for the site are provisioned on the device (based on Cisco best practices).

Create fabric

This step involves selection of fabric edge, fabric border and fabric control plane nodes. Additionally, pre-verification and post-verification checks are provided to verify the state of the devices in the fabric. A fabric is built with the following steps:

1. Add edge nodes to the fabric. Add fabric extended nodes (optional).

2. Choose one or more fabric border nodes.

   - You also need to provide the external and/or transit connectivity parameters, to connect to the outside networks.

3. Choose one or more fabric control plane nodes.

4. Add one or more wireless LAN controller to the fabric, for SD-Access wireless (optional).

5 Choose one or more fabric rendezvous point (RP), for SD-Access multicast (optional).

Note that it is possible to combine multiple roles on a single device or device stack. For example, small sites may only require a single control plane, border, edge and WLC node. This is commonly referred to as Fabric-in-a-box.

Guest authentication

Guest authentication in SD-Access fabric is fully automated and includes using workflows for automation of guest SSID and the policies in Cisco ISE.

Host onboarding

Host onboarding enables the attachment of endpoints to the fabric nodes. The host onboarding workflow allows you to authenticate, classify and assign an endpoint to a scalable group, and then associate to an IP pool and virtual network. Key steps to achieve this are as follows:

1 **Authentication template selection:** Cisco DNA Center provides several predefined authentication templates to streamline the process of applying authentication to your network. Selection of a template will automatically push the required configurations to the fabric edge.

2 **Virtual networks and IP pools selection:** Associate unicast or multicast IP address pools to the virtual networks (VN).

3 **Fabric SSID selection:** For integrating SD-Access wireless with a fabric-enabled IP address pool (and VN).

4 **Static port settings:** This allows custom IP address pool (and VN) and SGT settings at port level.

Pre-verification and post-verification checks

Each fabric creation step allows the administrator to do a pre-verification check which ensures that the selected network devices are capable and are correctly configured to accept the fabric provisioning. Likewise, post-verification checks allow the administra-

tor to verify the correct operation of the fabric by highlighting devices which might have reported errors during configuration. This step helps the administrator to find any explicit things that might not be working as expected.

Summary

Once the above tasks are completed to deploy SD-Access fabric, configuration changes to meet evolving use cases can be easily achieved through the Cisco DNA Center instead of through manual interaction.

SD-Access assurance

Assurance overview

As a network grows and evolves to accommodate new business requirements, complexity grows with it. Complexity is introduced in multiple dimensions: from ensuring that new functionalities can be successfully overlaid on top of existing network designs and architectures, to guaranteeing new features and functionalities co-exist and interoperate successfully with existing ones, to implementing appropriate lifecycle management of the infrastructure without disrupting the business.

In addition, as the role of the network becomes more prominent, any outage or degradation in the network performance is proportionally disruptive to the business. As a result, IT is constantly under pressure to ensure optimal connectivity and user experience for the applications, no matter where the user is connecting from, which device he/she is using, or which application he/she is trying to access.

In essence, IT is faced with the following challenges:

- **Reactive troubleshooting**: IT is often made aware of an incident after it occurs and must troubleshoot it reactively.

- **Disparate tools**: Enterprises have a plethora of tools which introduce complexity because each tool only provides a subset of the required functionality.

- **Lack of insights**: Today's tools often provide data that isn't correlated, and as a result, are unable to drive actionable insights.

What is assurance in the context of Cisco SD-Access?

Cisco DNA defines **Assurance** as the ability to quantify network availability and risk. This section will cover Cisco DNA Assurance as it relates to SD-Access.

For more detailed information on Cisco DNA Assurance as an overall solution, please refer to the Cisco DNA Assurance eBook *cs.co/assurancebook*

Cisco DNA Assurance as it relates to SD-Access provides analytics and insights for two main categories: network underlay and fabric overlay.

Analytics provided for the **fabric underlay** include the following:

- Network — traditional (non-fabric) LAN, WLAN and WAN protocols and tables

- Device — switch, router, wireless software and hardware (CPU, memory, temperature, etc.)

- Client — traditional (non-fabric) wired and wireless client status and statistics

- Application — traditional (non-fabric) wired and wireless flow status, statistics, and performance

Analytics provided for the **fabric overlay** include the following:

- Fabric reachability — connectivity checks between all the fabric nodes

- Fabric device — fabric nodes mapping entries, protocols, and performance

- Fabric clients — client onboarding and shared services (DHCP, DNS, AAA, RADIUS)

- Group-based policies — Cisco ISE (pxGrid, AAA) and border and edge node policy entries

Health and insights for SD-Access

SD-Access fabric health

Cisco DNA Center Assurance provides correlated, actionable insights based on a wide variety of telemetry data ingested from sources throughout the network. As it relates to SD-Access, Cisco DNA Assurance provides specifc health scores. These are referred to as SD-Access health scores, and fall into the following three categories:

- **System health:** Considers metrics such as CPU and memory utilization for switches, routers, APs, and wireless LAN controllers

- **Data plane connectivity:** Considers metrics such as link errors and uplink availability status with additional RF-related information for wireless networks

- **Control plane connectivity:** Considers metrics such as connectivity or reachability to the fabric control plane node

SD-Access fabric insights

Cisco DNA Assurance also provides specific insights for SD-Access fabric, as shown in the preceding graphic. These are referred to as SD-Access fabric insights and fall into the following categories:

- **Control plane insights:** Assurance provides insights into the fabric underlay and overlay control planes in a correlated manner by checking for reachability (link down, adjacencies flapping), measuring control plane response times (latency), and validating device configurations (MTU mismatches) that may lead to network problems

- **Data plane insights:** Assurance leverages IP SLA to proactively generate probes, both in the underlay and overlay from fabric borders towards shared services in

order to detect any issues that might arise. Assurance also leverages path trace functionality to provide additional context.

- **Policy plane insights:** Assurance gives visibility into the policy instantiations on network elements, such as whether SGACLs have failed to download or failed to instantiate due to a lack of TCAM resources.

- **Device insights:** Assurance monitors many individual resources such as CPU, memory, temperature, environment, fan, line cards, PoE power, and TCAM tables of the physical network devices, and can provide trending capabilities to help avoid issues from happening in the fabric.

DIAGRAM   Fabric network health
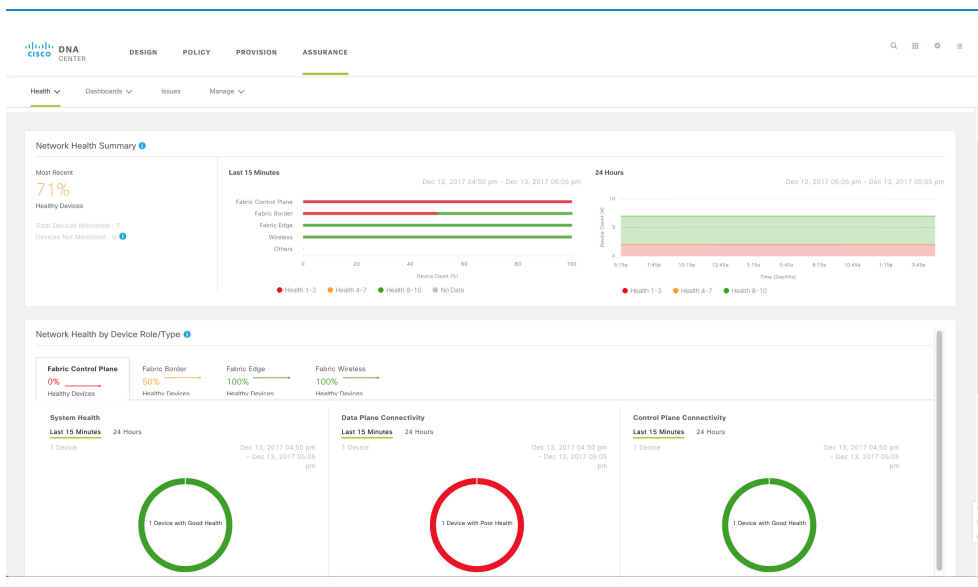
Reporting

In future releases, Cisco DNA Center will provide both pre-defined and customizable reporting capabilities to help plan capacity, detect overall baselines and pattern changes for client and infrastructure devices, and provide views into operational activities such as software upgrades and provisioning failures.

For more about the reporting capabilities offered by Cisco DNA Center Assurance, please refer to the Cisco DNA Assurance e-Book *cs.co/assurancebook*

Integration with partner ecosystems

Integration Overview

The value of SD-Access in terms of enhancing IT agility and efficiency is further augmented through ecosystem integrations with other Cisco and non-Cisco solutions and products. These integrations help customers manage the entire enterprise environment, including the network infrastructure as a single orchestrated entity, enabling faster introduction of new services and helping focus on business outcomes.

These integrations are built using the open northbound Cisco DNA Center APIs that expose a rich set of intent-driven business flows along with data-as-a-service. Cisco DNA Center provides integration capabilities and a development environment that can be leveraged by independent software vendors (ISV) partners, customers, and ecosystem partners. Examples include integrations with IT service management (ITSM), IP address management (IPAM), and others such as:

- API catalog and documentation

- Runtime monitoring and analytics

- API lifecycle management

- Sample code and script generation capability

DIAGRAM   Cisco DNA Center platform ecosystem

Business Apps

Business intelligence	License mgmt	Impact analysis	Financial costing

Network Apps

Advanced functions	End-user monitoring	Policy / compliance	UC performance monitoring

Northbound APIs

Network / IT process integrations

Connectors

Cisco DNA Center as a platform

Connectors

Cross-domain integrations

Identity services	IT service mgmt	Network services

Cloud	Data center	Security	Collab

Southbound APIs and collectors

Device and telemetry integration

Routers	Switch	Wireless	IoT

APIs and programmability

APIs in the SD-Access solution fall into the following three categories:

- **Device APIs** for direct access to configuration and operational functionality on individual network devices

- **Cisco DNA Center APIs** for network-wide automation and orchestration

- **Cisco ISE APIs** for rich contextual data about user and device access

These three categories are explored in further detail below.

Device APIs

Network infrastructure devices based on Cisco IOS® XE provide APIs that are based on IETF and OpenConfig YANG Models as well as Cisco native models and are exposed via NETCONF, RESTCONF and gRPC interfaces. These APIs support both configuration automation and operational models. Newer Cisco IOS XE devices (Catalyst 9000 family, ASR, and ISR) additionally support streaming telemetry, which allows analytics solutions to subscribe to low-latency datastreams based on YANG data-models for real-time data analytics at scale.

Cisco DNA Center APIs

Cisco DNA Center provides REST APIs for network-wide automation and orchestration. This helps provide an abstracted view of the network such that customers can drive network automation at scale. Additionally, Cisco DNA Center will provide REST APIs to expose operational data present in client 360, device 360 and application 360 views along with issues, trends and insights. This enables external entities to consume and create business-relevant, data-driven workflows.

Identity Services Engine APIs

Cisco Identity Services Engine (ISE) APIs expose contextual information about network behavior, including user and device access policies, authentication and authorization

events, and more. These APIs also allow for policy-related configuration automation, such as enabling external systems to trigger changes of authorization for users and endpoints for rapid threat containment, and provides a platform-exchange grid (px-Grid) for rich integration with more than 100 ecosystem partners.

Security technical alliance partners: *https://www.cisco.com/c/m/en_us/products/security/technical-alliance-partners.html*

Ecosystem integrations

Customers are able to derive enhanced value and functionality from the SD-Access architecture through open northbound REST APIs from Cisco DNA Center and integrations with other solutions that are part of the Cisco DNA ecosystem. The ecosystem integrations span various categories of solutions such as IP address management, IT service management, public cloud, external SDN orchestrators, security analytics and operations, security infrastructure such as firewalls, and more.

Some of the key capabilities that are achieved as a result of these integrations include:

- **Improved operations and security**: through orchestration between external systems and SD-Access, customers can programmatically define and deploy closed-loop systems that take insights and state information from the network, enrich it with external data and functionalities, and feed a response action back to the network. This helps customers operate multiple systems with independent automation and management platforms as a single orchestrated entity, thereby accelerating IT agility, efficiency and improving security posture.

- **Simplified best-of-breed deployments**: in an increasingly complex yet innovative world, customers are often looking for best-of-breed tools to incorporate into their IT environments which address key functionalities desired by the customer or end-user. Open APIs and validated interoperability ensure that an end-to-end orchestrated environment can be easily deployed.

- **Enhanced value from technology investments**: for customers who have already adopted and standardized on certain Cisco or non-Cisco solutions, the ecosystem integrations help ensure that customers gain maximum value from existing investment while adopting SD-Access.

The capabilities for each of the solution and vendor integrations is evolving rapidly. Below is a sample of the various integrated technology partners currently in development, and the joint functionalities with SD-Access.

IP address management (IPAM)

SD-Access offers a validated integration with leading IPAM solutions such as InfoBlox and Bluecat through IPAM-API integration. This allows the importation of IP address pools defined in an IPAM solution into Cisco DNA Center for use in SD-Access, and also allows IP address pools defined in Cisco DNA Center to be available in IPAM.

Cisco network services orchestrator

Some large enterprise customers have standardized on Cisco network services orchestrator (NSO) as their network orchestration solution, to have a single way to automate management of their WAN, data center and campus deployments using the NETCONF / RESTCONF interfaces of NSO and the infrastructure. SD-Access deployments can be orchestrated by Cisco NSO through APIs exposed from Cisco DNA Center, and used by Cisco NSO.

The key benefits of SD-Access integration with network orchestrators include:

- Consistency across networks: same solution for current and NFV/SDN networks

- Agile service management: reduces time to market for deploying new services

- Best-in-industry multi-vendor support

Firewalls

Stateful firewalls can be used to enforce policy driven inter-VN communication in an SD-Access deployment. Integration of Cisco ASA with Cisco SD-Access enables the firewalls to have full access to context information, including the security context of the endpoint's group classification. This enables simplified policy administration and monitoring of traffic through firewalls, such as ASA firewalls and Firepower Threat Defense firewalls, in a manner that is consistent with the policy model in Cisco DNA Center. This integration is also available on some non-Cisco firewalls such as CheckPoint.

Security analytics

Integrations between security analytics solutions such as StealthWatch and SD-Access empower network and security teams to accelerate incident response. Through the Rapid Threat Containment capabilities of Cisco ISE as part of SD-Access, hosts identified as being compromised or exhibiting suspect behavior through Stealthwatch can be easily quarantined or blocked on the network.

Summary, next steps and references

Summary

In today's world, the network connects everything. But with the growth of users and different devices, typical manual provisioning, monitoring and troubleshooting makes managing the network far more challenging. Cisco SD-Access has been designed from the ground up to address these challenges by enabling policy-based automation and assurance from the edge to the cloud, for both wired and wireless access. This is the difference between a network that needs continuous attention and one that simply understands what your organization needs and makes it happen. It's the difference between doing thousands of tasks manually and having an automated system that helps you focus on the needs of the business, not on the network — now and in the future.

Using Cisco SD-Access offers multiple benefits for an organization:

Greater speed and agility: SD-Access allows an organization to move faster and more nimbly — allowing for more rapid and seamless deployment of new network innovations that support business requirements and generate more successful outcomes.

Greater efficiency, deeper insights: SD-Access allows an organization to save money, by allowing for significantly more rapid and secure network deployment, driving greater efficiencies in network operation, and by providing insights into the ways users are leveraging the network and the operation of the applications that are in use.

Reduction of risk: SD-Access allows an organization to provide integrated security as an inherent property of the SD-Access fabric, reducing the attack surface of the network and providing granular user / device / application access controls, allowing an organization to rapidly roll out and maintain an inherently secure and flexible infrastructure.

All of this is made possible by the key elements within SD-Access fabric that we have examined: **automation, policy and assurance**, all enabled by **Cisco DNA Center**. SD-Access provides a substantially more powerful and flexible way to deploy enterprise networks that are at the same time simpler to design, deploy, and operate.

SD-Access gives IT time back by dramatically reducing the time it takes to manage and secure a network. It also improves the overall end-user experience.

DIAGRAM   Benefits of SD-Access

Network Provisioning	Threat Defence	Monitoring and Troubleshooting	End User Experience
67%	48%	80%	94%
Reduction in network provisioning costs	Reduction in cost impact of a security breach	Reduction in cost impact of a security breach	Reduction in costs to optimize policies

Source: Internal TCO analysis with major Cisco Global Enterprise customers

Next steps

So now that Cisco SD-Access and its benefits have been understood, what are the next steps a network administrator can take? How can an organization begin its journey towards a software-defined future?

SD-Access offers several value propositions for common challenges encountered in traditional enterprise networking, based on evolving trends around mobility, enterprise IoT and cloud; and providing for an agile, secure enterprise networking architecture. SD-Access can be deployed on existing networks and for new infrastructure projects.

Identify key objectives and use cases

In order to migrate an existing network to an SD-Access architecture, it is worthwhile to start with identifying the key objectives or use cases for the network environment, including anticipated and upcoming requirements that may need to be addressed. Based on the stated key objectives of the network architecture and the content covered in this book, one could map how the capabilities offered by SD-Access can be applied to the use cases and requirements of the organization.

Assess infrastructure readiness

Once the SD-Access deployment use case objectives have been defined, the next step is to identify parts of the network infrastructure that will be migrated to the SD-Access architecture to achieve project objectives. This includes planning for the deployment of Cisco DNA Center and Cisco Identity Services Engine (ISE) based on use case needs. It is also advisable to consider the hardware and software compatibility of the network infrastructure under consideration to meet the solution deployment needs.

Define policy objectives

Moving to SD-Access often entails discussing policy objectives for the environment with the appropriate stakeholders, based on business requirements. This is an important consideration since the policies in SD-Access are applied in terms of business functions on a network-wide basis versus traditional methods. As a result, design of

group structures and key criteria to map the organization's endpoints can greatly facilitate future policy definition and automation. Once appropriate policy objectives are defined, organizations can rapidly start to realize the benefits of centralized policy automation with SD-Access.

It is recommended that organizations start with a coarse grouping structure, such as employees versus partners, before moving on to more granular classifications such as finance employee versus HR employee. This ensures that the number of scalable groups can grow alongside the customer's operational familiarity with the SD-Access policy model.

Plan and execute the migration

Some customers, especially when trying to incorporate SD-Access into their brownfield projects, may choose to go down the path of installing a new parallel network infrastructure that is exclusively SD-Access-enabled while keeping the older network in place. Others will perform a phased migration by selecting a specific segment or segments of their network as initial migration targets. In either case, it is recommended that customers start with a limited area of deployment and pilot SD-Access, familiarize themselves with the operations and technology, and then over time, expand the deployment and use cases.

References

Additional sites which offer more detailed information about Cisco Software-Defined Access include:

www.cisco.com/go/sda - provides an overview and additional information on all components and aspects of SD-Access: automation, assurance, supported platforms, customer references and testimonials, and a wealth of the most up-to-date information on SD-Access.

www.cisco.com/go/dnacenter — provides an overview and additional information on Cisco DNA Center.

cs.co/sda_tech_paper — SD-Access solution white paper. It provides a mid-level technical overiview of all of the major Cisco DNA and SD-Access components and their relationships. It's a great place to continue your journey!

www.cisco.com/go/cvd — includes the Software-Defined Access Cisco validated design (CVD) document covering SD-Access design options, operational capabilities, and recommendations for deployment. It provides direct insight into the best practices for the design, operation, and use of SD-Access in customer network deployments.

Key Cisco Live 365 Sessions:
Cisco SD-Access — *A Look Under the Hood* — BRKCRS-2810
(*https://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2810*)

Cisco SD-Access — *Building the Routed Underlay* — BRKCRS-2816
(*http://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2816*)

Cisco SD-Access — *External Connectivity* — BRKCRS-2811
(*https://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2811*)

Cisco SD-Access — *Extending Segmentation and Policy into IoT* — BRKCRS-2817
(*http://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2817*)

Cisco SD-Access — *Migration — BRKCRS-2812*
(https://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2812)

Cisco SD-Access — *Monitoring and Troubleshooting — BRKCRS-2813*
(https://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2813)

Cisco SD-Access — *Assurance — BRKCRS-2814*
(https://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-2814)

Cisco SD-Access — *Policy — BRKCRS-3811*
(https://www.ciscolive.com/global/on-demand-library.html?search=BRKCRS-3811)

Cisco SD-Access — *Wireless Integration — BRKEWN-2020*
(https://www.ciscolive.com/global/on-demand-library.html?search=BRKEWN-2020)

Cisco SD-Access — *Data Center Integration — BRKDCN-2489*
(https://www.ciscolive.com/global/on-demand-library.html?search=BRKDCN-2489)

Cisco SD-Access YouTube channel:

www.youtube.com/user/Cisco/search?query=SD-Access

Acronyms

AAA — Authentication, authorization, and accounting

ACI — Application-centric infrastructure

ACL — Access control list

AD — Active directory

AP — Access point

API — Application programming interface

APIC — Application policy infrastructure controller

AWS — Amazon Web Services

BGP — Border gateway protocol

CAPWAP — Control and provisioning of wireless access points

CDB — Cisco digital building

Cisco DNA — Cisco Digital Network Architecture

Cisco ISE — Cisco Identity Services Engine

CDP — Cisco discovery protocol

CLI — Command line interface

CP — Control plane

CPP — Cloud policy platform

CUCM — Cisco Unified Communications Manager

DDI — DHCP, DNS, and IPAM

DHCP — Dynamic host configuration protocol

DMVPN- Dynamic multi-point VPN

DNS — Domain name system

EPG — Endpoint group

ERSPAN — Encapsulated remote switch port analyzer

eWLC — embedded wireless LAN controller

FiaB — Fabric-in-a-box

GUI — Graphical user interface

HA — High availabilty

HIPAA- Health Insurance Portability and Accountability Act

HSRP — Hot standby routing protocol

IE — Industrial ethernet

IP — Internet protocol

IPAM — IP address management

IoT — Internet of things

IS-IS — Intermediate system to intermediate system

ISV — Independent software vendor

IT — Information technology

ITSM — IT services management

LAN — Local area network

LISP — Locator/identity separation protocol

MAC — Media access control

MPLS — Multi-protocol label switching

NETCONF — Network configuration protocol; used to install, manipulate, and delete the configuration of network devices

NGFW — Next-generation firewall

OMP — Overlay management protocol

OSPF — Open shortest path first

OT — Operational technology

PCI — Payment card industry

PIM — Protocol independent multicast

PIM-ASM — PIM any source multicast

PIM-SSM — PIM source specific multicast

PNP — Plug-n-Play

PxGRID — Platform exchange grid

REP — Resilient ethernet protocol

RESTCONF — YANG network management protocol specification to a Representational State Transfer (REST) interface

RP — Rendezvous point

SD-Access — Software-Defined Access

SGACL — Scalable group access control list

SGT — Scalable group tag

SGT-aware NGFW — Scalable group tag-aware next-generation firewall

SIEM- Security information and event manager

SNMP — Simple network management protocol

SPAN — Switch port analyzer

STP — Spanning tree protocol

SWIM — Software Image Management

SXP — SGT exchange protocol

TCP — Transmission Control Protocol

TLV — Type length value

UADP — Unified Access Data Plane

UI — User interface

VN — Virtual network

VNI — Virtual network instance

VPN — Virtual private network

VRF — Virtual routing and forwarding

VRRP — Virtual router redundancy protocol

VXLAN — Virtual extensible local area network

WLC — Wireless LAN controller

WLAN — Wireless local area network

YANG — Yet another next generation; data modeling language for the definition of data sent over network management protocols such as the NETCONF and RESTCONF