Cisco Software-Defined Access for Industry Verticals

From Design to Migration

cisco

Cisco Software-Defined Access for Industry Verticals

Icons used in this book

Border Node - Switch

Control Plane Node - Switch

Border Node - Router

Control Plane Node - Router

Edge Node

Extended Node

Transit Control
Plane Node - Switch

Transit Control
Plane Node - Router

Fabric Wireless
LAN Controller

Fabric WLC HA SSO

Border Node
Switch Stack

Border Node and Control
Plane Node Switch Stack

Edge Node Switch Stack

Colocated Border Node and Control
Plane Node with Layer 2 Handoff

Cisco DNA Center

Cisco Identity
Services Engine

Cisco vManage

Cisco vBond

Colocated Border Node and
Control Plane Node - Switch

Colocated Border Node and
Control Plane Node - Router

Policy Extended Node

Fabric in a Box

Fabric Access Point

Fabric in a Box
Switch Stack

Fabric Site

DHCP, DNS, and AD

Cisco vSmart

Embedded Wireless
LAN Controller

Multisite Remote Border Node

Multicast Rendezvous Point

Layer 3 Switch

Layer 2 Switch

Layer 3 Switch Stack

Layer 2 Switch Stack

Access Point

Access Point

Wireless LAN Controller

WLC HA SSO

User

User Group

Laptop

Servers

Machinery

Machinery

IP Phone

Mobile Phone

Security
Camera

Credit Card
Reader


Smart Waste


Badge Reader


Smart City


IT Professional


OT Professional


Medical Devices


Traffic Light


Medical Devices


Street Light


Router


Firewall


Services Block Switch


Internet


User


Building Management Systems


Lights


Security Group Tag


WAN Edge

Preface

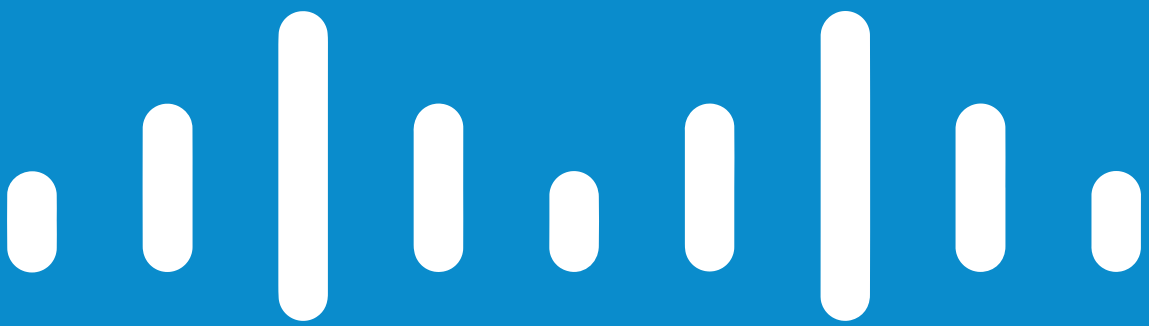

Authors

This book represents an intense collaboration between Technical Marketing, Engineering, Sales, and CX during a week-long comprehensive session at Cisco Headquarters in San Jose, CA.

Devi Bellamkonda | CCIE (DC, SP)
Technical Marketing Technical Leader, Cisco SD-Access

Dhrumil Prajapati | CCIE (R&S, SP), CCDE
Senior Multi-Domain Architect – CX GES Architectures

Jonathan Cuthbert
Technical Marketing Engineer, Cisco SD-Access CVD Author, UI Architecture

Kedar Karmarkar
Principal, SD-Access Technical Marketing Engineer

Keith Baldwin | CCIE (R&S, Wireless), CCDE
Senior Technical Solutions Architect, Campus Automation Center of Excellence

Mahesh Nagireddy | CCIE (R&S)
Technical Marketing Technical Leader, Cisco SD-Access

Parthiv Shah
Principal Engineer, Cisco SD-Access Enterprise Networking Engineering

Pete Kavanagh
Solution Architect, Industrial IoT

Prakash Jain
Principal Engineer, Cisco SD-Access Enterprise Networking Engineering

Prashanth Kumar Davanager Honneshappa
Technical Marketing Engineering, Cisco SD-Access | SD-WAN CVD Author

Raja Janardanan
Principal Engineer, Cisco Enterprise Solutions Engineering

Sanjay Hooda
Distinguished Engineer, Cisco SD-Access Design and Architecture

Scott Hodgdon
Technical Marketing Technical Leader, Cisco SD-Access

Acknowledgment

A tremendous thank you to Cisco's Enterprise Networking Technical Marketing, Product Management, Engineering, Sales, and Customer Experience teams who recognized the need for this book and supported its development. A special thanks to Jeff Scheaffer, Jeff McLaughlin, and Paul Nguyen for supporting the efforts of the authoring team. We would also like to thank Shannon Chavez for her incredible resource organization throughout this process, and the many amazing Cisco resources who provided information and clarifications throughout the process.

We would also like to extend our sincerest appreciation to our Book Sprints team (*www.booksprints.net*):

- Karina Piersig – Book Sprint Facilitator
- Raewyn Whyte – Copy Editor
- Christine Davis – Copy Editor
- Lennart Wolfert – Illustrator
- Henrik van Leeuwen – Illustrator
- Manu Vazquez – Book Designer

Karina and the team were outstanding in creating an environment that allowed our thoughts and ideas to collaboratively flourish which resulted in producing this publication.

Feedback

The team is proud to have come together for this special occasion to share our collective experience with the Cisco SD-Access solution with you. This team came from a diverse background in technical marketing, development, design, sales, implementation, and support.

We are already looking forward to writing the next book and would love to collaborate with you on the next topics!

Your feedback will help drive the right focus and direction to the solution, product, and portfolio.

We encourage you to share your thoughts, ideas, and general comments about this book at the following link:

http://cs.co/sda_book_feedback

Introduction

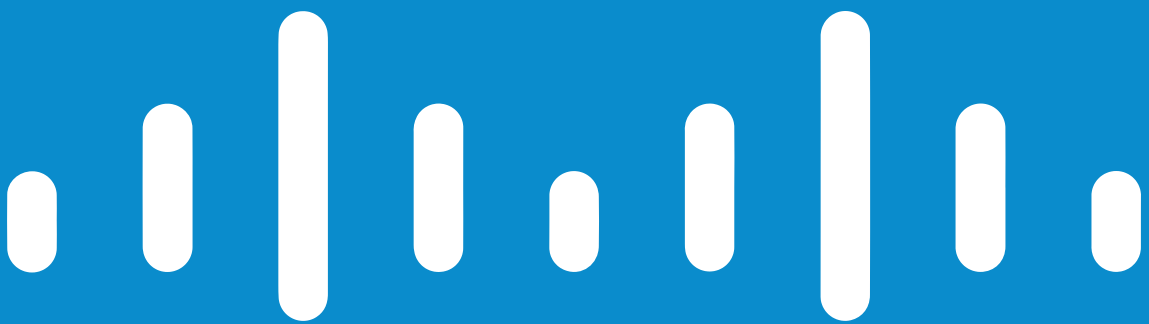

Executive Summary

Digital transformation is creating new opportunities in every industry. In healthcare, doctors can monitor patients remotely and leverage medical analytics to predict health issues. In education, technology is enabling connected campuses, providing personalized and equal access to learning resources. In retail, shops can provide a seamless, engaging experience in-store and online using location awareness. In the world of finance, digital technology enables users to securely bank anywhere, anytime, using the device of their choice. In today's world, digital transformation is essential for businesses to stay relevant.

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

What is Cisco Software-Defined Access?

Cisco Software-Defined Access (SD-Access), a solution within the Cisco Digital Network Architecture (Cisco DNA), is built on Intent-based networking principles. This solution provides visibility-based, automated end-to-end segmentation to separate user, device, and application traffic without redesigning the underlying physical network.

Cisco SD-Access automates user access policy so organizations can ensure the right policies are established for any user, device, or application across the network. With unified access policies across LAN and WLAN, a consistent user experience is created without compromising security. By combining security and networking operations, SD-Access enhances visibility by defining access policies and automating network configurations to implement these policies.

Why Cisco SD-Access?

Challenges in Traditional Networks

Organizations deploying traditional network architectures are facing mounting challenges as the number of users, devices, and types of devices continue to proliferate. These challenges began with the introduction of Bring Your Own Device (BYOD) networks and has accelerated as the Internet of Things (IoT) trend has been adopted by increasingly more organizations.

Identifying, grouping, and analyzing the traffic of all of these users and devices is a significant concern for organizations that want to ensure that they do not impact their corporate infrastructure should a BYOD device become compromised. The ability to define and enforce granular access policies in an automated way is nearly impossible if these devices cannot be classified.

Figure 1.1: Traditional network challenges

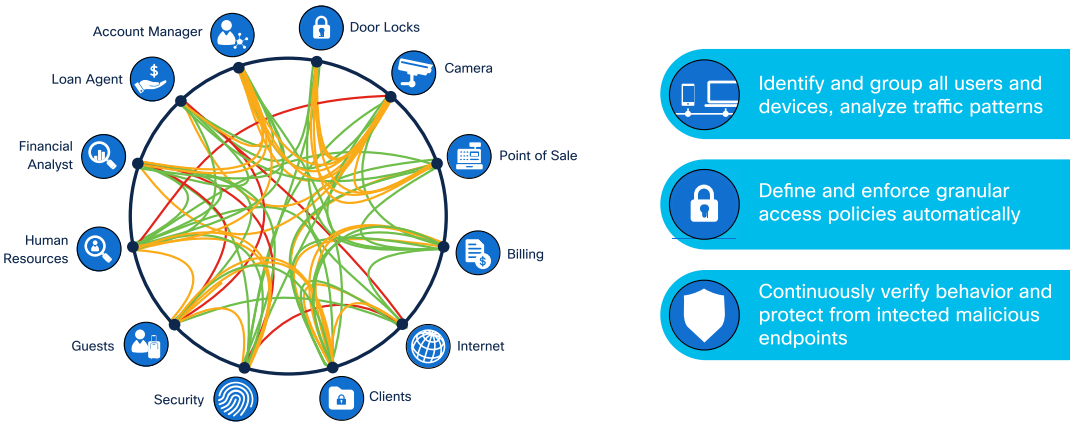

In traditional networks, the need for a significant number of VLANs and manual Access Control Lists (ACLs) across multiple and often disparate devices becomes a recipe for manual misconfiguration disasters. As time goes on and the business expands, more devices and locations are added, increasing complexity and the possibility for errors. New and more complex security rules must be manually updated across the enterprise.

For example, consider ACL 2 in Figure 1.2. When the organization adds Branch A to the enterprise, the network operations team may need to update the ACL in both the HQ and Branch 1 locations. If a manual mistake is made during that update, then the security policy will be inconsistent and could result in a security breach.

Figure 1.2: Traditional network security challenges

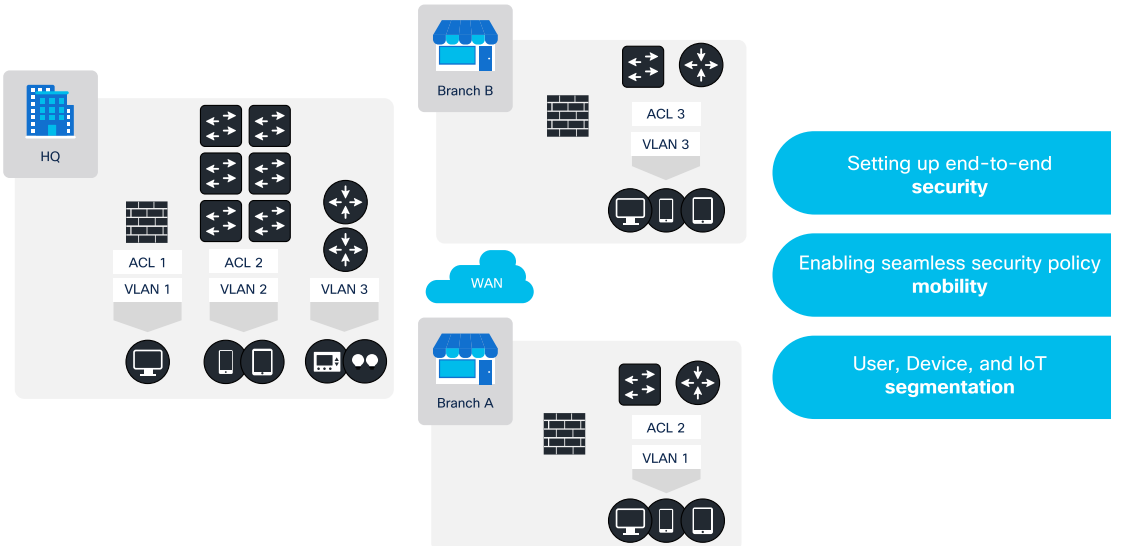

Figure 1.2 also outlines three of the main security obstacles that customers face when utilizing traditional networking methodologies in today's digital world:

- Setting up end-to-end security
- Enabling seamless security policy mobility
- User, device, and IoT segmentation

Setting up end-to-end security is a challenge as the network grows. An even larger issue is ensuring that security policy follows a user or device from location to location. In the traditional networking world, security policies are based on IP addresses. If a user or device moves from one location to another, then security policies must exist based on all of the different IP addresses that the user or device will use in the different locations. This becomes extremely difficult to manage as the network evolves.

Networks can contain corporate users and devices, user BYOD devices, and IoT devices. Without a dynamic method to identify, onboard, and secure users and devices, network administrators must spend significant and arduous time planning and configuring network changes to ensure that every device is securely onboarded onto the network using the proper network segment. If the users and devices move, those changes must be manually made in anticipation of the move.

How Cisco SD-Access Resolves These Challenges

Cisco SD-Access is built on an Intent-based Networking foundation that encompasses visibility, automation, security, and simplification. Using Cisco DNA Center automation and orchestration, network administrators can implement changes across the entire enterprise environment through an intuitive, GUI-based interface. Using that same controller, they can build enterprise-wide Fabric architectures, classify endpoints for security grouping, create and distribute security policies, and monitor network performance and availability.

SD-Access secures the network at the macro- and micro-segmentation levels using Virtual Routing and Forwarding (VRF) tables and Security Group Tags (SGTs), respectively. This is called Multi-Tier Segmentation, which is not optimal in traditional networks. With this SD-Access, this segmentation happens at the access port level. This means the security boundary is pushed to the very edge of the network infrastructure for both wired and wireless clients.

With Multi-Tier Segmentation, network administrators no longer have to undertake configurations in anticipation of a user or device moving locations as all of the security contexts associated with a user or device are dynamically assigned when they authenticate their network connection. SD-Access provides the same security policy capabilities whether the user or

device is attached via a wired or wireless medium, so secure policy consistency is maintained as the user or device changes attachment type.

Instead of relying on IP-Based security rules as in a traditional network, SD-Access relies on centralized group-based security rules utilizing SGTs that are IP address-agnostic. This means that as a user or device moves from location to location and changes IP addresses, their security policy will remain the same because their group membership is unchanged regardless of where they access the network. This reduces pressure on network administrators since they do not have to create as many rules or manually update them on different devices. This, in turn, leads to a more dynamic and stable environment for network consumers.

How can a network be both dynamic and stable at the same time? When a rule does have to be created or changed, it can be done for all users of a group in Cisco DNA Center. Those rules are then dynamically populated to all relevant network devices that need that rule, ensuring both accuracy and speed for the update.

Looking back at the original question of Why Cisco SD-Access? There are three primary reasons which make it superior to traditional network deployments:

- Complexity reduction and operational consistency through orchestration and automation
- Multi-Tier Segmentation which includes group-based policies
- Dynamic policy mobility for wired and wireless clients

Cisco SD-Access for Industry Verticals

Cisco SD-Access has been deployed in thousands of customer networks across all major industry verticals. This book focuses on five of those specific verticals: Operational Technology (OT), Healthcare, Large Enterprise, Finance, and Universities.

The book provides a brief review of the components of the SD-Access architecture and will introduce commonly used features from these verticals. The final section focuses on migration, and provides the next steps to evolve from the traditional network architecture to the SD-Access architecture.

This book is intended for practitioners of all architectural levels. The technical nature of this book is best suited for network architects in our customer and partner community.

Chapter 2: Architecture Overview

This chapter provides an overview of the Fabric components and feature information that is needed to understand the designs and topics discussed in later chapters. It also covers common Cisco and third-party application integrations used in SD-Access deployments.

Chapter 3: OT Integration with Cisco SD-Access

This chapter focuses on deployments of Cisco SD-Access in Operation Technology (OT) environments. An OT network is typically a dedicated and physically separate network from an enterprise network. OT networks are often deployed by manufacturing, industrial, and utility organizations.

Chapter 4: Cisco SD-Access in Healthcare

Significant changes have been taking place in the Healthcare industry, such as exponential growth in telehealth and virtual care, sudden increases in remote workforces, and fast-evolving primary care models. The purpose of this chapter is to provide design guidance for a typical healthcare deployment profile using Cisco DNA Center and SD-Access.

Chapter 5: Cisco SD-Access for Large Enterprises and Government

Enterprise Networks owned by multinational corporations and governments are some of the largest networks in existence. This chapter explains how Cisco SD-Access addresses the unique scale and connectivity requirements faced by organizations with these large-scale networks.

Chapter 6: Cisco SD-Access in Universities

University networks are unique because they are built to provide the ultimate sharing and collaborative experience in the pursuit of learning. This chapter explores typical design caveats that university networks must overcome when adopting new technology and methods in which SD-Access design inherently overcomes these challenges.

Chapter 7: Cisco SD-Access in Financial Verticals

This chapter addresses common SD-Access use cases in Financial verticals. This vertical faces unique challenges in the areas of regulatory compliance, service quality, site scale, and resiliency, just to name a few.

Chapter 8: Migration to Cisco SD-Access

This chapter focuses on migration options, how to migrate your existing network, and steps to evolve a traditional network into an SD-Access network.

Architecture Overview

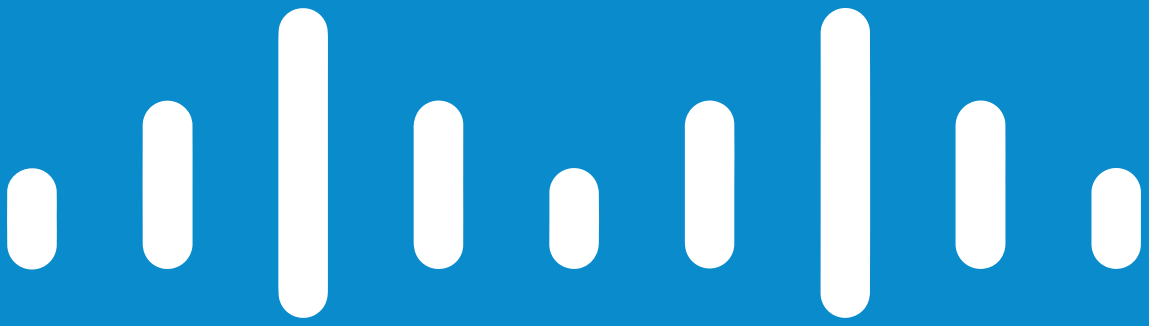

If you are reading this book, there is a good chance you already have some knowledge about Cisco SD-Access from an architectural or design perspective. The purpose of this chapter is to provide you with an overview of the Fabric component, design, and feature information that you will need to understand the designs and topics discussed in later chapters.

As every reader will have a different level of experience with Cisco SD-Access, the chapter contains information for new to advanced users. We have tried to identify common components across all verticals so that you do not have to see them repeated multiple times.

We recommend that you refer to the Table of Contents to determine the appropriate starting point for you or find a specific topic or feature of interest.

Solution Components

There are three fundamental pillars of the Cisco SD-Access solution. These pillars are Cisco DNA Center, the Cisco Identity Services Engine, and the wired and wireless device platforms that support Fabric functionality.

Cisco DNA Center

Cisco DNA Center is a foundational component of Cisco SD-Access, enabling automation of device deployments and configurations into the network to provide operational efficiency and the speed and consistency required. Through its automation capabilities, the control plane, data plane, and policy plane for the network devices are easily, seamlessly, and consistently deployed. In addition, Cisco DNA Center Assurance provides visibility and context into the managed infrastructure devices and connected endpoints.

Identity Services Engine

Cisco Identity Services Engine (ISE) is a secure network access platform enabling increased management awareness, control, and consistency for users and devices accessing an organization's network. ISE is an integral component of Cisco SD-Access for implementing a network access control policy. ISE performs policy implementation, enabling dynamic mapping of users and devices to security groups and simplifying end-to-end security policy enforcement.

Network Infrastructure

The Cisco SD-Access solution infrastructure includes routers, switches, access points, and Wireless LAN Controllers. On these devices, Cisco DNA Center will deploy the various Fabric Site roles based on the network administrator's choices in the User Interface (UI).

Fabric Overview

A Fabric is simply an Overlay network. Overlays are created through encapsulation, which adds one or more additional headers to the original packet or frame. An Overlay network creates a logical topology which virtually connects devices built over an arbitrary physical Underlay topology.

In an idealized, theoretical network, every device would be connected to every other device. In this way, any connectivity or topology imagined could be created. While this theoretical network does not exist, there is still a technical desire to connect all these devices in a full mesh. This is where the term Fabric comes from: it is a cloth by which everything is connected. An Overlay (or tunnel) provides this logical full-mesh connection in networking.

Underlay

The Underlay network is defined by the physical switches and routers that are used to deploy the Cisco SD-Access network. All network elements of the Underlay must establish IP connectivity via the use of a routing protocol. Instead of using arbitrary network topologies and protocols, the Underlay implementation for Cisco SD-Access uses a well-designed, Layer 3 foundation inclusive of the Access Layer switches, known as a *Layer 3 Routed Access* design. This ensures performance, scalability, resiliency, and deterministic convergence of the network.

In Cisco SD-Access, the Underlay devices support the physical connectivity for users and endpoints. However, end-user subnets and endpoints are not part of the Underlay network and become part of the automated Overlay network.

Overlay

An Overlay network is a logical topology used to virtually connect devices and is built over an arbitrary physical Underlay topology. The Cisco SD-Access Overlay network is created on top of the Underlay network through virtualization, creating Virtual Networks (VNs). Data, traffic, and control plane signaling are contained within each Virtual Network, maintaining isolation among the networks and independence from the Underlay network. Multiple Overlay networks can run across the same Underlay network through virtualization.

Fabric Network Overview

Fabric Site 

A Fabric Site is composed of a unique set of devices operating in a Fabric role along with the intermediate nodes that connect those devices. At a minimum, a Fabric Site must have a Border Node and a Control Plane Node, but most often, it will also have Edge Nodes. A Fabric Site has an associated Fabric Wireless LAN Controller (WLC) and potentially an ISE Policy Service Node (PSN).

Transits

Transits can connect multiple Fabric Sites or can connect a Fabric Site to non-Fabric domains such as a data center or the Internet. Transits are a Cisco SD-Access construct that defines how Cisco DNA Center will automate the Border Node configuration for the connections between Fabric Sites or between a Fabric Site and an external domain. There are two types of Transits: IP-Based and SD-Access.

IP-Based Transit

With IP-Based Transits, the Fabric VXLAN header is removed, leaving the original native IP packet. Once in native IP, packets are forwarded using traditional routing and switching protocols between Fabric Sites. Unlike an IP-Based Transit, an SD-Access Transit is an Overlay that rides on top of a WAN/MAN network much like SD-WAN and DMVPN do. IP-Based Transits are provisioned with VRF-Lite connection to an upstream peer device. IP-Based Transits most often connect to a data center, WAN, or Internet. An IP-Based Transit can also be used to connect to shared services using a *VRF-Aware Peer*.

SD-Access Transit

An SD-Access Transit uses VXLAN encapsulation and does not rely on a VRF-Lite connection to an upstream peer. Like an IP-Based Transit, packets are forwarded using traditional routing and switching protocols between Fabric Sites. Unlike an IP-Based Transit, an SD-Access Transit is an Overlay that rides on top of a WAN/MAN network much like SD-WAN and DMVPN. There are key considerations when using an SD-Access Transit:

- Connections should accommodate the recommended MTU settings used for Cisco SD-Access in the Campus Network.
- IP reachability must exist between Fabric Sites. Specifically, there must be a known Underlay route between all Fabric Nodes (the default route cannot be used for this purpose).

Virtual Networks

Cisco SD-Access provides Layer 3 and Layer 2 connectivity across the Overlay using Virtual Networks.

Layer 3 Overlays emulate an isolated routing table and transport Layer 3 frames over the Layer 3 network. This type of Overlay is called a Layer 3 Virtual Network. A Layer 3 Virtual Network is a virtual routing domain that is analogous to a Virtual Routing and Forwarding (VRF) table in a traditional network.

Layer 2 Overlays emulate a LAN segment and transport Layer 2 frames over the Layer 3 network. This type of Overlay is called a Layer 2 Virtual Network. Layer 2 Virtual Networks are virtual switching domains that are analogous to a VLAN in a traditional network.

Security Group Tags (SGTs)

Security Group Tags (SGTs) are metadata values that indicate the privileges of the source within the entire network. There are several methods to propagate SGTs. Within the SD-Access Fabric, SGTs are propagated in the header of the VXLAN encapsulated packets.

With identity services provided through ISE, users and devices connected to the Fabric are dynamically mapped to an SGT. This simplifies end-to-end security policy management and enforcement at a greater scale than traditional network policy implementations, which rely on IP access lists.

Segmentation Capabilities

Cisco SD-Access creates two layers of segmentation. The first layer of segmentation is referred to as macro-segmentation and is achieved through the use of Virtual Networks. Users, devices, and applications can be put into different Overlay networks, enabling isolation between them.

The second layer of segmentation is referred to as micro-segmentation and is achieved through the use of SGTs. SGTs are used to segment inside the Virtual Network. By default, users and devices in the same Virtual Network can communicate with each other. SGTs can be used to permit or deny communication within a given Virtual Network.

Fabric Components

Control Plane Node 

An SD-Access Control Plane Node runs the Locator ID Separation Protocol (LISP) to register endpoints and provide next-hop forwarding information for traffic generated by those endpoints.

Border Node 

An SD-Access Border Node is an entry and exit point to the Fabric Site. In effect, it speaks two languages: SD-Access Fabric on one link and traditional routing and switching on another.

Edge Node 

An SD-Access Edge Node provides first-hop services and hosts the Anycast Layer 3 Gateways needed for users, devices, and applications to connect to the network.

Transit Control Plane Node 

The role of a Transit Control Plane Node is to learn which prefixes are associated with each Fabric Site and to direct traffic to these sites across an SD-Access Transit using control plane signaling.

When traffic from an endpoint in one site needs to send traffic to an endpoint in another site, the Transit Control Plane Node is queried to determine to which site's Border Node this traffic should be sent.

Fabric Access Point 

Fabric-Mode Access Points (APs) are APs associated with a Fabric Wireless LAN Controller (WLC) that have been configured with Fabric-enabled SSIDs.

Fabric Wireless LAN Controller 

A Fabric Wireless LAN Controller connects Fabric APs and Wireless endpoints to the SD-Access Fabric. The Fabric WLC registers wireless clients with the Control Plane Node.

Embedded Wireless LAN Controller 

The Embedded Wireless LAN Controller on Catalyst 9000 Series switches allows easy deployment of wireless in the SD-Access network without having to manage or deploy a separate physical device.

Fabric-Mode APs continue to support the same wireless media services that traditional APs support, such as applying Application Visibility and Control (AVC), Quality of Service (QoS), and other wireless policies.

Extended Node 

Extended Nodes offer a Layer 2 port extension to an Edge Node while providing segmentation and group-based policy to the endpoints connected to these switches. Endpoints, including Fabric APs, can connect directly to the Extended Node. VLANs and SGTs are assigned using host onboarding as part of Fabric provisioning.

Policy Extended Node 

A Policy Extended Node supports enhanced security capabilities compared to an Extended Node. In addition to the operation and management provided by a classic Extended Node, Policy Extended Nodes directly support SGTs. This local SGT support provides direct east-west traffic enforcement on the device.

Intermediate Node 

An Intermediate Node is part of the Layer 3 Underlay network used for interconnections among the devices operating in a Fabric Site. For example, if a Three-Tier Campus deployment provisions the Core switches as the Border Nodes and the Access switches as the Edge Nodes, the Distribution switches are the Intermediate Nodes. Intermediate Nodes are not limited to a single layer of devices.

Fabric in a Box 

Fabric in a Box is an SD-Access construct where the Border Node, Control Plane Node, and Edge Node run on the same Fabric Node. This may be a single switch, switch with hardware stacking, or StackWise Virtual deployment.

Peer Device (Fusion) 

Border Nodes are connected to next-hop devices, providing access to the network outside the Fabric Site. There are several common configuration options for the next-hop peer device. This device may peer (have IP connectivity and routing adjacency) with the Border Node using VRFs. This next-hop device may even continue the VRF segmentation extension to its next hop. This next-hop may not be VRF-aware and peer to the Border Node using the global routing table. The term *Fusion* has commonly been used in existing collateral for each of these next-hop device deployment types. Throughout the book, Peer will be used in place of *Fusion*.

Fabric Features and Capabilities

LISP Publisher/Subscriber (Pub/Sub)

LISP Pub/Sub was purpose-built for SD-Access and is extensible, and adaptable. It vastly simplifies network operations by natively publishing all Fabric Site prefixes to subscribed Border Nodes, utilizing native LISP rather BGP. This simplification will make the deployment and maintenance of an SD-Access solution less complex for operations and support teams.

LISP Pub/Sub enables several new capabilities that increase an SD-Access deployment's resiliency and flexibility, including Dynamic Default Border, Backup Internet, and SD-Access Extranet.

Dynamic Default Border

With the SD-Access Dynamic Default Border, the Fabric Overlay converges quickly in the event of uplink failures or upstream device failures on Border Nodes which result in the loss of the default route.

SD-Access Backup Internet

In multisite SD-Access Transit deployment, several Fabric Sites may have access to the Internet. Using the SD-Access Backup Internet functionality, Fabric Sites can use each other as a backup path to the Internet if their site-local Internet access is lost.

Cisco SD-Access Extranet

With Cisco SD-Access, users, devices, and applications are onboarded to join a Virtual Network in the Fabric. While these Virtual Networks are isolated from one another, users, devices, and applications require shared services such as DHCP, DNS, Active Directory, and ISE. Access to the Internet is usually needed as well. However, shared services and the Internet are often in dedicated VRFs or in the Global Routing Table (GRT).

Using SD-Access Extranet, Cisco DNA Center automates the configuration necessary to provide secure communication between Fabric Virtual Networks to shared services and the Internet. SD-Access Extranet does not need additional hardware, such as a Peer (*Fusion*) device, to provide this capability.

Wireless in Cisco SD-Access Fabric

Cisco SD-Access provides a unique differentiator by integrating the wireless control plane with the Overlay control plane of the wired world. Cisco SD-Access Wireless offers a centralized control and management plane via the WLC with a distributed data plane providing the best of both worlds – centralized and distributed wireless designs. The WLC integrates with the Control Plane Node, registering endpoints as they are onboarded and updating their location as they roam. This is the first instance where there is synergy between the wireless and the wired control planes.

This unique integration of wired and wireless brings several benefits to network users and the operations teams that support them:

- Simplification – networks can have a single subnet for both wired and wireless clients
- Consistency of Policy – the rich set of wired policies are extended to wireless traffic, and they are both applied at the Edge Node

- Improved performance – wireless roams are Layer 2 and do not require any form of anchoring

Wireless Over-the-Top

Cisco SD-Access has the flexibility to support a centralized wireless deployment called Wireless Over-the-Top (OTT). Wireless OTT support is crucial because there are several situations where it could be required:

- Existing Cisco WLCs and APs are not SD-Access Wireless-capable
- Third-party wireless presence in the network
- Wired and wireless asymmetric migration pace

In Wireless OTT deployments, wireless control, management, and data plane traffic traverse the Fabric in a Control and Provisioning of Wireless Access Points (CAPWAP) tunnel between the APs and WLC. This CAPWAP tunnel uses the Cisco SD-Access Fabric as a transport. Other vendor's wireless equipment may use a different tunneling protocol other than CAPWAP, but the concept of using the SD-Access Fabric as a transport is the same.

FlexConnect Over-the-Top

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

End-to-End Segmentation

Proper end-to-end network segmentation is a foundational measure to address various security and scalability issues that are seen in networks today. Cisco SD-Access provides the ability for network architects to create and deploy secure, network-wide segmentation for their users, devices, and applications across their networks, making networks more robust and

scalable. Macro- and micro-segmentation contexts are carried across the SD-Access Fabric and can be further extended beyond the Fabric network using methods described in this book's various vertical chapters.

Multisite Remote Border

The Multisite Remote Border feature enables the configuration of a single IP Address Pool across multiple Fabric Sites. It allows the same IP Address Pool to exist at multiple Fabric Sites by anchoring the IP Address Pool to specific Border Node(s) and Control Plane Node(s). Preserving the subnet across multiple sites provides a scalable strategy for conserving address space.

Multisite Remote Border specifies that a designated Fabric Site will be the ingress and egress location for all traffic in a designated Virtual Network even though that Virtual Network stretches across multiple Fabric Sites. A common use case for Multisite Remote Border is Guest traffic isolation.

In a traditional network, a wireless Guest Anchor Controller is usually located in the DMZ. Multisite Remote Border provides the same capabilities while also providing the segmentation and automation benefits of an SD-Access network.

Fabric Zones

Typically, in Cisco SD-Access deployments, all the Overlay subnets are available across all the Edge Nodes in that Site. Fabric Zones are a way to localize certain subnets to certain sections of the Fabric Site. For example, a Fabric Site may include ten buildings. There may not be a need to have all the subnets across all the ten buildings. Fabric Zones ensure that some IP subnets are only available in some buildings and not on the entire Campus.

Critical VLAN

The Critical VLAN functionality allows network access to users, devices, and applications using a specific VLAN if the ISE PSN is unreachable from the SD-Access network. In addition, currently, authentication users, devices, and applications will remain in their presently assigned VLANs, as periodic re-authentication is paused until the ISE PSN is reachable again.

Gateway Outside of the Fabric

In SD-Access, a default gateway is present on all Edge Nodes for each subnet in a Virtual Network within a given Fabric Site. Traffic destined to a remote subnet is processed by the default gateway on the Edge Node and then routed to the appropriate destination.

In many OT networks, there is a need for the default gateway to be on an external firewall and not on the local Edge Node. Firewall traffic inspection is a common security and compliance requirement in many networks.

By enabling Gateway Outside of the Fabric functionality, the default gateway is not provisioned on the Edge Nodes. The gateway can be provisioned on an external device such as a firewall, where traffic to that gateway can be inspected.

Layer 2 Flooding

By default, there is no flooding of traffic in the Cisco SD-Access network. ARP resolution is done in an optimized fashion by default in Fabric. However, some endpoints need broadcast support, particularly to support silent hosts. Layer 2 Flooding is a feature that enables the flooding of broadcast, link-local multicast, and ARP traffic for a given Overlay subnet. Layer 2 Flooding requires the support of native multicast in the Fabric Underlay network.

LAN Automation

LAN Automation helps enterprise IT administrators prepare, plan, and automate SD-Access Underlay networks. This simplifies network operations, frees IT staff from time-consuming and repetitive network configuration tasks, and creates a standard error-free Underlay network.

Multicast in Fabric

Multicast is supported in both the Overlay Virtual Networks and the physical Underlay networks in SD-Access, each achieving different purposes.

The multicast source can either be outside the Fabric Site (commonly in the data center) or can be in the Fabric Overlay, directly connected to an Edge Node, Extended Node, or associated with a Fabric AP. Multicast receivers are commonly directly connected to Edge Nodes or Extended Nodes, although multicast receivers can also be outside of the Fabric Site if the source is in the Overlay.

PIM Any-Source Multicast (PIM-ASM) and PIM Source-Specific Multicast (PIM-SSM) are supported.

Multicast Forwarding in SD-Access

SD-Access supports two different transport methods for forwarding multicast. One uses the Overlay, referred to as Head-End Replication, and the other uses the Underlay and is called Native Multicast.

Head-End Replication

Head-End Replication (or Ingress Replication) is performed either by the multicast first-hop router (FHR) when the multicast source is in the Fabric Overlay, or by the Border Nodes when the source is outside the Fabric Site.

Native Multicast

Native multicast does not require the ingress Fabric Node to do unicast replication. Rather, the whole Underlay including Intermediate Nodes is used to do the replication. To support Native multicast, the FHRs, Last-Hop Routrers (LHRs), and all network infrastructure between them must be enabled for multicast.

Quality of Service (QoS) in Cisco SD-Access Fabric

Cisco SD-Access Fabric encapsulation preserves the Differentiated Services Code Point (DSCP) value in the IP header of the incoming packet by copying the value to the DSCP in the outer IP packet. Thus, QoS DSCP values are preserved end-to-end across the Overlay.

Multiple Cisco DNA Center to Single ISE

The multiple Cisco DNA Center capability allows multiple Cisco DNA Center clusters to integrate with the same single ISE deployment. Using the concept of Author clusters and Reader clusters, this feature creates a single management point for policy definitions in the deployment. Replication of these definitions is then propagated from the Author Node to the Reader Nodes.

Figure 2.1: Multiple Cisco DNA Center architecture

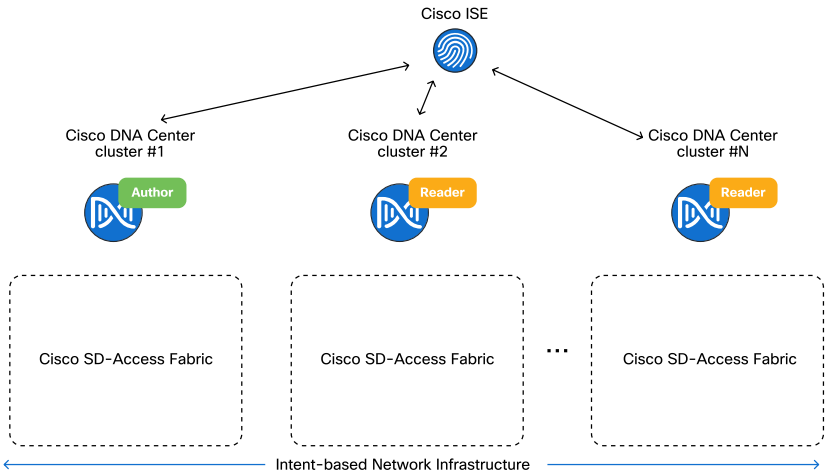

Ecosystem

Cisco DNA Center has an ecosystem with a variety of parallel solutions and third-party apps. This section describes those ecosystem solutions in the context of Cisco SD-Access.

Wide Area Bonjour

Bonjour is a zero-configuration solution that simplifies network configuration and enables communication between connected devices, services, and applications. Bonjour is designed for single Layer 2 domains such as small, flat networks.

The Cisco Wide Area Bonjour application on Cisco DNA Center eliminates the single Layer 2 domain constraint and expands the scope to larger Layer 3 domains which are used in SD-Access wired and wireless networks.

Cisco Secure Network Analytics (StealthWatch)

Just as it is important to secure endpoints in the network, it is just as important to have visibility into the traffic on the network and a method by which to audit it. Cisco Secure Network Analytics (StealthWatch) leverages NetFlow data from network devices throughout all areas of the network to provide a concise view of traffic patterns. This visibility allows a StealthWatch database record to be maintained for every communication that traverses a network device.

The StealthWatch Security Analytics application on Cisco DNA Center automates the provisioning of network devices so they send NetFlow data to StealthWatch. This provides visibility and real-time monitoring of all network traffic from the Cisco Secure Network Analytics Manager.

ThousandEyes

The ThousandEyes application is hosted on Catalyst 9000 Series Switches and is provisioned through a workflow in Cisco DNA Center. ThousandEyes provides a way to monitor and observe devices and applications in the network.

ThousandEyes Agent on Edge Nodes provides network and application visibility from client subnet to services. ThousandEyes Agent on Border Nodes provides network and application visibility from the Border Node to services outside of the Fabric.

Cisco DNA Center Plug-and-Play

Cisco DNA Center helps automate and onboard Cisco network devices with built-in Plug-and-Play (PnP) functionality. Plug-and-Play is a software agent on the network devices that calls home to Cisco DNA Center and downloads the required software and device configuration.

Return Material Authorization (RMA) Workflow

Cisco DNA Center RMA workflow makes replacing devices a simpler and zero-touch process. Customers flag a failed device in Cisco DNA Center, physically install the new device, and run the basic zero-touch workflow to bring up the device through the Plug-and-Play process. Using this process, Cisco DNA Center automates software image upgrades, installs appropriate licenses and certificates, and applies the basic configuration. Once a device

is detected by Cisco DNA Center, it will configure the replacement device with the old device configuration. Cisco DNA Center supports RMA for switching device platformsin both Fabric and non-Fabric Devices.

Cisco AI Endpoint Analytics

Cisco AI Endpoint Analytics is a solution that detects and classifies endpoints and IoT devices into different categories based on Endpoint Type, Hardware Model, Manufacturer, and Operating System type. The Cisco AI Endpoint Analytics engine and user interface runs on Cisco DNA Center and assigns labels to endpoints by receiving telemetry from the network infrastructure.

Fabric Assurance

Cisco SD-Access Assurance provides visibility into the Underlay and Overlay, and it enables reachability into critical network services in the Fabric Infrastructure. Fabric network devices are provisioned with model-driven streaming telemetry which monitors the status of certain protocol states. Any change in the protocol state is reported by the network device to Cisco DNA Center. In the Assurance dashboard, suggestive actions help network operators to narrow down the issue domain and eventually remediate the issue.

SD-Access Assurance provides visibility into the health and state of the Fabric Site, Virtual Networks, and SD-Access Transit.

3rd-Party Integrations

Cisco DNA Center integrates with various other IT applications such as IPAM services and ServiceNow IT Service Management (ITSM) suite. Two integrations are explained below.

IP Address Management systems (IPAM)

With third-party IPAM integration, all aspects of IPAM management, such as DNS and DHCP, can be done using one integrated platform. This integration eliminates manual processes and patchwork tools, increasing IP address management efficiency. Cisco DNA Center supports the ability to integrate third-party IPAM systems Infoblox and BlueCat.

Service Now

The Cisco SD-Access integration with ServiceNow monitors and publishes Fabric events that include Fabric Role updates. This provides security and other operational triggers to the IT Service Management system.

OT Integration with Cisco SD-Access

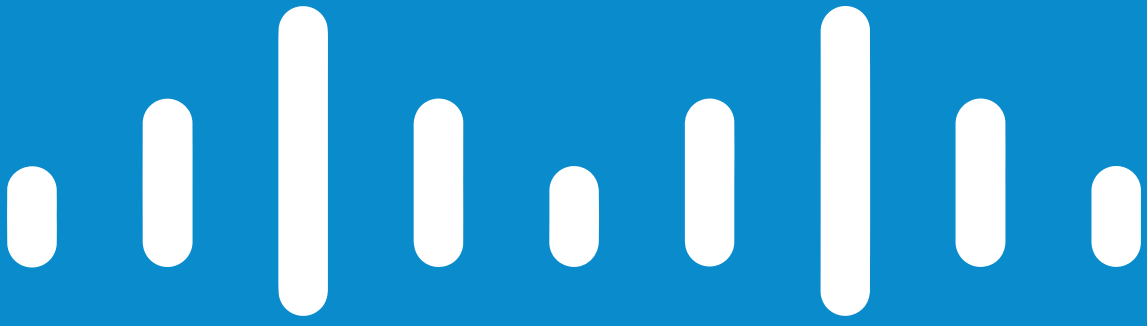

Introduction

Chapter Overview

The Purdue Model was released in 1990, and it is the standard by which Operational Technology (OT) networks have designed their networks. However, significant changes have occurred in the networking industry over the last 30 years. These innovations are being leveraged in the OT industry today.

This chapter provides design guidance for the manufacturing space, focusing on how the new innovations in Cisco SD-Access are being utilized to create simple, secure, and flexible OT networks. The chapter starts with the description of the OT Network and follows with requirements and solutions for the OT space. Finally, the chapter ends with deployment options.

What is an OT Network?

An OT network is typically a physically separate network from a customer's IT Corporate network. It can be tied to a division or department within the organization and is often in direct support of the customer's line of business. OT networks are best described as a spectrum of technologies, where the spectrum includes traffic types, flows, topology mandates/restrictions, segmentation, and so on. These components can be introduced if the organization is part of a regulated industry or if there are existing norms within certain industries that have been in place for many years.

For example, an Electrical Utility High-Voltage Substation could be at one end of the spectrum. In this case, standards such as IEEE 1613 and IEC 61850 must be adopted. Other unique requirements include lossless LAN

duplication, Supervisory Control, and Data Acquisition (SCADA) protocols, and having elements in-sync to a sub-microsecond accuracy.

Somewhere in the middle of the spectrum could be Manufacturing, which includes Industrial Automation. Here we typically see the Profinet protocol and a strongly segmented hierarchical approach, usually in line with IEC 62443.

Roadways could be at the least strict end of the spectrum. Here we see acceptable recovery times of less than 500 milliseconds (ms) and mostly regular IPv4 TCP and UDP unicast traffic.

Different OT networks can be categorized in the spectrum of quadrants at a high level as shown in the following chart. Globally, interpretations might vary, but the chart describes how these interpretations are categorized.

Figure 3.1: Traffic types, resiliency requirements, and industry framework adherence across various OT examples

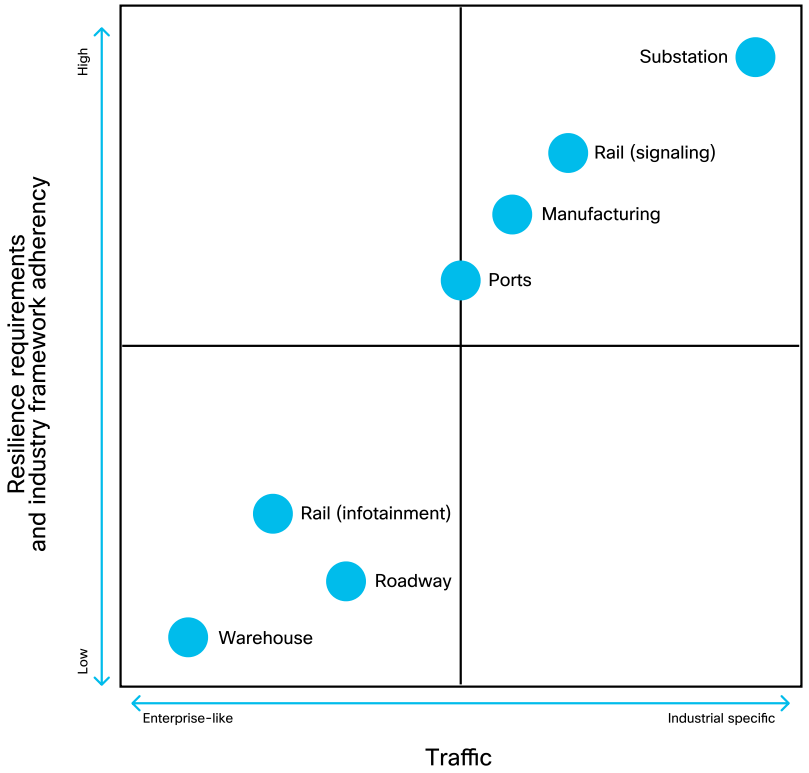

All of these OT networks may have a lot of differences, but they also have some commonalities. We will discuss how Cisco SD-Access helps in the implementation and integration of OT networks with IT networks and how Cisco SD-Access can be used to build OT networks.

Organization and Infrastructural Separation

IT and OT are often two different departments within a single organization. OT networks often have their own IT infrastructure and, depending on the size of the OT department, staff usually have varying networking skill sets. In most cases, OT and IT separation are due to security or compliance requirements, and the two environments are separated by a firewall.

Firewall separation is common between IT and OT networks. This firewall separation has several benefits in the OT/IT environment:

- Policy enforcement
- User control
- Threat containment
- Change control

Having this firewall separation ensures that events such as an accidental change on the IT or OT network should not impact other networks.

Figure 3.2: IT and OT networks separated by a firewall

The OT Network often uses the IT Network as a backbone to traverse to the Internet and access external resources. This connectivity requires a firewall between the IT Network and the outside world to ensure only authorized users and applications are allowed to access OT networks.

The Purdue Model, outlined in Figure 3.3, shows how these different zones are connected and separated using firewalls.

Figure 3.3: Purdue Model

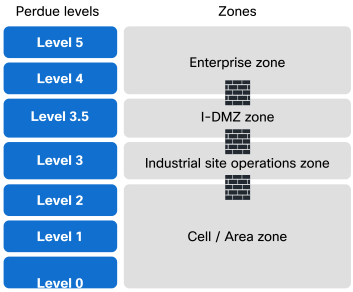

Traditional Deployment Environment and Environmental Conditions

OT networks can be indoors, outdoors, or both. When indoors, this is often in an uncarpeted space. When outdoors, they are either in an open area with IP67-rated conditions or housed in enclosures with environmental protection.

For both indoor and outdoor models, we assume that the location will not be temperature-controlled, and so, typically, we see the need for industrial-grade network devices. Depending on the specific industry, requirements and standards vary when it comes to temperature ranges, ability to withstand vibration, passive cooling, and so on. Cisco's Enterprise portfolio is often not up to specification for such harsh environments. To address these harsh environments, Cisco has the Cisco Industrial Ethernet (IE) family of switches covering 19" rack and DIN-rail form-factors.

It is also common in OT networks to have elements far more geographically distributed compared with IT networks, as there can be several miles between neighboring switches. Due to these geographically dispersed

installations, it becomes impractical or too costly for customers to use a star topology since they are not able to “home run” every access switch back to the Distribution Layer.

While OT networks often contain networks with Star Topology, we also see others such as:

Linear daisy chain: Here, switches are arranged linearly due to a lack of available fiber cores. To a network architect this may look like a bad design, with a compounded set of single-points-of-failure. This might be the intent as customers often do their best with their fiber resources and the physical world they are working in.

Figure 3.4: Linear daisy chain of Layer 2 Access switches

Ring: Here, unlike the linear daisy chain, the last Access switch effectively connects back to the Distribution switch, forming a ring. This provides an advantage of an alternate path in the event of connectivity loss or equipment failure. This topology requires loop avoidance technology, such as Resilient Ethernet Protocol (REP), which is critical for its function.

Figure 3.5: Ring of Layer 2 Access switches

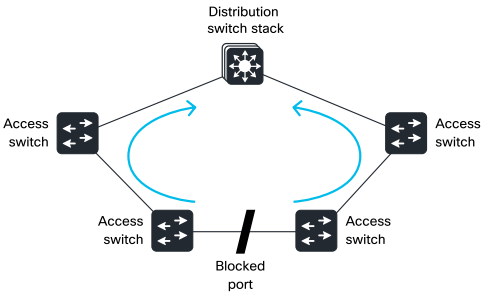

Resiliency and Uptime

In OT networking, rings may often cover large distances, but they can also be used within relatively short distances such as 100 feet. Especially within Industrial Automation and Utilities, there is often the requirement for a lossless Layer 2 mechanism. Protocols such as Parallel Redundancy Protocol (PRP) or High-availability Seamless Redundancy (HSR) are often used together. If a failure occurs in one of the paths, traffic continues to flow over the other path uninterrupted, and zero recovery time is required.

Uptime is incredibly important for manufacturing customers, utilities, and other critical systems. Precision Time Protocol (PTP) protocol helps with OT devices exchanging data in a Programmable Logic Controllers (PLC) control loop, providing a highly synchronized environment down to a sub-microsecond level. Production lines shut down if small amounts of data are lost or devices become out of sync. Avoiding these shutdowns is essential in keeping the national infrastructure running smoothly at the macro-economic level and, of course, for safety.

As a general rule for manufacturing, zero-loss and sync are required for motion applications. They are less stringent for applications such as I/O,

which could withstand up to 100ms of Layer 2 convergence times. In most manufacturing industries, this level of resilience and precision are not mandated. Here, OT networks are commonly deployed with REP rings, with variants to the Spanning Tree Protocol (STP) still providing resiliency to an extent.

Bringing all this back to line-of-business, you need to think about critical data and application flows between Cell/Area Zone and backend components in the Industrial Site Operations and Enterprise Zones. Maximum uptime is needed across all of these, and if line-of-business systems like SAP cannot communicate and function correctly, this likely means that manufacturing or some other industrial process needs to be halted. Compliance with the ISO 9000 family of standards is considered mandatory by most customers across the OT landscape, especially in Manufacturing. This includes tracking components from the supply chain to production and all the factors that go into a quality management system. Network architectures that address the necessary resiliency requirements and business uptime are critical to OT customers being successful.

Challenges

With the evolution of IT networks adapting newer technology, the OT network has the opportunity to utilize the flexibility provided by the enhancements in IT networking. Before delving into how the OT network can leverage the new functionality provided to it, let us look at the pain points from the perspective of both business and technical requirements.

Business Requirements

Critical Business Continuity, Resiliency, and High Availability

Regardless of size or industry, companies have clearly-defined processes and should be able to consistently meet the expectations of their customers. To certify the quality of a product, OT uses a quality management system to maintain and track the production of each component and part in the manufacturing process. As ISO 9001 is a globally recognized way to demonstrate the implementation of an optimal standard of quality in your company, the application suite which tracks this is not only mandatory but also a critical business application.

Business Agility

This is the requirement for increased efficiency during deployment and upgrade of the network and endpoints, such that the OT team can add new network hardware easily and RMA network hardware, when required, with minimal business disruption. More forward-looking industrial automation customers are looking into technology changes that allow them to virtualize Programmable Logical Controllers (PLC), which will gain them a better level of agility.

Network Performance and Agility

With any new architecture, overall network performance should be the same or better, but the speed of operations change should be enhanced. Performance in this context is measured as an increase in productivity and user/endpoint experience. Agile methodologies to change or drive network changes are crucial. In a Manufacturing plant, uptime is critical.

Organization Considerations

Providing an integration point between IT and OT teams and networks is critical, ideally with a shared view of policy but with a separation of duties. Network device software/firmware for OT devices is upgraded on a different cadence than the IT network. Often OT personnel do not have access to IT tools such as Identity Services Engine (ISE), and typically there is a separate Policy Services Node (PSN) from the ISE cluster within the OT space. This PSN communicates through the firewall to the IT network's data center. Thus, there is always a dependency on IT for network changes.

Purdue Model Alignment

To address and minimize cyber security exposures for industrial control and automation systems, many customers align to an OT architecture based on the Purdue Model.

This model articulates the layers and zones of separation for ISA99, specifically levels 0, 1, 2, 3, 3.5, 4, and 5.

Technical Requirements

Resiliency and Redundancy and High Availability

As OT networks have become extremely critical for the functioning of an organization, and in many cases an OT network is the core of a company’s standing, it is critical to provide resiliency all the way to full redundancy options for such networks. If the quality management system like SAP for a

manufacturing floor is down, no matter whether that is Tier-1, Tier-2, or Tier-3 manufacturing, then the manufacturing line is down. Thus, if the network is impacted, chances are extremely high that the plant will be down, costing millions of dollars per second.

Replication Requirements

A small failure in OT networks can lead to big losses. For example, stopping an auto manufacturing line can run into millions of dollars per second per stoppage. As a result, it can be mandatory to have resiliency in the form of packet replication in the OT networks to ensure that packet loss on failures is minimized or even zero.

Segmentation

Today's OT network occasionally provides limited segmentation capabilities. Mostly these networks are VLAN based. When deeper segmentation is needed, many organizations create physically separate OT networks.

There is a need to solve this by providing logical segmentation with the strength of physical separation but the simplicity of a single network.

Segmentation is sometimes used as a means to help address cyber security. If segmentation is instantiated by IP ACLs, it can become difficult to scale, troubleshoot, and maintain over time in all but the most simple deployments.

Security on Host Ports of OT Network

In most of the OT deployments, the host-facing ports are “open.” This is a security risk as anyone accessing the physical ports can get into the network and create havoc.

To minimize the damage due to this, many OT networks are using a firewall at the Edge, but it is not a watertight solution. For east-west traffic, there is no protection, leading to an enhanced security risk for OT assets. If customers want to address this and move from default-open to default-closed, it is an extremely large undertaking fraught with challenges.

Due to the open nature of ports on the OT network, there are a lot of cases where third-party unmanaged hubs/switches and other untrusted devices are connected to the OT network, thus increasing the attack surface. The open nature of ports can lead to various failures in OT networks from both the security and network resiliency perspectives. Open physical ports can access the devices at network layers, allowing an intruder to exploit any known vulnerabilities and/or launch attacks.

Visibility

One of the biggest problems in the OT network is the variety of endpoints on the manufacturing floor. There is limited visibility into the connected devices. Some devices are unmanaged, some remain untouched for years, and others can be geographically far away.

Silent Hosts

OT networks have different types of devices (hosts) connected to them, such as:

- Dynamically addressed hosts which get their IP addresses from a DHCP server
- Hosts with statically assigned IP
- Hosts that do not speak until some other host tries to reach them (Silent Hosts)
- Hosts that do not respond to pings and other network triggers and only respond to specific packets destined for them (Silent Hosts)
- Hosts that come online but then go to sleep if no one is speaking to them (Silent Hosts)
- Even hosts that are Layer 2-only, having no IP stack

Therefore, in terms of troubleshooting and general visibility, silent hosts can be problematic on an OT network.

Environmental Factors

Depending upon the location of OT equipment being connected to the network, the need for ruggedized network devices is increased compared to IT networks.

Large Layer 2 Networks

OT networks may require large Layer 2 networks for specific use cases or business reasons. The current network limits the size of the Layer 2 domains due to the issues with STP and other Layer 2 protocols. There is a requirement to solve this Layer 2 network size issue while allowing the flexibility for the OT device to be part of the same Layer 2 without the downside of Layer 2 protocols such as STP.

Ring Topologies

Ring topologies are prevalent in OT networks, presenting Layer 2 loop avoidance protocol challenges.

Layer 2 Network Address Translation

Machine builders usually reuse the same IP address scheme when delivering a process to a manufacturer, so Layer 2 Network Address Translation (L2NAT) is required when adding that process to the plant network.

Customer Solutions

The business and technical requirements of OT networks being agile, separating layers with the Purdue Model, moving easily towards closed mode authentication, etc., can be met by introducing the Cisco SD-Access Fabric solution to the OT network. The following sections will show how Cisco SD-Access solutions can meet these business and technical requirements.

This section will review solutions and how they address the previously stated requirements. Each solution will be based on real customer OT environments. In each use case, we will call out the requirements in the next sections of the chapter.

Based on how the OT network is deployed and how Cisco SD-Access Fabric can help meet the requirements, we propose the following four solution examples. These examples can be used in the OT space to enhance security, segmentation, availability, and increase OT network deployment agility. They are ordered from least to most in terms of segmentation capability, business, and technical resiliency:

- A shared IT-OT Cisco SD-Access Fabric, with one or more Virtual Networks created for OT, macro-segmentation, micro-segmentation, and a Peer (*Fusion*) firewall (for simple, non-regulated deployments).
- A dedicated OT Cisco SD-Access Fabric Site, in addition to the IT Cisco SD-Access Fabric Site, separated by a Peer (*Fusion*) firewall but with a shared Cisco DNA Center/Cisco ISE instance (for more complex, departmental separation, e.g., roadways, but still non-regulated deployments) (Separate OT and IT).

- Industrial demilitarized zone and Industrial security zones have their own Cisco SD-Access Fabric Site. Each of the OT sites is connected to the IT Fabric using a common Peer (*Fusion*) firewall device with a dedicated Cisco ISE Policy Service Node (PSN), enabling survivability zones. A Cisco DNA Center deployment is shared with site-based RBAC (two different OT Zones).
- Dedicated OT Cisco SD-Access Fabric Sites with dedicated SD-Access Transits (multiple, dedicated SD-Access Transit).

Shared IT-OT Cisco SD-Access Fabric

This is the most straightforward and simplistic design where the Enterprise Cisco SD-Access IT network has one or more dedicated Virtual Networks for OT networks. For this example, let us name this Virtual Network VN-OT. In this design, the firewall function between the IT and OT network is served by the Common Enterprise firewall, where the virtual network VN-OT is handed off to the Peer (*Fusion*) device as a firewall Zone for OT. The traffic between the VN-OT and corporate Virtual Networks is inspected by the Peer (*Fusion*) firewall device and based on the policy, the traffic may be:

- Dropped by the firewall
- Permitted by the firewall
- Selectively forwarded by the firewall. based on the L4-L7 inspection rules

Figure 3.6: Shared IT/OT Cisco SD-Access Fabric – Logical view

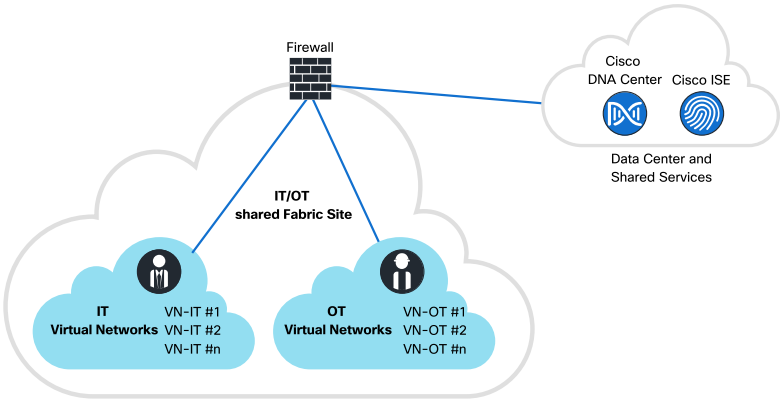

Figure 3.6 depicts the topology of Virtual Networks in the IT-OT shared Fabric Site

- OT Virtual Networks (VN-OT#1 to VN-OT#n) in IT/OT Shared Fabric
- IT Virtual Networks (VN-IT#1 To VN-IT#n) in IT/OT Shared Fabric

The IT/OT Shared Fabric Site hands off the IT and OT Virtual Networks to the Peer Node firewalls. Firewalls based on the rules allow/deny communication across these Virtual Networks. Additionally, the firewall also provides shared services access to the data center and other services, including DHCP, DNS, Internet, etc.

Now that we understand the logical topology of this design let us look at the physical topology. Figure 3.7 depicts the physical topology of the design.

Figure 3.7: Shared IT/OT Cisco SD-Access Fabric – Physical view plus logical Virtual Network mapping

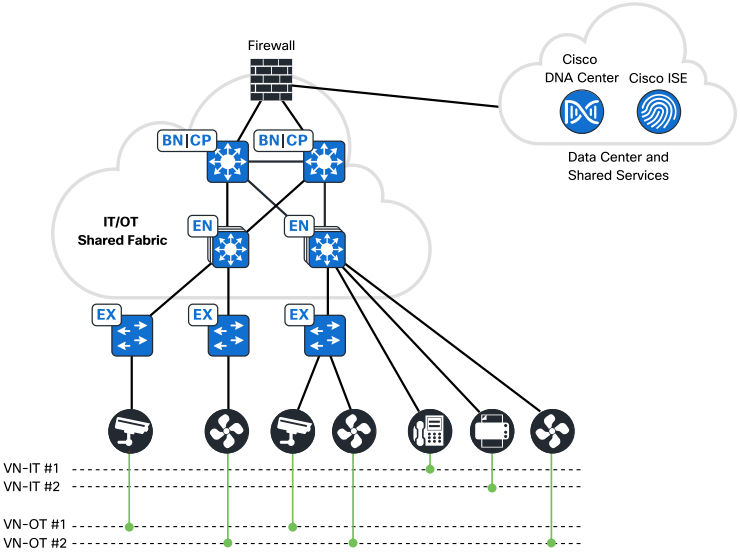

You can see that both IT and OT devices can connect to Access Ports of either Extended Nodes or Edge Nodes, which are mapped to respective Virtual Networks. Device/host onboarding can be achieved in three ways:

- Ports on Edge Nodes and Extended Nodes can be assigned to a particular VLAN of a Virtual Network using the Host Onboarding workflow in Cisco DNA Center (no-auth case).
- Ports on Edge Nodes and Extended Nodes can be set for dynamic authentication, where the endpoint can be admitted to the relevant Virtual Network/VLAN based on defined authentication methods (Open Auth and Closed Auth).

- Ports on Edge Nodes and Extended Nodes can be assigned to a particular Virtual Network via SD-Access REST APIs. This can be especially powerful when driven from a Line-of-Business (LoB) process.

Given OT devices are likely to be in non-carpeted/rugged locations, you would use Cisco Industrial Ethernet (IE) switches as Extended Nodes. Please refer to product data sheets and to the Cisco SD-Access Compatibility Matrix on Cisco's website to see relevant Cisco IE switches and the Cisco SD-Access Fabric roles they can perform.

As a reminder, Edge Nodes are connected to the rest of the Cisco SD-Access Fabric via a routed Layer 3 Underlay network. Extended Nodes are connected via a Layer 2 802.1Q trunk, which is also part of a PAgP port channel (even if the group only has one member/uplink).

The onboarding of Extended Node devices is OT-friendly. The switches can be unboxed, patched in, powered up, and auto-configured only if the device has the No Authentication Template enabled for the Fabric Site. If you enable any other template, create a port channel manually, using the Cisco DNA Center GUI of the Extended Node or the Policy Extended Node to which the new Node is connected for onboarding.

Extended Nodes must connect back to a single Edge Node. For better redundancy, Extended Node should be connected to an Edge Node stack (either physical StackWise on certain Cisco Catalyst 9000 family members or StackWise Virtual capable switches). This helps maintain the upstream connectivity even during the failure of the upstream switch or link.

Figure 3.8: Extended Node via Cisco Catalyst 9000 StackWise

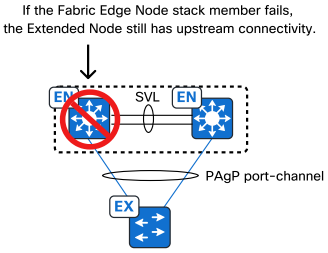

Cisco SD-Access is an effective solution in this case because an entirely parallel OT network is not required. IT and OT share a physical network but are logically separated and segmented.

It is important to understand that while this solves some problems, from an OT perspective, additional caveats may require you to take a different approach for some services. When considering shared services, always consider function over form. In OT, functions that rely on shared services that could take down the environment, whether in this area or across the WAN, should be reasons to pause and consider alternative approaches.

Dedicated OT Cisco SD-Access Fabric Site

In this design, the OT network has a dedicated OT Fabric Site in addition to the Enterprise Cisco SD-Access IT network. The OT Fabric Site is separated from the Enterprise Cisco SD-Access Fabric Site with a firewall. The firewall has at least two zones, one connecting to the OT Fabric Site and another one connecting to the Enterprise Cisco SD-Access Fabric Site. The firewall can be a Layer 2 (transparent) firewall or a Layer 3 (routed) firewall. If needed, the firewall could NAT the traffic coming from the OT site towards the Enterprise Cisco SD-Access site.

Having separate Fabric Sites for IT and OT might be an organizational decision driven due to security policies or departmental boundaries. This design has a major advantage for both IT and OT networks, as any changes they implement on one side will not impact the other. This is particularly useful when the OT network supports something such as a Smart City or Connected Roadway. Examining the case of a roadway customer, they would want to keep network traffic separate between device types such as:

- Virtual Network for road toll equipment
- Virtual Network for cameras
- Virtual Network for intersections
    - SGT for Traffic Signal Controllers
    - SGT for LiDAR
    - SGT for Edge Compute
    - SGT for V2X radios
- Virtual Network for weather stations

At an intersection, we can allow OT devices to communicate at Layer 2, with very low latency, without hairpinning the traffic via a Peer (*Fusion*) router/firewall. Using SGTs, we can still enforce via micro-segmentation since Cisco SD-Access gives us a highly automated way to deploy a micro-segmentation policy.

Cisco Cyber Vision is a useful tool in OT networks. Having the capability of running on the Extended Nodes and Policy Extended Nodes, Cyber Vision helps to identify and classify OT devices in conjunction with Cisco ISE. Cisco Cyber Vision is particularly important because Cisco Endpoint Analytics often cannot be leveraged to recognize endpoints connected to Cisco IE switches.

Another advantage of this design is that the Enterprise Cisco SD-Access site and dedicated OT Cisco SD-Access site share common components such as Cisco DNA Center and ISE. By using site-based RBAC, this design provides a true separation from the logical perspective but does share the

physical Cisco DNA Center appliance/cluster between IT and OT Cisco SD-Access sites.

Figure 3.9: Dedicated OT Cisco SD-Access Fabric

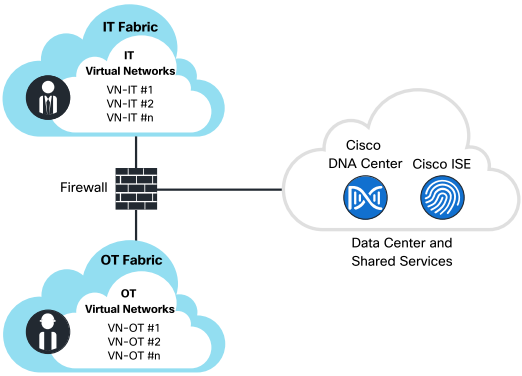

The logical topology depicted in Figure 3.10 applies to Virtual Networks in different Fabric Sites

- OT Virtual Networks (Virtual Network-OT#1 to Virtual Network-OT#n) in OT Fabric Site
- IT Virtual Networks (Virtual Network-IT#1 To Virtual Network-IT#n) in the IT Fabric Site

Each site hands off its respective Virtual Networks to the firewalls. The firewall will allow/deny communication across these Virtual Networks based on the firewall rules. Additionally, the firewall provides shared services access to the data center and other services, including DHCP, DNS, Internet, etc.

Now that we understand this design's logical topology let us look at the physical topology of the design, shown in Figure 3.10.

Figure 3.10: Dedicated OT Cisco SD-Access Fabric – Physical view plus logical Virtual Network mapping

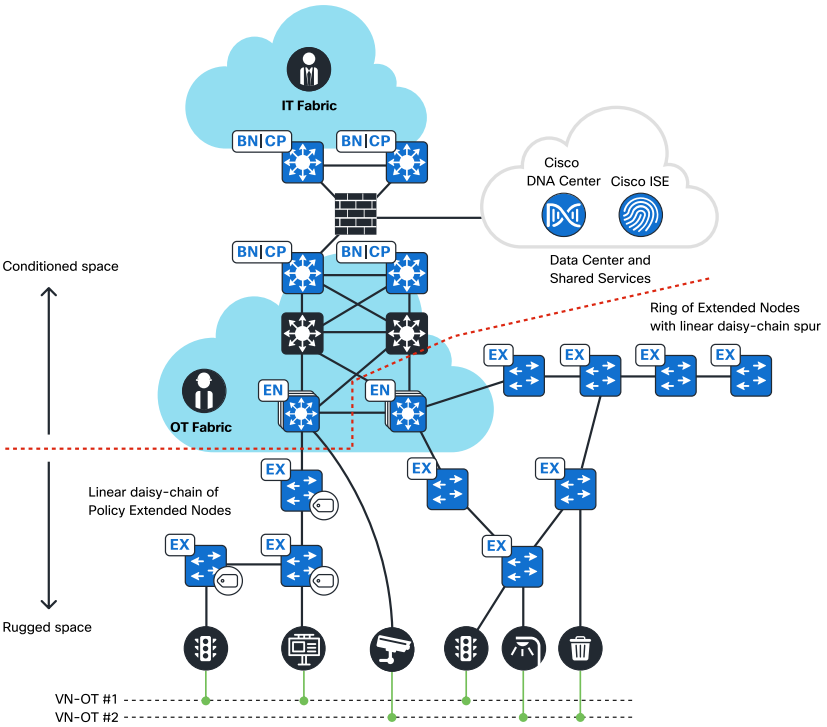

As shown in Figure 3.10:

- IT devices connect to access ports of the Edge Nodes and Extended Nodes, which are part of the IT Cisco SD-Access Fabric Site (IT Fabric) – not shown on the diagram.
- OT devices connect to access ports on both the Edge Nodes and Extended Nodes, which are part of the OT Cisco SD-Access Fabric Site (OT Fabric).

This design provides full physical separation between the IT and OT networks while maintaining common segmentation constructs such as Virtual Networks and Security Group Policy constructs across IT and OT Fabric Sites.

A ring topology for IE switches can be automated in Cisco DNA Center, converting an RSTP ring to Resilient Ethernet Protocol (REP). This typically results in an approximately 50% improvement in Layer 2 recovery time in the event of a failure. Also, you can see that it is possible to spur off a ring with a linear daisy chain.

For the choice of appropriate enterprise switches, please refer to product data sheets and to the Cisco SD-Access Compatibility Matrix on the Cisco website to see relevant Cisco Enterprise switches and the Cisco SD-Access Fabric roles they can perform.

For OT devices that are likely to be in non-carpeted/rugged locations, the IE switch platforms are the same as in example #1.

It is important to note that the Border Node and Control Plane Node in the OT Fabric are “not” in non-carpeted/rugged locations. They are part of the Conditioned space, the same as the Enterprise IT switches and routers supported by Border Nodes and Control Plane Nodes.

Industrial DMZ and Industrial Security Zone with Dedicated Fabric Sites

This design separates the OT Fabric into two Fabric Sites, the first for the Industrial Demilitarized Zone (DMZ) and the second for the Industrial Security Zone. The OT Fabric Sites are connected to the Peer firewall. The Peer firewall has multiple firewall zones, one for each OT Fabric Site and another connecting to the Enterprise SD-Access Fabric. The OT Fabric Sites bring all the traffic towards the Peer (*Fusion*) firewall device. The Peer firewall can be a Layer 2 transparent firewall or a Layer 3 routed firewall.

This model is driven by maintaining separation between the Purdue levels while saving costs by using a common multi-zone firewall between the different layers versus using dedicated firewalls. This architecture also helps to maintain separate micro-segmentation and macro-segmentation within the layers while allowing firewalls to do a stateful inspection for any cross-layer traffic. This design has the benefit of introducing segmentation on demand as well as the ability to deploy dynamic authentication on a selective set of OT ports connecting to the hosts. In OT networks, IoT devices such as robots, measuring instruments, or certain PLCs are often used. When such endpoints are silent, we can leverage the Layer 2 Flooding functionality in Cisco SD-Access to provide connectivity in the network. This feature can also cover Layer 2-only protocols.

In this model, a common Cisco DNA Center is used for the OT and IT SD-Access Fabrics but dedicated ISE PSNs are used for OT and IT Fabric Sites, respectively.

Figure 3.11: Industrial DMZ and Industrial Security Zone with dedicated Fabric Site

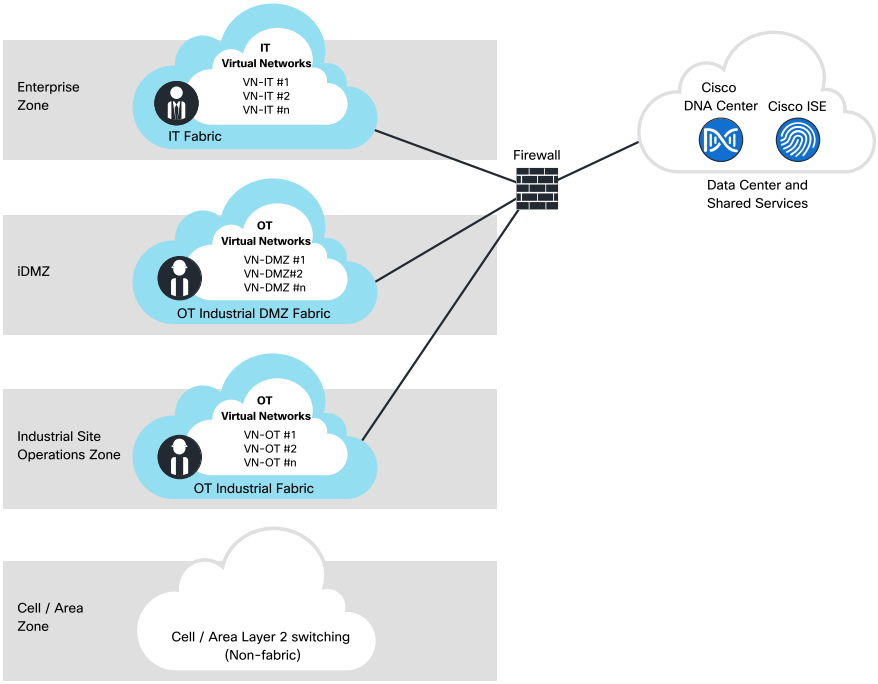

Figure 3.11 depicts the logical topology of the following Virtual Networks in different Fabric Sites

- OT Virtual Networks (VN-OT#1 to VN-OT#n) in OT Industrial DMZ Fabric Site
- OT Virtual Networks (VN-DMZ#1 to VN-DMZ#1n) in OT Industrial Fabric Site
- IT Virtual Networks (VN-IT#1 To VN-IT#n) in the IT Fabric Site

Each of these sites hands off their respective Virtual Networks to the firewalls. Based on the firewall rules, the firewall allows/denies

communication across these Virtual Networks. The firewall also provides shared services access to the data center and other services, including DHCP, DNS, Internet, etc.

The Peer firewall could NAT the traffic from the OT sites towards the Enterprise SD-Access site, as required, but Layer 2 NAT is often done at the edge of the Cell/Area Zone. You should also think carefully when deploying firewalls in and between the various zones. Enabling features such as inspection and others may have a performance/latency impact, so you should consider how enabling these features could impact an Industrial Automation network.

One of the key requirements of the Industrial DMZ (IDMZ) is that no traffic flows directly between Enterprise and Industrial zones. Firewall rules are set up to ensure Enterprise Zone traffic can reach the IDMZ, and the IDMZ can be used to access the Industrial zone, and vice-versa.

The IDMZ allows specific and limited access into the Industrial Site Operations Zone and the Cell/Area Zone. Access is typically extremely limited, locked-down, and audited, with SGTs applied and enforced by Cisco SD-Access becoming an effective method to support these security requirements.

Now that we understand this design's logical topology let us look at the physical topology of the design, depicted in Figure 3.1.

Figure: 3.12: Dedicated OT Cisco SD-Access Fabric Site with dedicated Industrial DMZ Fabric Site – Physical view

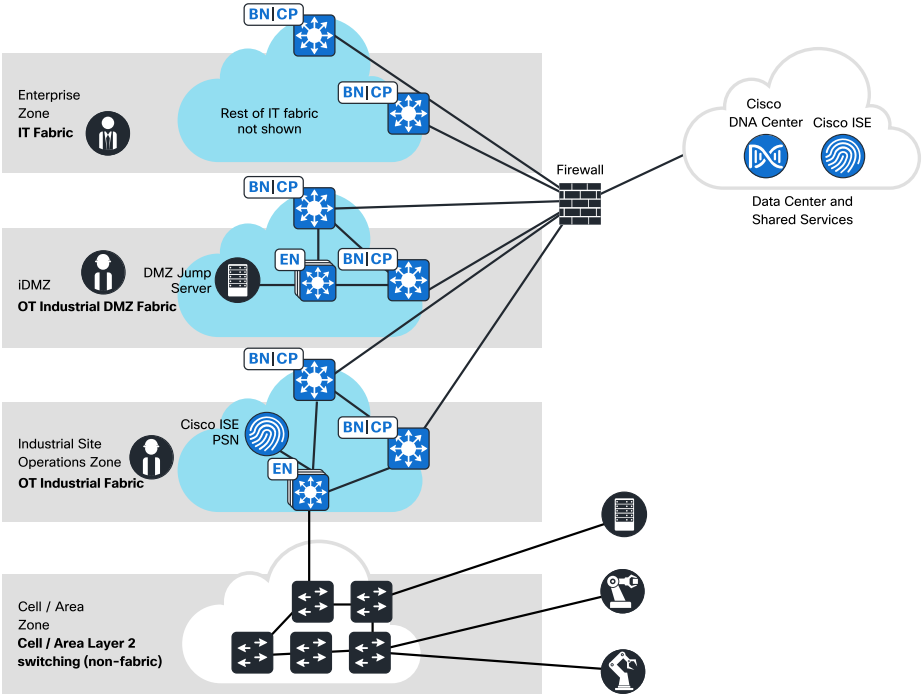

You might have already guessed because of the zones referenced in Figure 3.12, but this solution applies to Manufacturing and Industrial Automation, with Fabric separation aligned to the Purdue Model. For the control loops themselves (levels 0 to 2), we are keeping this outside of the Fabric as an External Layer 2 Switching Domain. The Edge Node, between the Cell/Area Zone and the Industrial Site Operations Zone, can sometimes take the place of an Industrial Firewall, with enforcement utilizing both macro-segmentation and micro-segmentation – this can simplify security at the Cell/Area Zone level.

SD-Access can be looked at as a toolset, and while there is no one way to construct a segmentation scheme for Industrial Automation, we see design patterns with customers:

- Controlling segmentation within a Cell/Area Zone
- Controlling segmentation across more than one Cell/Area Zone

Segmentation on demand – where OTs are empowered to create, update and destroy virtual networks – is a real differentiator for SD-Access versus traditional LAN options. All of this control is available via REST APIs, meaning it can be fully integrated into the OT line-of-business processes and applications, truly executing on business intent.

Dedicated OT Cisco SD-Access Fabric Sites with Dedicated SD-Access Transit

This design brings constructs of separate SD-Access Transits for the OT SD-Access sites of IDMZ and Industrial security zones, with the Peer firewall having at least three zones. The first zone connects to the IDMZ Fabric, the second zone connects to the Industrial security zone Fabric, and the third zone connects to the Enterprise SD-Access Fabric. The Peer firewall directs traffic between the zones and also filters the traffic based on the firewall rules. The Peer firewall can be a Layer 2 (transparent) firewall or a Layer 3 (routed) firewall. The Peer firewall could NAT the traffic between different Cisco SD-Access networks.

Figure 3.13: Dedicated OT Cisco SD-Access Fabric Sites with dedicated SD-Access Transit

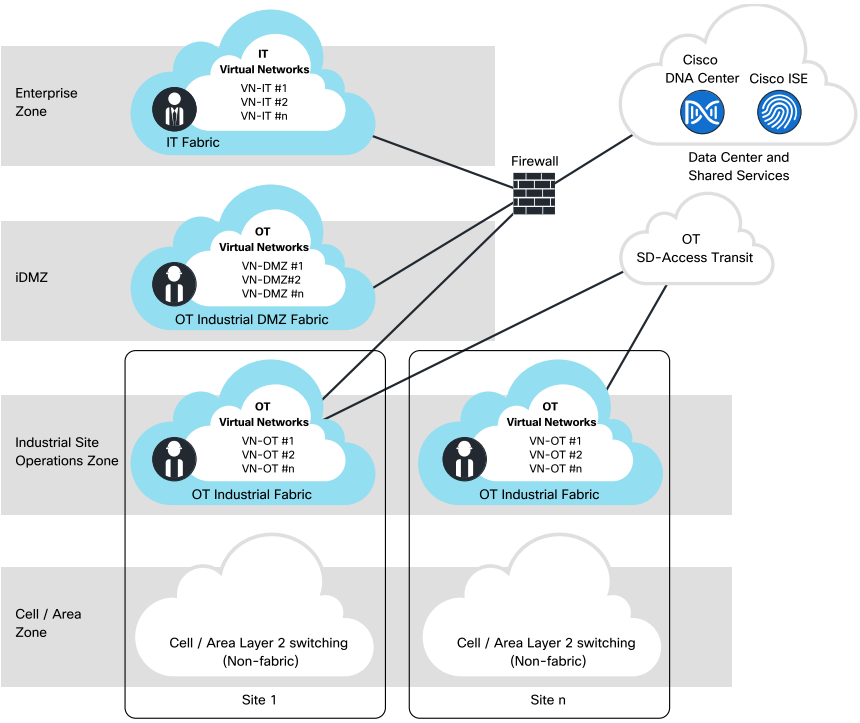

Figure 3.13 logical topology depicts the following Virtual Networks in different Fabric Sites:

- OT Virtual Networks (VN-OT#1 to VN-OT#n) in OT Industrial Fabric Sites. To keep the consistency of the Virtual Networks in the Industrial Site Operations zone, the OT Virtual Networks can be consistent across the multiple Cisco SD-Access OT Industrial Fabric Sites.
- OT Virtual Networks (VN-DMZ#1 to VN-DMZ#n) in OT Industrial DMZ Fabric Site
- IT Virtual Networks (VN-IT#1 To VN-IT#n) in the IT Fabric Site

Each site hands off its respective Virtual Networks to the firewalls. Based on the firewall rules, the firewall allows/denies communication across these Virtual Networks. Additionally, the firewall also provides shared services access to the data center and other services, including DHCP, DNS, Internet, etc.

Now that we understand this design's logical topology, let us look at the physical topology.

Figure 3.14 depicts the physical topology of this design.

Figure 3.14: Dedicated OT Cisco SD-Access Fabric Sites with dedicated SD-Access Transit – Physical view

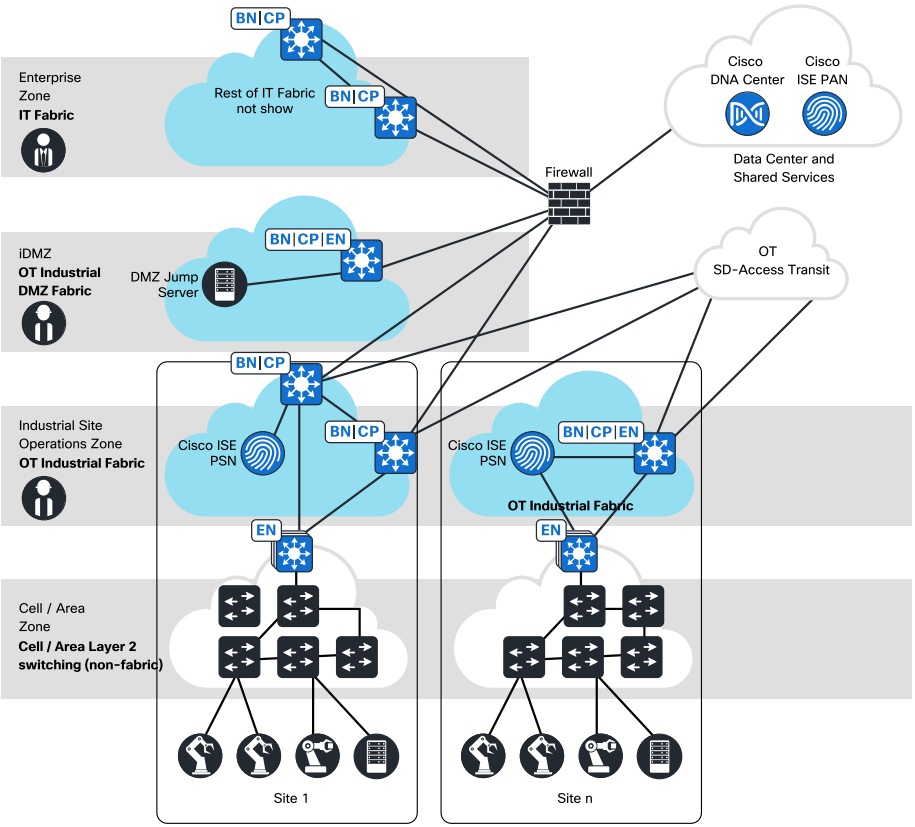

Cisco SD-Access Fabric in a Box can be a useful building block where a smaller scale is required, and resiliency derived from stacking is sufficient. Here we have used a Fabric in a Box in the IDMZ, as this typically houses a modest number of servers and services. Similarly, for Site N, we can assume this is a smaller site, and so an Edge Node connected to Fabric in a Box is appropriate. For Site 1, we have deployed with Colocated Fabric Border Node and Control Plane Node, with separate Edge Nodes.

Site 1 through Site N could be different sections of the same physical factory or different factories spread across geography. Such is the flexibility of using multiple Fabric Sites.

We also note that separating into multiple Fabric Sites affords more isolation and better site survivability. For example, if the Fabric at Site 1 is somehow impacted through operator error misconfiguration, Site N is unaffected.

As more forward-looking industrial automation customers look into technology changes that allow them to virtualize PLCs, Cisco SD-Access can provide a framework to support this in ways traditional and static networking does not. SGT-based policy across industrial zones and sites, SDN to reconfigure the network allowing the customer to be agile and responsive to changes in demand – these things can help lower a customer's operational costs and increase productivity.

Deployment Options

Fabric Site Sizes – Design Strategy

In a Cisco SD-Access deployment design, you can create fewer, larger Fabric Sites rather than multiple, smaller Fabric Sites. The design strategy is to have a larger Fabric Site size while minimizing total site count, which can help reduce management overhead. Larger Fabric Sites can also help in deploying common policy across the Fabric. Though having a Large Site has multiple benefits, your business requirements may necessitate having a Small/Medium Site. The multi-dimensional factors of survivability, high availability, endpoint count, services, and geography are all factors that may drive the need for multiple, smaller Fabric Sites instead of a single Large Site. The following reference models can be used to help you design Fabric Sites of different sizes.

Fabric Site Reference Models

In deployments with physical locations, customers use different templates for each site type, such as a large branch, a regional hub, headquarters, or a small remote office. The underlying design challenge is to look at the existing network, wiring, and deployment and propose a method to layer SD-Access Fabric Sites in these areas. This process can be simplified and streamlined by templating designs into reference models. The templates drive understanding of common site designs by offering reference categories based on the multi-dimensional design elements like endpoint count, number of Fabric Devices, and number of users to provide guidelines for similar site size designs.

The numbers are used as guidelines only and do not necessarily match maximum specific scale and performance limits for devices within a reference design.

Each Fabric Site includes a set of Wireless LAN Controllers, Border Nodes, Control Plane Nodes, and Edge Nodes, sized appropriately from the listed categories. The ISE PSNs are also distributed across the sites to meet survivability requirements.

The models which are deployed widely in the OT environment are:

1. Very Small Site (Fabric in a Box)
2. Small Site
3. Medium Site
4. Large Site
5. Extra-Large Site

Very Small Site (Fabric in a Box Site)

You can use Fabric in a Box to cover a single Fabric Site, with resilience supported by switch stacking or StackWise Virtual to support up to 200 endpoints and 40 APs.

For Fabric in a Box deployments, SD-Access Embedded Wireless provides site-local WLC functionality. The OT switches can be connected in a Fabric in a Box site as Extended Nodes or a ring connected to an SVL or stack switch. The site may contain an ISE PSN depending on the resiliency requirements of the OT cell.

Small Site

The Small Site Reference Model covers a single office or building with single wiring closets, usually up to 4,000 and up to 100 APs. The Border Node function is colocated with the Control Plane Node function on one or two devices and usually uses embedded wireless with the option of hardware WLCs.

Medium Site

The Medium Site Reference Model usually covers an OT network with multiple wiring closets with a physical topology consisting of a Two-Tier Collapsed Core/Distribution with an Access Layer.

The Medium Site is designed to support less than 25,000 endpoints and less than 2,000 APs. The Border Node function is colocated with the Control Plane Node function on one or two devices or a highly-resilient single device, and a separate wireless controller is ideally deployed in an HA configuration.

Large OT Site

Figure 3.15: Large OT Site design

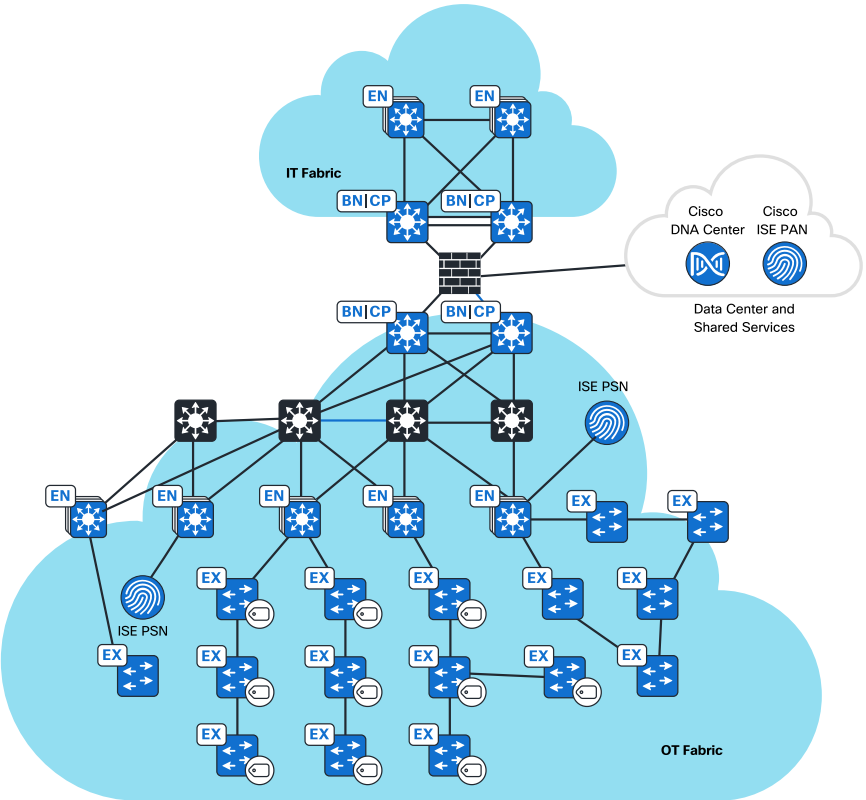

The Large Site Reference Model covers a factory floor with a large Industrial security zones or stretched across a large manufacturing floors or buildings. The physical network could compromise multiple tiers, switch rings, or daisy chains of switches and is designed to support less than 100,000 endpoints. This network is large enough to require dedicated services, including dedicated ISE PSNs with an ISE cluster dedicated to this large OT network.

Multiple Border Nodes are distributed from the Control Plane Node function on redundant devices from the OT network towards the Peer (*Fusion*) firewalls or Enterprise IT SD-Access network; a separate wireless controller has an HA configuration.

Extra-Large Site

The Extra-Large Site Reference Model covers an OT network with multiple wiring closets or multiple facilities stretched across a large Campus. The physical network is a Three-Tier network with Core, Distribution, and Access Layers and may sometimes have a Super Core in a Fourth-Tier. An Extra-Large network requires dedicated services exit points such as a dedicated data center, shared services block, and Internet services.

The Extra-Large Site supports up to 200,000 endpoints and 10,000 APs Multiple Border Nodes are distributed from the Control Plane Node function on redundant devices, and a separate wireless controller in an HA configuration.

Summary

We learned from this chapter that Enterprise OT Networks are expansive and complex. We saw how Cisco SD-Access can help solve these complex challenges and take an organization's OT network on the digital transformation journey. We learned how SD-Access enabled OT Enterprise networks to transform static OT networks without security and segmentation into a highly agile OT network with flexible designs to introduce security for the OT networks. We also went from a shared IT-OT segmented network to multiple types of dedicated IT-OT networks following the Purdue Model of computer integrated manufacturing, providing separation of layers and authenticated access ports to further enhance security. With the introduction of SD-Access to an OT network, the ability to do segmentation on demand, and its ability to be integrated into the OT line-of-business processes and applications using REST APIs, is a game-changing innovation that can be utilized by OT networks.

To round out the conversation and to review, the methods above are various solutions that help solve problems, but it is important to always remember that uptime is critical for what is a lossless environment in the case of most Industrial Automation organizations. Unlike most other verticals, outages here can have catastrophic consequences. Reliance on line-of-business systems like SAP for compliance with the ISO 9000 family of standards is mandatory, and thus any impact on the tracking of manufactured components within the production to supply chain is heavy. During this chapter, we shared various methods and use cases and how they play in the manufacturing space.

The main takeaway from the chapter is that Cisco SD-Access helps implement and integrate OT networks with IT networks. Additionally, Cisco SD-Access can be used to build dedicated OT networks. It is important to remember that not all approaches are suitable for all OT verticals. For

example, Utility Substation and Rail Signaling are not addressed by any of the four examples, and for some elements of Manufacturing/Industrial Automation, it is not suitable to bring Cisco SD-Access right down to the plant floor as we showed in examples 3 and 4. Finally, we discussed Small and Large deployment models for the OT networks with Multi-Tier, ring, and daisy chain physical topologies. Lastly, remember that OT networks are many and varied, and the four different solution examples are best applied to different industries.

Cisco SD-Access for Healthcare

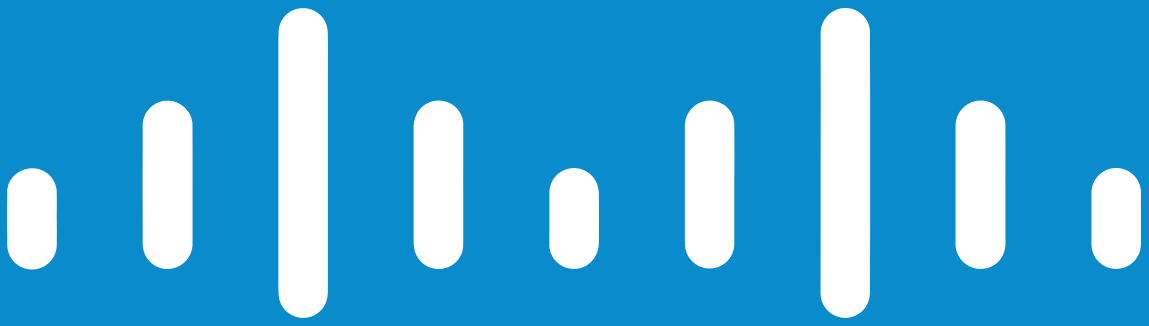

Introduction

Chapter Overview

Significant changes have been taking place in the healthcare industry, such as exponential growth in telehealth and virtual care, sudden increases in remote workforces, increased security concerns, fast-evolving primary care models, shifts in care delivery sites, and the prioritization of worker safety and wellness.

This chapter provides design guidance for a typical healthcare deployment using Cisco DNA Center and Cisco SD-Access. The following sections describe the key considerations for a large, evolving healthcare network that needs to meet today's healthcare requirements.

Traditional Network Architecture for Healthcare

Healthcare networks are often comprised of multiple sites of varying sizes and needs. They can vary from multi-hundred-bed facilities with critical and emergency care units down to small or mobile clinics servicing a single specialty or general care.

Regardless of their size, certain requirements such as security, mobility, resiliency, and device and patient tracking are crucial. These services must be available to users and devices regardless of whether they connect to the network via wired or wireless media.

Challenges

Given the critical nature of the services delivered by healthcare providers, the network architecture challenges they face are often quite unique compared to customers in other verticals. The following sections outline several of the most critical capabilities required by healthcare providers when deploying a network.

Security Compliance Mandates

Healthcare systems need to protect highly sensitive medical records and financial information of patients and are strictly bound by government regulations (for example, HIPAA in the United States and GDPR in the European Union). Hospitals, clinics, doctor's offices, and other healthcare-related entities must provide regulation-compliant wired and wireless networks. These networks must provide complete and constant visibility for network management and monitoring. Sensitive data and medical devices such as Electronic Medical Records (EMR) servers and vital sign monitors must be protected so that malicious devices cannot compromise the network.

The proliferation of the Internet of Things (IoT) devices and Bring Your Own Device (BYOD) policies is adding to the need for segmentation between different groups and also within the groups between different types of users and devices. This means that macro-segmentation and micro-segmentation capabilities must be top priorities for any healthcare network architecture.

Finally, given that staff, medical devices, and visitor traffic all share the same network infrastructure, every group must be isolated from one another and restricted to only the resources they are permitted to access.

Mobility of Staff and Devices

Within the healthcare environment, staff and equipment must often move between rooms, floors, and buildings in a rapid amount of time. It is estimated that, on average, a nurse walks anywhere between 4-5 hours in a shift. It is also quite common for devices to have static IP addresses resulting in subnets must be stretched across the facility to allow the equipment to move anywhere it is needed.

To increase operational efficiency, device and staff mobility should be completely dynamic without needing network staff to constantly reconfigure networking infrastructure. Mobility must be available for devices and users whether they attach to the network via wired or wireless media.

Increased Reliance on IoT Endpoints

Healthcare networks are populated by a wide variety of devices in multiple locations. This makes locating and identifying all of the devices in a network time-consuming and tedious. In some cases, these devices will sit silently on the network and will not send any traffic unless they are first contacted by another device. In other cases, the network staff may not even know all of the types of devices in their network which makes onboarding an extremely difficult task.

Modern security threats seek vulnerable points of entry, such as undetected or silent hosts, to exploit the network's valuable resources. These devices must be identified and tracked to meet the secure framework of the network.

Resiliency and High Availability

Healthcare providers demand levels of high availability and resiliency beyond most other enterprises. The network is a key component in delivering access to key information and services, such as patient record availability, and it must be part of a wider resiliency and high availability strategy across IT.

Any network solution must provide user and device access to critical resources in the event of loss of communication to the authentication infrastructure. To minimize service impact, strict network-level and application-level resiliency must exist.

Patient Safety and Asset Tracking

Given the steadily increasing reliance on IoT endpoints in healthcare providers, asset management is becoming more difficult as increasingly more diverse devices become mobile and are shared across locations. It is estimated that 20-30% of time spent by nurses is searching for equipment and people. Additionally, there are times when patient tracking is crucial to ensure that a patient has not left a specific area – for example, newborns from the nursery. The ability to track both patients and devices in a single system is a critical requirement for every modern healthcare provider.

IT Network Operational Challenges

The combination of ever-increasing devices and security requirements on the network is putting immense operational pressure on healthcare IT staff. These are manifested in the complexity of traditional IP-Based rule management and the struggle of managing those policies across large enterprise footprints. Timely troubleshooting is becoming more critical and difficult, especially in smaller and/or mobile clinics without onsite staff.

Faster Response Times to Attend Patients

Hospitals need to provide patients with the ability to connect with nurses and doctors and for immediate needs, including emergencies. Therefore, healthcare networks need to facilitate sending alarms and alerts to people in their proximity.

Solutions

Cisco SD-Access can address and resolve the challenges outlined in the previous section. SD-Access can be implemented across all areas of healthcare, providing consistency across the enterprise.

Security and Compliance

Macro-segmentation

One of the preferred segmentation methods for healthcare customers is to isolate the patient user endpoints (such as PCs, Tablets, IP Phones, etc.) from the doctor, nurse, and enterprise IoT endpoints by placing them in different virtual routing and forwarding instances (VRFs). SD-Access offers the flexibility for macro-segmentation of endpoints in different VRFs, all of which can be provisioned in the network from Cisco DNA Center.

Micro-segmentation

Within the scope of a single VRF, customers tend to have further segmentation needs for use cases such as:

- Placing medical devices and video surveillance of patient floors in different groups
- Placing visiting doctors and resident doctors in different groups

For such requirements, in the traditional network architecture, the only means to segment was by placing groups in different subnets enforced by IP ACLs. In Cisco SD-Access, in addition to providing the flexibility of using different subnets, we provide the flexibility of micro-segmentation, i.e., using the same subnet in a more user and endpoint-centric approach.

Referring to the visiting and resident doctor example, each group can still be placed in the same subnet. However, by leveraging dynamic authorization, they can be assigned different Security Group Tags (SGTs) by ISE based on their authentication credentials. These SGT-based rules can then be enforced by Security-Group Access Control Lists (SGACLs).

MACsec

Media Access Control security (MACsec) is a standardized solution based on IEEE 802.1AE and MKA based on 802.1X that can add another layer of security to a Cisco SD-Access Fabric-based architecture. Customers can choose to deploy MACsec in 3 different places within the Fabric:

- Endpoint to Edge Node
- In the Underlay connecting different Fabric Devices
- A manual IP handoff from the Border Nodes connected to WAN/Peer devices

Cisco SD-Access offers the flexibility to enable MACsec in any of the options above by using the Template Editor application within Cisco DNA Center to provision these configurations onto the network devices.

Separation of Guest Endpoints

Cisco DNA Center and SD-Access provide various levels of flexibility in the separation of patient and visitor endpoints connecting to the network due to security reasons by providing:

- Separate Border Nodes and Control Plane Nodes located in the DMZ by using the Multisite Remote Border functionality
- Separate ISE Policy Service Nodes (PSN) to manage Guest users
- A Virtual Network to isolate Guest user traffic, segregating them from other VRFs serving healthcare providers and other critical IoT endpoints in the healthcare network.

Staff and Device Mobility

Anycast Gateway

The Cisco SD-Access Fabric architecture is built on the prevalence of the same Anycast Gateway subnet on every Edge Node within a Fabric Site. This facilitates seamless roaming for wired and wireless endpoints. Clients roaming between APs connected to different Edge Nodes perform a Layer 2 roam instead of a Layer 3 roam. This is achieved by having wireless data plane traffic distributed in the Fabric Overlay instead of encapsulating all traffic to be dispatched to a centralized wireless controller.

Figure 4.1 illustrates how a simple Cisco SD-Access Fabric can be set up within a hospital building. When wireless clients roam from a Fabric AP connected to one Edge Node to a Fabric AP connected to a different Edge Node, they will retain their IP address and do a seamless Layer 2 roam thanks to the presence of the Anycast Gateway.

Figure 4.1: Cisco SD-Access Fabric wireless architecture

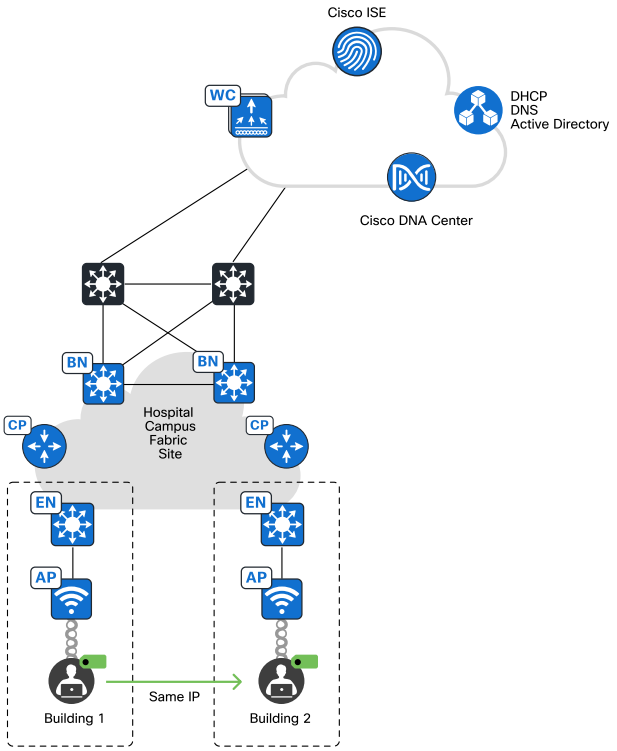

Multisite Remote Border

For healthcare customers who require roaming for their providers between different buildings, Cisco SD-Access provides the flexibility to anchor such users into a separate Virtual Network that can be centrally managed by a common Border Node and Control Plane Node. This feature is referred to as Multisite Remote Border, and an example is shown in Figure 4.2.

Figure 4.2: A depiction of seamless roaming using Multisite Remote Border

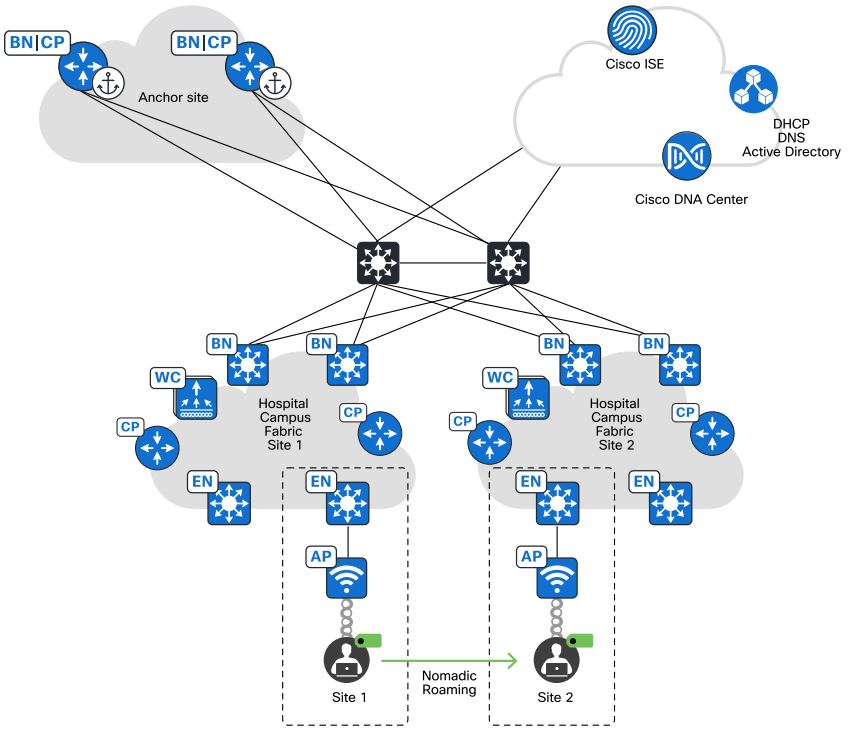

The Multisite Remote Border feature enables the configuration of a single IP Address Pool across multiple Fabric Sites, as shown in Figure 4.2. It allows the same IP Address Pool to exist at multiple Fabric Sites by anchoring the IP Address Pool to specific Border Node(s) and Control Plane Node(s). Preserving the subnet across multiple sites provides a scalable strategy for conserving address space.

As we can see here, the two hospital buildings are configured as independent Fabric Sites. The Virtual Network that contains the endpoints, is

anchored to a common Multisite Remote Border Node at the Anchor Site. This same architecture supports medical devices, such as infusion pumps or heart monitors, which need to be moved wirelessly between buildings as patients are moved.

Having a common Border and Control Plane Node ensures that endpoints are anchored to the common site as they roam between the buildings, across Fabric Sites. The WLC that is local to each site updates the Anchored site's Control Plane regarding the location of these wireless endpoints.

Note The roaming in this scenario is not seamless.

Automated Mobility Groups

Another means of deployment for healthcare customers is the ability to provide seamless roaming when large buildings are managed by multiple WLCs within a given Fabric Site. This is common when large healthcare enterprises have different floors in a building exclusively managed by a specific WLC. Cisco DNA Center supports the automated deployment of Wireless Mobility Groups when customers configure multiple WLCs in a given Fabric Site.

Figure 4.3 shows how such a hospital building looks and how a Cisco SD-Access architecture can be carved for such buildings:

Figure 4.3: Multiple WLCs managing different floors in a single Fabric Site

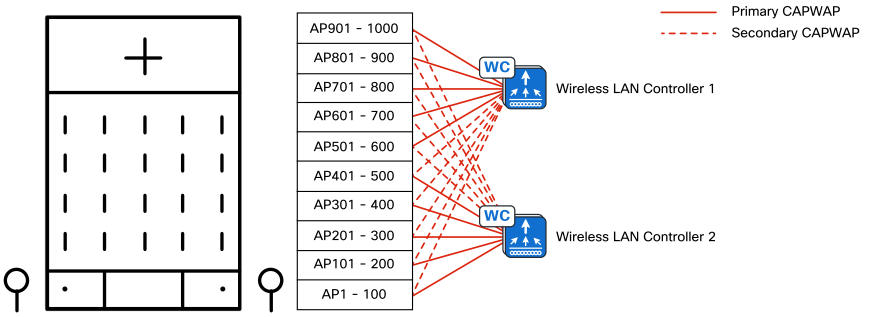

Figure 4.4: Large Hospital Fabric Site

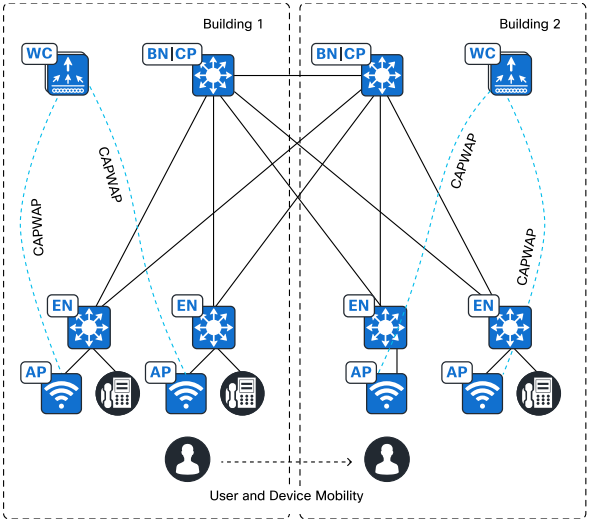

Hospitals have IoT endpoints like heart monitors connected to the network as wired and wireless devices. In the patient's room, these heart monitors are a wired endpoint, but when moving the patient around for other tests, these monitors are wireless endpoints. Further, these monitors must maintain the same IP address in the same Layer 2 segment regardless of whether they are wired or wireless. SD-Access provides the flexibility to use the same subnet for wired and wireless clients, allowing a given endpoint to always retain its IP address irrespective of how it is connected to the network.

Onboarding IoT Endpoints

Faster Onboarding Using Micro-Segmentation

Cisco SD-Access enables accelerated deployment of new endpoints connecting to an existing Fabric network. A healthcare environment has numerous critical IoT endpoints that may constantly need to get refreshed or new ones ordered as the hospital scales its needs. In a traditional network, before the endpoints can connect to the network, a customer has a long list of manual steps that must be completed on a box-by-box basis. These include assigning new access port configurations, authentication rules, access control lists, and more. Cisco DNA Center automation in the SD-Access architecture provides unmatched flexibility for new endpoint onboarding in several ways:

- Sites are already pre-configured with Authentication Templates as prescribed by the customer
- Cisco Endpoint Analytics has already facilitated the classification of similar endpoints in the network
- SGT assignment is set up with Cisco DNA Center, with the corresponding authorization rules being created automatically in Cisco ISE.

Using this endpoint onboarding methodology, any new endpoint can be onboarded simply by connecting it to the network and letting the endpoint be automatically placed into the right subnet and security group. The corresponding access policies are automatically downloaded to all pertinent network devices as the endpoint comes online.

Endpoint and Policy Analytics

Healthcare customers benefit immensely from the ability to classify an innumerable number of IoT devices based on the kind of data traffic that passes through the network. Once the endpoint is onboarded, customers can then group them into various security groups. For example, infusion pumps from different vendors will still share the same attributes to be grouped in the same security group.

When moving from a traditional network to SD-Access, many healthcare customers can leverage the insights Endpoint Analytics provides to profile novel, unclassified devices into specific endpoint groups. These endpoint groups can be further separated into meaningful security groups based on the classification.

For example, a customer may choose to group infusion pumps and heart monitors with a similar security group because they share a similar security profile. This facilitates a smooth transition to embrace a group-based policy enforcement strategy that will work across multiple sites. This scales quickly and more accurately than utilizing a traditional box-by-box access control list strategy.

In addition to Endpoint Analytics that help classify, Group-Based Policy Analytics helps customers identify the right security group policies that need to be configured in the network. GBPA assists the customer when migrating from cumbersome IP-address-based rule management to more flexible group-based management.

Figure 4.5: Cisco DNA Center Group Based Policy Analytics (GBPA)

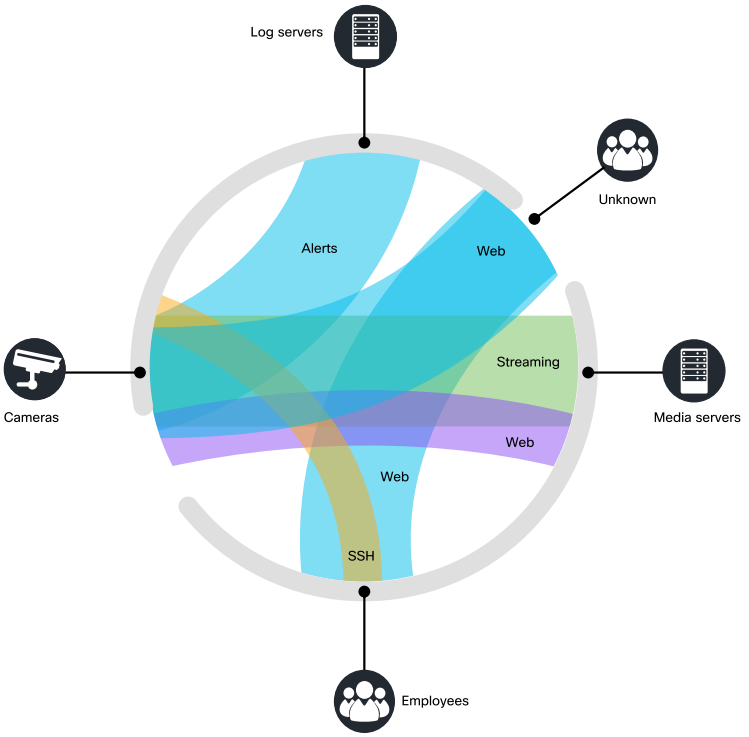

In Figure 4.5 there are several flows from a given security group to destinations for whom a security group has not been assigned yet (Unknown). For example, the Cameras group and Employees group both have web-based flows to the Unknown destination group. Double-clicking into what those flows are can help identify those Unknown destinations and can facilitate grouping them into common IP-SGT mappings in Cisco ISE.

The network operations team, in consultation with the IoT team, identifies the basis of these flows and can enforce the following outcomes as a result of these audits:

- Block illegitimate flows that were previously allowed in the network
- For legitimate flows, group those with similar patterns and assign access policies to help customers migrate their traditional network to a micro-segmentation-based network.

Silent Hosts

Some IoT endpoints on healthcare networks exhibit unique communication traits. These types of devices are sometimes referred to as Silent Hosts due to their unique communication. They are broadly categorized into two categories based on their communication characteristics:

- Once the endpoints are onboarded onto the network, they do not send any further packets or frames. These endpoints are essentially hibernating and only respond to specific frames or packets directly addressed to them. These devices may be DHCP-capable or statically addressed.
- Endpoints that are cabled and have not registered to the Fabric Overlay because of a static IP address

For the first set of endpoints, even though the endpoint is active, the Edge Node that is connected to that endpoint may have removed that endpoint from the Edge Node's host database. Since the endpoint is hibernating, the Edge Node is not receiving any response to periodic Internet Control Message Protocol (ICMP) probes, so the host is considered as no longer present on the Edge Node. Any other endpoint in the network with a previous record of communicating with the silent host will continue to communicate using the Silent Host's MAC address.

Cisco SD-Access supports handling unknown unicast frames by flooding them to every Edge Node within an IP Address Pool that is enabled for Layer 2 Flooding in the Fabric Site. While not responding to ICMP, hibernating devices will respond to specific payloads of traffic that are more directly relevant. As a result, the Silent Host will respond and rejoin the Fabric

network if the unknown unicast frames are relevant to it. Flooding to the endpoint will also happen if the endpoint is part of a Layer 2 VNI, which requires the Layer 2 Flooding feature to be enabled. Once the endpoint is registered in the Control Plane Node, any further communication to the endpoint can continue without concern.

They have never registered with the Control Plane Node for the second set of endpoints. The only way the Fabric can even know of the presence of such an endpoint is by configuring a manual IP Device Tracking (IPDT) entry on the Edge Node to which such a Silent Host is connected. This status mapping permanently ties the host's MAC address to the Device Tracking database. Once configured, such hosts will never lose traction from the Control Plane and can continue to attract traffic.

High Availability Capabilities

Critical VLAN

Critical care units must be assured of uninterrupted connectivity to the larger enterprise network. For clients that are already onboarded, if connectivity to the ISE Policy Service Node is lost, periodic re-authorization is paused to ensure disruptions in the authentication path do not impact the data plane.

For clients that are not yet onboarded into the network, the Critical VLAN feature places the client on a particular VLAN if connectivity to ISE is lost. This VLAN provides limited access to the network. Critical VLANs can utilize static micro-segmentation to enforce the policy in the absence of ISE.

Rolling AP Upgrades

Healthcare customers can upgrade wireless networks without network downtime. Using the Rolling AP Upgrade feature, APs can be upgraded in a staggered manner while still being connected to the same controller. This method can be used to avoid upgrade downtime even for N+1 networks by using the N+1 Hitless Rolling AP Upgrade feature and a spare controller.

Hospitals with Server Nodes on Campus

Healthcare networks are faced with the challenge of hosting confidential patient records on-premises to ensure there is survivability and access to such records in case of fiber cuts between buildings. Hospitals typically need access to Picture Archiving and Communication System (PACS) infrastructure, which contains vital imaging records of patients. Storing these records locally to the campus, rather than fetching all of the data from a remote data center, significantly lowers the possibility of communication disruption between the PACS storage and the user needing this information.

To address challenges like this, Cisco SD-Access provides the flexibility of supporting PACS servers connected directly to Edge Nodes in the Fabric. In addition to supporting trunk ports on Edge Nodes, SD-Access allows the PACS system to utilize a Layer 2 connection to a Border Node, thereby making Layer 3 Border Nodes also function like a Layer 2 terminating device.

Device and Building Level Resiliency

Device redundancy in Border Nodes can be managed with a StackWise Virtual (SWV) based deployment. A recommended design for a highly resilient Border Node is to deploy SWV Border Node pairs, thereby providing both intra-chassis and inter-chassis resilience.

To provide building-level resiliency, one of the models that healthcare customers tend to embrace is to deploy Border Nodes in multiple buildings to terminate connections from various hospital buildings on a large campus. Figure 4.6 shows an example of such a resilient design.

Figure 4.6: Border Resilience across buildings

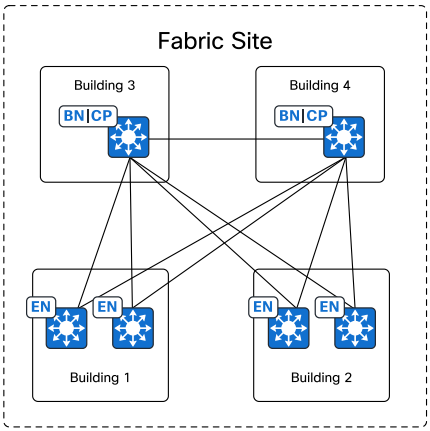

As can be seen in Figure 4.6, Buildings 1-4 are in the same Fabric Site, with the Colocated Border Nodes and Control Plane Nodes being located in separate buildings. Cisco SD-Access provides the flexibility to assign priorities to these border deployments such that one Border Node can be prioritized more or even become the only active border carrying traffic, so long as it is active. When a building fails, the Fabric Border Node in the other building can automatically take over all traffic from the Edge Nodes.

Endpoint Level Resiliency

Often, we encounter healthcare customers requiring critical endpoints to be hosted on servers that are dual-homed into the network infrastructure.

Cisco SD-Access supports the ability to dual-home endpoints in an active/standby mode into different Edge Nodes. Failover can happen based on two failure mechanisms, NIC bonding and Layer 2 Flooding.

When utilizing NIC bonding, the endpoint hosting server runs Link Aggregation Control Protocol (LACP) in active mode to two back

stacked/SVL paired Edge Nodes. LACP is set up with a maximum of 1 link to ensure only one of the links is active and the other is on standby, and any failover kicks in based on LACP failures.

Figure 4.7: Server NIC bonding between Edge Nodes

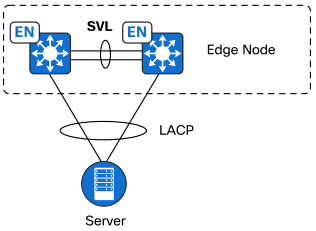

In the other model, the endpoint server application has its own means to detect which link needs to be active by sending heartbeats into the network to a Layer 2 multicast address. Only the active link will send the heartbeat, while the standby link keeps receiving it. As long as the application continues to receive the heartbeat on the standby link, it will be status quo. In Cisco SD-Access, this is achieved by using Layer 2 Flooding for Broadcast, Unknown Unicast, and Multicast (BUM) traffic which sends a multicast frame from one Edge Node to be received by clients on another Edge Node, as shown in Figure 4.8.

Figure 4.8: Using Layer 2 Flooding for application resiliency

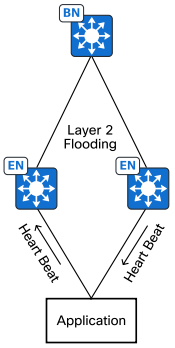

If the active heartbeat link fails, the standby heartbeat link stops receiving the heartbeat via the Layer 2 Flooding in the Fabric Site. The application will automatically switch to using the standby heartbeat link to send all its traffic, thereby providing failover resiliency.

Patient Safety and Asset Tracking

Cisco DNA Spaces

With Cisco DNA Spaces, you can leverage your existing access network to glean insights into the location of users, connected medical equipment, and IoT devices, allowing you to deliver data-driven care and improved clinician workflows. With our rich ecosystem of partner applications and devices, you can transform your facilities to offer patients an enhanced experience and provide clinicians and staff instant access to records and equipment.

Utilizing integration with Cisco partners such as Stanley Healthcare, healthcare providers can leverage their Wi-Fi infrastructure and Cisco DNA Spaces for use cases like asset management, environmental and temperature monitoring, Staff, Patient, and Infant Protection, as well as

various applications such as workflow and hand-hygiene compliance. With its robust Firehose API and 24/7 monitoring, Cisco DNA Spaces reliably provides data to the Stanley Healthcare platform, all while complying with user privacy.

Within hospitals, staff can track the location of medical devices and other critical equipment and be notified proactively if any equipment is misplaced or enters a prohibited area. For medications, vaccines, and other environmentally sensitive assets, staff can set temperature thresholds for monitoring and get immediate notifications when conditions deviate. Healthcare providers can implement real-time monitoring of infants, patients, and staff and trigger notifications in the event of an abduction, entry into prohibited areas, or staff duress. All these contribute to time and cost savings, higher productivity, and safer patient and staff experiences.

Healthcare providers can deliver these use cases while increasing value in the following ways:

- Simplified deployment: Reduce complexity and increase uptime through a lightweight cloud and connector model with active/active HA. Significantly reduce time to deploy, maintain and upgrade compared to traditional on-premises architecture.
- High Scalability: Deliver Real-Time Location Service (RTLS) capabilities and use cases at scale across multiple locations, using a lightweight deployment model. The DNA Spaces connector can support up to 10,000 Access Points (APs) and 350,000 devices with High Availability.
- Enhanced monitoring and support: Enable faster time to issue resolution and reduce downtime through proactive end-to-end monitoring. This includes monitoring the partner app, the Firehose API, and the data flow between the controllers, connectors, applications, and DNA Spaces.

Cisco DNA Spaces integrates with Stanley Healthcare RTLS to provide critical patient safety and asset tracking services. The Cisco SD-Access Wireless architecture acts as the backbone for delivering this functionality.

Reducing Operational Complexity

Centralized Troubleshooting using iCAP and Sensors

Healthcare customers often lack an onsite network support team in smaller clinics which makes troubleshooting that much more difficult.

Cisco DNA Center's Intelligent Capture capability supports a direct communication link between Cisco DNA Center and APs in Cisco SD-Access Wireless environments. Cisco DNA Center can receive packet capture data, AP and client statistics, and spectrum data. With the direct communication link between Cisco DNA Center and APs, Intelligent Capture allows you to access data from APs that is not available from wireless controllers.

Replacing IP-Based Rules with Group-based Rules

Cisco DNA Center allows healthcare customers to group similar endpoint groups that have been classified using Endpoint Analytics. These endpoint groups can be placed into common security groups based on Cisco DNA Center's visibility of traffic flows from each group. This, in turn, can be used to create security policies to allow or disallow data traffic to different SGTs based on simplified SGACL rules instead of complex IP-Based access list rules.

SGT-based rules are more self-explanatory and will help eliminate possibilities of unnecessary configurations, thereby alleviating the burden on newer NetOps engineers even though the policies themselves may have been established by a different team of engineers in the past.

Single Touchpoint for Policy Management

For healthcare organizations that scale beyond the limits of a single Cisco DNA Center cluster in their SD-Access environment, the option exists to deploy multiple Cisco DNA Center clusters that can integrate with a single Cisco ISE deployment. This reduces the complexity of policy management in a large healthcare enterprise by supporting the same policy framework to be established across all Cisco DNA Centers, managing different regions of the network. For more details, refer to the Multiple Cisco DNA Center to Single ISE section in the Architecture Overview chapter.

Deployment Options

The examples given in this section are based on Cisco's observations of healthcare customer deployments. Any numbers given do not represent a hard limit or maximum. The scale should always be determined by the device being deployed and the information presented in the Cisco DNA Center datasheet.

If the provider cloud is WAN or MAN, then either IP Transit or SD-Access Transit are viable options depending on WAN/MAN characteristics and other factors such as traffic profile and type.

Figure 4.9: Large Healthcare Campus, Branch LAN, and WLAN view

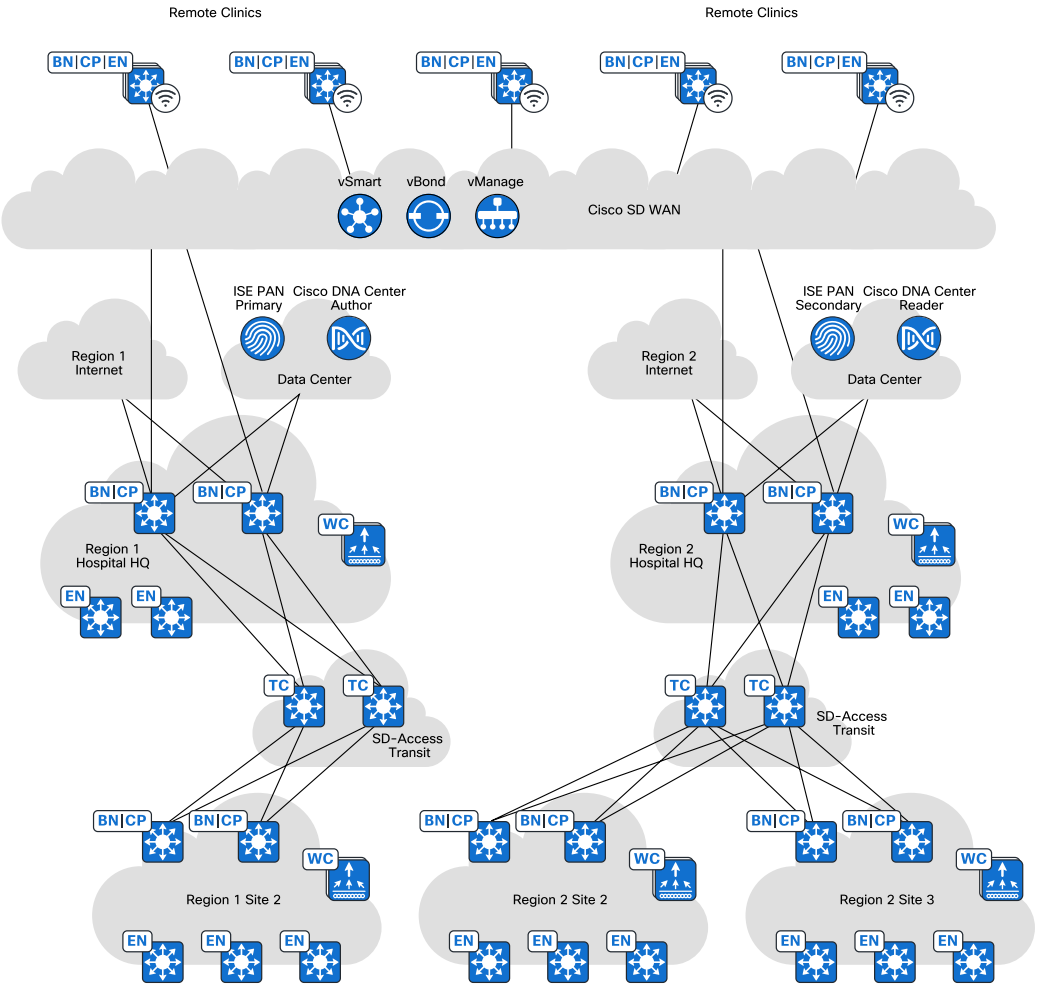

If the provider cloud is SD-WAN enabled, then IP Transit will be the most common option.

Micro/Mobile Clinic

Figure 4.10: Micro/Mobile Clinic design

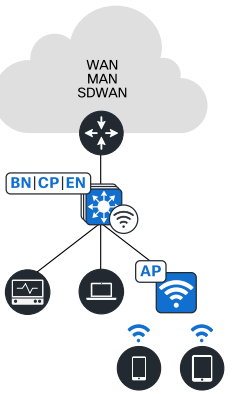

These sites typically support fewer than 50 users/devices, especially in mobile clinic deployments.

Small Clinic

Figure 4.11: Small Clinic design

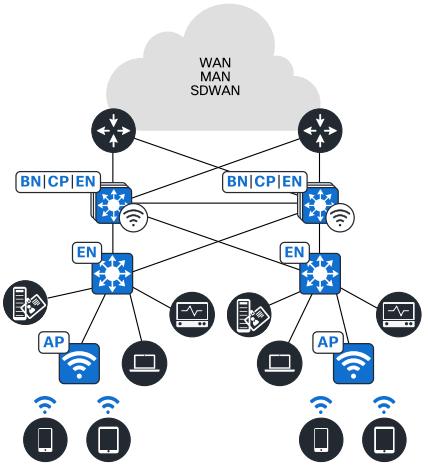

These sites typically support fewer than 1,000 users/devices and 50 APs.

Medium Size Clinic

Figure 4.12: Medium Size Clinic design

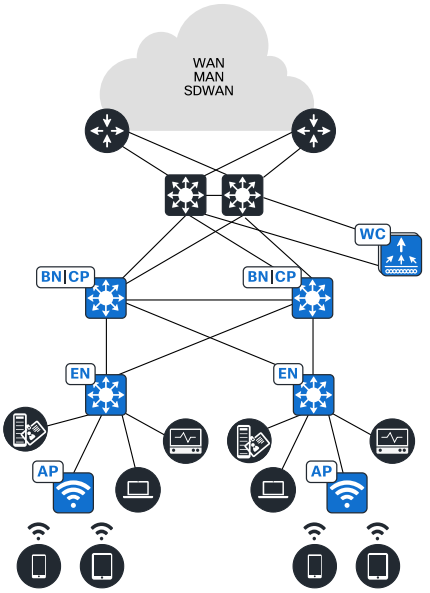

These sites typically support fewer than 5,000 users. Fabric-enabled AP scale will usually require dedicated WLCs instead of embedding on the Border Nodes. Edge Node count will often require dedicated Fabric Border Nodes instead of Fabric in a Box. It is most common to have Colocated Border Nodes and Control Plane Nodes.

Large Clinic

Figure 4.13: Large Clinic design

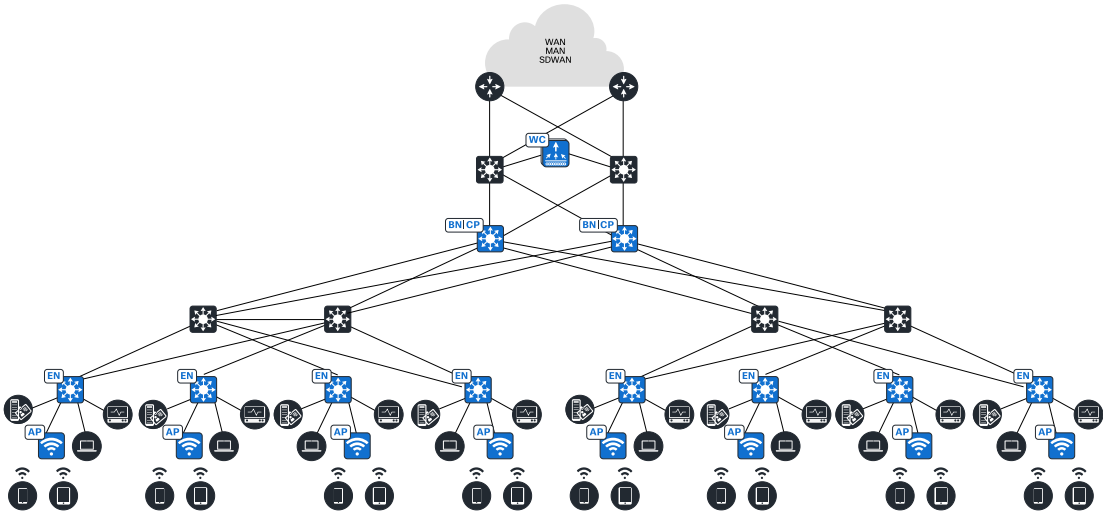

These sites typically support fewer than 25,000 users. Fabric-enabled AP scale will require dedicated WLCs instead of embedding on the Border Nodes. Edge Node count will often require a Distribution Layer that is just intermediate nodes connected to dedicated Border Nodes. It is most common to have Colocated Border Nodes and Control Plane Nodes. Shared services can be used where patients, doctors, and devices connected through Virtual Networks need to talk to each other.

Summary

Given their environments' criticality and privacy requirements, healthcare organizations face unique challenges when deploying a network infrastructure. When designing their networks, they must be extra vigilant to ensure that data integrity, worker and patient safety, and operational resiliency are maintained at the highest levels.

In this chapter, we showed how Cisco SD-Access provides a secure, dynamic, resilient solution that addresses the most demanding use cases and challenges faced by healthcare organizations. With the addition of Cisco DNA Center capabilities such as Endpoint Analytics and Ecosystem Partner applications such as Cisco DNA Spaces, Cisco SD-Access ensures that healthcare organizations can deliver the critical services demanded by their users and customers.

Cisco SD-Access for Large Enterprises and Governments

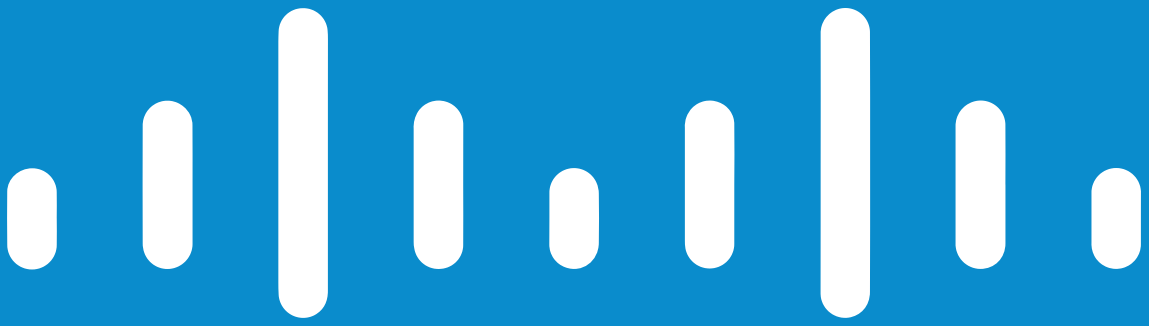

Introduction

Enterprise Network Introduction

Computer networks are like fingerprints: each is different and built to address the organization's business needs. Enterprise Networks owned by multinational corporations and governments are some of the largest networks in existence. They are composed of hundreds of switches and multiple Campus locations consisting of thousands of users and endpoints.

In most cases, the following characteristics describe an Enterprise Network:

- The Enterprise Network is a mix of multiple large, medium, and small campuses
- Individual Campus locations consist of multiple buildings with multiple floors connected typically through dark fiber
- Dedicated leased lines connect large campuses, data centers, and Carrier Neutral Facilities (CNFs) across geographic locations
- Internet access is centralized for these Campuses, and traffic egress a security stack

Within the vertical of Enterprise Networks, there are three typical types of networks: Government, Large Enterprise, and Managed Services Provider (MSP). Each network may have many features in common, but there are specific challenges that each type of network will need to solve.

Challenges

Some of the common organizations and the challenges seen in the Enterprise Networks include the following:

Government Agencies

- Government agencies such as federal, state, and local governments need end-to-end network segmentation for security reasons.
- Government agencies have security mandates requiring networks to be air gapped.
- Federal agencies have stringent requirements for their VLANs which are fixed quantity resources. Overlapping VLAN IDs are avoided unless there are no other options available.

IT as a Business (ITaaB)

- Large Enterprise Networks are moving to a *cost model* which allows them to internally cross-charge their departments to recover the expenses involved in network support.
- Individual tenant consumers of the IT network must be isolated, though each tenant needs access to a common set of shared services.
- Enterprise Networks have significant organic growth. IP address space must be preserved and maintained for security and compliance purposes as the network expands.

Managed Service Providers

- Managed Service Providers (MSPs) manage their customer's IT infrastructure and end-user systems.
- MSPs provide a set of day-to-day management services such as network and infrastructure management along with security and monitoring services.
- MSPs must have visibility into their network to know what users and devices are connected to the network and how applications are performing in the network.
- MSPs must address potential major network issues proactively.

Other Enterprise Networks

- Network-level and service-level resiliency must exist across geographically dispersed networks.
- Latency round-trip-time (RTT) between management systems and managed devices has created a need for multiple management systems for different regions of the globe.
- Bring Your Own Device (BYOD) modalities allow employees to bring personal devices. Mobile Device Management (MDM) platforms that onboard these BYOD devices must be incorporated into the design.
- Specialized Application servers and audio/video devices support Access Layer switches instead of just data centers.

Customer Solutions

The following section will describe some of the most common use cases of Large Enterprise Networks.

Government Agencies

Government Departments Isolation via Virtual Networks

Many Government agencies have segmented their networks to achieve consolidation while maintaining separation and isolation. Deploying multiple Virtual Networks is also an effective way to increase the Return on Investment (ROI) of assets and reduce the overall cost of implementation.

IP Address Management and Conservation

Traditionally, network best practices recommended using an IP subnet using 24 network bits (/24) for users and endpoints. Most networks for Enterprises are designed with multiple subnets of this size across their access network. This results in a large number of Layer 3 Switched Virtual Interfaces (SVIs) on Distribution Layer switches.

The driving force behind networks using subnets of this size (/24) was to limit broadcast traffic in that Layer 2 Flooding domain. To address this challenge, an SD-Access network, by default, does not flood broadcast traffic. The result is that the total number of IP subnets for users and endpoints can be reduced by using subnets with a smaller number of network bits. This, in turn, provides more host addresses, such as using a /20.

As the total number of subnets is reduced, the performance of the Global Routing Table is increased. In addition, this provides a fresh start in

standardizing the IP address schema for the global deployment. In most cases, organic growth in an organization leads to non-summarizable IP addresses across the network.

Another major advantage of this flexibility is the Custom VLAN capability in Cisco DNA Center. Custom VLAN allows you to migrate brownfield networks that require IP portability using the same VLAN IDs which are configured in a legacy network. Having flexibility is crucial as it scales the migration of endpoints with static IP addresses. This means the same networks and static IP addresses on endpoints can continue to be used.

Air Gap Support

Many government agencies require support for air gap deployment – especially for Cisco DNA Center. Today, if we have to deploy Cisco DNA Center in an air gap environment, the deployment of any SD-Access Fabric managed by that Cisco DNA Center can also be supported in that mode. An air gap is a highly requested feature for government agencies owing to their need to be completely disconnected from the Internet for security reasons. An air gap network is a top requirement in Defense, National Security Agencies, and many others.

Inter-Agency Audits

Figure 5.1: Inter-Virtual Network route leaking with Peer device

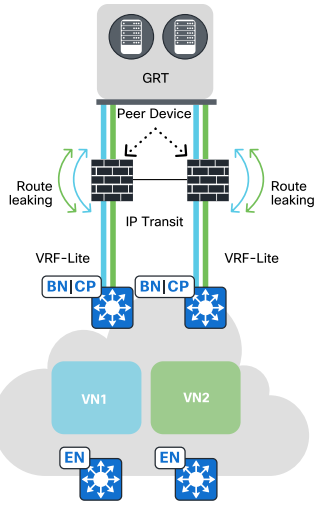

By hosting different agencies in different Virtual Networks (VNs)/Virtual Routing and Forwarding (VRF) tables, isolation is achieved inside the Fabric. In selective cases, there is a need to have communications between endpoints in different VRFs. For example, an engineer in a Contractor VRF might need to reach out to the IoT VRF to upgrade software remotely on a sensor via a Firewall. This can be easily achieved in the Cisco SD-Access solution, where the Border handoff of each VRF connects to a Firewall. Using either dynamic routing or default routing, traffic can be routed through the Firewall for inspection.

Access to Shared Services

Figure 5.2: Inter-Virtual Network route leaking with SD-Access Extranet

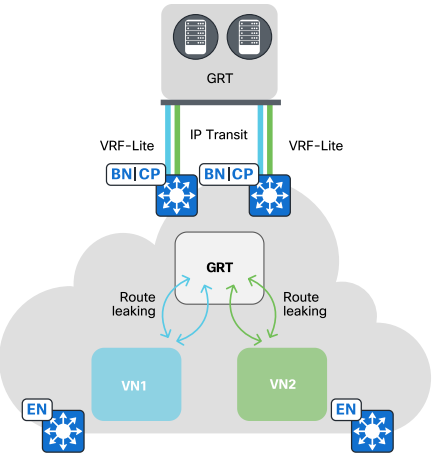

There is always a need to access Shared Services like DHCP, DNS, and Internet, among others. Traditional methods of route-leaking result in duplicate routes being installed in each VRF. This leads to an increase in the utilization of the Ternary Content Addressable Memory (TCAM) on network devices.

Cisco SD-Access provides an efficient method of inter-VRF routing without increasing the TCAM utilization. This feature is called SD-Access Extranet. You do not need a peer device to perform the route-leaking; the route-leaking is done internally within the Fabric. In such scenarios, SD-Access provides the flexibility for traffic to be selectively forwarded, enforced, and audited by a firewall, as shown in Figure 5.2.

IPv6 mandates in the Underlay from Government Agencies (Reference to future solutions coming)

IPv6 is one of the most debated topics and has two different versions depending on the organization. For most multinational companies, using IPv4 addresses is not an issue as they have a large allocation of IPv4 addresses for consumption. However, some government agencies have a mandate to start using IPv6 addresses. SD-Access supports IPv6 in the Overlay for the clients and endpoints. Recent mandates with various governments are requiring infrastructure IPs to be IPv6 and reduce dependency on IPv4 addressing.

Note: IPv6 in the Fabric Underlay is a feature still under development. Availability of this feature comes after the publication of this book.

IT as a Business (ITaaB)

Shared Infrastructure using IT as a Cost Center

Historically, Information Technology (IT) has always been seen as a cost center for an organization. With high resource and operational costs, IT had always been slow in transformation, leading to an enterprise's slow adoption of digitization. Cisco DNA Center, along with SD-Access, helps with not just the standardization of all IP service offerings across the enterprise but also with a single pane of glass that provides a view into managing different departments within a Large Enterprise Network.

ITaaB works by offering standard services to different business units within an organization. For example, if a regional site in a city is used as a sales office, they would be provided a catalog of standard physical architecture depending on the number of users or criticality of the site. The sales organization picks the network design, and the IT team deploys and manages them with SD-Access providing high resiliency and consistent security policy for a recurring cross charge towards the sales department. If that same building is used for other departments such as Research and Development (R&D) in the future, the cost will be shared across two

departments. New departments can be rapidly deployed using the automated workflows from Cisco DNA Center (refer to topology in VRF isolation section).

Technology IT Services

More industries across many different verticals are leveraging the technological expertise of third-party IT services technology companies in developing and maintaining their IT applications and systems. Such IT services companies are Large Enterprises by themselves, providing service to all of their clients by employing a workforce that is dedicated to different customer accounts. Such customers seek a high level of security and isolation within the network infrastructure of the IT services companies to ensure that there is no compromise to the security of their applications and data.

SD-Access provides flexibility to such IT services companies by providing isolation at two different levels.

Network Isolation

The most straightforward and common model to provide isolation for tenants of IT services companies is to segregate them into different Virtual Networks (VNs). SD-Access supports up to 256 VRFs to be set up within a given Fabric Site to provide complete isolation between tenants within each Virtual Network.

Figure 5.3: IT as a Business Example

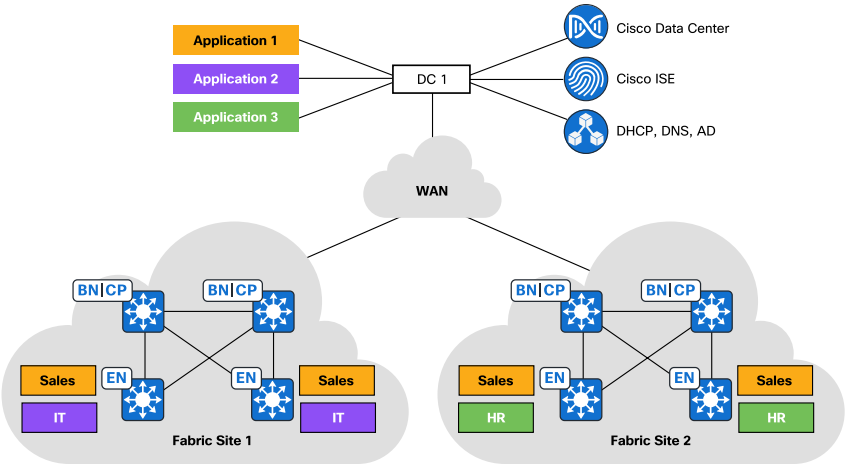

Fabric Zones

Large campuses are comprised of networks across many buildings. The IT teams have an IP schema where they map subnets to some of the buildings and use it as a geographic identity within the campus. Cisco SD-Access configures all the Overlay subnets on all Edge Nodes by default. This might not be the desired behavior in this case. Using the Fabric Zones capability, you can constrain certain subnets in certain buildings or floors, preserving the previous IP schema. Using Fabric Zones, you can assign different Edge Nodes to be in separate zones so that no other network device on the same site is enabled with the same set of subnets.

Figure 5.4: Fabric Zones

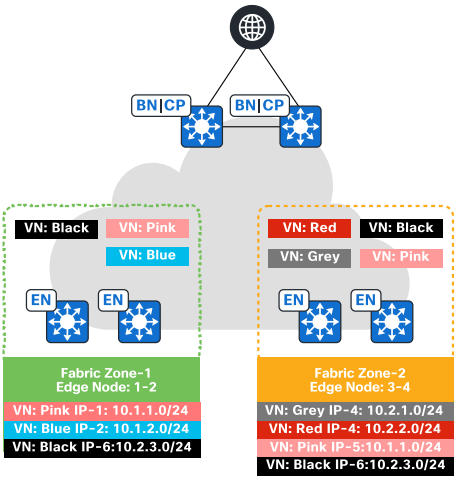

In the Fabric Site design shown in Figure 5.4, you use a single large Fabric Site with multiple Edge Nodes. The Fabric Site uses 5 Virtual Networks – Gray, Red, Black, Blue, and Pink with IP Subnet: 10.1.1.0/24, 10.1.2.0./24, 10.2.1.0/24, 10.2.2.0/24, 10.2.3.0/24. Two Fabric Zones – Green and Orange – are configured within this Fabric Site. With this design, Fabric Edge switches in the Green Fabric Zone will be configured with only Pink, Blue, and Black Virtual Network and related subnet configurations. Fabric Edge switches in the Orange Fabric Zone will be configured with Gray, Red, Pink, and Black Virtual Networks and related subnet configurations.

You can extend the same Fabric Zone concept for different types of Medium and Large Fabric Sites to have better management control.

Multisite Remote Border

For certain other tenants of IT Services companies, it is not enough to just zone their subnets into exclusive switches. Tenants may also seek the

deployment of exclusive Border and Control Plane Nodes for their Virtual Networks for full isolation. SD-Access supports Multisite Remote Borders whereby specific Virtual Networks can be terminated on specific Fabric and Control Plane Nodes without sharing those with any other tenants running on the same infrastructure. This Multisite Remote Border feature can also be used to terminate VRFs across sites as well. The most common application of this feature is in providing Guest access across sites. This allows the organization to have a common Guest subnet across many individual Fabric Sites and funnels all the Guest traffic to a central site where the Multisite Remote Border is located. Figure 5.5 depicts the Tenant use case.

Figure 5.5: Multi-Tenant Fabric with Multisite Remote Border

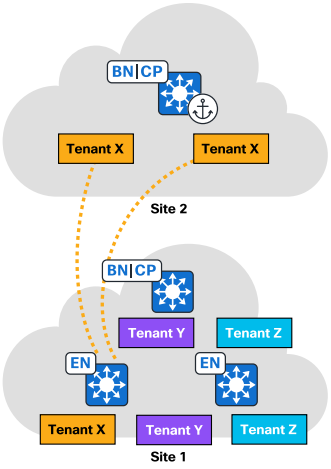

Fabric as a Tenant – Managed Service Providers (MSP)

Managed Service Providers (MSPs) are another set of Enterprise Networks with a little different use case than ITaaB. For MSPs, their customers use the MSP's SD-Access managed infrastructure as a means of transport to access centralized applications within the MSP's network. The SD-Access infrastructure is completely managed by the MSP, restricting management access to their customers. The Fabric as a tenant design warrants a network that can be flexible enough to be deployed on demand with a security policy to ensure no two customer networks can communicate with each other. In addition, large MSPs can share high-performance hardware for multiple customers without taking any downtime in provisioning new services.

An example of an enterprise acting as an MSP is a county IT team managing the infrastructure that supports agencies such as police, fire, ambulance, and schools. Each of these will have its own unique policy requirements, but they want to share the same hardware to avoid the cost of deploying a dedicated network for each.

By using Cisco SD-Access, Large Enterprises working as MSPs can more easily and flexibly deliver the services required by their customers while using Cisco DNA Center as a single touchpoint for network operations.

Figure 5.6: Managed Service Providers (MSP) – Fabric as Tenant

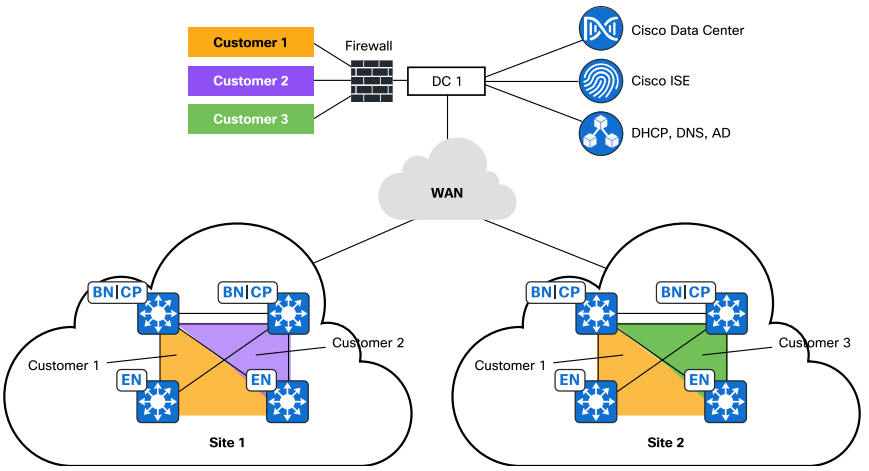

Figure 5.7: Managed Service Provider – Shared infrastructure with VRF-based isolation

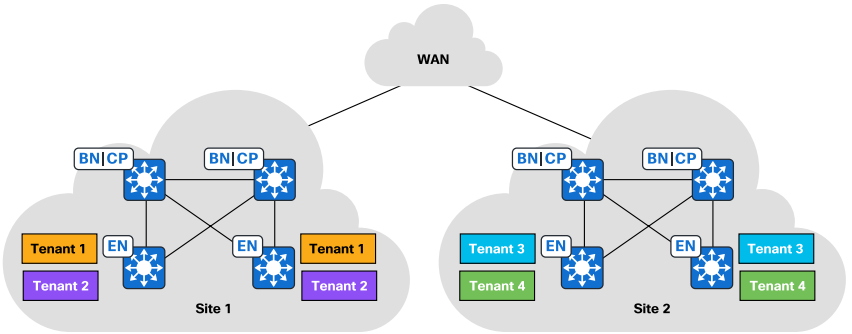

Other Enterprise Use Cases

Large enterprise deployments spread across the globe require multiple Cisco DNA Centers to be deployed to meet scale, network latency, and/or global compliance requirements. Multiple Cisco DNA Center clusters can be integrated into a single Cisco ISE, providing a centralization and standardization of Authentication, Authorization, and Security Group Policy across the enterprise.

Multiple Cisco DNA Center

Cisco Identity Service Engine can scale up to 2 million endpoints. Cisco DNA Center systems can support an endpoint and network device scale based on different Cisco DNA Center models, which range from 25K to 300K endpoints and 500 to 8000 network devices. You can integrate up to 5 Cisco DNA Center systems (standalone or clusters) with one Cisco ISE system, but you cannot integrate Cisco DNA Center with more than one Cisco ISE system.

If you have an extensive network that extends beyond a single Cisco DNA Center-supported endpoint/network device scale, the Multiple Cisco DNA Center to Single Cisco ISE feature can assist you in taking full advantage of Cisco DNA Center by integrating multiple Cisco DNA Center systems (standalone or clusters) with a single Cisco ISE. The Multiple Cisco DNA Center to Single Cisco ISE feature enables standard policy and other Cisco ISE features across all integrated Cisco DNA Centers and related Client Endpoint/Network devices.

Inter-Geographical-Region Communication

In Large Enterprises, when you deploy the Multiple Cisco DNA Center to Single Cisco ISE solution, each Cisco DNA Center manages its own network devices and Fabric Sites. To ensure and maintain end-to-end propagation of Secure Group Tags (SGTs) for traffic from one Cisco DNA Center-controlled SD-Access Fabric to another Fabric, the best option is to use a Shared SD-Access Transit. A Shared SD-Access Transit Control Plane Node can be

created with one Cisco DNA Center, which can add/delete/modify the site, but other Cisco DNA Centers can use it as a read-only shared Transit Control Plane Node. To ensure Fabric Sites created on different Cisco DNA Centers can communicate with each other, the different Fabrics must have the same segment assigned to the Virtual Network.

Figure 5.8: Global SD-Access Deployment with Cisco Multiple DNA Center to single Cisco ISE

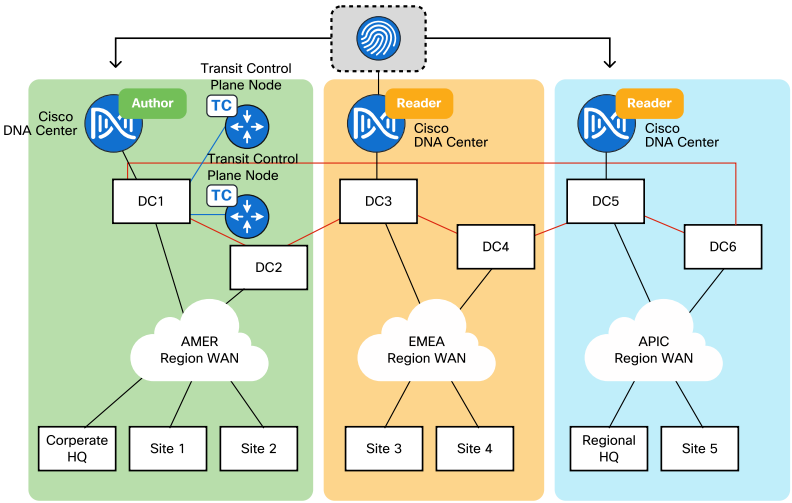

Bring Your Own Device/Guest Access

Guest or Bring Your Own Device (BYOD) wireless access provides differentiated access and user policy to known/unknown devices. Guest or BYOD access can be achieved in an SD-Access design by using various options such as Static URL Redirect and RADIUS-Based Change of Authorization (CoA). Cisco DNA Center can configure Central Web Authentication (CWA), External Web Authentication (EWA), and hotspot SSIDs for Cisco WLCs in the Fabric design. With Cisco ISE integrated, customers can use Cisco DNA Center and ISE BYOD workflows to onboard their endpoints by provisioning Certificate Authority (CA) signed endpoint

certificates and configuring the network interface and OS native supplicant to utilize the provisioned certificate for network access.

Telepresence Devices

User computers and IP Phones are the most common devices connected to the Edge Nodes. These devices can be profiled, postured, and onboarded into the Fabric.

Certain Audio/Visual (AV) and telepresence devices require a Layer 2 VLAN-only infrastructure with no Layer 3 connectivity. In this case, a Layer 2 VXLAN Network Identifier (L2VNI) functionality in SD-Access can be leveraged where a VLAN can be provisioned without any Layer 3 IP Address Pool. The provisioned VLAN can operate in Layer 2-only mode, satisfying a scenario where an AV vendor usually inserts a small 8-port unmanaged Layer 2 switch to connect multiple AV components to communicate with the controller.

Servers at the Edge

In many enterprise networks, you can have local servers like FTP servers, web servers, database servers, ESXi servers, and local application servers. These servers often need to be accessed by users from different sites, and require High Availability connected at the Access Layer. In such scenarios, Cisco DNA Center provides the functionality of configuring trunk ports on Edge Nodes, which will allow customers to deploy servers that can carry traffic from multiple subnets.

Deployment Options

Fabric Site Sizes – Design Strategy

In a Cisco SD-Access deployment design, you can create fewer, larger Fabric Sites rather than multiple, smaller Fabric Sites. The design strategy is to have a larger Fabric Site size while minimizing total site count, which can help reduce management overhead. Larger Fabric Sites can also help in deploying common policy across the Fabric. Though having a Large Site has multiple benefits, your business requirements may necessitate having a Small/Medium Site. The multi-dimensional factors of survivability, high availability, endpoint count, services, and geography are all factors that may drive the need for multiple, smaller Fabric Sites instead of a single Large Site. To help you in designing Fabric Sites of different sizes, the following reference models can be used.

Fabric Site Reference Models

Customers use different templates for each site type in deployments with physical locations, such as a large branch, a regional hub, headquarters, or a small, remote office. The underlying design challenge is to look at the existing network, deployment, and wiring and propose a method to layer SD-Access Fabric Sites in these areas. The design process can be simplified and streamlined by templating designs into reference models. The templates drive understanding of common site designs by offering reference categories based on multi-dimensional design elements like endpoint count, the number of Fabric Devices, and the number of users to provide guidelines for similar site size designs.

The numbers presented in this section are used as guidelines only and do not necessarily match maximum specific scale and performance limits for devices within a reference design.

Each Fabric Site includes a supporting set of Control Plane Nodes, Edge Nodes, Border Nodes, and Wireless LAN Controllers, sized appropriately from the listed categories. ISE Policy Service Nodes are also distributed across the sites to meet survivability requirements.

Major Fabric Site model:

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

Very Small Site

Figure 5.9: Very Small Site design

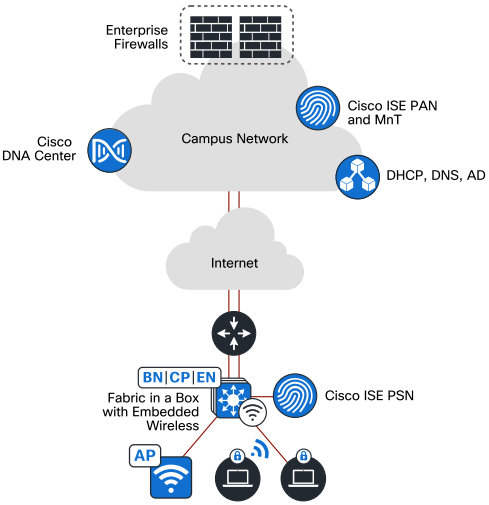

You can use Fabric in a Box to cover a single Fabric Site, with resilience supported by switch stacking or StackWise Virtual to support up to 200 endpoints and 40 APs.

For Fabric in a Box deployments, SD-Access Embedded Wireless is used to provide site-local WLC functionality. The site may contain an ISE PSN depending on the WAN/Internet circuit and latency.

Small Site

Figure 5.10: Small Site design

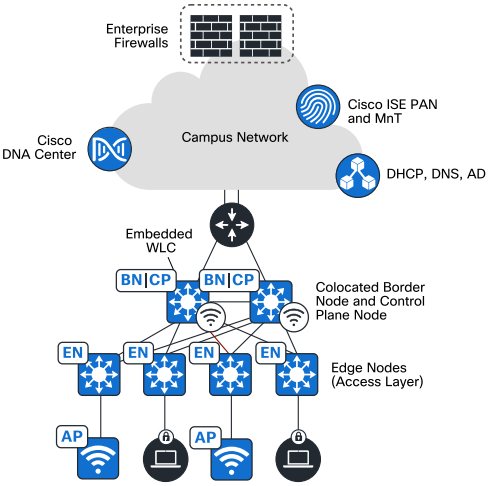

The Small Site Reference Model covers a single office or building with single wiring closets, usually up to 4,000 and up to 100 APs. The Border Node function is colocated with the Control Plane Node function on one or two devices and usually uses embedded wireless with the option of hardware WLCs.

The physical network is usually a Two-Tier Collapsed Core/Distribution with an Access Layer servicing several wiring closets. Rather than colocating all roles in one device, the Small Site Reference Model provides added resiliency, redundancy, and a more significant number of endpoints by separating the Edge Node role onto dedicated devices in the Access Layer. The Border Node and Control Plane Nodes are colocated in the Collapsed Core Layer. For SD-Access Wireless, the embedded WLC is provisioned on

the Colocated Border Node and Control Plane Node. Optionally, a virtual- or hardware-based WLC is used.

Medium Site

Figure 5.11: Medium Site design

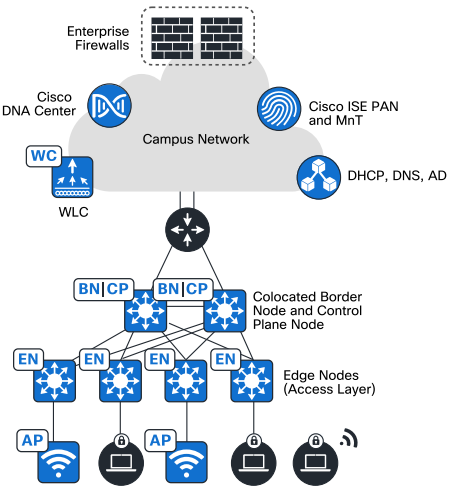

The Medium Site Reference Model usually covers a building with multiple wiring closets with a physical topology consisting of a Two-Tier Collapsed Core/Distribution with an Access Layer.

The Medium Site is designed to support less than 25,000 endpoints and less than 2,000 APs. The border function is colocated with the control plane function on either two devices or a highly resilient single device, and a separate wireless controller is ideally deployed in a High A configuration.

Large Site

Figure 5.12: Large Site design

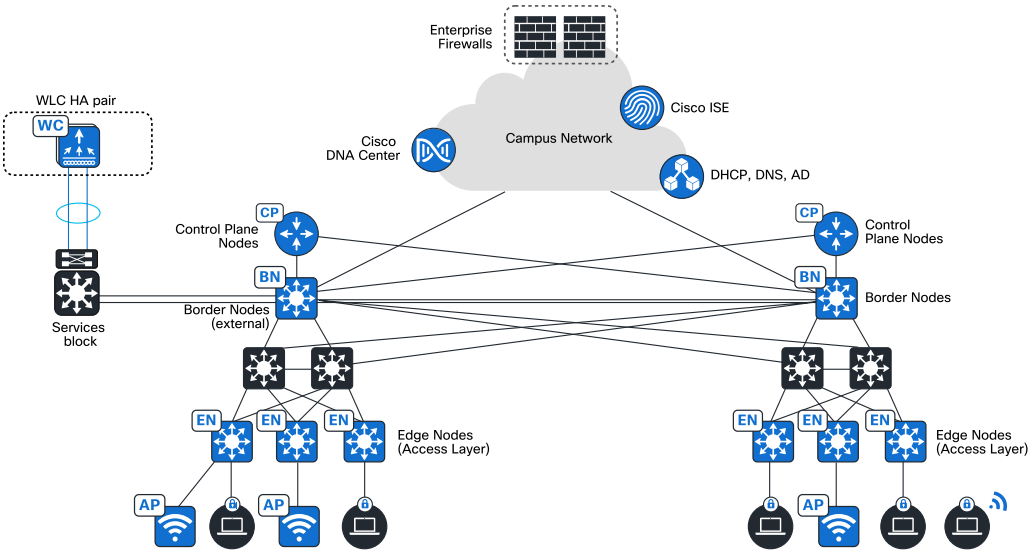

The Large Site Reference Model covers multiple buildings or a building with multiple wiring closets. The physical network is usually Three-Tier with Core, Distribution, and Access Layers. It may even contain a routed Super Core that aggregates numerous buildings and serves as the network egress point to the WAN and Internet. The Border Plane and Control Plane functions are provisioned on separate devices rather than colocating.

The Large Site supports up to 100,000 endpoints and 6,000 APs. Border Nodes are distributed from Control Plane Nodes using redundant devices, and a separate wireless controller is in a High Availability configuration.

Extra-Large Site

Figure 5.13: Extra-Large Site design

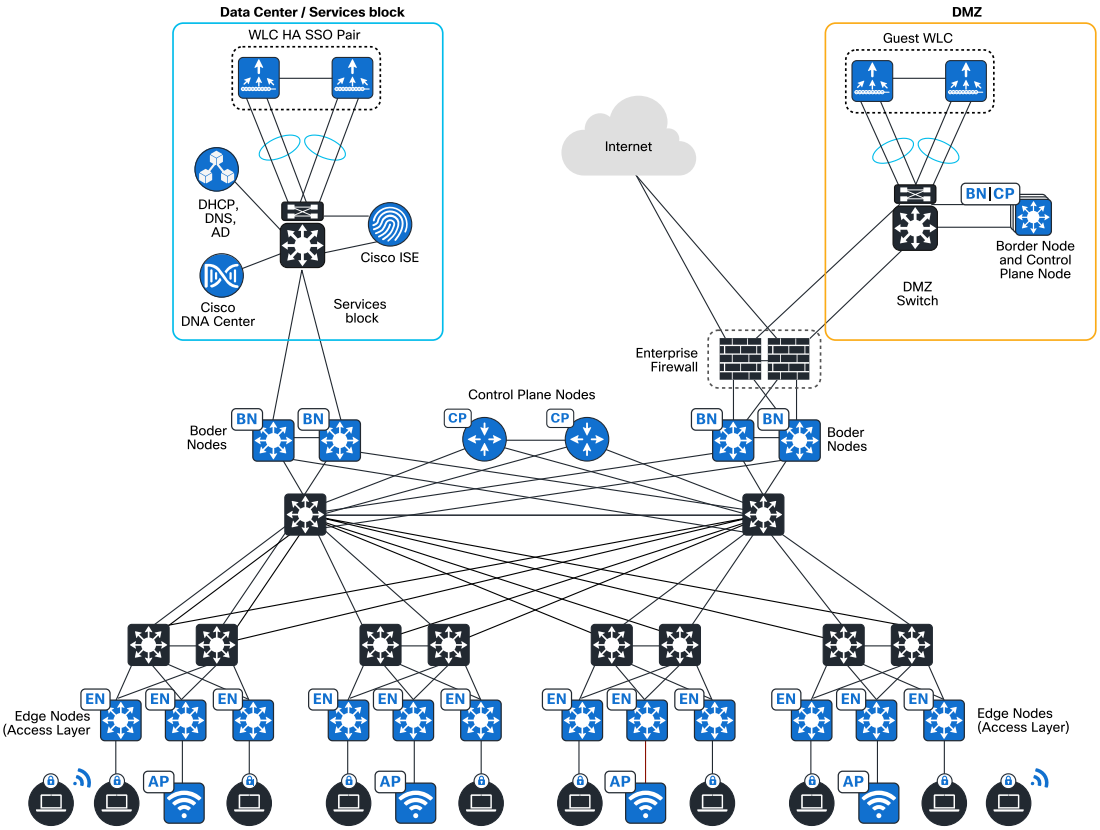

The Extra-Large Site Reference Model covers a building with multiple wiring closets, or multiple facilities stretched across a large campus. The physical network is a Three-Tier network with Core, Distribution, and Access Layers and may sometimes have a Super Core in a Fourth-Tier. An Extra-Large

network requires dedicated services exit points such as a dedicated data center, shared services block, and Internet services.

The Extra-Large Site supports up to 200,000 endpoints and 10,000 APs with Multiple Border Nodes distributed from the Control Plane Node on redundant devices, and a separate wireless controller in an HA configuration.

Summary

This chapter showed that Large Enterprise and Government Networks are expansive and complex. You saw how Cisco SD-Access can help solve complex challenges such as global end-to-end segmentation and how it can take an organization's network on a digital transformation journey. We learned how Cisco SD-Access enabled Large Enterprise Networks to transform from IT as a Cost Center to IT as a Business. We examined how some Large Enterprises will require multi-tenancy capability, making them essentially Managed Service Providers (MSPs). We demonstrated how Cisco SD-Access can provide scalable, secure networks for their customers.

Finally, we showed how SD-Access allows Large Enterprise organizations to support comprehensive security segmentation strategies in their complex, global, end-to-end networks.

Cisco SD-Access in Universities

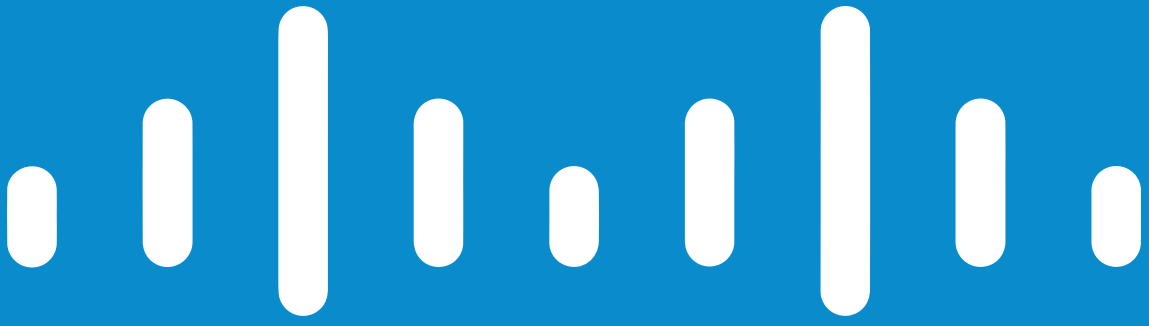

Introduction

University networks are unique because they are built on the premise of providing the ultimate sharing and collaborative experience in the pursuit of learning. To that end, they are divided into parts to accommodate various needs within the university, all with that common goal of sharing. Universities are subdivided into campuses of various sizes where specific needs of the student and faculty are addressed.

These campuses are typically all interconnected via high-speed optical links to centralized facilities such as data centers where access to shared resources is offered. Typically, the interconnecting networks are owned solely by the university, but occasionally, various partners or providers may be involved in small ways.

Figure 6.1: University Campus Network reference topology

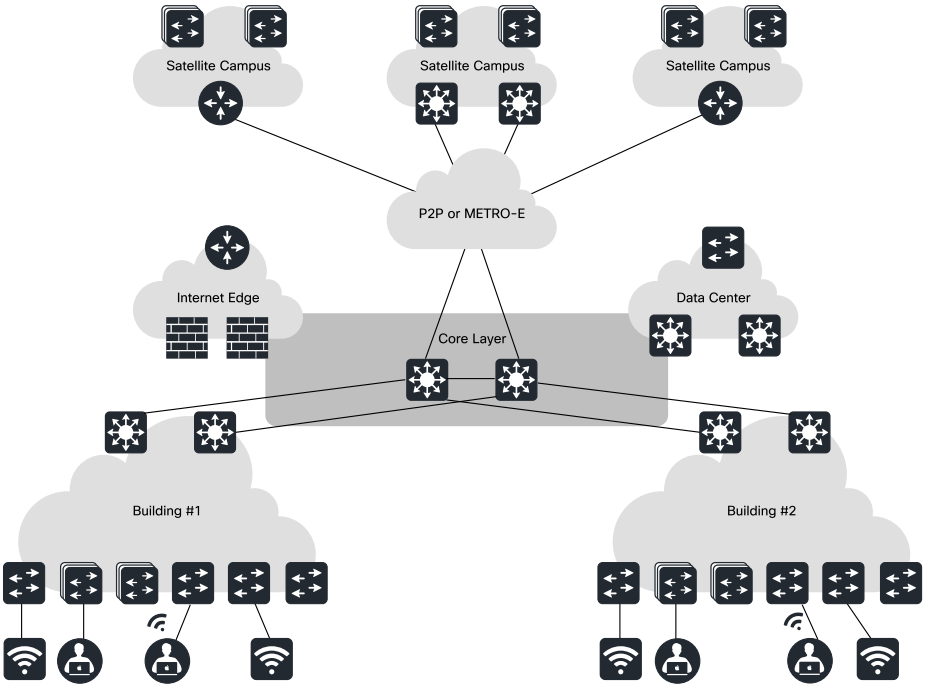

Universities functionally are divided into departments that take on centers of learning and provide instruction to students. In that regard, each department has varied needs, some more physical than others, but all with the idea of sharing some resources across the network to be used in instruction. Typically, these shared resources are housed on servers or specific machines within the campus or in the data centers, but it does vary from university to university. Research labs require massive amounts of data transfers from one end of the campus to the other. Frequently, the requirements for specific applications or services offered to students or faculty require a stretched Layer 2 broadcast environment. At times, simple Internet Protocol (IP) connectivity is all that is required, but the emphasis is always on connectivity and sharing of information.

Unlike every other vertical for which networks are designed, it is crucial to understand that everything moves in this vertical every 90 minutes. When classes change, the load and stresses on the infrastructure are pushed to their maximum. Association, authentication, and roaming clients within the campus push identity services to the maximum as well as the wireless environment. The aggregation of traffic across main conduits where students congregate increases load in areas you would not expect, such as hallways and cafeterias. To that end, all of these concerns must be addressed in the form of scale.

The network must provide various functions in the classroom that benefit faculty and students. The primary function it needs to address is access to the network, and that access should be similar to whatever media is offered, wired or wireless.

The second function that is critical is the speed of the network, and universities are often graded on this metric. Slow or poor access to resources, or poor onboarding experiences, can result in reputational loss as students vent their frustrations on social media. This could have an impact on future enrollment.

Once the devices are on the network within the classroom, they typically use some shared resource or application, and each has its own requirements. Typically, teaching applications like CANVAS are multicast-based and allow the teachers and students to interact in a collaborative classroom setting.

While that is one example of such an application, some universities have developed their own respective applications to facilitate instruction, and these applications usually involve multicast in some form or fashion.

The other application widely used is screen sharing for presentation purposes. Long gone are the projectors of the past, as today there is a high percentage of adoption of Apple TV or Chromecast, which together allow for a student or faculty member to share wirelessly to the class. Whatever the network is, these needs would need to be addressed as they have specific challenges and constraints when discussing how to let them operate while conforming to policy needs.

Similarly, there are common services that all departments will utilize to onboard clients onto the network. These include Dynamic Host Control Protocol (DHCP), Domain Name System (DNS), and Identity Based Services such as Microsoft Active Directory (AD) or Authentication, Authorization, and Accounting (AAA). These services are essential for the student's clients that they bring with them to gain access to the various common and department resources.

Additionally, it is important to understand that while there are the direct learning needs of the student and faculty, there are also the needs of the supporting departments. Those supporting departments are typically named facilities, and there are many ranging from catering, janitorial, vending, health and wellness, medical, sports, housing, finance, and administration. Each facility has a goal, and each goal has specific requirements. In this area, the network is a transit to allow the supporting facility to achieve its goal. Each application within that facility or client used will have needs, and the network needs to address these needs.

The last major area of need in this evolving world is a well-defined and planned campus security requirement. This specific requirement has seen increasing needs for physical badging, IP cameras, panic stations, telephone, and public announcement requirements, as well as those of a physical nature such as door locks and Internet of Things (IoT) devices similar to those used in the hospitality industry. The common goal is to provide a safe space for the faculty and students while allowing for a usable pleasant environment for learning.

To build the Campus Networks today, the Three-Tier network model is very much in use, and as such, modularization of infrastructure for simplicity is preferred, with High Availability and redundancy very much designed into the common university networks. Within these Three-Tier networks, the focus again is on the speed of processing and sharing. Once a client is onboarded in the Access Layer, the emphasis is on providing whatever service they need quickly and most expediently with as few policies as possible. While this is the general rule, there is an emphasis in most universities to protect specific elements from threats, such as faculty devices, intellectual property, and data. While policies are reduced, there will always be protections in place that must be adhered to and enforced.

Because of the scale of most university networks, both in size and the number of endpoints, there is a need for the intent, which we would define as policies and access, to be automatically assigned as much as possible to reduce strains on support staff and to provide robustness of user and application experience. To deal with this in the current network, automation and orchestration platforms are beginning to be heavily used, along with centralized AAA services for policy. Integration points for these common services are typically housed within large data centers and shared as common services to the network as a whole.

As there are a varied number of multicast use cases that must all be addressed in a scalable way so that the main core need of fast network access to all for shared learning is not impeded.

Remote learning and access have been of importance as well in the past few years, and while the focus of this chapter is on the Campus Network environment, there are specific applications and network services that are used to address and facilitate remote learning, like video endpoints and meeting applications, which run on the student or faculty laptops or devices and which are crucial to providing an inclusive environment for all. To that end, the network must also be able to deal with the prioritization of applications across the network to ensure a robust user experience for these real-time applications.

The dormitories where students are housed on campus also have specific network needs. While universities try to ensure the student has access to all their learning needs, they are essentially the students' home away from home, and so the emphasis in these environments shifts from just the academic to more of a social network experience. In this setting, the network must accommodate the students' devices, including personal computers, TVs and platforms where streaming video is utilized, and even gaming environments. To this end, most universities outsource cable TV and keep these network requirements, but the rest must be accommodated.

In the following sections, we will dive into the typical design caveats that network architects must strive to overcome when adopting new technology and methods in which Cisco SD-Access inherently overcomes the challenges by design.

Challenges

University networks are shared platforms whereby once a client is authenticated, it is allowed to access the resources in the quickest and easiest method possible. The university network architecture's challenges are often unique compared to customers in other lines of business. The following sections outline several of the most critical capabilities required by University networks.

Fully Redundant Network Design

University network design has to be remarkably similar to most enterprise network designs which cannot afford to have any downtime. The network should be fully redundant to eliminate any single point of failure, resilient to meet requirements for high network availability, and follow guidelines of hierarchical network design.

Bring Your Own Device (BYOD)

In the past 10 years, wireless devices have become an integral part of personal and professional life. The education sector is no exception in making the learning environment more flexible for students. Bring Your Own Device (BYOD) trends also come with their fair share of challenges, caveats, and concerns. Two such challenges/concerns are:

1. Ungoverned devices accessing critical faculty resources
2. Profiling student endpoints to provide the right level of access

Silent Host

Silent hosts are devices, such as vending machines and third-party IoT devices, that do not announce their presence to the network, which poses a challenge in onboarding such clients onto the network. Modern security threats seek vulnerable entry points, such as these unknown or silent hosts, to exploit a network's valuable enterprise information. These devices must be identified and brought into the secure framework of the network.

Onboarding Clients with Poor Supplicants

As the education sector embarks on a digital transformation journey, they are adding increasingly more devices to their network. With the start of the new academic year, every new student brings in new devices which need network access on campus. Profiling plays a very key role in onboarding these new devices onto the network. The first step in securing devices is knowing what devices you have in your network. Hence the solution should be able to profile all the different device types which are getting connected to the network.

Isolate Research Partners from Main Campus

Universities typically have multiple partner organizations whose users are onsite within the university and need access to services or their own networks via the SD-Access environment. A simple use case that needs support is doctors who teach for medical schools by partnering with neighboring hospitals. In such a case, back-to-back networks extend various services through a Demilitarized Zone (DMZ) environment to the university, and guarded access must be provided within the network.

Guest Services with Firewall as Gateway

Universities abound with students and friends bringing all kinds of endpoints and all access levels into the network, including Guest access. In such scenarios, a next-generation network should be able to support traffic inspection for guest traffic which means the first hop for all guest traffic needs to be a firewall in DMZ.

Faculty, Students, and Building Management Systems

With exponential growth in the number of endpoints connecting to the network, there is a need for the IT Team to protect student and faculty information, research data, financial data, and so on. Since all of this information runs over the same network infrastructure, it is crucial to establish a segmented environment to prevent compromised clients in one group from impacting operations of another group.

Video Surveillance and Digital Signage

Video surveillance has been a key component of ensuring student and faculty safety and security in an educational institution. Video surveillance has become more important as security risk increases for universities that need to visually monitor or record events. Traditionally, several separate systems are deployed for disparate applications such as video surveillance, fire and smoke detection, and digital signage. As a result, IT administrators lack consistency, interoperability, and capabilities, translating into higher capital and operational cost. Hence the newer architecture should meet the use case requirements outlined above through the functionality of the IP network infrastructure and still enable secure access to live and recorded video.

Screen Sharing and Lecture Presentations

Universities tend to have a high proliferation of wireless endpoints for varied uses, from AirPrinting to screen sharing and lecture presentations. Networks should be capable of fulfilling the following requirements:

- Flexible enough to share students' creations with their classmates
- Enable faculty to share information to an Apple TV in a classroom, but not students
- Allow faculty to temporarily allow students to project their work in the classroom

- Support an IT policy regarding which printers can be used by students and which cannot.
- Advertise local service availability based on location

Wireless Challenges

High-density Wireless

A university has a lot of buildings spread out over a larger area of land, and there are a couple of challenges to be met by wireless networks. Students are always moving across campus and do not just want access inside the buildings, so they will also need coverage outdoors. Other than this coverage, there may be many students in an area for a single AP to manage, such as a lecture hall or auditorium. Density is even more of a concern when we think about modern/smart education using technology. These challenges are demanding and could cause dissatisfaction among students and faculties.

eduroam (Education Roaming) Support

University students and faculty members need instant access to networks to access research materials over free, secure, reliable Wi-Fi wherever they travel or study. Users should be able to connect to the eduroam Wi-Fi network anywhere in the world and use their home institution credentials for authentication.

Guest Services Using Multisite Remote Border

Universities must be able to support thousands of students who are all simultaneously connecting to the network. Besides the students and faculty, guest users also need to access the network. Guest access is needed for visitors, event attendees, and so on.

Customer Solutions

There are many use cases that SD-Access can address within the University vertical. In a large campus deployment, the most important functions that the Campus Network needs to solve are connectivity, user, and application experience. To reiterate, students typically are not tolerant of network latency or disruptions to access, and some of these events have been covered in the news media.

SD-Access Design and Automation

SD-Access by design inherently includes the capabilities to address an available and resilient network. It is important to understand that the ability to build a network based on roles allows for the high availability of those discrete roles within the Fabric, as an example. Border Nodes are built out redundantly to not have a single point of failure.

Figure 6.2: SD-Access for Distributed Campus

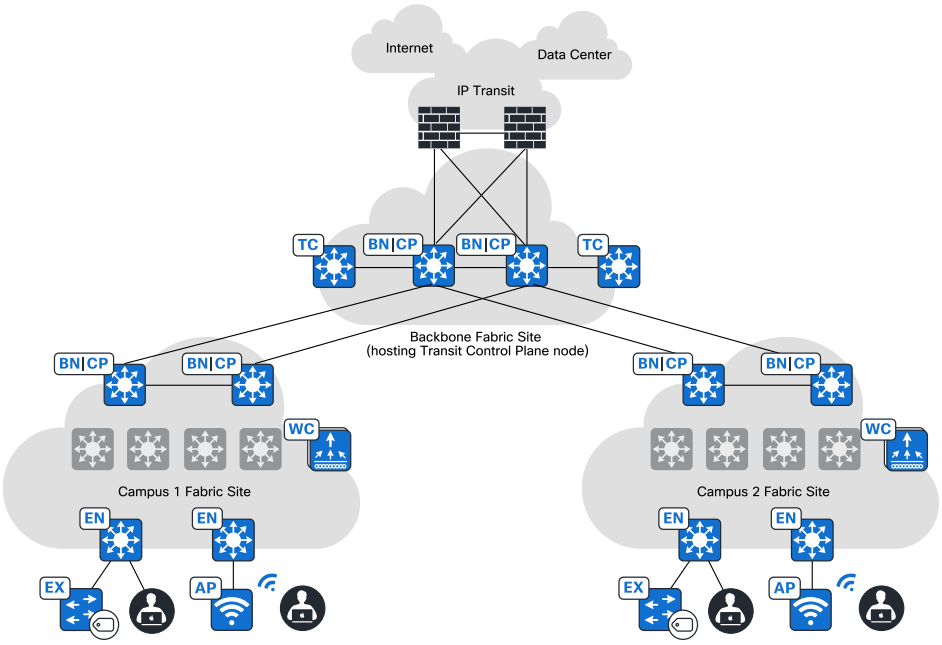

Additionally, those nodes are all routed at Layer 3 from an Underlay perspective. This allows for Equal-cost Multipath (ECMP) and the ability to service network load from a bandwidth perspective over multiple diverse links. Considering the load on a Campus Network with hundreds of thousands of endpoints, the ability to offload data throughput over multiple optics, fiber, transceivers, and ASICs allows us to address the latency and load of the network. Border Nodes must carry the ingress traffic load to transit traffic and the route prefixes from the environment into which they connect. To that end, use appropriate hardware within the SD-Access Compatibility Matrix and design for both bandwidth and prefix concerns, as we have multiple switch and router platforms that can accommodate the requirements.

From a Control Plane Node perspective, when considering redundancy, high availability, and load, we must consider the number of Endpoint Identifiers (EIDs) within the space. Separation of Fabric roles and responsibilities is a capability of SD-Access, so while the border handoff is being managed for various use cases, we can further separate the Control Plane Node from the Border Node and have that function on a device that can carry the table sizes for the specific overall design.

Closer to the Client from an access perspective, while we want to address a flexible mobile environment as much as possible, we have to build the network to suit the physical environment of the university. To that end, we must conform to the nature of the physical building where the users are located. Sometimes we have more Access Points than switch ports, or the cable plant would exceed the IEEE UTP specification. In those cases, we need additional access switches to support the large wireless environments with hundreds of Access Points in large buildings. Due to the nature of cabling and location of the switching, we will need multiple distribution switches to address the high availability and resilience concerns where we want to have these Edge Nodes multi-homed. In situations like this, we can use the Intermediate Node, which inherently does not participate in the Fabric.

Multicast at Scale in the University Campus

Due to technology changes and the proliferation of new capabilities within the client endpoints that users are bringing into the network, Cisco SD-Access must address multicast at scale within the university space. The largest concern in the last couple of years has been the multicast environment and how the changes to various endpoints by certain vendors have impacted the University vertical.

Figure 6.3: ASM design with External RP

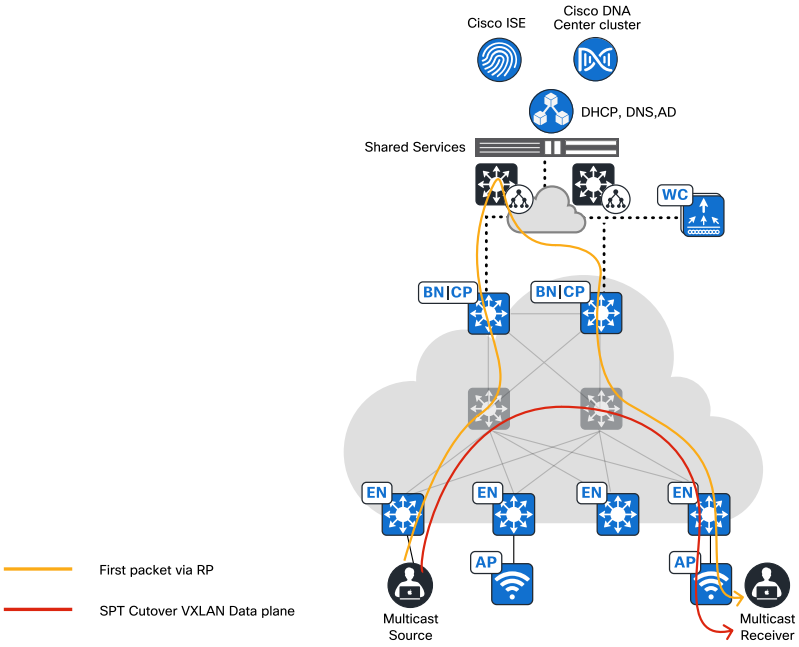

In the SD-Access Fabric, multicast forwarding still uses Any-Source Multicast (ASM) design aspects. The Rendezvous Point (RP) outside the Fabric is selected within each Virtual Network as the RP the clients should use to join the tree. The concepts of Shared Tree and Source Tree are still in operation.

Wide Area Bonjour (WAB) and Scale in the University Campus

As you may recall, the depth and breadth of clients, which can now function as screen-sharing devices, have multiplied exponentially in the past few

years, with various major vendors providing the capability of screen sharing on laptops and mobile devices such as tablets. As a result, we have multiplied the number of multicast sources, which in turn has put pressure on the table sizes within the network. To rein that in, we need to have a way of controlling the multicast environment while allowing for us to control which devices may be sources and which devices may be receivers. To this end we can incorporate two methods of operating:

1. Ensure our multicast network is designed and built with resilience and connected to Rendezvous Points in an available manner, using ACLs to limit stream subscription.
2. Leverage Wide Area Bonjour (WAB) capabilities within Cisco DNA Center to appropriately control and scale the traffic within the network.

As the clients will be in virtual networks, the multicast traffic for those clients will be carried in the Overlay within VXLAN.

While deploying Wide Area Bonjour services in a network, it is important to understand the concepts involved and understand how to scale the solution for the end-state. The Cisco IOS software on network devices introduces new and advanced Wide Area Service Discovery Gateway (SDG) functionality that performs the Agent role in the overall solution.

The SDG gateway switch provides a single gateway solution at the LAN or wireless distribution block in the controller-less Bonjour solution. The SDG switch communicates with Multicast DNS (mDNS) clients to build and manage the service information.

Figure 6.4: Wide Area Bonjour reference topology

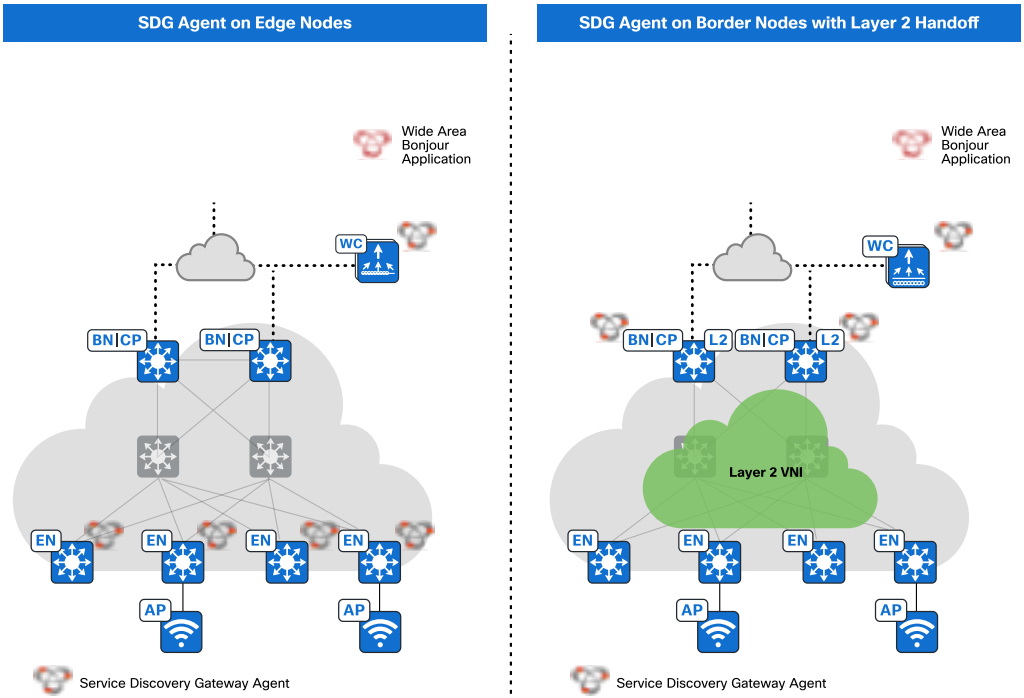

An important consideration while deploying WAB services is the support of 150 TCP sessions. Depending on how many Edge Node devices are in the Fabric, attention will need to be paid to how we support the WAB service. If we need to scale for large Layer 2 environments, then one option is converging the SDG functionality at a Layer 2 Border Node and using a Layer 2 Pool supported across a Virtual Network. Alternatively, if the Edge Node counts across Campus are lower, then we can support this with SDG at the Routed Access Edge Node via Layer 3.

It is important to understand first the scope of the problem, from a client scale perspective, the number of nodes in the network, and then to adjust your design to accommodate the caveats.

SD-Access and Meraki Wireless

Many organizations, including universities, have deployed Cisco Meraki cloud-managed wireless for their Campus Networks. While we could roll out SD-Access for both wired and wireless, we may also need to augment the network with various equipment to be more flexible with the SD-Access architecture. Cisco Meraki wireless differs from traditional OTT technology, where data plane traffic terminates into Fabric Overlays for communication with the Fabric instead of being tunneled outside of the Fabric. Cisco Meraki APs use an 802.1Q trunk with management traffic being untagged and mapped to the AP Management VLAN, and user-tagged traffic maps to a unique VLAN on the switch for wireless endpoint, maintaining separation of traffic.

Figure 6.5: Cisco Meraki Control and data integration

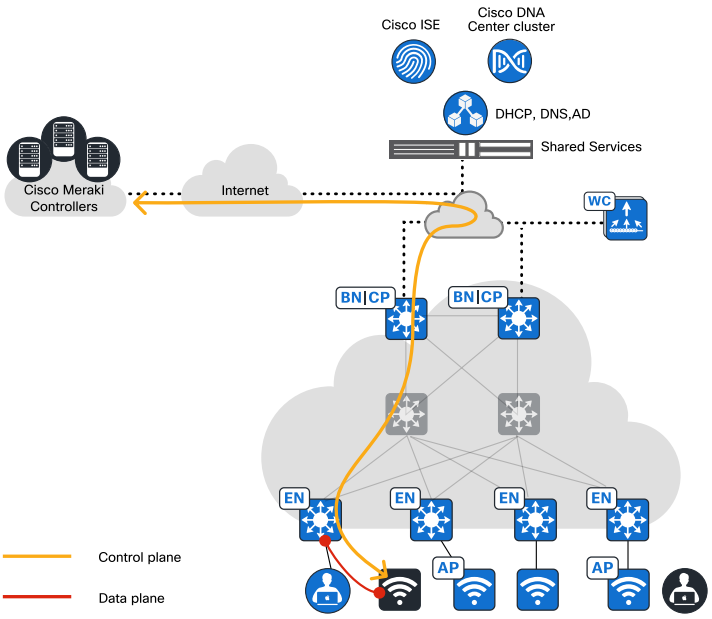

SD-Access Macro-segmentation

There are multiple departments and varied clients and use cases within universities. The nature of the university space is that they can bring any device at any time and connect it to the network if it aids the student and faculty in achieving the goal of collaboration. In a collaborative environment, it is important to deal directly with these varied clients securely. SD-Access, by design, incorporates the ideas of macro-segmentation and micro-segmentation. Within this section, we will discuss areas where macro-segmentation can be leveraged. While this is not an authoritative list, it is here to show you the flexibility of the design and to allow us to extrapolate beyond this book to solve varied use cases.

Figure 6.6: Macro-segmentation via Virtual Networks

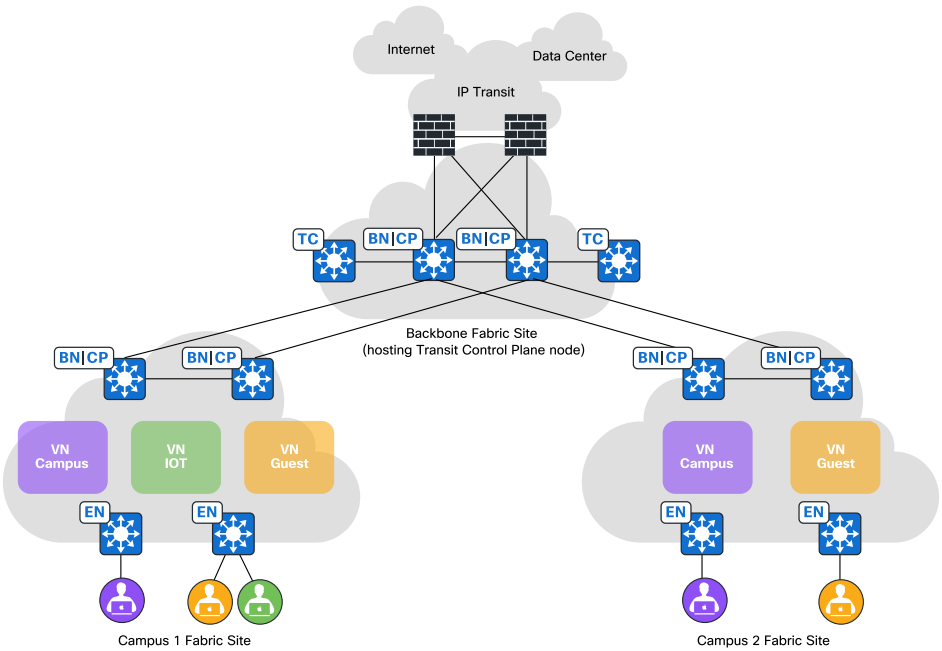

Universities have many personal devices which are brought onto campus. As a solution, Cisco SD-Access allows us to onboard these specific devices via Bring Your Own Device (BYOD) workflows. Faculty who are university employees may have phones or other devices which need to be controlled for security reasons by a Mobile Device Manager (MDM) to make specific applications of supporting services available to help them facilitate learning within the university. To that end, those typical flows may ensure that the device is appropriately categorized by the network when onboarding and tracked for anti-virus signatures and anomalies to ensure that it conforms to the university policies. These categorizations can be used within the policies to force a Change of Authorizations (CoA) to occur in the event of a breach of policy or virus infection. Additionally, the client may be placed in a

separate virtual network entirely from the student population, thus achieving the aim of separating student and faculty traffic. This is one method of dealing with the situation, and later in the chapter we will see how to deal with this challenge in a different way.

As clients are brought onto the network within the university space, we want to allow as much traffic as we can for collaboration. As a result, a default Permit List methodology is utilized. Using a default Permit policy, we would create rules that deny behavior we do not want while everything else would be permitted.

As with any network, there are good operating systems and poor ones, and the wireless supplicant which is used on a device may be robust but in some cases is not. A good example of this is personal devices students bring on campus, such as Microsoft Xbox or other gaming systems. In situations like these, we ensure that we profile the device and, from categorization, can use endpoint logical policies to ensure the device is onboarded securely, but we" with "securely. We can also utilize a BYOD portal to allow the user to add their devices to the network manually. This allows for not only the endpoint to be onboarded appropriately but also for the profiling to ensure that type of endpoint ends up in the correct segment of the network.

Lastly, we may need to segment partner organizations from each other while allowing them to utilize the network for transit to their own networks. An example of this is a university medical school that is partnered with one or more large healthcare providers. In situations like this, a back-to-back service provider-type connection is made between the existing organizations to facilitate the exchange of AAA traffic and data traffic. These two flows may be logically separated from one another. In some designs, a Layer 2 segment is used very much like a handoff in Carrier Supporting Carrier (CSC) networks. This design uses pure IP to forward traffic and connectivity, limiting route exchange.

It is, therefore, possible to connect a doctor who teaches at the university to the Campus Network using his own credentials and segment based on university policies by either proxying RADIUS across the Business-to-

Business (B2B) link or by direct integration with the foreign Identity Store. This also can be utilized in reverse within the hospitals for student interns.

SD-Access Micro-segmentation

Within universities, there are times when we will want the various clients to be contained within the same macro-segment but, at the same time, micro-segmented from each other. In Cisco SD-Access, in addition to providing the flexibility of using different subnets, we provide the flexibility of micro-segmentation, i.e., using the same subnet in a more user and endpoint-centric approach. This allows you to onboard clients into different groups within the same Virtual Network while controlling access between those groups. In this way you have achieved segmentation within segmentation, or what we call micro-segmentation.

Figure 6.7: Micro-segmentation with Enforcement

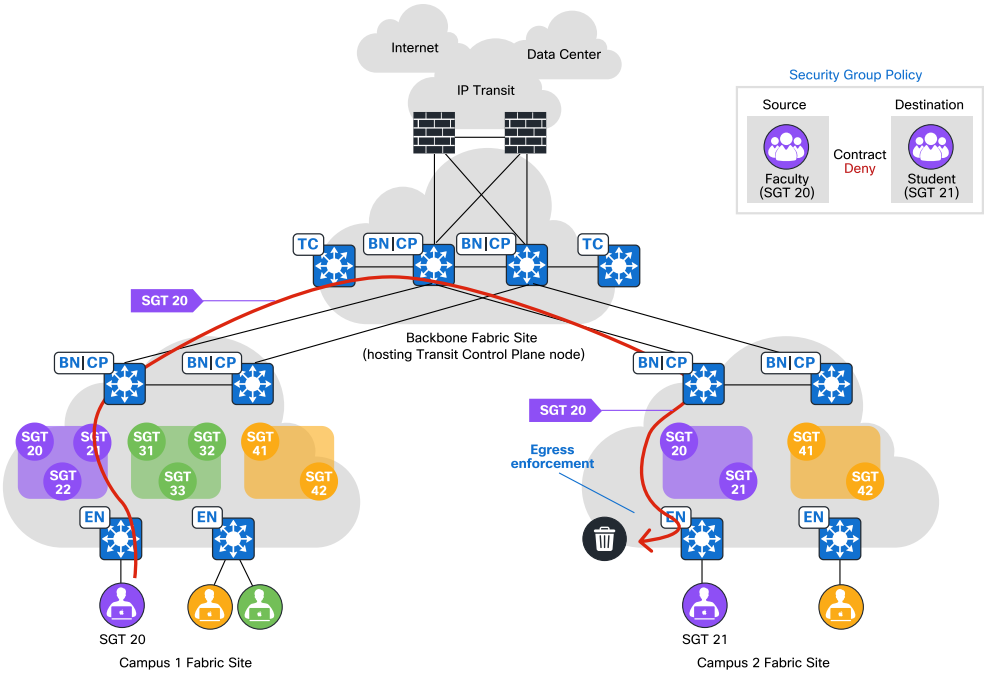

Micro-segmentation, in this case, would be delivered by SGTs assigned via Cisco ISE at authorization, and an egress policy in the form of a Security Group Access Control List (SGACL) would enforce and restrict client-to-client traffic. The SGACL is programmed into the Fabric when clients are authorized by Cisco ISE via 802.1x or MAB, which enables the dynamic nature of the network. It also allows the network to use only those policies it needs when clients exist, freeing up valuable resources for forwarding traffic when they are not needed.

There are various other cases where micro-segmentation may be required to secure specific streams from other traffic. Those may include, but are not limited to, video surveillance feed from wireless cameras as well as a

dashboard or digital signage devices connected wired or wirelessly for display requirements at administrative buildings, large campus buildings, dormitories, or catering locations. These devices may have critical feeds in the event of security incidents, which may be utilized to help inform the student body of a specific event.

Cisco SD-Access Supporting Guest Services

University campus administrators often need to manage extensive guest services with a common subnet across all their campus sites. To address this challenge, SD-Access provides the Multisite Remote Border solution using Virtual Network Anchors. This solution allows traffic from a given Virtual Network at multiple dispersed sites to be aggregated back to a central location, sometimes referred to as an Anchor Site, where it uses a single common subnet rather than having to define and use per-site subnets for the guest Virtual Networks.

With a simplified and centralized common subnet structure, Virtual Network Anchor Sites significantly simplify the guest service deployments across sites and provide consistent and secure segmentation for guest traffic in university environments.

Using anchored services, traffic for the endpoints that belong to the Anchored Virtual Network at each site are aggregated and tunneled back to the remote anchor border at the Anchor Site over VXLAN. An Anchor Site functions very much like a traditional Fabric Site, but it forms a virtual Fabric Site serving a particular Virtual Network.

This virtual Fabric Site has its own site Border and Control Plane Nodes in the Anchor Site. What is special about the Anchor Site is that its edges and wireless controllers are dispersed across multiple Fabric Sites, referred to as anchoring sites.

Figure 6.8: Guest Traffic Flow with Anchored Virtual Networks

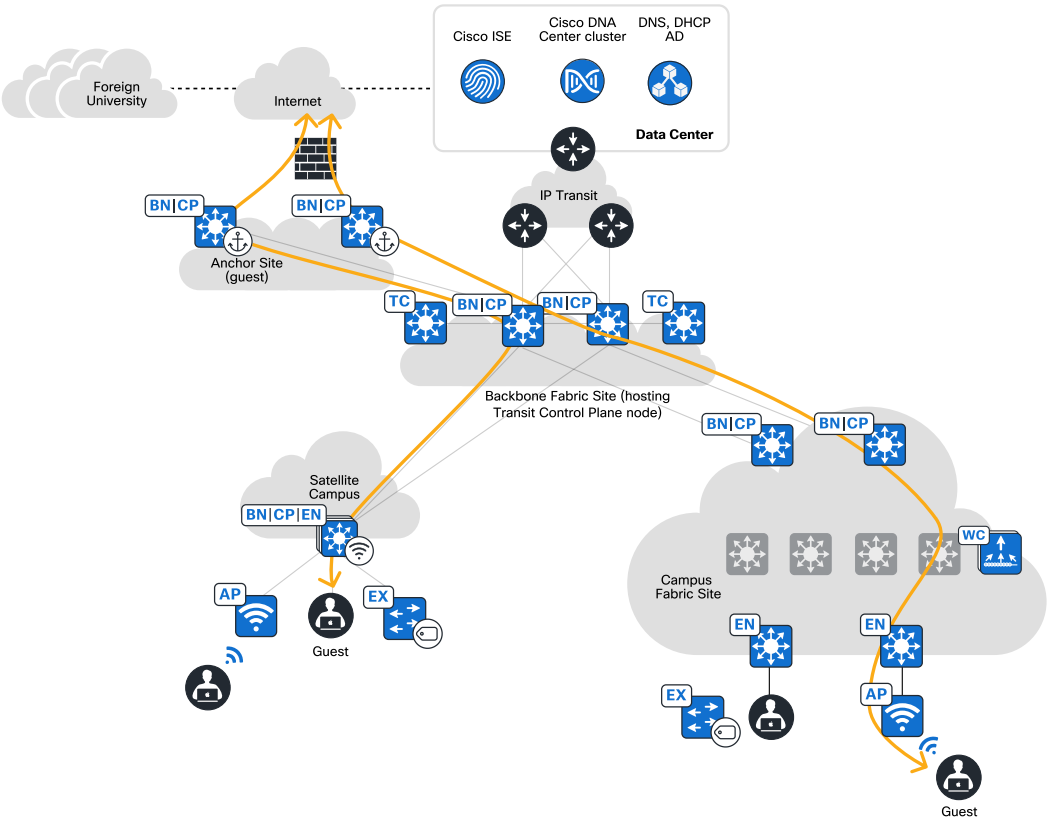

Multisite Remote Border is enabled on a per-Virtual Network basis. For an Anchored Virtual Network, all edges in the anchoring sites use the anchor Border and Control Plane Nodes for data plane and control communication. Wireless controllers in the anchoring sites communicate with the Anchor Control Plane Node for wireless endpoint registration specific to the Anchored Virtual Network.

Since the anchor border reachability may traverse multiple IP networks, special attention must be paid to the MTU across the entire path to

accommodate the VXLAN header overhead of 50 bytes.

Anchored Virtual Network is configured to use the Anchor Site. After a guest endpoint joins the guest SSID and passes the Central Web Authentication using Cisco ISE, it is associated with the anchored guest Virtual Network. Guest traffic is tunneled to the Anchor Site Border Node and egresses to the Internet through a firewall.

Additionally, while the guest traffic is VXLAN-encapsulated and passes through the Fabric, the first hop or gateway for guest traffic can be outside the Fabric and connected at Layer 2 to the Multisite Remote Border. Such a device can be a firewall for inspection purposes, should it be a design requirement.

Cisco SD-Access for Shared Services

Typically, there are shared services within every network which are housed or connected to a data center. Each case depends on policy and authorization regarding who or what is entitled to access a specific resource. Every network deployment has a common set of resources needed by every endpoint:

- Identity services (e.g., AAA/RADIUS)
- Domain name services (DNS)
- Dynamic host configuration protocol (DHCP)
- IP address management (IPAM)
- Data collectors (e.g., Netflow, Syslog)
- Internet Access
- IP voice/video collaboration services
- Other infrastructure elements

These common resources are often called "shared services." These shared services will generally reside outside of the SD-Access Fabric and must be

deployed correctly to preserve the isolation between different Virtual Networks accessing those services. In most cases, such services reside in the global routing table (GRT) of the existing network OR may be in a specific Virtual Network or VRF.

Figure 6.9: Shared Services access

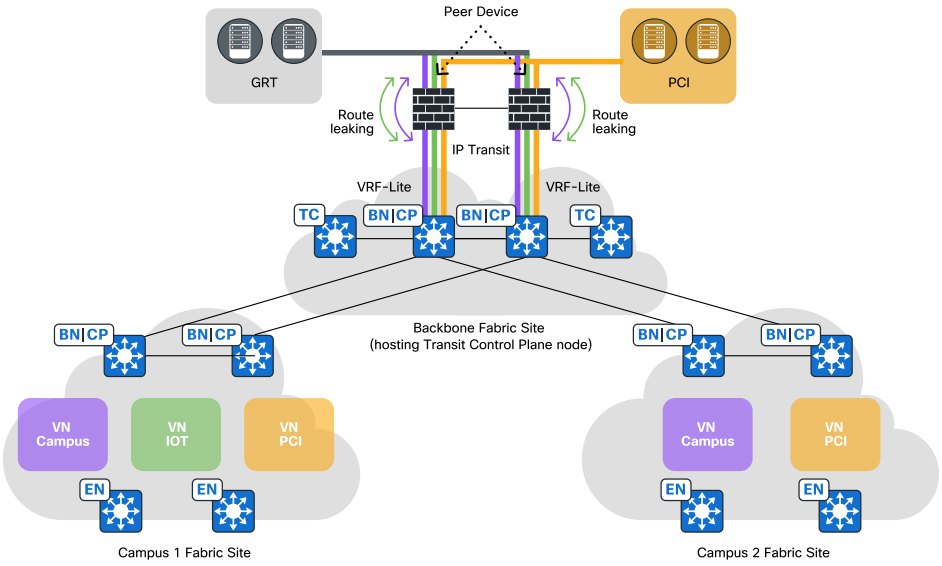

Fabric clients in the SD-Access network operate in Overlay Virtual Networks. Thus, if the shared services are part of the global routing space or a specific Virtual Network or VRF, some method of Inter-VRF route leaking is required.

There are two ways to do inter-VRF route leaking:

- Use of a Peer (*Fusion*) device (Switch/Router/Firewall)
- Use of the SD-Access Extranet feature

The use of a VRF-Aware Peer (*Fusion*) device directly attached to the Fabric Border Node provides a mechanism for route leaking of shared services

prefixes across multiple networks. The use of firewalls provides an additional layer of security and monitoring of traffic between Virtual Networks or between Virtual Network and GRT.

If firewall inspection is not a requirement, then SD-Access Extranet is an innovative way to access shared services from other VRFs without requiring a peer device for route-leaking. For more information about SD-Access Extranet, please refer to that section in the Architecture Overview chapter.

Wireless Solutions

Fabric-enabled Wireless and Over-the-Top Wireless Design

In this day and age, almost no one physically connects their personal devices to a network unless they have to, and so wireless networks must be highly resilient and easily available. From a wireless perspective, when discussing the Fabric, the wireless Access Point (AP) will be dispersed amongst various line cards of chassis-based switches or across various members of a StackWise deployment. This ensures a higher level of resiliency within the wireless environment, allowing not only for the switch network to be upgraded but also to spread the load of the wireless users across multiple switches ASICs to facilitate client load and bandwidth concerns. To take this even further, you can also utilize power priority on the switches so that in the event of power starvation, specific APs will go down to create a brownout scenario instead of a complete blackout.

There are multiple ways to architect a wireless solution within a Fabric. The first and leading primary method would be to utilize SD-Access Wireless which allows for the handoff of client data traffic from the AP directly into the Fabric by encapsulating client traffic within VXLAN on the AP itself. A secondary supported method is Over-the-Top (OTT) Wireless, whereby the client traffic (data) is encapsulated within CAPWAP with the management (Control) traffic and sent to the Wireless LAN Controller (WLC) for forwarding. With the SD-Access solution, you have the flexibility to have a mix of SD-Access Wireless and OTT SSIDs if needed.

Figure 6.10: SD-Access Wireless/OTT Wireless Architecture

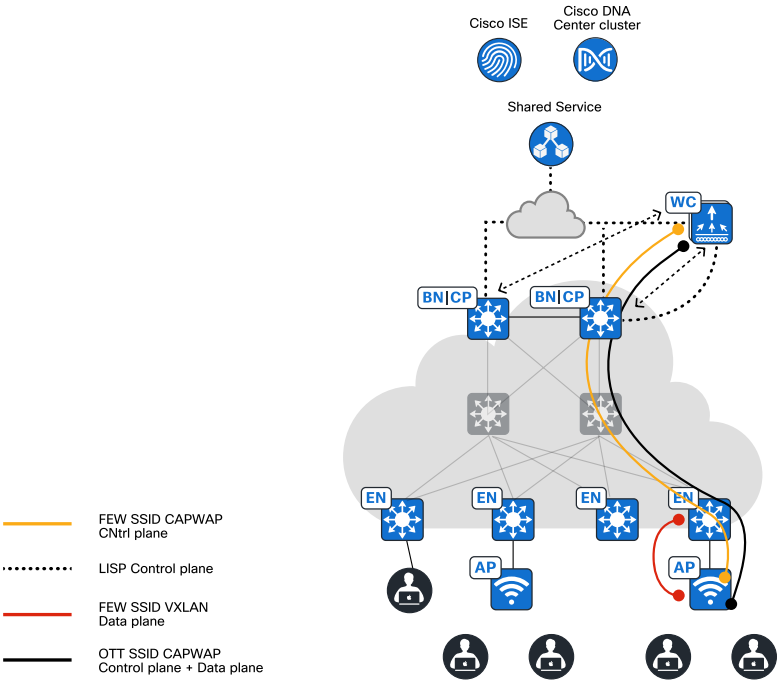

Cisco SD-Access again helps facilitate these large networks by delivering automation as a foundational core component, allowing for the quick deployment of new hardware with the same policy and configuration. This helps from a deployment perspective and prevents misconfiguration as the network configuration is the same from switch to switch and AP to AP across the network. In large universities with hundreds of Edge Nodes and thousands of APs, this operational consistency can save the network team countless hours of troubleshooting due to human error in a configuration.

Cisco SD-Access Support for eduroam

For universities supporting the eduroam solution, there is a need to be able to accommodate wireless connectivity within the guest environment across the SD-Access Fabric. To expand on this specific use case, SD-Access can be used to extend wireless services to any visitor from other universities as long as that visitor has an eduroam account.

Figure 6.11: eduroam Authentication Workflow

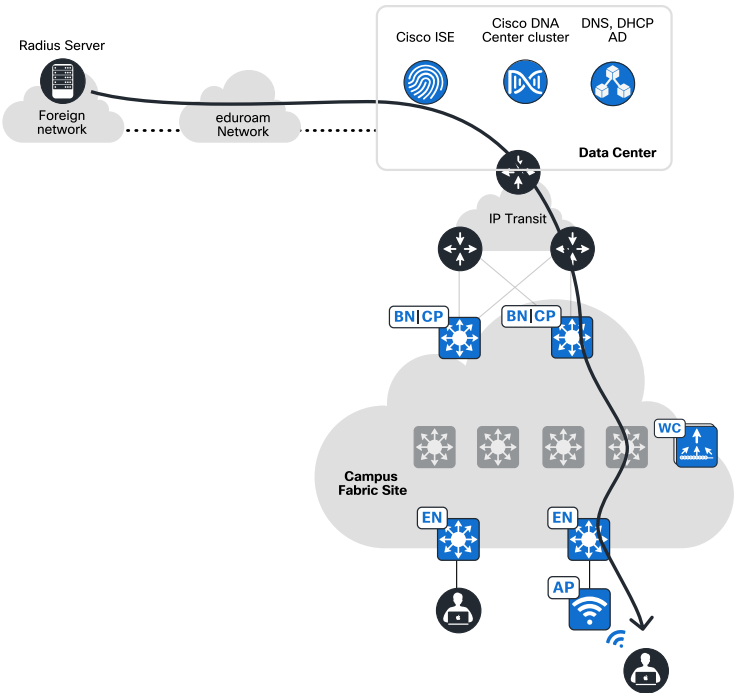

Anyone with an eduroam account can connect to the eduroam SD-Access Wireless SSID, which was designed by large universities to address cross-university authentication for university-to-university roaming. In this scenario, we would anchor the eduroam client data traffic to a set of Border Nodes and Control Plane Nodes utilizing the Multisite Remote Border (see Figure

6.8) within the DMZ. Wireless clients would utilize the eduroam infrastructure, located off-campus in most cases, for client authentication as shown in Figure 6.11. In this scenario, the client's data session is terminated on the Guest Anchor Border Node and not the Foreign Controller.

The authentication request is proxied to the eduroam servers, which in turn forwards the request to the user's home institution. If the authentication is successful, the foreign RADIUS server sends an "Access-Accept" response to the eduroam server, and the eduroam server forwards this response to the university where the user is currently located. The user is authenticated and gains network access at the foreign university.

Stadiums, Auditoriums, and High-Density Wireless:

On university campuses, there are often areas of extremely high-density wireless. These can be large auditoriums, used for conferences and large classes, or even massive stadiums used for sports such as football. To that end, the university needs to be able to ensure the availability of wireless connectivity given the massive number of wireless clients present during these events.

Cisco SD-Access is ideally suited to deal with a high-density wireless environment. The wireless AP control traffic is highly available to the WLC, allowing for optimal control of the WI-FI channels and RF spectrum and delivering a better user experience for wireless clients. The distributed data plane capability of SD-Access Wireless means that there are no bandwidth bottlenecks for the wireless clients, even if there are thousands online all at once.

Deployment Options

All universities range in size, but it is typical to have large numbers of endpoints supported across a Campus Network. For the academic year of 2023 alone, enrollment at some universities has a combined student population of over 50,000 for both undergraduate and graduate students. If each student brings only two devices on campus, that would be conservatively 100,000 endpoints. If you peel back the onion with all the IoT devices, faculty devices, printers, scanners, and other devices on the network, you can hit the scale limits of a single standalone DNA Center XL appliance easily.

As a result, it is common for universities to split roles for various parts of the network into various clusters, allowing Cisco SD-Access to be supported in multiple environments. A large central campus environment is on one Cisco DNA Center cluster, while multiple smaller satellite campus locations are on another cluster. In such cases, Multiple Cisco DNA Center to single ISE support is required to meet the scale requirements of the University network. In addition to meeting the scale required, the Multiple Cisco DNA Center to single ISE deployment maintains security policy consistency by assigning one Cisco DNA Center cluster with the policy author role while the other cluster is in a read-only role.

Cisco DNA Center offers both High Availability and Disaster Recovery options which can be used independently or in conjunction with one another. One or both of these options are recommended in any network where SD-Access is deployed. Which one you use depends on your organization's resilience requirements regarding network management and analytics.

While we have used high-scale universities with potential geographic dispersion as the example in this section, it is certainly possible that other smaller universities will not need these same kinds of capabilities. Always make sure to understand your scale requirements at the network and management levels and adjust accordingly.

Design Best Practices for Universities

When considering a network design for a university, we recommend that you follow the following design principles and guidelines:

- Design with scale in mind at every layer of the environment, including the management layer
- Select the appropriate hardware platform to support the expected load, providing overhead for growth
- Understand the required hardware table sizes needed at each layer to support the features required, providing overhead for growth
- Keep in mind the previously mentioned design aspects concerning a Cisco SD-Access architecture

Summary

We have seen that University environments are unique due to their challenging use cases related to their goal of providing the ultimate sharing and collaborative experience in pursuing higher learning. When designing their networks, they must be extra vigilant to ensure that they account for the sheer scale of their environment, and where possible, they should limit the size of fault domains to ensure that convergence and operational resiliency are maintained at the highest levels. While we have covered a number of the various design caveats and use cases in this chapter, you might encounter a similar use case in another vertical. If you encounter such a use case, please refer to that other vertical chapter.

In this chapter, we showed how Cisco SD-Access provides a secure, dynamic, resilient solution that addresses the most demanding challenges faced by universities. With the addition of Cisco DNA Center capabilities such as Endpoint Analytics and Ecosystem partner applications such as Cisco DNA Spaces and eduroam, Cisco SD-Access ensures that universities can continue to deliver the shared network environment most conducive to learning that both students and faculty alike have enjoyed for many years.

Cisco SD-Access for Financial Customers

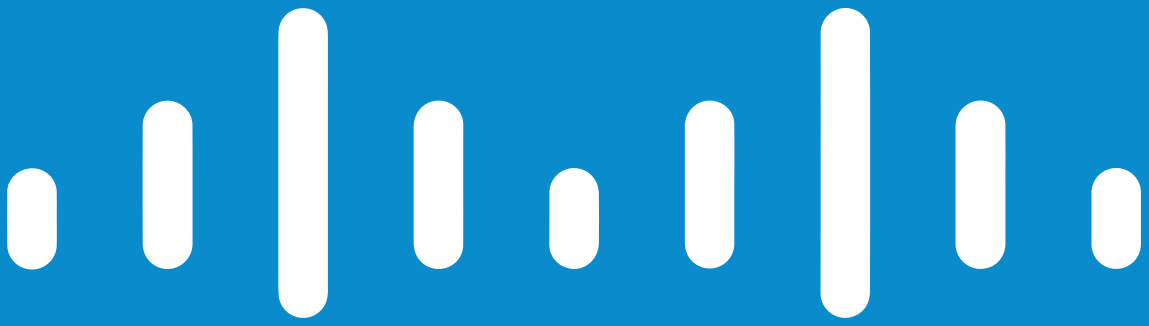

Introduction

Financial organizations such as retail banks, credit unions, insurance companies, and capital firms are shifting to digital network architectures to securely build and operate 1000s of locations across the world. While each location may have special requirements, the financial vertical needs standardized and secure network connectivity, simplifying network operations and maintenance, failure-resilient systems, and consistent policy implementation across the entire organization.

This chapter provides design guidance for the financial space, focusing on how the new innovations (especially Cisco SD-Access) are being utilized in this industry to have a simple, secure, and flexible network. The chapter starts with a description of the existing Financial Network architecture, followed by challenges and solutions for financial organizations. Finally, the chapter ends by reviewing some standard deployment options for financials.

Challenges

Financial Networks are complex and support wired, wireless, WAN, and security infrastructures deployed across a country, multiple countries, or worldwide. Today the network infrastructure in these organizations can be broadly divided into Corporate Headquarters (HQ) and Remote Branch business segments. The network requirements can vary depending on the network segments to support various business-critical use cases.

HQ sites are corporate buildings connected by Metropolitan Area Network (MAN) for administrative buildings and by Wide Area Network (WAN) for remote site connectivity.

Figure 7.1: Traditional network

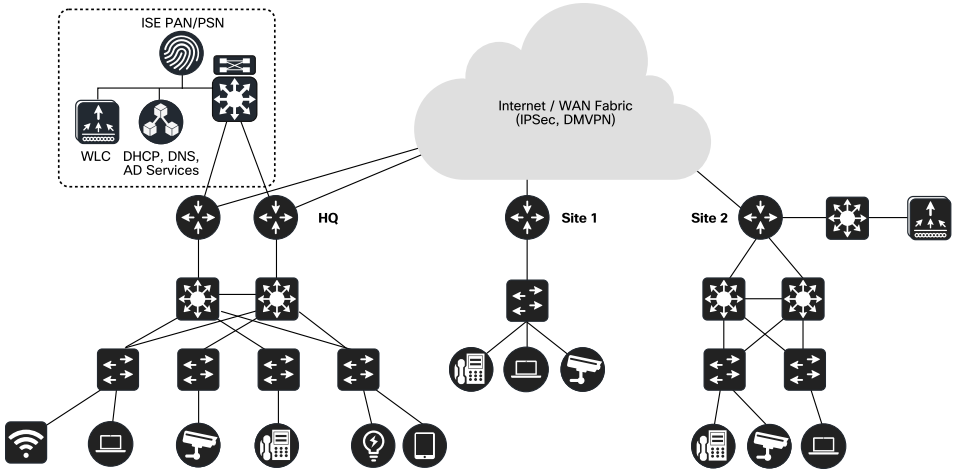

The topology shown in Figure 7.1 depicts a generic traditional network for Financials with Core, Distribution, and Access Layer switches in the Head Quarters (HQ) site serving wired, wireless, and IoT endpoints. This network provides Internet, data center, and WAN connectivity for the endpoints connecting to it.

The Branch locations in the topology, Site1 and Site 2, have different scale requirements by embracing either a Two-Tier, Three-Tier, or single device network hierarchy for the LAN Infrastructure.

Requirements in Financial Customers

Financial Networks are usually geographically dispersed, complex networks with unique challenges due to the high levels of security, resilience, and regulatory compliance required in the financial industry. The following sections outline several of the most critical capabilities required by financial organizations when deploying a network.

Security Challenges

Cybersecurity threats are the biggest threat that keeps a financial company's Chief Information Security Officer (CISO) awake at night. The rapid transition to hybrid work and the evolution of digital business services to customers have massively increased attack surface vectors available to cyber criminals. If not kept in check, these bad actors can take advantage of vulnerabilities and cause massive losses in both financial and reputational terms. The CISO group periodically reviews basic security fundamentals and security processes.

Compliance Mandates

Financial systems need to protect highly sensitive financial records and information of customers, and they are strictly bound by government regulations. For example, the PCI-DSS (Payment Card Industry) standards

include mandates such as data encryption in flight, security mandates for storing customer data, and tracking and monitoring network resources and cardholder data. Financial Networks must provide complete and constant visibility for network management and monitoring.

Finally, given that different departments and guests all share the same network infrastructure, every group must be isolated from one another and restricted to only the resources they are permitted to access. At the same time, these diverse groups of users and devices will need access to shared services.

Availability (Five-Nines)

Due to the criticality of financial systems in our daily lives, 100% availability is the goal. A five-nines availability comes close to 100% availability and guarantees that the network will be up for 99.999% of the time (roughly 5 minutes and 16 seconds of outage per year). Automation, Monitoring, Load-balancing, and failover schemes can allow financial firms to meet or beat the goal of five-nines availability.

Large Scale Multisite Deployments

Financials are typically large multisite deployments often running thousands of sites across vast geographic regions. It is an enormous challenge to deploy and manage such large networks box-by-box or site-by-site with onsite network management teams. If you ask any network engineer in a Financial, they will tell you that automation is the need of the hour to be able to do complex site deployments with minimal manual work.

Complete Isolation of Guest Users

Customers visiting financial branches may bring different endpoints that could be compromised, so the traffic from such guest devices must be isolated and inspected property. In such scenarios, the next-generation network should be able to support traffic inspection for wired and wireless

guest traffic, which means the first hop for all guest traffic needs to be a firewall in a demilitarized zone (DMZ).

Centralized Policy Management

With exponential growth in the number of endpoints connecting to the network and the spanning of large financial organizations across the globe, there is a need to manage security policy in different geographical regions. This leads to management complexity as the rules may be driven by the local laws. There is a crucial need to simplify the grouping of users and devices so that security policies can be managed intuitively.

High Sensitivity to Quality of Service (QoS)

In addition to worrying about security, compliance, and availability, a slow and QoS-disparate network can lead to poor customer satisfaction and monetary losses. When it comes to trading floors, which are highly sensitive to delays, low latency and consistent QoS are mandatory to fulfill the needs of the organization.

Acquisitions Integration

To reap the full synergy and benefits of acquiring an organization, the systems of the two merging organizations must be integrated to optimize redundant operations. The network is usually the first area of integration, but this has to be done securely so that the organization is not exposed to new vulnerabilities.

Ticker Streaming and Video applications

As the trading business for financial organizations has grown rapidly, working with real-time information has become a primary focus area. Financial traders need access to real-time data to remain competitive in the marketplace. Banks want to showcase to visiting customers the various programs and strategies available for them by utilizing IPTV-based advertisements and videos constantly running in all financial branches.

Lean IT staff

Given the global footprint of large Financials, a solution is required so that any location can be brought up faster and managed remotely so that financial organizations can maintain a lean IT staff. The requirement is to have automation and assurance in the network to reduce the complexity and time spent for both deployments and troubleshooting.

Customer Solutions

Cisco SD-Access

This section considers many specific use cases that are solved with Cisco SD-Access solutions automated from Cisco DNA Center. The Cisco SD-Access application runs on the Cisco DNA Center controller for designing, provisioning, applying policy, and providing the campus wired and wireless network with the context that enables an Intent-based intuitive network.

Fabric is an integral part of SD-Access and introduces Overlays that enable easy-to-deploy network virtualization across the wired and wireless campus. In addition to network virtualization, Fabric provides software-defined segmentation and policy enforcement based on user identity and group membership. Software-defined segmentation is seamlessly integrated using Cisco Identity Service Engine (ISE), providing micro-segmentation through the use of security groups within a virtual network.

Cisco DNA Center automates the creation of Virtual Networks, reducing operational expenses and risk with integrated security and improved network performance provided by assurance and analytics capabilities.

Security Challenges

Macro-segmentation

One of the preferred segmentation methods for Financial customers is to isolate the user endpoints (CORP, GUEST, IoT) by placing them in different virtual routing and forwarding instances (VRFs). Cisco SD-Access offers the flexibility for macro-segmentation of endpoints in different VRFs, all of which can be provisioned in the network from Cisco DNA Center.

Micro-segmentation

Within the scope of a single VRF, customers tend to have further segmentation needs for use cases such as:

- Placing loan department and credit card department in different security groups
- Placing video surveillance of branches and IoT devices like printers in different security groups

For such requirements, in the traditional network architecture, the only means to segment was by placing groups in different subnets enforced by IP ACLs. In Cisco SD-Access, in addition to providing the flexibility of using different subnets, we provide the flexibility of micro-segmentation, i.e., using the same subnet in a more user and endpoint-centric approach.

Referring to the loan and credit card department example, each group can still be placed in the same subnet. However, by leveraging dynamic authorization, they can be assigned different Security Group Tags (SGTs) by ISE based on their authentication credentials. Traffic between or among these groups can then be enforced by Security-Group Access Control Lists (SGACLs) based on group membership instead of endpoint IP address.

Figure 7.2: Macro-segmentation and Micro-segmentation

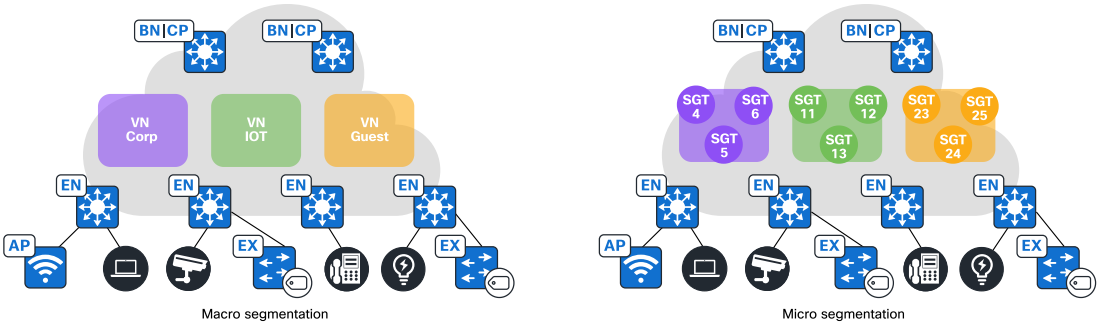

Secure Onboarding

Device Onboarding

In Financial Networks, Zero Trust is finding more and more adoption. It is crucial to securely onboard users and network devices. Cisco SD-Access can force network devices such as APs and switches to be securely onboarded using IEEE 802.1x mechanisms. This protects the network from the attachment of unauthorized devices by maintaining closed authentication on all Edge Node access ports. Extended Nodes onboarded securely using closed authentication are called Supplicant-Based Extended Nodes (SBEN).

Supplicant-Based Extended Nodes (SBEN) are provisioned as Policy Extended Nodes by Cisco DNA Center to have a supplicant with EAP-TLS authentication on their uplink to the Edge Node. The EAP-TLS certificate is provisioned by Cisco DNA Center using Cisco DNA Center Certificate Authority (CA). After successful onboarding, access to the port is purely based on authentication status. If the device/port goes down, the authentication session is cleared, and traffic is not allowed on the port. When the port comes back, it goes through dot1x authentication to regain access to the Cisco SD-Access network.

Secure AP onboarding is done by authorizing the Access Point on a closed authentication port by allowing limited access to DHCP/DNS and Cisco DNA Center for the PnP workflow. The PnP workflow on the Cisco DNA Center is enhanced to enable a dot1x supplicant on the AP, and the AP uses this supplicant to authenticate with Cisco ISE.

Wired Client Onboarding

The north star goal for most financial customers is to slowly migrate their networks into completely closed authorization networks, fully leveraging the secure access benefits a Cisco SD-Access architecture provides. A preferred model we recommend to all such financial customers is to slowly migrate all parts of their network to a dot1x-based closed authentication mode. All traffic is dropped before authentication, including DHCP, DNS, and ARP in this mode. For clients that do not support a dot1x client, you can allow dot1x to fail before MAC Authentication Bypass (MAB) can authenticate and succeed, or you can set MAB as the primary authentication method. While user endpoints are typically onboarded using dot1x, IoT endpoints are typically onboarded using MAB.

Wireless Client Onboarding

Multiple Wireless Service Set Identifiers (SSID) are deployed with different security – Enterprise Dot1x, Identity/Pre-Shared Key, MAC Filtering – to support an array of wireless endpoint connectivity. Critical endpoints and infrastructure access can be tightly secured by mandating dot1x-based onboarding for enterprise SSIDs for all users. Guests can be onboarded securely through a DMZ Fabric Border Node deployment that will be covered in detail in the Guest Isolation section of this chapter.

Data Loss Prevention using Secure Network Analytics

Cisco Secure Network Analytics (SNA) (formerly StealthWatch) leverages NetFlow data from network devices throughout all areas of the network – access, distribution, core, data center, and WAN Edge – providing a concise view of traffic patterns. This visibility allows an SNA database record to be

maintained for every communication that traverses a network device. The StealthWatch Security Analytics (SSA) service on Cisco DNA Center automates the provisioning of network elements so that they send data to Cisco SNA, enabling you to gain more visibility and providing real-time monitoring of all network traffic. In addition to this, Cisco SNA can be used by customers to identify flows that are perceived to be a threat, and those flows can be marked for isolation within the network by quarantining them.

A quarantine operation results in the specific user or device being assigned a quarantine SGT in Cisco ISE by leveraging the pxGrid communication that Cisco SNA maintains with Cisco ISE. Customers are expected to have preset Quarantine policies for the Quarantine SGT group so that endpoints that have been identified for quarantine in Cisco SNA can automatically be enforced by all the security group policies that will restrict them to minimal to zero connectivity without any user intervention on Cisco ISE.

Rogue Detection using aWIPS (Adaptive WIPS)

Financial customers benefit immensely from the native capabilities of Cisco DNA Center in collaboration with Cisco wireless to monitor and detect rogues in the network and mitigate such risks through containment. Cisco DNA Center supports honeypot detection of rogue APs and various forms of DoS attacks that rogue wireless clients can perform on a network, from flooding of Auth requests to probes to association/disassociation requests.

End-to-End Encryption

Financial customers require data traversing their network to be encrypted, and certain deployments have privacy, data confidentiality, and/or regulatory requirements. Financials often demand encryption at Layer 2 within the corporate network and require encryption at Layer 3 across the WAN. Encryption inside the Fabric Site, specifically encryption at the MAC layer, is achieved with MACsec.

Figure 7.3 MACsec Encryption

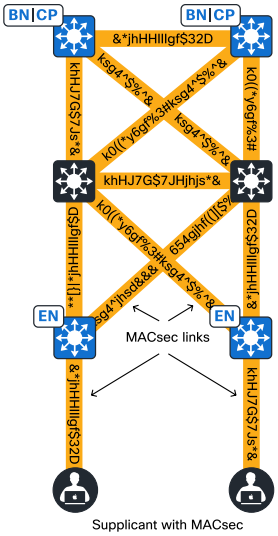

To provide encryption on these links, as shown in Figure 7.3, we can use Media Access Control security (MACsec). MACsec is the IEEE 802.1AE standard that enables devices on Ethernet networks to provide confidentiality, integrity, and authenticity for data in Transit. Switch-to-host and switch-to-switch MACsec can be enabled in the Cisco SD-Access network on compatible switches via templates in Cisco DNA Center. Encryption at the WAN is taken care of by IPsec tunnels between each WAN Edge router.

Availability (Five-Nines)

Disaster Recovery for Cisco DNA-Center

Financial institutions have a low tolerance for management, control, or data plane failure. Cisco DNA Center supports both intra-cluster and inter-cluster resiliency. Cisco DNA Center's Disaster Recovery implementation consists of

three components: the Main Site, the Recovery Site, and the Witness Site. At any given time, the Main and Recovery Sites operate in either active or standby roles. The active site manages your network while the standby site maintains a continuously updated copy of the active site's data and managed services. Whenever an active site goes down, Cisco DNA Center automatically initiates a failover, completing the tasks necessary to designate the former standby site as the new active site.

Highly Resilient SD-Access Network Architecture

Apart from the conventional means of providing resiliency using stacking and StackWise Virtual, regional hubs and Campus headquarters in Financial Networks often need protection from building failures. This ensures that connectivity to data centers for payment gateways and other critical applications is always available. Cisco SD-Access supports a flexible deployment architecture where deployments can take advantage of positioning Fabric borders in different physical sites while tying them together under the ambit of a single Fabric Site, as shown in Figure 7.4.

Figure 7.4: SD-Access Fabric resiliency

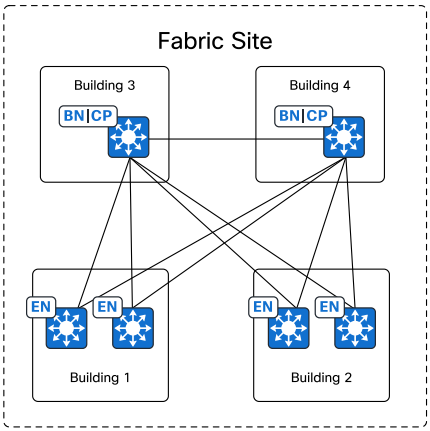

As can be seen in Figure 7.4, Buildings 1-4 are in the same Fabric Site, with the Colocated Border Nodes and Control Plane Nodes being located in separate buildings. Cisco SD-Access provides the flexibility to assign priorities to these Border Node deployments such that one Border Node can be prioritized more or even become the only active border carrying traffic, so long as it is active. When a building fails, the Border Node in the other building can automatically take over all traffic from the Edge Nodes.

Critical VLAN

Cisco SD-Access supports a Critical VLAN capability to provide minimum required network connectivity options to endpoints when the connectivity to their ISE server is lost due to outages such as a WAN outage.

For clients that are already onboarded, if connectivity to the ISE Policy Service Node is lost, periodic re-authorization is paused to ensure disruptions in the authentication path do not impact the data plane. For clients that are not yet onboarded into the network, the Critical VLAN feature

places the client on a particular VLAN if connectivity to ISE is lost. This VLAN provides limited access to the network. Critical VLANs can utilize micro-segmentation to enforce the policy in absence of ISE.

Compliance and Audits

Financials need an always-on visibility into what goes on in the network and a strict compliance policy regarding who is authorized to manage and provision the network. Cisco DNA Center has a host of features built in that help with all such compliance and audit mandates in an organization.

Audit Logs

Audit logs capture information about the various applications running on Cisco DNA Center. Audit logs also capture information about device public key infrastructure (PKI) notifications. The information in these audit logs can be used to assist in troubleshooting issues, if any, involving the applications or the device PKI certificates. Audit logs also record system events, when and where, and which users initiated them. With audit logging, configuration changes to the system get logged in separate log files for auditing.

Configuration Compliance

Compliance helps identify any intent deviation or out-of-band changes in the network that may be injected or reconfigured without affecting the original content.

With Cisco DNA Center, a network administrator can conveniently identify devices that do not meet compliance requirements for the different aspects of compliance such as Configuration, Software Image, Product Security Incident Response Team (PSIRT) exposure, Network Profile, Fabric, and more.

Compliance checks can be automated or performed on demand, as shown in the following:

- Automated compliance check: Uses the latest data collected from devices in Cisco DNA Center. This compliance check listens to the traps and notifications from various services such as inventory, Software Image Management (SWIM), and so on to assess data.
- Manual compliance check: Enables users to manually trigger the compliance in Cisco DNA Center
- Scheduled compliance check: A scheduled compliance job is a weekly compliance check that runs on a set schedule, for example, every Saturday at 11 pm.

Cisco DNA Center supports the following types of compliance checks at the time of this writing:

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

Configuration Drift

Financial organizations are often subject to compliance mandates that dictate that organizations maintain archives of configurations for all network devices. Cisco DNA Center supports configuration drift that highlights the current configuration of every single device with the flexibility to go back up to a month to view how configurations have changed on a specific device.

Role-Based Access Control

Cisco DNA Center supports the flexibility to assign permissions to users either based on a local or external RADIUS/TACACS database. Users can be assigned one of the following roles and can also be assigned access to specific applications within Cisco DNA Center:

- **Administrator (SUPER-ADMIN-ROLE)**: Users with this role have full access to all of the Cisco DNA Center functions. They can create other user profiles with various roles, including those with the SUPER-ADMIN-ROLE.
- **Network Administrator (NETWORK-ADMIN-ROLE)**: Users with this role have full access to all of the network-related Cisco DNA Center functions. However, they do not have access to system-related functions, such as backup and restore.
- **Observer (OBSERVER-ROLE)**: Users with this role have view-only access to the Cisco DNA Center functions. Users with an observer role cannot access any functions that configure or control Cisco DNA Center or the devices it manages.

Lean IT Staff

Centralized Troubleshooting using iCAP and Sensors

Financial customers often lack an onsite network support team in smaller branches, making troubleshooting much more difficult.

Cisco DNA Center's Intelligent Capture capability provides support for a direct communication link between Cisco DNA Center and APs in Cisco SD-Access Wireless environments. Cisco DNA Center can receive packet capture data, AP and client statistics, and spectrum data. With the direct communication link between Cisco DNA Center and APs, Intelligent Capture allows you to access data from APs that is not available from wireless controllers.

Power of Automation

Financials are often lean on IT staff and hence benefit immensely from the automation that Cisco DNA Center brings to the table. Cisco DNA Center supports Zero Touch Provisioning (ZTP) by way of automating the bring-up of the entire Underlay network in an SD-Access Fabric Site. Even a single network engineer can bring up 100s of branches from a central location once the network devices are physically connected. This accelerates the onboarding of new locations to productivity and reduces site rollout time to weeks or days from months.

Monitoring and Troubleshooting

Cisco DNA Center has a powerful and robust Assurance application that provides next-generation visibility into the telemetry that is returned from the network infrastructure. The Assurance suite allows for issues to be alerted, automatically correlated, and triaged via advanced machine analytics. This allows for intelligence to be gathered for a wide range of network and client-related problems, which can be solved either through the Cisco DNA Center user interface or remotely through advanced integration via Rest-API with Service Now or other ITSM applications. Additionally, issues can be prioritized in a severity list and reported through robust exportable reports.

Complete Isolation of Guest Users

Financial Network administrators often need to manage extensive guest services with a common subnet across all their campus and branch sites. To address this challenge, Cisco SD-Access provides the Multisite Remote Border solution using Virtual Network Anchors. This solution allows traffic from a given Virtual Network at multiple dispersed sites to be aggregated back to a central location, sometimes referred to as an Anchor Site, where the Virtual Network uses a single common subnet rather than having to define and use per-site subnets for guest traffic.

Virtual Network Anchor Sites significantly simplify the guest service deployments across SD-Access Fabric Sites and provide consistent and secure segmentation for guest traffic in financial environments with a simplified and centralized common subnet structure.

Using anchored services, traffic for guest endpoints that belong to the Anchored Virtual Network at each site are aggregated and tunneled back to the Multisite Remote Border at the Anchor Site over VXLAN. An Anchor Site functions much like a traditional Fabric Site, but it forms a virtual Fabric Site serving a particular Virtual Network.

This virtual Fabric Site has its own Border Nodes and Control Plane Nodes which are dedicated to the Anchor Site. What is special about the Anchor Site is that its Edge Nodes and Fabric WLCs are dispersed across multiple Fabric Sites, referred to as Anchoring Sites.

Figure 7.5: Guest Traffic Flow with Anchor Virtual Networks

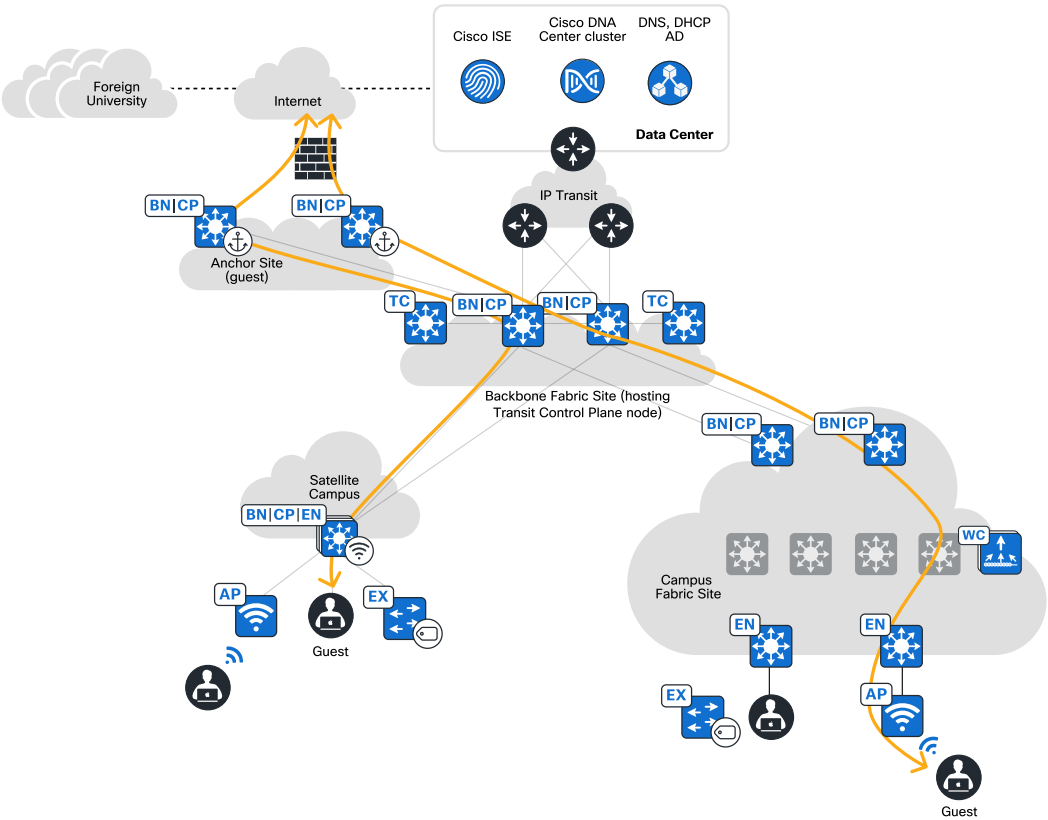

Multisite Remote Border is enabled on a per-Virtual Network basis. For an Anchored Virtual Network, all edges in the anchoring sites use the Anchor Site Border Nodes and Control Plane Nodes for data plane and control communication. Fabric WLCs in the anchoring sites communicate with the Anchor Site Control Plane Node for wireless endpoint registration specific to the Anchored Virtual Network.

Since the Anchor Site Border Node reachability may traverse multiple IP networks, special attention must be paid to the MTU across the entire path

to accommodate the additional VXLAN header overhead of 50 bytes.

After a guest endpoint joins the guest SSID and passes the Central Web Authentication using Cisco ISE, it is associated with the anchored guest Virtual Network. Guest traffic is tunneled to the Anchor Site Border Node and egresses to the Internet through a firewall.

Additionally, while the guest traffic is VXLAN-encapsulated and passes through the Fabric, the first hop or gateway for guest traffic can be outside the Fabric and connected at Layer 2 to the Multisite Remote Border. Such a device can be a firewall for inspection purposes, should it be a design requirement.

From a Cisco ISE perspective, the following are the suggested options for the financial customers that need isolated Guest environments:

- Separate PSN hosted in DMZ for Guest users
- Dedicated PSN for Guest users

By having separate authentication, authorization, and Multisite Remote Border for the guest user, the guest network is now fully isolated from the Control Plane, Data Plane, and Policy Plane perspective, ensuring the Guest network is isolated from the other financial users, devices, and resources.

Trading Floor and Video Applications Streaming

Multicast will be required in the Cisco SD-Access campus Fabric for various trading floor and video applications that exist in a financial environment. Video applications include various advertisements that can be streamed to various branches across various regions explaining new product offerings or programs to visiting customers.

Such multicast data can be streamed from a variety of sources, including regional data centers, corporate data centers, etc. The Cisco SD-Access architecture provides the flexibility of end-to-end seamless multicast data traffic to flow from anywhere in the larger enterprise network to anywhere in the globe. Such multicast traffic can be streamed over SD-Access Transit or IP-Based Transit networks. Cisco SD-Access supports all flavors of Multicast, with flexibility in allocating Multicast RP nodes either inside the SD-Access Fabric or outside of it.

Multicast forwarding in SD-Access Fabric uses two methods for distributing the traffic on the Underlay: Head-End Replication and Native Multicast.

Head-End Multicast Replication (Overlay)

Head-End Multicast replication is usually deployed in smaller SD-Access Fabric Sites with <100 receivers and just a handful of nodes needing multicast streaming.

Head-End replication (or ingress replication) is performed either by the multicast first-hop router (FHR) when the multicast source is in the Fabric Overlay or by the Border Nodes when the source is outside of the Fabric Site.

Native Multicast (Underlay)

Native multicast is our recommended means to do Multicast in Medium and Large Sites with 100s of devices and 1000s of clients.

This type of deployment significantly improves the efficiency, latency, and overall scale of multicast traffic handling. This reduces the replication load at the Head-End Node (Border), allowing traffic to be replicated by the Underlay infrastructure to all of the Edge Nodes. This approach allows each multicast group in the Overlay to be mapped to a corresponding Source-Specific Multicast (SSM) Underlay Group. To enable this Native Multicast deployment, it is required to configure Protocol Independent Multicast (PIM) in the Fabric Underlay.

Multicast over SD-WAN

Cisco SD-Access native multicast relies on Source-Specific Multicast (SSM) in the Underlay. Therefore, SSM configurations are configured on the Fabric Nodes and Intermediate Nodes. The Multicast Rendezvous Points (RPs) can be external to the SD-Access Fabric in the data center or even local to Fabric. Typically, since most of such streaming happens from the data center, it may make more sense to place the RP nodes closer to the sources. Receivers placed across Fabric Sites can subscribe to multicast traffic over Cisco SD-WAN. Multicast (SSM/ASM) is enabled on the service VPN subinterfaces of the Cisco SD-WAN Edges towards the Overlay Virtual Network on the Fabric Border Nodes. This ensures complete end-to-end native multicast configuration throughout the Fabric Nodes and SD-WAN Edges.

Figure 7.6: Multicast across SD-WAN

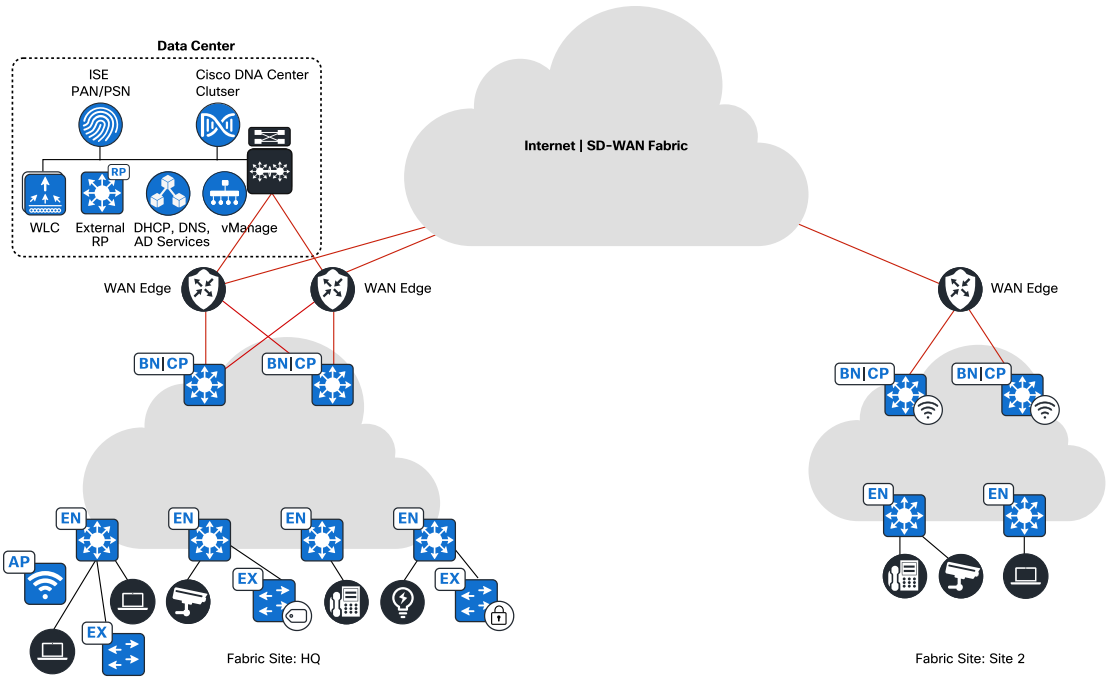

Global deployment with Multi-DNA Center to centralized ISE

Cisco ISE can be deployed either in a standalone or distributed model. For multiple site enterprise deployments, Cisco ISE is deployed in a distributed model with distributed Policy Administration Nodes (PAN) and Monitoring Nodes (MNT), as well as active and standby Platform Exchange Grid (pxGrid) Nodes and Policy Service Nodes (PSNs). PSNs can be spread across various sites to provide a low latency Network Access Control (NAC) solution. PSNs

synchronize with the ISE PAN to provide consistent policy and SGT assignment based on client authentication.

Global deployments spread across various locations require multiple Cisco DNA Centers to be deployed due to scale considerations, latency network requirements, and Global compliance reasons. Multiple Cisco DNA Center clusters can be integrated into a single Cisco ISE, providing centralization and standardization of Authentication, Authorization, and Security Group Policy across the global enterprise.

Large Cisco ISE deployments, such as in Financial Networks, can benefit by integrating multiple Cisco DNA Center clusters with a single Cisco ISE. Cisco DNA Center supports multiple Cisco DNA Center clusters per Cisco ISE deployment to utilize Cisco ISE better and provide a centralized policy management plane for multiple Cisco DNA Centers. For more information, see *support for Multiple Cisco DNA Center Clusters with a Single Cisco ISE System*.

Figure 7.7: Global Multi DNA Cluster deployment for consistent Group Based Policy

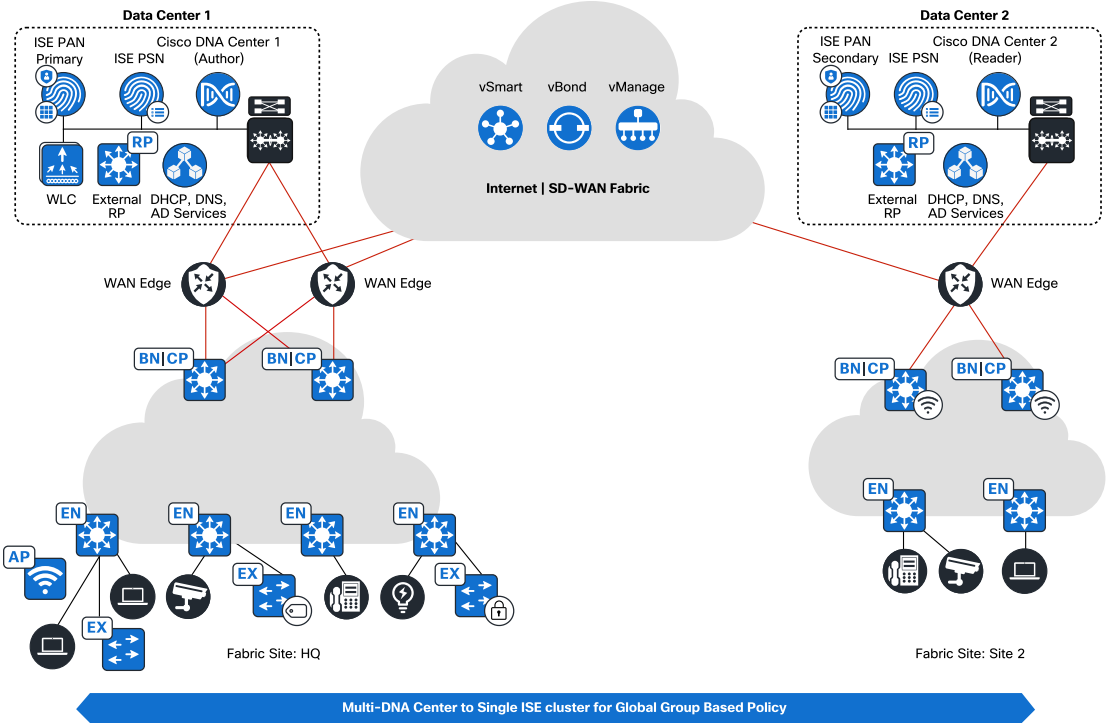

Note At the time of the writing of this book, this feature is currently a Limited Availability offering. It is a temporary, initial approach to the problem in response to urgent business requirements. A more comprehensive General Availability offering will be available in the future. Check with your Cisco account team on the latest status of this capability before deploying it in your production network.

Network Visibility

Cisco DNA Center not only manages your network by automating network devices and services but also provides network assurance and analytic

capabilities. Cisco DNA Center collects telemetry from network devices, Cisco ISE, users/endpoints, applications, and other integrations across the network. Cisco DNA Center Network Analytics correlates data from various sources to help administrators/operators to provide comprehensive Network Insight into:

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

Centralized Policy Management

To deal with the problems associated with managing the exponential growth in the number of endpoints connecting to the network, we need to design a solution that integrates Bring Your Own Device (BYOD) capabilities. Globally we need to be able to build a set of security policies that span the globe but are easy to manage and propagate.

We solve these problems with Cisco DNA Center integrated with Identity Services Engine (ISE), both of which are foundational elements of a Cisco SD-Access network. Within the network, we enable secure onboarding for endpoints using Cisco SD-Access and ISE, utilizing a redirect to a Mobile Device Manager (MDM). Once the endpoint is onboarded appropriately, it will have all the security policies, security clients, and signatures matching

the organization's security policies. With Cisco SD-Access, we can enforce compliance with the security policy based on posture assessment.

Once devices are authorized on the network, we can scalably embed a security tag based on the authorized session and apply a security policy that is end-to-end enforced anywhere in the network.

When we think globally, Cisco DNA Center's powerful capability of authoring policy in one cluster and then absorbing that on multiple Cisco DNA Center clusters around the world allows for the same policies to be propagated globally. This allows us to flexibly build a policy that makes sense and simultaneously scales across the entire global enterprise.

Finally, from an observation perspective, with Cisco DNA Center and Cisco SD-Access, we can evaluate group interactions with Security Group Policy Analytics. Security Group Policy Analytics is a Cisco DNA Center application that allows you to determine which users/devices are interacting within the network, anywhere in the world. If you notice behavior that must be immediately corrected, you can build a contract in Cisco DNA Center and deploy it globally in a matter of minutes.

High Sensitivity to Quality of Service (QoS)

Financial organizations have stringent application service level agreements, which affect their application availability. A slow and QoS-disparate network can lead to poor application experience and consequently impact operations and trading floors, which are highly sensitive to delays.

Cisco DNA Center has major advancements which help large and global financial organizations manage and scale application policies for deployment. When building the Cisco SD-Access Fabric, we can deploy a tiered Cisco-recommended or Customized QoS policy to both the wired and wireless infrastructures. This Application Policy uses the DiffServ model, where applications are categorized in application sets and are divided into Business Relevance based on administrator preference.

Cisco DNA Center's advancement into Controller-Based Application Recognition utilizes the NBAR 2.0 Deep Packet Inspection (DPI) capabilities that are utilized for Endpoint Analytics and allows for applications to be learned by Cisco DNA Center directly from packets in Transit across the network. This allows for seamless learning of applications from Cisco's NBAR2.0 Cloud, Microsoft's 0365, and Infoblox. The application's Fully Qualified Domain Names (FQDN) and URL information can be learned and used to update application sets within QoS Policies. This allows for the immediate deployment of those application policies at scale anywhere in the network.

Cisco DNA Center also can gather telemetry based on how those Application Policies are operating, which allows for administrator intervention to modify and redeploy policies at scale across the network. The ability to see how an application operates is crucial to understanding the user experience and solving issues before being notified by an end user. This goes a long way to ensure that we have a method of proving mean time to innocence.

Lastly, as all application traffic is now correctly classified and marked by Application QoS policies, we can then deal with them appropriately by any device end-to-end in the network, across the WAN, in a SD-WAN cloud, or even in the data center.

Acquisitions Integration

As in any other industry, financial organization mergers and acquisitions occur from time to time. One of the benefits of Cisco SD-Access is the ability to deal with these interesting challenges seamlessly. This allows for merging two organizations to integrate and optimize redundant operations.

Once the Business-to-Business (B2B) connection is enabled, merging the two networks can begin. The integration of Cisco DNA Center and Cisco ISE enables the integration of authorization across the Cisco SD-Access Fabric from one side of the network to the other. By connecting ISE to a foreign

Active Directory environment, we can connect directly or proxy authentications to onboard clients from the newly merged network. This allows for the application of SGTs from end to end, which allows for the application of a seamless end-to-end security policy.

Initially, the newly-acquired/newly-merged organization's endpoints may be automatically placed into a Virtual Network within Cisco SD-Access and enforcement configured to utilize pxGrid integration with firewalls that can enforce policies.

Figure 7.8: Cisco SD-Access Layer 2 Fabric for Merger and Acquisition

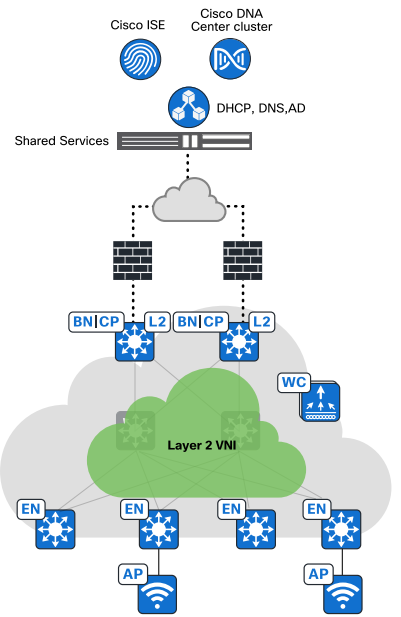

As the acquisition progresses, more and more integrations can be leveraged, allowing for the slow progression of change of ownership of assets and the complete onboarding of the newly-acquired/newly-merged organization. Initially, we can deploy a Cisco SD-Access Fabric as a Layer 2 Fabric Site

and allow for the central onboarding of clients within the Fabric Site. Those clients would then use a gateway hosted on a pxGrid-integrated firewall which would be initially responsible for traffic inspection. As confidence improves and the acquisition continues, the client endpoints from the newly acquired organization will be onboarded properly within the organization network.

Deployment Options

Cisco SD-Access provides flexibility to deploy multiple Fabric Sites interconnected with SD-Access Transit or IP Transit. Individual Fabric Site size is based on several metrics, including:

- Fabric Site client count scale
- Fabric Site device count
- Number of IP Address Pools
- Virtual Networks
- Security Groups
- Security Group Policies
- High-Availability/Redundancy requirements

The intent of this vertical is to design and support multiple sites with site survivability and support end-to-end macro-segmentation and micro-segmentation with consistent policy across the enterprise.

Major Fabric Site model:

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

Very Small Site

Figure 7.9: Very Small Site design

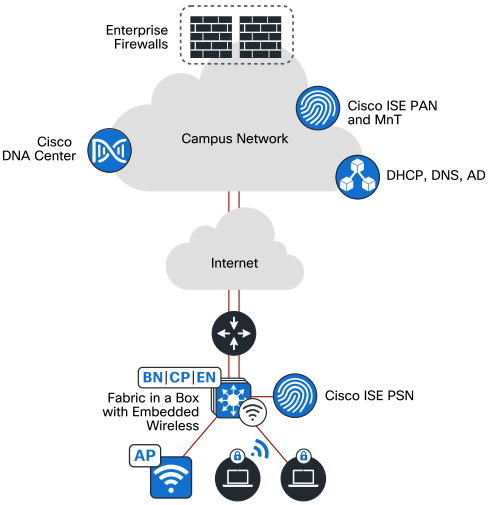

You can use Fabric in a Box to cover a single Fabric Site, with resilience supported by switch stacking or StackWise Virtual to support up to 200 endpoints and 40 APs. For Fabric in a Box deployments, SD-Access Embedded Wireless provides site-local WLC functionality. The site may contain an ISE PSN depending on the WAN/Internet circuit and latency.

Small Site

Figure 7.10: Small Site design

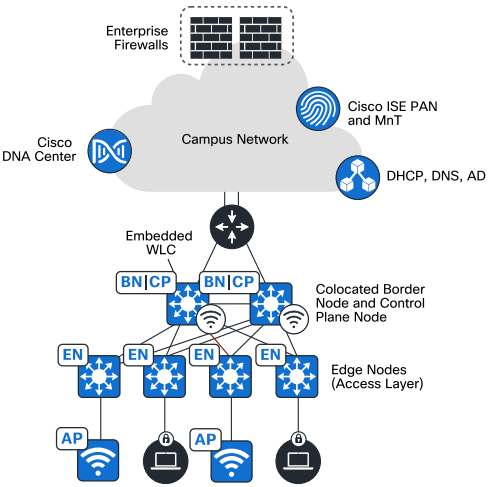

The Small Site Reference Model covers a single office or building with single wiring closets, usually to 4,000 and up to 100 APs. The border function is colocated with the control plane function on one or two devices and usually uses embedded wireless with the option of hardware WLCs.

The physical network is usually a Two-Tier Collapsed Core/Distribution with an Access Layer servicing several wiring closets. Rather than colocating all roles in one device, the Small Site Reference Model provides added resiliency, redundancy, and a more significant number of endpoints by separating the Edge Node role onto dedicated devices in the Access Layer. The Border Node and Control Plane Nodes are colocated in the Collapsed Core Layer. For SD-Access Wireless, the embedded WLC is provisioned on the Colocated Border Node and Control Plane Node. Optionally, a virtual- or hardware-based WLC is used.

Medium Site

Figure 7.11: Medium Site design

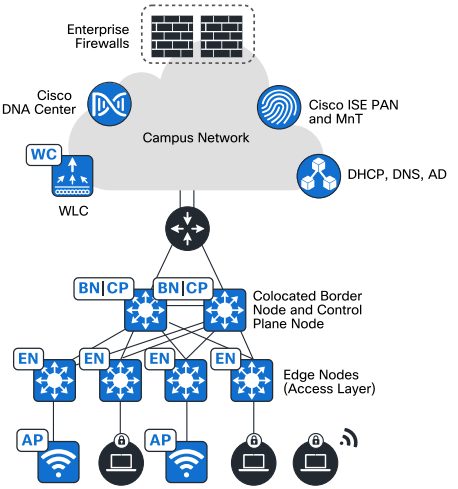

The Medium Site Reference Model usually covers a building with multiple wiring closets with a physical topology consisting of a Two-Tier Collapsed Core/Distribution with an Access Layer.

The Medium Site is designed to support less than 25,000 endpoints and less than 2,000 APs. The Border Node function is colocated with the Control Plane Node function on one or two devices or a highly-resilient single device, and a separate wireless controller is ideally deployed in a HA configuration.

Large Site

Figure 7.12: Large Site design

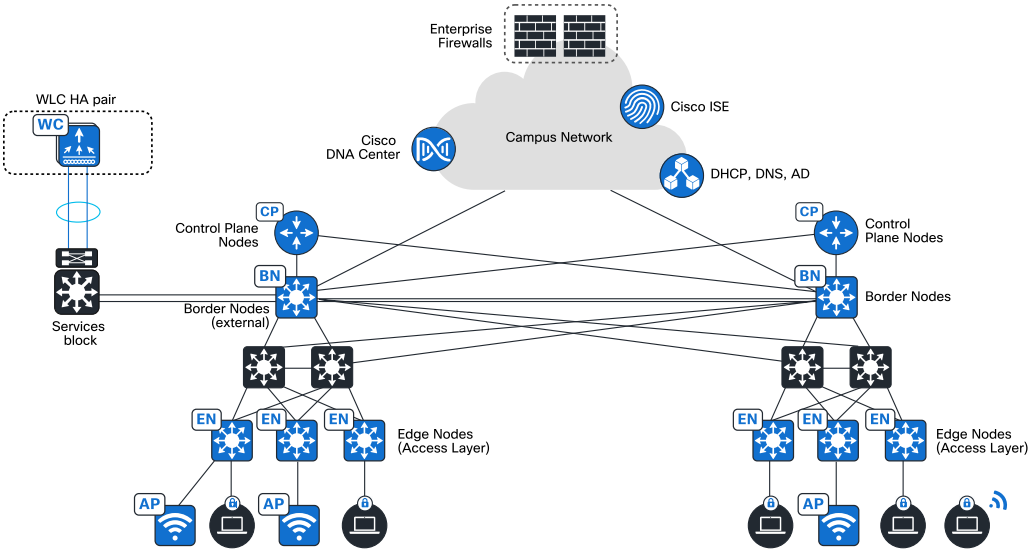

The Large Site Reference Model covers multiple buildings or a building with multiple wiring closets. The physical network is usually Three-Tier with Core, Distribution, and Access Layers. It may even contain a routed Super Core that aggregates numerous buildings and serves as the network egress point to the WAN and Internet. The Border and Control Plane Nodes functions are provisioned on separate devices rather than colocating.

The Large Site supports up to 100,000 endpoints and 6,000 APs. The border is distributed using redundant devices from the control plane function, and a separate wireless controller is in an HA configuration.

Summary

The main IT goal of Financial organizations is to provide maximum value and desired business outcomes to users through industry-leading technologies such as Cisco SD-Access. We have seen during this chapter how SD-Access provides maximum value and desired business outcomes while providing security, flexibility, and resiliency in parallel by solving use cases with ease.

In this chapter, we demonstrated how Cisco SD-Access solves the challenges faced by organizations in the Financial vertical. You learned how you can securely onboard network devices, users, and endpoints. We showed you how to encrypt your network while preserving security group information across global sites via WAN Infrastructure. We also demonstrated how the automation, orchestration, and assurance capabilities of Cisco SD-Access and Cisco DNA Center can empower IT staff to support global networks, even when staffing is lean. All of these exemplify the beauty of software-defined architecture and controller-based networking.

Migration to Cisco SD-Access

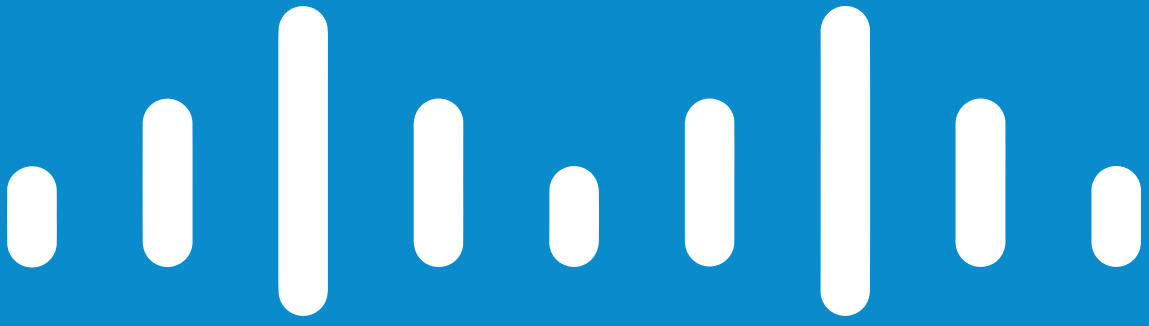

After reading any of the previous chapters or hearing about Cisco SD-Access in different forums, you are now convinced that this is the right solution for your environment. So, the next immediate question is, “How can I migrate my network to Cisco SD-Access? What migration options do I have?” Well, you have come to the right chapter in this book. This chapter focuses solely on network migration and not any other scenarios. This chapter does not focus on how to install and configure Cisco DNA Center, set up network hierarchy, network environments, discover switches, or integration of Cisco DNA Center with Cisco Identity Services Engine (ISE).

In the case of an existing Cisco ISE installation, we recommend that you use a new ISE to migrate to Cisco SD-Access. Why a new ISE? The integration with Cisco DNA Center might need a software version upgrade and might impact the policies in the existing network. Once the entire network is migrated to Cisco SD-Access, you can decommission the existing installation and make the new one the primary ISE instance.

Considerations

We will start off with a Considerations section with some of the “little things that matter much.”

Selecting A Great Start

The network can be divided into different portions from a risk/criticality factor. Introducing any new technology into the network has its challenges. It should be done in a manner that results in a smooth and positive experience for all involved, including the end users. The wise approach is to start planning the migration in a small, low-risk/non-critical (IT-friendly if possible) part of the network while creating a path for the end users to fall back on if problems arise. For example, if you want to migrate a healthcare facility or a factory with a carpeted part of the network, start with the outpatient clinics in the former case or the non-real-time administrative area in the latter scenario. This will give you the right experience to migrate the rest of your critical network infrastructure.

Approaches to Migration

Let us look at the two approaches to migration.

Table 1: Comparison between Parallel and Incremental Approach

Parallel	Incremental
Best for Branch (small scale) deployments	Best for Campus (any size)
Requires cable runs to create a new parallel network	Requires a couple of cables from new access and distribution switches
Power and outlets for the parallel network	Incremental power and outlet requirement
Legacy hardware in existing network	Legacy hardware in existing network
Upgrade most of the network infrastructure	Upgrade most of the network infrastructure
Clean slate (leaving behind any complexity in the old design)	Will need to carry forward the constraints of the old design in the Underlay
Test users in a completely new network	Test of functionality is partial
Easy Rollback of migrated users	Easy Rollback of migrated users

Parallel Approach

With the Parallel Approach, you build the new network alongside the existing network. Benefits to this approach include:

- You can ignore the complexities of the existing network
- You can upgrade the entire network with new hardware and build the solution realizing the full new capabilities
- You do not have to deal with a mixture of old and new hardware and software, which can often force you away from the new capabilities you want in your new environment.
- You can test use cases in an entirely new network before migrating the users
- You can roll back the users into the old environment if any issues occur in the new network

While the Parallel Approach has many appealing benefits, there is one "gotcha" that you must consider. You will potentially need a lot of new Layer 1 – i.e., space, power, cooling, cables – to support a parallel installation. This approach might not apply to large Campus Networks with a lot of buildings or a single building with a lot of floors. Depending on the situation, you might be able to use this approach in a small branch-type network where there are available Layer 1 resources to execute this.

Incremental Approach

An Incremental Approach is an approach where you start small and then increase the footprint of the new solution within your existing network. In this case, as Table 1 shows, you do not require a lot of new Layer 1 resources. However, you must live with the complexities of the old network and sometimes operate in a mixed environment of different hardware and software capabilities. With this approach, you test use cases in a mix of old and new environments and roll users back to the old network if problems occur.

Maximum Transmission Unit (MTU)

Figure 8.1: IP Packet of 1500 bytes + 50 bytes of VXLAN Encapsulation

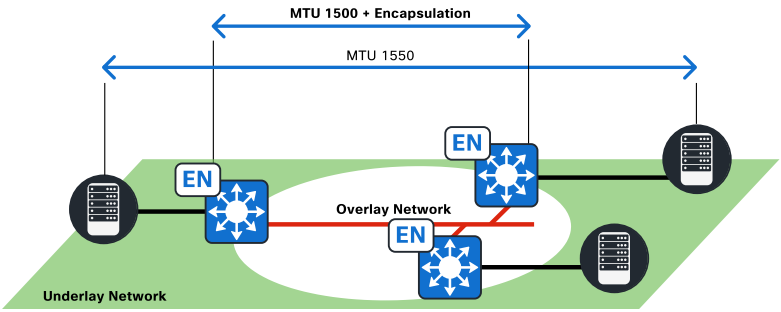

At the network device level, the primary thing to consider is MTU (Maximum Transmission Unit) support for not only Fabric Nodes but also for any other devices over which Fabric traffic must be transmitted. Cisco SD-Access uses a VXLAN (Virtual eXtensible Local Area Network) encapsulation for all data traffic, adding an additional 50 bytes to the packet.

At a minimum, devices carrying SD-Access VXLAN-encapsulated data traffic should support an MTU large enough to accommodate the extra 50 bytes. Cisco strongly recommends enabling `system mtu 9100` on the Cisco Catalyst 9000 Series switches to accommodate all scenarios.

There are scenarios where the Underlay network does not support more than 1500-byte packets – for example, if the Fabric Sites are connected via a Cisco SD-Access Transit over a WAN that does not support more than 1500-byte packets. In these cases, TCP MSS (Transmission Control Protocol Maximum Segment Size) can be specified to constrain the packet size,

considering the VXLAN header's encapsulation overhead. The best practice recommended value of setting MSS is 1250 bytes.

TCP MSS is configured on all the endpoint-facing SVIs (Switched Virtual Interfaces) – Routed VLAN interfaces (Virtual Local Area Networks) on all Edge Nodes. However, this method works only on TCP applications and not any other applications.

Underlay Transformation into Routed Access

Layer 2 Access designs are the most common network design across all industry verticals. Traditionally, the Layer 2 Access design has provided the most flexibility for subnet availability across the network, enabling plug-and-play simplicity. However, Layer 2 Access designs rely heavily on Spanning Tree Protocol (STP) for loop prevention. Network teams always live in fear of a network meltdown caused by STP due to a misconfiguration or improper connection, and STP has poor support for delay-sensitive applications, such as voice traffic, during a failover.

Figure 8.2: Diagram shows a Layer 2 Access Underlay transformed to Routed Access

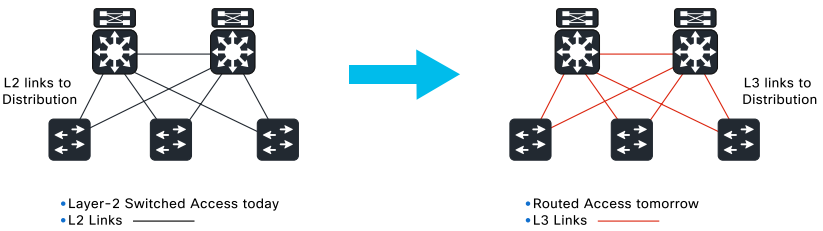

The Fabric provides the best of both worlds – plug-and-play connection capability over a stable, reliable Routed Access Underlay network. Hence, it is mandatory that the Underlay access design be converted from Layer 2 Access to a Routed Access design.

Support for Different Topologies

We acknowledge that Campus Network topologies are dictated by the physical structure of buildings and hence have constraints in design. In most cases, Campus Networks are either standard Three-Tier or Two-Tier networks. In a Three-Tier design, the Access Layer switch connects to a Distribution Layer switch that then connects to a Core Layer switch. In a Two-Tier design, the Access Layer switch connects to a Collapsed Distribution/Core Layer switch.

In some cases, the design could be a star topology where the Access Layer switch connects to the Distribution switch or Collapsed Distribution/Core switch and to other Access Layer switches. In other cases, the Access Layer switches might be daisy-chained to each other, or there might be small compact switches connected to the Access Layer switches.

Regardless of the topology, Cisco SD-Access is supported as far as the Underlay is Routed Access and each element is IP reachable to the Cisco DNA Center, to the other Fabric Nodes, and to the Wireless LAN Controller (WLC) in that site. If the topology was built based on a Cisco Validated Design document, there should be few concerns regarding the topology.

IP Addressing in the Underlay and Overlay

Figure 8.3: Different IP addressing schema for Underlay and Overlay

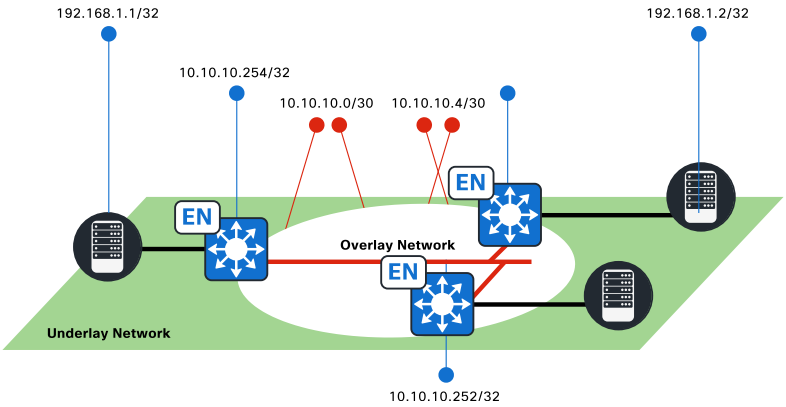

Cisco highly recommends planning for a different IP schema to be used in Underlay and Overlay networks. While this is not mandatory, it is certainly a best practice. The importance of having two separate schemas in the IP addressing is especially realized when troubleshooting problems in the two different network areas. Having the same IP addressing schema would get confusing, and you might troubleshoot the Underlay inadvertently when there is a problem in the Overlay.

Movement of Feature Application

Figure 8.4: Movement in point of application of features

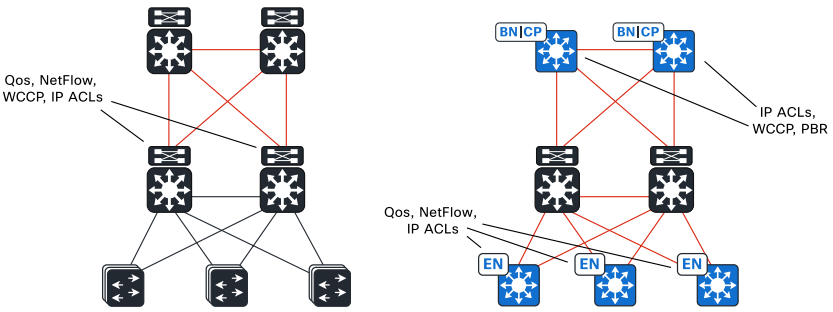

In a Layer 2 Access design, the Layer 3 boundary is typically at the Distribution Layer and is where all the features and policies are applied. Within an SD-Access Fabric network, the Underlay transforms into a Routed Access design with the Layer 3 boundary at the Access Layer. In the Fabric network, some features such as Quality of Service (QoS), NetFlow, and IP Access Lists (ACLs) are applied at the Access Layer, while some features such as Policy-Based Routing (PBR) or Web-Cache Control Protocol (WCCP) are applied at the egress interface of the Fabric Border Node.

Migrating a Layer 2 Access Network with New Subnets

Figure 8.5: Layer 2 Access with Layer 3 boundary at each Distribution block

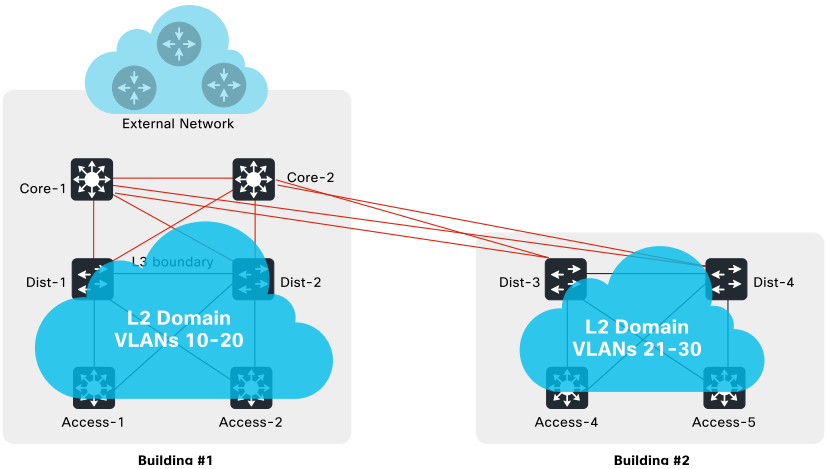

Consider a Campus Network with a typical Layer 2 access design, where the Layer 3 boundary is at each distribution. All VLANs span below each distribution block in this type of network, as seen in Figure 8.5. This section focuses on how to start an incremental migration to Cisco SD-Access in this type of network design using new IP subnets.

The reasons why we start with new IP subnets are compelling. One of the advantages of having a Fabric network with Anycast Gateways is that a smaller number of larger subnets can be used compared to a traditional network that uses a large number of smaller subnets.

The traditional best practice for designing IP subnets is to design using /24 subnets to avoid large broadcast domains. In the Fabric, broadcasts are disabled by default and are only limited to each Edge Node. One can safely have a /16 or a /20 subnet in the network with broadcast disabled. This also improves efficiency in utilizing DHCP (Dynamic Host Configuration Protocol) scopes and works well for endpoints amenable to changing IP addresses via DHCP.

The high-level concept of the migration process can be seen in Figure 8.6.

Figure 8.6: High-level conceptual view of migration with new IP subnets

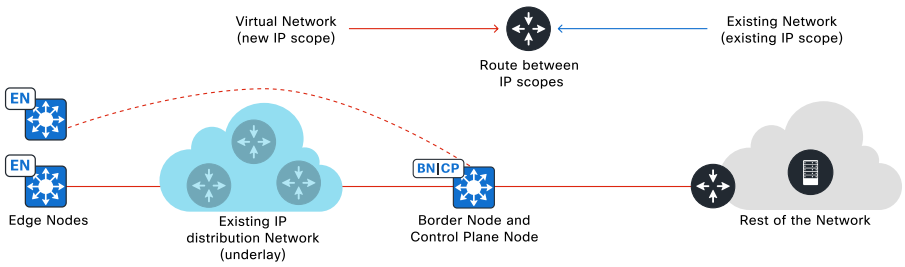

We start a New Subnet migration by introducing a new Access switch in the network. This will play the role of an Edge Node and will be connected via Routed Access links to each of the existing distribution switches, as shown in Figure 8.7. This forms the ingress into the Fabric from an endpoint perspective.

Figure 8.7: Connect a new switch in the Access Layer with a Routed Access design

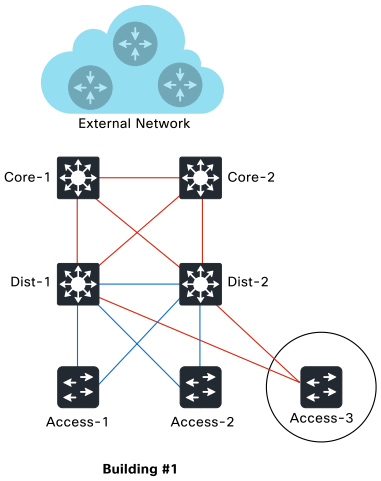

We need a Control Plane Node and a Border Node to function as the Overlay control plane for the Fabric and exit point, respectively. This Node can either be a new Node or the control plane and Fabric border functionality can be added to an existing Core if the hardware/software platform supports the functions.

Configure a new VN (Virtual Network) with the new IP subnet on these two Nodes. An IP handoff for this Virtual Network must be configured from the Border Node to the upstream Layer 3-capable Peer Device (*Fusion*). This Peer Device (*Fusion*) can be any Layer 3-capable switch, router, or firewall. You have effectively built a Fabric with just two switches on top of the existing network at this stage. The communication happens between endpoints in the new Fabric and those in the existing network via the IP handoff at the border.

As part of this example migration, we assume all the considerations like MTU have been applied in the existing network as they are a prerequisite before starting the process.

Along with configuring the new switch with the normal management parameters like Network Time Protocol (NTP), Simple Network Management Protocol (SNMP), and access VTY/Console lines, do not forget to configure the system MTU to 9100 bytes. You will also need to configure a Loopback0 interface as this is the preferred management interface of this switch by Cisco DNA Center. The Loopback0 interface also acts as the Routing Locator (RLOC) for the VXLAN encapsulation.

For the Control Plane Node and Border Node, as previously mentioned, you can add the functionality on top of the Core functionality if the platform supports it. For example, if Core-2 in Figure 8.9 is configured with the added Fabric role functionalities, it still provides Core redundancy to the existing network. In other words, Core redundancy is not lost by configuring Core-2 as a Control Plane Node and Border Node.

Figure 8.8: Alternate method of connecting Control Plane Node and Border Nodes to existing network

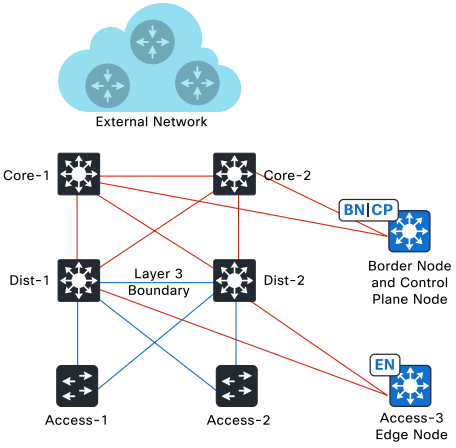

If the Core platform does not support Fabric functionality, or if the preference is that the Core is not modified, another switch can be connected to the existing Core and can be configured with the Fabric functionalities, as seen

in Figure 8.8. In this case, traffic from the Edge Node to the external or existing network will be hairpinning. This is not a big issue since it is directly connected to switches on high speed, high bandwidth Ethernet on the Campus and because it is a small Fabric at this point.

For simplicity's sake, let us assume that the existing Core-2 supports Fabric functionality and is configured as a Control Plane Node and Border Node. The network then looks like Figure 8.9.

Figure 8.9: Fabric network on top of existing network

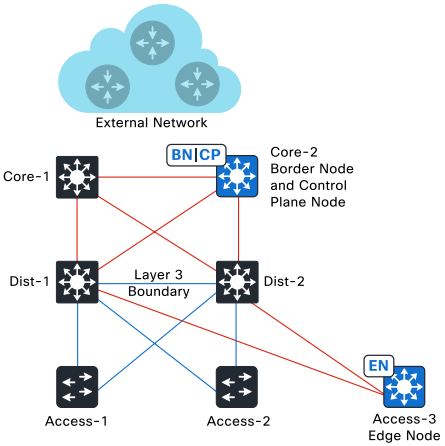

Configure an IP Transit handoff from the Border Node, Core-2, to the Peer Device (*Fusion*) to advertise the Fabric Overlay VNs and the GRT outside the Fabric. The traffic between existing endpoints and endpoints in the Fabric will traverse as shown Figure 8.10.

Figure 8.10: Communication between endpoints in Fabric and existing network

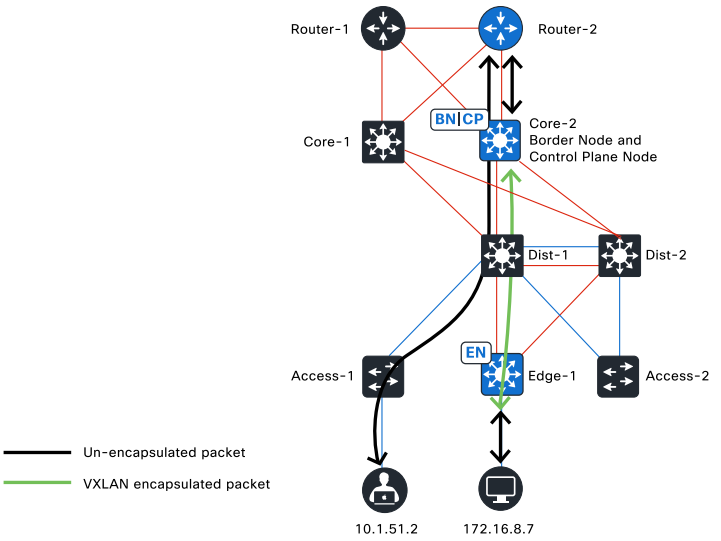

At this point, you can statically assign the client interfaces on the Edge Node to the respective VLANs to start. This will allow you to evaluate connectivity and performance without dealing with 802.1X authentication. If you already have a wired 802.1X infrastructure in place, you can also utilize that as part of the first round of migration tests.

You are ready to move users from existing Access switches to the Edge Node and assess connectivity and performance. If everything looks good and there are no issues, perform a rolling migration by configuring another Access switch as an Edge Node starting with Routed Access connectivity to the Distribution and following the same steps taken to onboard the first Edge Node.

Figure 8.11: Rolling migration increasing Fabric footprint in the network

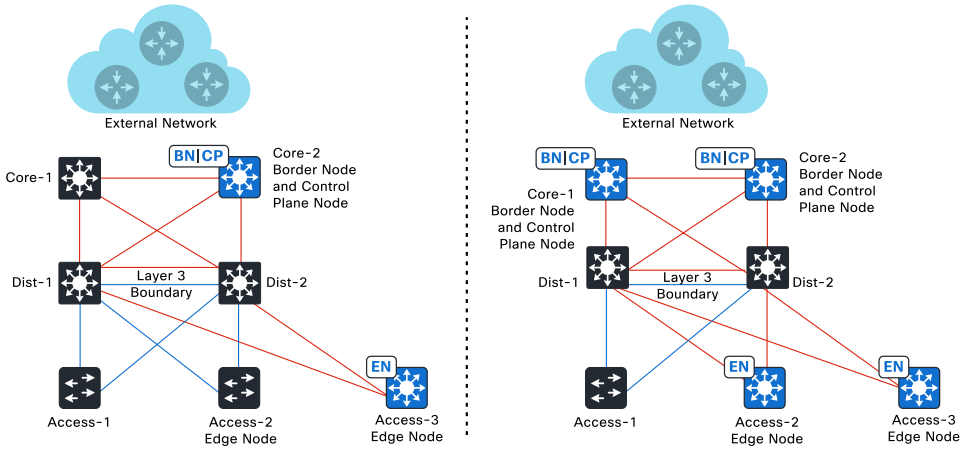

You can add the Control Plane Node and Border Node functions on Core-1 for redundancy purposes as well, as shown in Figure 8.11.

Figure 8.12: Rolling migration of Edge Node

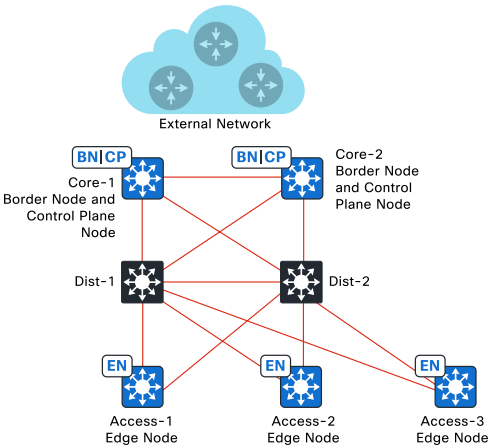

Once you complete the last Access switch in the Fabric migration, you can clean up configurations in the intermediate switches in the network and Core switches. You will now have a fully Fabric-enabled infrastructure, as shown in Figure 8.12.

Migrating a Layer 2 Access with Existing Subnets with Edge Node at Distribution

In this section, we will review the available options for migrating an existing Layer 2 Access network with existing subnets into Cisco SD-Access. Migrations keeping the existing subnets and, in this case, the existing Access switches are a little tricky.

The first option is for you to enable the Fabric at the Distribution Layer instead of the Access Layer. The reasons for doing this are:

- The Access Layer is not compatible with Fabric functionality
- The Access Layer and endpoints underneath it are sensitive to large changes
- Quick win for the IT team to get Fabric segmentation going in the network

Figure 8.13: Sample Layer 2 Access network design

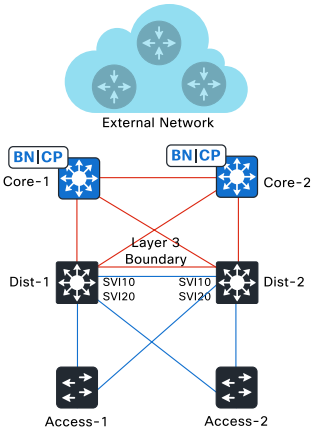

Consider the sample Layer 2 Access as shown in Figure 8.13. VLAN10 (10.1.1.0/24), and VLAN20 (10.1.2.0/24) have Layer 3 boundaries SVI10 and SVI20, respectively, on the Distribution Layer, as shown. Also note that the Border Node and Control Plane Nodes are already built on Core-1 and Core-2. For this example, we used the same guidelines to determine Border Node and Control Plane Node placement as we used in the previous section.

Once the Distribution switches have been discovered in Cisco DNA Center, you begin with shutting down SVI10 and SVI20 on the Distribution switches. Next, you convert the link between Dist-1 and Dist-2 to a Layer 3 link running the Interior Gateway Protocol (IGP) of the existing network. Configure Dist-1 as an Edge Node. Use the Custom VLAN feature within the Cisco DNA Center workflow to configure VLAN10 and VLAN20 as the subnets in the Overlay Virtual Network.

Figure 8.14 shows what the network should look like at this point.

Figure 8.14: Initial steps for migration

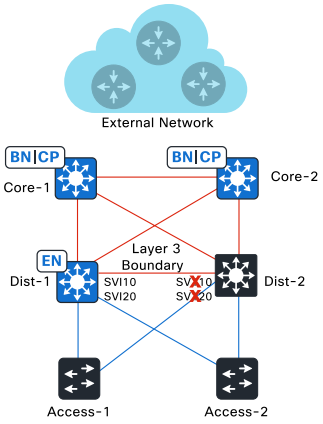

The redundancy at Dist-1, which is now the Edge Node, is provided by in-system redundancy options like a modular Cisco Catalyst 9400 Series chassis with redundant supervisors, Cisco Catalyst 9500 Series switches with Cisco StackWise Virtual (SWV), or Cisco Catalyst 9300 Series switches with Cisco StackWise stack.

The Cisco Catalyst 9000 Series switches (9300 onwards) support a scale of 16,000 endpoints (IPv4). This is a large enough number to onboard four /20 size subnets, or eight /21 size subnets, resulting in a total of 16,000 hosts from the downstream network.

If you want to load-balance and increase the scale of endpoints being onboarded into the Fabric by this method and want to continue with the migration, configure DIST-2 as the second Edge Node and proceed with load-balancing VLAN10 and VLAN20 as shown in Figure 8.15.

Figure 8.15: Load-balance of VLANs

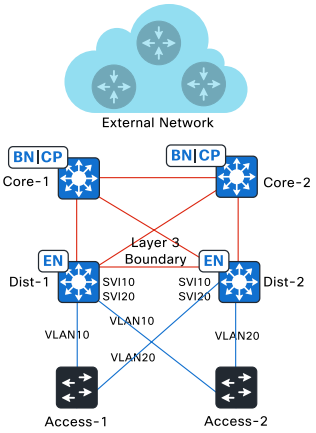

Please ensure that no VLAN is multi-homed into two different Edge Nodes from the Access Layer switches. As seen in Figure 8.15, VLAN10 terminates on Dist-1, and VLAN20 terminates on Dist-2.

In this case, the Fabric segmentation will only start from the Distribution Layer, and East-West traffic below the Distribution Layer is controlled by the Access Layer switches. The micro-segmentation enforcement point is the Edge Nodes at the Distribution Layer.

To proceed with the migration, configure one of the Access Layer switches with routed links upstream to the Distribution Layer switches and then convert it into an Edge Node. Proceed with a rolling migration of Access Layer switches as shown in Figure 8.16.

Figure 8.16: Rolling migration of Access switches

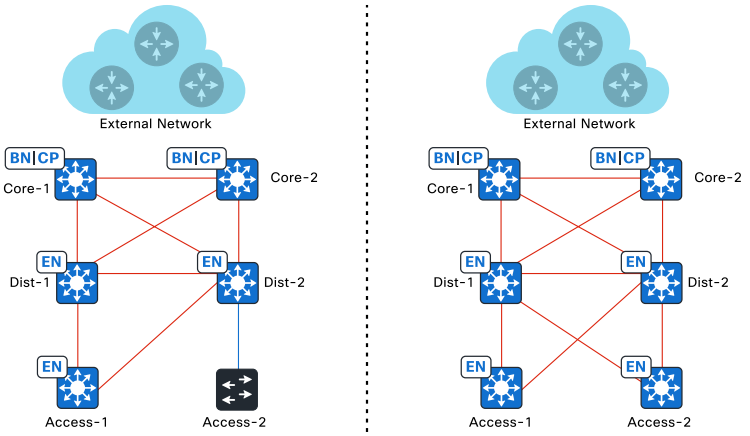

Once all Access Layer switches are configured as Edge Nodes, remove the Edge Node function from the Distribution Layer switches and convert them to normal Intermediate Nodes to complete the migration.

Figure 8.17: Remove Edge Node functionality from the Distribution Layer

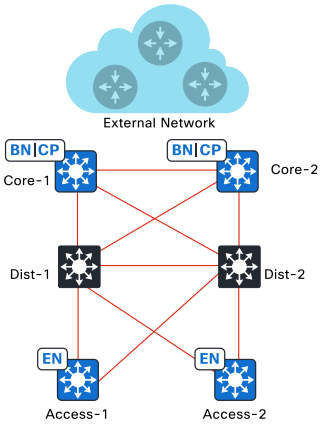

If you have a second building, repeat the same steps in Building #2 to complete the entire Campus migration to Cisco SD-Access, as shown in Figure 8.18.

Figure 8.18: End state of Campus migrated to Cisco SD-Access

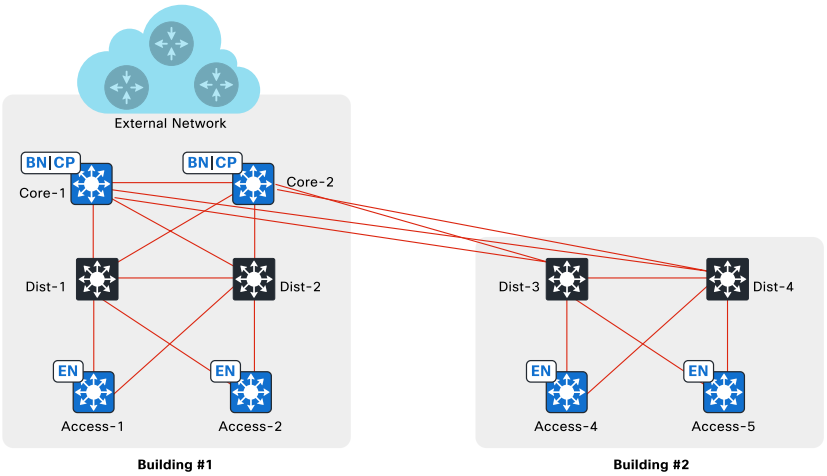

Migrating Layer 2 Access using Layer 2 Handoff

Let us look at the other approach of migration using the existing subnet schema and a Layer 2 Handoff. In this approach, you have a scenario with the same subnet inside and outside the Fabric. To support this, you will use a Layer 2 Handoff on a Border Node to connect the Layer 2 network inside the Fabric with the Layer 2 network outside the Fabric. Consider the scenario shown in Figure 8.19.

Figure 8.19: Sample Layer 2 Access design with Layer 3 boundary at the Core

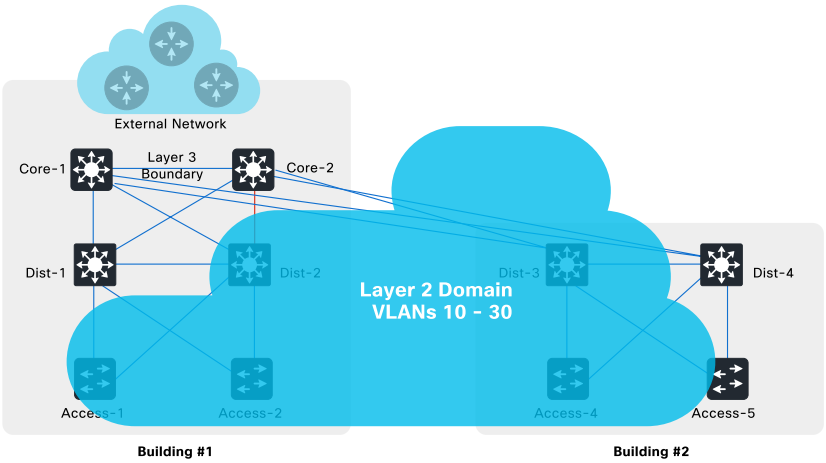

As seen in Figure 8.19, the Layer 3 Boundary is at the network's core. In this type of network, the VLANs span the entire Campus as shown.

We start this migration by introducing a new Access switch in the network. This will play the role of an Edge Node and will be connected via Routed Access links to each existing Distribution switch. This forms the ingress into the Fabric from an endpoint perspective.

You will need to configure a routed path within the Layer 2 access network upstream of this new Access switch to the Layer 3 boundary on the Core Layer. The Core needs to route packets between the Loopback0 interface, the Routed Access interfaces of this new switch, and the management infrastructure upstream of this entire Campus Network. The subnets used in these point-to-point VLANs can later be re-used once that link can be converted to a Layer 3 routed link to achieve Routed Access in the Underlay.

We need the Control Plane Node and Border Nodes to function as the Fabric Overlay control plane and exit point, respectively. This Node can either be a new Node or the Control Plane Node and Border Node functionality can be added to an existing Core if the hardware/software platform supports the functions.

Figure 8.20: High-level concept of migration with switching between existing IP scopes

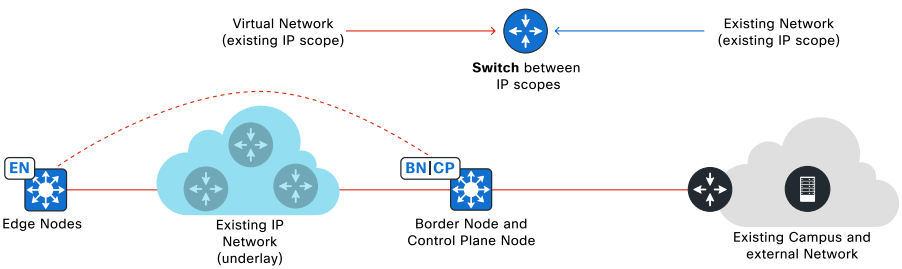

Configure a new Virtual Network with the existing IP subnet on these two Nodes. A Layer 3 IP Transit handoff for this Virtual Network must be configured from the Border Node to the upstream Peer Device (*Fusion*). This Peer Device (*Fusion*) can be any Layer 3-capable switch, router, or firewall.

You will also need to configure a Layer 2 Handoff on the Border Node to support your existing Layer 2 domain. This ensures that communication between the same subnet in the existing network is switched via the Layer 2 Handoff at the border. At this stage, you have effectively built a Fabric with just two switches on top of the existing network. The communication happens between endpoints in the new Fabric and those in the existing network via the IP Handoff at the border.

Figure 8.21: Sample network with VLAN10

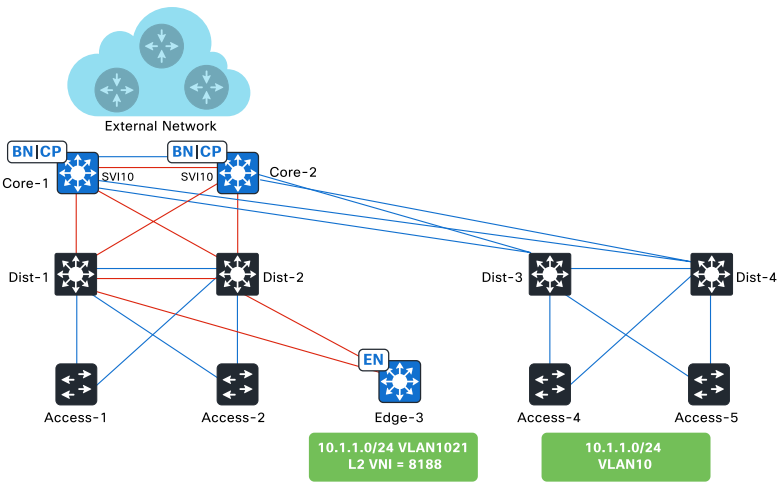

Consider the network shown in Figure 8.20. Existing subnet 10.1.1.0/24 (VLAN 10) is terminated at Layer 3 by SVI10 at the Core-1 and Core-2 switches which are already configured with Control Plane Node and Border Node functionality.

Insert a new switch (Edge-3) and undertake the necessary steps to discover it via Cisco DNA Center. Configure the Edge Node functionality on Edge-3 and use a different VLAN number for the same existing IP subnet 10.1.1.0/24. In the example in Figure 8.21, we see that the 10.1.1.0/24

subnet uses VLAN1021 on the Edge Node with a Layer 2 VNID of 8188. The Layer 2 VNID is automatically assigned when subnet 10.1.1.0/24 is added to the Virtual Network.

Figure 8.22: Initial steps to begin migration to Cisco SD-Access

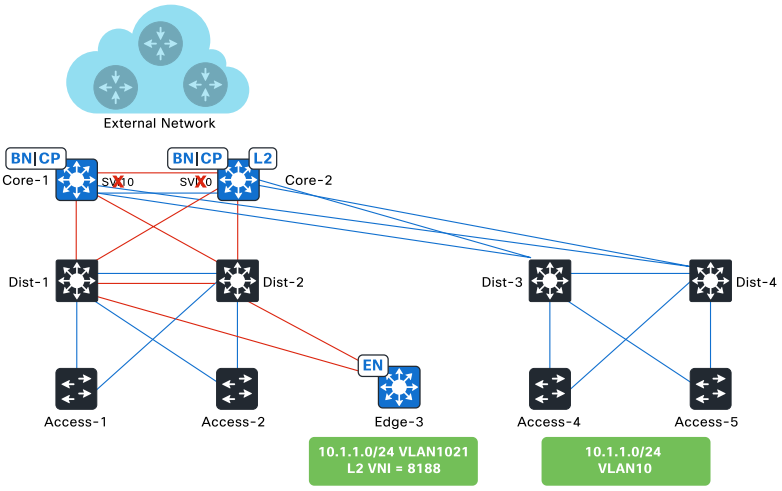

Next, remove SVI10 on Core-1 and Core-2. Configure a Layer 2 Handoff on Core-2, mapping subnet 10.1.1.0/24 to VLAN10. This will place SVI10 in the correct Virtual Network on Core-2 and internally map Layer 2 VNID 8188 to VLAN10 in the configuration.

The redundancy at Core-2, which is now also acting as an Edge Node for clients outside the Fabric Site, is provided by in-system redundancy options such as a modular Cisco Catalyst 9400/9600 Series chassis with redundant supervisors or Cisco Catalyst 9500 Series switches with Cisco StackWise Virtual.

The Cisco Catalyst 9000 Series switches (9300 onwards) support a scale of at least 16,000 endpoints (IPv4). This is a large enough number to onboard

four /20 size subnets or eight /21 size subnets resulting in a total of 16,000 hosts from the Layer 2 Handoff.

Re-configure an existing access switch for Routed Access and perform rolling migration on the Access switches as outlined earlier, moving the endpoints as the Edge Nodes are configured.

Figure 8.23: Rolling migration of Access switches to Edge Nodes with routed links

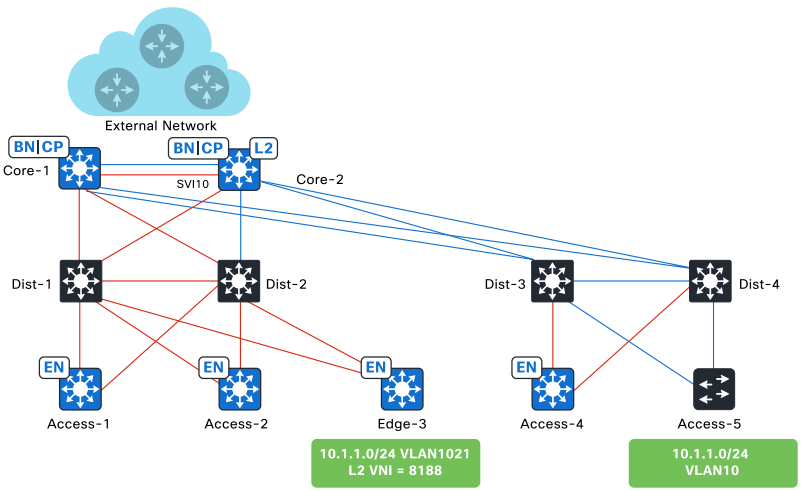

As the last of the Access switches are migrated to Cisco SD-Access, you can remove the Layer 2 Handoff from Core-2 to complete the migration of the Campus Network as shown in Figure 8.24.

Figure 8.24: End state of the network migrated to Cisco SD-Access

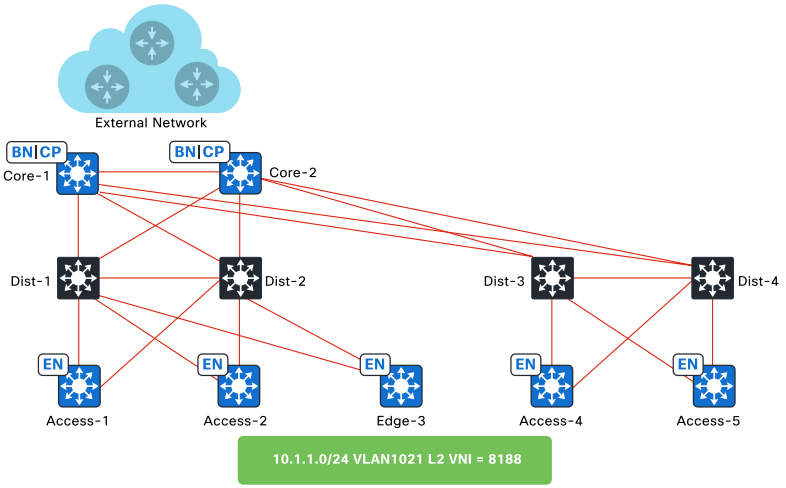

Migrating Existing Routed Access Deployments

Routed Access brownfield networks are easier to migrate than Layer 2 Access designs. Also, you can migrate with existing subnets or new subnets easily.

Consider the sample network topology in Figure 8.25, where VLAN10 is 10.1.1.0/2 and VLAN20 is 10.1.2.0/24.

Figure 8.25: Sample Routed Access topology

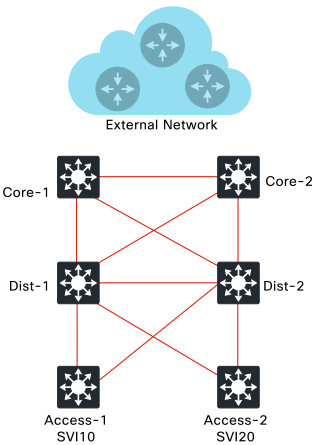

In the topology shown in Figure 8.25, you cannot span VLANs across the Access Layer, hence VLAN10 is localized at Access-1 and VLAN20 is localized at Access-2, respectively. You can begin the migration by

discovering the switches in Cisco DNA Center and using the Fabric workflow to begin configuring Fabric capabilities on the switches.

Before you begin the addition of the Edge Node capability on Access-1, shut down the existing SVI10 on the switch. Proceed with configuring Access-1 as an Edge Node and Core-1 as the Control Plane Node and Border Node as shown in Figure 8.26.

Figure 8.26: Initial steps in migrating Routed Access design

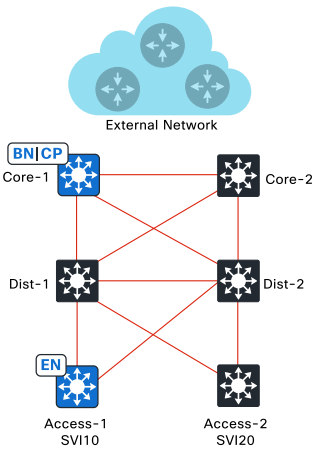

Configure the same existing subnet using a different VLAN ID on the Edge Node. For example, in Figure 8.26, we see that subnet 10.1.1.0/24 now maps to VLAN110, instead of VLAN10. Statically allocate the test client ports to VLAN110 to move the clients into the Fabric and test for connectivity. If you experience issues, you can move the users back to VLAN10 by mapping the client ports to VLAN 10 and un-shutting SVI10. When you have successfully completed the migration of the users to VLAN 110, you can remove SVI10.

Figure 8.27: End state of the migrated network

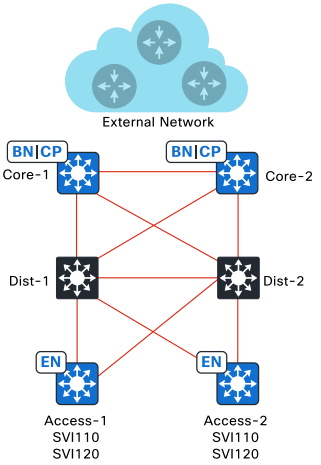

Configure the Control Plane Node and Border Node functionality on Core-2 and enable Edge Node functionality on Access-2. Repeat the same steps for adding subnet 10.1.2.0/24 to the Fabric as you did for subnet 10.1.1.0/24, adjusting the VLAN information accordingly. You will now see VLAN110 and VLAN120 on both the Access switches – a feature you could not get in traditional Routed Access networks.

Integrating Cisco Wireless in Cisco SD-Access Networks

Figure 8.28: Sample wireless network

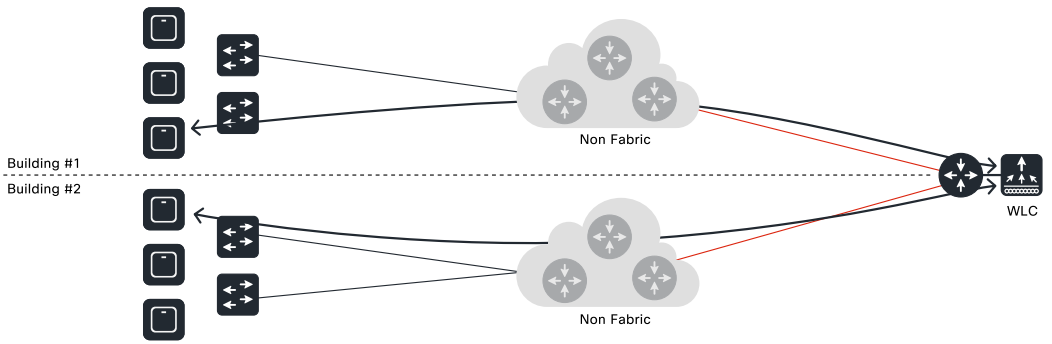

Consider the sample network of two buildings running Cisco centralized wireless with an existing WLC located on the campus, as shown in Figure 8.28. The assumption is that there is no requirement for seamless roaming between the buildings. The wireless integration is per seamless roaming domain. The Access Points (APs) connected to the Access switches build Control and Provisioning of Wireless Access Point (CAPWAP) tunnels to the centralized WLC over which all wireless management, control, and data traffic flows, as shown by the gray arrows from the APs to the WLC.

Figure 8.29: Centralized wireless over a wired Fabric network

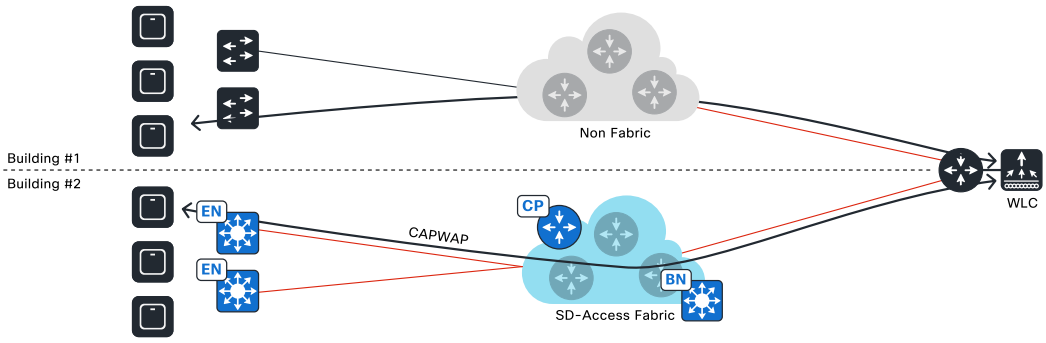

The first step to integrating wireless into Cisco SD-Access Fabric is to create a wired Fabric on the existing network Underlay in Building 2. You can still operate a centralized wireless network on top of the Fabric, as shown in Figure 8.29. There is minimal to no impact on existing wireless networks as you migrate the wired portion of the network into the Fabric.

Cisco recommends a separate WLC be used to integrate wireless into the Fabric in Building 2. This is shown as the blue WLC in Figure 8.30. This recommendation is because we do not want to impact existing wireless on the whole Campus, as a move to Fabric might involve software version changes or even a hardware refresh. Having a new WLC constrains the impact domain to just Building 2 and not the whole Campus. However, if you cannot deploy for a new WLC, you can integrate the existing WLC into the Fabric using Cisco DNA Center workflows. For Fabric Sites, you must create new SSIDs that will be part of the Fabric network, while the older SSIDs will continue to service the non-Fabric portion of the network.

Once you discover the WLC via Cisco DNA Center, you can integrate it into the Fabric in Building 2. You can configure the same SSID on both the new SD-Access WLC and the existing legacy WLC. You can also configure the same RF groups, so APs in the Building 2 wireless network are not seen as

rogue APs in the existing WLC and vice versa. There is no reason to configure the two WLCs in the same mobility group because there is no seamless roaming supported between centralized wireless and SD-Access Wireless today.

The next step in the migration is to schedule a maintenance window and change the primary legacy WLC to the SD-Access WLC in the configuration of APs in Building 2. The APs in Building 2 change their association to the SD-Access WLC. Upon joining the SD-Access WLC, the APs will download the code for Fabric support and will become Fabric-enabled APs. The SSIDs on these APs are locally switched via VXLAN or CAPWAP as per configuration. In Figure 8.30, you can see that the data plane of the wireless network terminates at the Edge Node while the CAPWAP tunnel carries the management and control plane traffic of the wireless network centrally back to the SD-Access WLC.

Figure 8.30: Inter-operation of Integrated and Centralized wireless in Campus

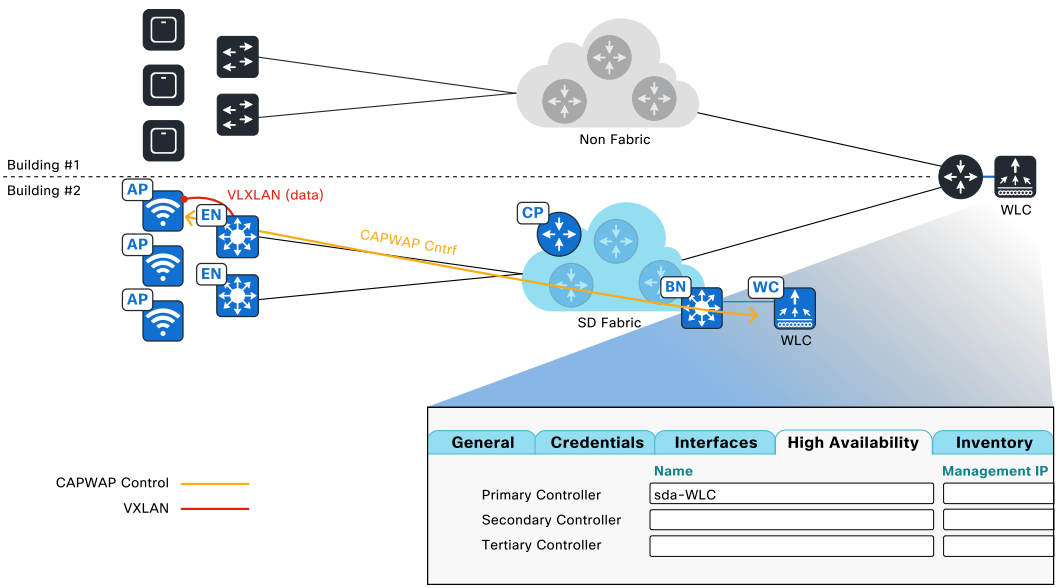

To complete the integration, convert the Building 1 wired network to Cisco SD-Access Fabric and repeat the same steps to integrate the APs into the SD-Access WLC. In this way, you will achieve full wireless integration for the Campus.

What Next?

To familiarize yourself with the Cisco SD-Access solution, we recommend setting up a small lab or a Proof of Concept (PoC) in your network with a couple of Edge Nodes and a couple of Control Plane Nodes and Border Nodes. Go through the motions of configuring a Fabric in Cisco DNA Center by utilizing the Fabric workflow, and make sure to get familiar with Cisco ISE for micro-segmentation. Plan for the border handoffs to integrate the PoC into the existing network.

Also, since things behave a little bit differently from a broadcast perspective in the Fabric compared to traditional Layer 2 Access, it is essential to set up all endpoints in the PoC network and test them out in the Fabric. Use the PoC as a testing ground to test all the endpoints in the network and confirm they behave as expected in the Fabric network before going to production.

Figure 8.31 shows how to connect the PoC to the existing network.

Figure 8.31: Connect PoC network to existing network

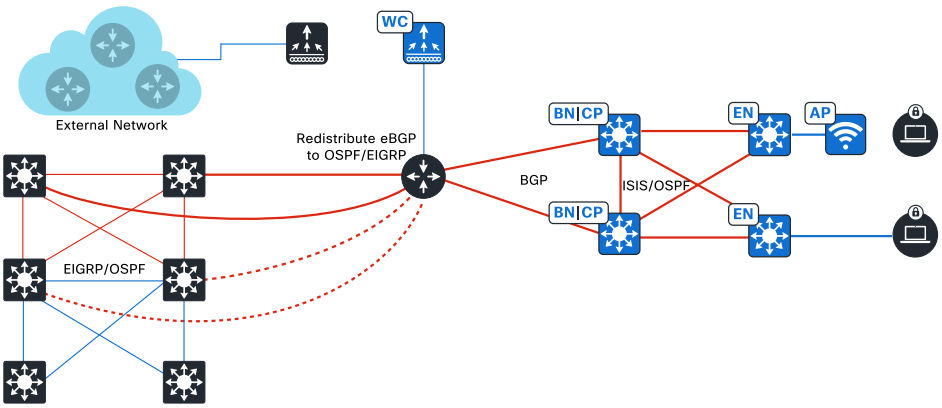

You could connect the border directly to the switches in the existing network, but that will mean introducing a new protocol such as Border Gateway Protocol (BGP) in the existing network. Instead, use a Peer Device (*Fusion*) to peer with the Border Nodes via BGP and redistribute the routes into the IGP of the existing network. This is a minimal change since the peer device is just a new IGP neighbor for the existing network.

As you start to plan your migration, keep these tips in mind:

- Pre-plan and execute management components and their integration, for example, Cisco DNA Center and Cisco ISE
- Backup device configurations and any other information that you need
- Plan an extended maintenance window for the initial days of migration to have a buffer to recover
- Start in low-risk areas, with a small footprint, and preferably IT-friendly or even in the IT department itself
- Be ready with a rollback plan if you face any challenges and have a deadline to recover connectivity for the end users

Summary

After reading this section, you are now convinced that migrating to Cisco SD-Access is realistically possible with as few as two switches connected to the existing network with minimal impact. As you have now seen, you can leverage Cisco DNA Center automation and orchestration to migrate diverse topologies with diverse options such as:

- Migrating Layer 2 Access designs
- Migrating Routed Access designs more easily compared to Layer 2 Access designs
- Migrating with new or existing subnets
- Integrating wireless into Cisco SD-Access

Appendix

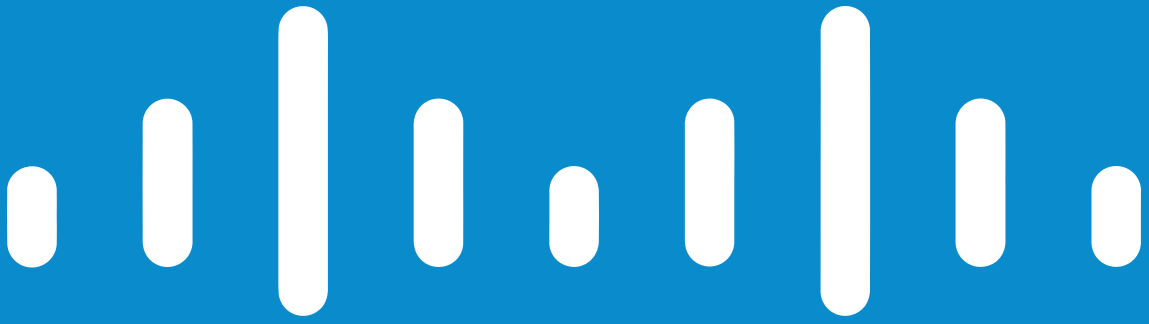

Further Resources and Materials

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

http://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.

http://www.cisco.com/go/dnacenter – Provides an overview and additional information on Cisco DNA Center.

http://cs.co/sda_tech_paper – SD-Access solution white paper. It provides a mid-level technical overview of the Cisco DNA and SD-Access components and their relationships. It is a great place to continue your journey.

http://cs.co/sda-cvd – Includes the Software-Defined Access Cisco validated design (CVD) document covering SD-Access design options, operational capabilities, and recommendations for deployment. Provides direct insight into best practices for designing, operating, and using SD-Access customer network deployments.

http://cs.co/sda-resources – Community resource portal consolidating all collaterals related to the support of Cisco SD-Access.

http://cs.co/sda-youtube – Check out our Cisco SD-Access YouTube Channel, for the latest video updates of features, use cases, demos, and more. Do not forget to subscribe, like, and click the notification on our videos!

https://cs.co/sda-macro-pdg - This guide is intended to provide technical guidance to design, deploy, and operate macro-segmentation across Software-Defined Access Fabric.

https://cs.co/sda-wireless-pdg – Guides integrating wireless access into the SD-Access architecture to gain all the advantages of Fabric and Cisco DNA Center automation.

http://cs.co/gbpa-dg – Provides technical guidance for deploying Group-Based Policy Analytics, covering design topics, deployment best practices, and how to get the most out of the technology operation.

http://cs.co/sda-ml-pdg – A guide for deploying Medium and Large Fabric Sites with Cisco Software-Defined Access.

http://cs.co/sda-distributed-pdg – Guides SD-Access customers in deploying a unified and automated policy across multiple physical locations in a metro-area network.

https://cs.co/independent-domain – Provides design and deployment steps to use the Cisco SD-Access and Cisco SD-WAN solutions to achieve end-to-end segmentation and consistent policy for the enterprise and branch.

https://cs.co/sda-infra-pdg – Provides the deployment guidance of Cisco DNA Center and Cisco Identity Services Engine (ISE) within a services block or data center network connected to a Cisco SD-Access Fabric or traditional Three-Tiered Campus topology.

http://cs.co/mdnac-to-ise – Provides technical guidance to design, deploy, and operate the Software-Defined Access Solution using Multiple Cisco DNA Center Clusters with a single Cisco Identity Service Engine (ISE) deployment.

https://cs.co/sda-segment-sdg – This guide is intended to provide technical guidance to design, deploy, and operate macro-segmentation across Software-Defined Access Fabric.

https://cs.co/en-cvds – A consolidated list of simple, modular, use case-based design and deployment guidance to provide you with Validated designs and best practices. Prescriptive, easy-to-follow deployment guides all with the intent to give you Confidence as you transform your network to meet your business goals.

https://cs.co/sda-el2sd – Recommends practices to be implemented on the Edge Nodes when connecting an External Layer 2 Switching Domain. These best practices are designed to minimize the common issues that arise in the Layer 2 part of the network while protecting the Fabric network against Layer 2 loops and Layer 2 misconfiguration.

http://cs.co/ciscolive-sda – Cisco Live 2022 learning map of Cisco's SD-Access sessions covering comprehensive overview, best practices, design, deployment, migration, and monitoring of the architecture.

http://cs.co/sda_tech_notes – Provides essential technical troubleshooting notes for Cisco SD-Access such as LAN and Guest automation, Fabric provisioning, multicast, wireless, and many others.

http://cs.co/sda-cm – Cisco SD-Access solution compatibility matrix showing what is supported for specific hardware and software versions.

http://cs.co/sda-design-tool-desk – Leverage Cisco SD-Access Design Tool to start your journey into SD-Access solution and Cisco SD-Access Consulting Design Desk for design queries.

Acronyms

AAA–Authentication, Authorization, and Accounting

ACI–Cisco Application Centric Infrastructure

ACL–Access Control List

AD–Microsoft Active Directory

AP–Access Point

API–Application Programming Interface

APIC–Cisco Application Policy Infrastructure Controller (ACI)

ASM–Any-Source Multicast (PIM)

ASIC–Application-Specific Integrated Circuit

ASR–Aggregation Services Router

Auto-RP–Cisco Automatic Rendezvous Point protocol (multicast)

AVC–Application Visibility and Control

BFD–Bidirectional Forwarding Detection

BGP–Border Gateway Protocol

BUM–Broadcast-Unknown Unicast and Multicast

BYOD–Bring Your Own Device

CA–Certificate Authority

CAPWAP–Control and Provisioning of Wireless Access Points Protocol

CMD–Cisco Metadata

CMDB–Configuration Management Database

CNF–Carrier Neutral Facility

CoA–Change of Authorization

CVD–Cisco Validated Design

CWA–Central Web Authentication

DC–Data Center

DHCP–Dynamic Host Configuration Protocol

DMVPN–Dynamic Multipoint Virtual Private Network

DMZ–Demilitarized Zone (firewall/networking construct)

DNA—Cisco Digital Network Architecture

DNS—Domain Name System

DSCP—Differentiated Services Code Point

EA—Cisco Endpoint Analytics

EAP-TLS—Extensible Authentication Protocol-Transport Layer Security

ECMP—Equal-cost Multipath

eduroam—Education Roaming (wireless)

EID—Endpoint Identifier

EIGRP—Enhanced Interior Gateway Routing Protocol

EMR—Electronic Medical Records

EPG—Endpoint Group

ETR—Egress Tunnel Router (LISP)

FEW—Fabric Embedded Wireless

FHR—First-Hop Router (multicast)

FTP—File Transfer Protocol

GBAC—Group-Based Access Control

GBPA—Group-Based Policy Analytics

GOOSE—Generic Object-Oriented Substation Events

GRE—Generic Routing Encapsulation

GRT—Global Routing Table

GUI—Graphic User Interface

HA—High-Availability

HQ—Headquarters

HSR—High-availability Seamless Redundancy

HSRP—Cisco Hot-Standby Routing Protocol

iCAP—Cisco Intelligent Capture

ICMP—Internet Control Message Protocol

IDF—Intermediate Distribution Frame; a wiring closet

IDMZ—Industrial Demilitarized Zone

IEEE—Institute of Electrical and Electronics Engineers

IGP—Interior Gateway Protocol

IoT—Internet of Things

IP—Internet Protocol

IPAM—IP Address Management

IPSec—Internet Protocol Security

ISE—Cisco Identity Services Engine

ITaaB—IT as a Business

L2 VNI—Layer 2 Virtual Network Identifier; as used in SD-Access Fabric, a VLAN

L3 VNI—Layer 3 Virtual Network Identifier; as used in SD-Access Fabric, a VRF

LACP—Link Aggregation Control Protocol

LAN—Local Area Network

LHR—Last-Hop Router (multicast)

LISP—Location Identifier Separation Protocol

LiDAR—Laser Imaging Detection and Ranging

MAB—MAC Authentication Bypass

MAC—Media Access Control Address (OSI Layer 2 Address)

MAN—Metro Area Network

MDF—Main Distribution Frame; the central wiring point of the network

MDM—Mobile Device Management

mDNS—Multicast DNS

MKA—MACsec Key Agreement

MNT—Monitoring Nodes (Cisco ISE persona)

MPLS—Multiprotocol Label Switching

MSDP—Multicast Source Discovery Protocol (multicast)

MSP—Managed Services Provider

MSS—Maximum Segment Size

MTU—Maximum Transmission Unit

NAC—Network Access Control

NAD—Network Access Device

NAT—Network Address Translation

NIC—Network Interface Card

NSF—Cisco Nonstop Forwarding

NTP—Network Time Protocol

OT—Operational Technology

OTT—Wireless Over-the-Top

PAgP—EtherChannel Port Aggregation Protocol

PAN—Primary Administration Node (Cisco ISE persona)

PBR—Policy-Based Routing

PCI DSS—Payment Card Industry Data Security Standard

PIM—Protocol-Independent Multicast

PLC—Programmable Logic Controllers

PnP—Plug-n-Play

PRP—Parallel Redundancy Protocol

PSN—Policy Service Node (Cisco ISE persona)

PTP—Precision Time Protocol

pxGrid—Platform Exchange Grid (Cisco ISE persona and publisher/subscriber service)

QoS—Quality of Service

RADIUS—Remote Authentication Dial-In User Service

RBAC—Role-Based Access Control

REP—Resilient Ethernet Protocol

REST—Representational State Transfer

RFC—Request for Comments Document (IETF)

RIB—Routing Information Base

RLOC—Routing Locator (LISP)

RMA—Return Material Authorization (Cisco DNA Center workflow)

RP—Rendezvous Point (multicast)

RRM—Radio Resource Management (wireless)

RSTP—Rapid Spanning Tree Protocol

RSTP—Real-Time Location Service

RSTP—Supplicant-Based Extended Nodes

SCADA—Supervisory Control and Data Acquisition

SD—Software Defined

SDA—Cisco Software-Defined Access

SDG—Cisco Service Discovery Gateway

SDN—Software-Defined Networking

SGACL—Security-Group ACL

SGT–Security Group Tag

SNMP—Simple Network Management Protocol

SSID—Service Set Identifier (wireless)

SSM—Source-Specific Multicast (PIM)

SSO—Stateful Switchover

STP—Spanning-tree protocol

SV–StackWise Virtual

SVL–Cisco StackWise Virtual

SXP–Security Group Tag Exchange Protocol

Syslog–System Logging Protocol

TCP–Transmission Control Protocol (OSI Layer 4)

TLOC–Transport Locator

UDP–User Datagram Protocol (OSI Layer 4)

URL–Uniform Resource Locator

UTP–Unshielded Twisted Pair (Cable)

VLAN–Virtual Local Area Network

VN–Virtual Network, analogous to a VRF in SD-Access

VNI–Virtual Network Identifier (VXLAN)

VPN–Virtual Private Network

VRF–Virtual Routing and Forwarding

VXLAN–Virtual Extensible LAN

WAB–Wide Area Bonjour

WAN–Wide-Area Network

WCCP–Web-Cache Control Protocol

Devi Bellamkonda
Dhrumil Prajapati
Jonathan Cuthbert
Kedar Karmarkar
Keith Baldwin
Mahesh Nagireddy
Parthiv Shah
Pete Kavanagh
Prakash Jain
Prashanth Kumar Davanager Honneshappa
Raja Janardanan
Sanjay Hooda
Scott Hodgdon