FACULTY
OF INFORMATION
TECHNOLOGY
CTU IN PRAGUE

Assignment of bachelor's thesis

Title:	Návrh migrace reálné velké firemní sítě na softwarově definovanou síť
Student:	Diana Prokopisina
Supervisor:	Ing. Alexandru Moucha, Ph.D.
Study program:	Informatics
Branch / specialization:	Computer Networks and Internet 2021
Department:	Department of Computer Systems
Validity:	until the end of summer semester 2024/2025

Instructions

Software-Defined Networks (SDN) are the cutting-edge trend in networking technologies as they offer unparalleled flexibility, services, central monitoring and management.
The aim of the thesis is to study the migration process of an anonymized existing large classic enterprise computer network, called CNX, to the SDN technology

- Analyze the current design of CNX and identify its drawbacks.
- Redesign both the internal local network architecture of CNX and the WAN interconnecting the different locations of CNX into SDN. Since CNX is implemented using Cisco devices, reuse as much as possible the current network devices in order to minimize the migration costs and base your SDN redesign also on Cisco technologies, mainly Cisco SD-WAN (for the interconnection) and Cisco SD-Access (for the local networks).
- Design a technical plan to migrate CNX into SDN.
- Perform a partial implementation of the SDN redesign in the EVE-NG simulator at FIT CVUT for the SD-WAN proposal, using either Cisco SD-WAN virtualized devices or Viptela virtualized technologies.
- Assess the capability of the EVE-NG simulator to simulate the proper functioning of SDN and its suitability as a teaching tool used to demonstrate the principles of SDN.
- Analyze the computational and memory requirements of the SDN simulation.

Electronically approved by prof. Ing. Pavel Tvrdík, CSc. on 8 January 2024 in Prague.

Bachelor's thesis

DESIGN OF MIGRATION FOR A REAL LARGE ENTERPRISE NETWORK TO SOFTWARE-DEFINED NETWORK TECHNOLOGIES

Diana Prokopisina

Faculty of Information Technology
Department of Computer Systems
Supervisor: Ing. Alexandru Moucha, Ph.D.
May 16, 2024

Czech Technical University in Prague
Faculty of Information Technology

Citation of this thesis: Prokopisina Diana. *Design of Migration for a Real Large Enterprise Network to Software-Defined Network Technologies.* Bachelor's thesis. Czech Technical University in Prague, Faculty of Information Technology, 2024.

Contents

List of Figures

I would like to thank the supervisor of this thesis, Ing.Alexandru Moucha, Ph.D., a Cisco expert and enthusiast, for inspiration and support. I am also grateful to my boyfriend for immense help with perfecting this work, especially for providing meaningful comments and suggesting corrections. Last but not least, I would like to thank Bc. Josef Zápotocký for providing guidance, support and help with practical part of the thesis.

Declaration

I hereby declare that the presented thesis is my own work and that I have cited all sources of information in accordance with the Guideline for adhering to ethical principles when elaborating an academic final thesis.

I acknowledge that my thesis is subject to the rights and obligations stipulated by the Act No. 121/2000 Coll., the Copyright Act, as amended, in particular that the Czech Technical University in Prague has the right to conclude a license agreement on the utilization of this thesis as a school work under the provisions of Article 60 (1) of the Act.

In Prague on May 16, 2024

Abstract

This Bachelor thesis focuses on the concept of Software-Defined Networking (SDN), that is aimed at addressing challenges of traditional networks. Theoretical description of Software-Defined and traditional networks is followed by presenting two proprietary Cisco SDN solutions: Cisco SD-Access and Cisco Catalyst SD-WAN. Their application to improve functionality of real networks is demonstrated on an example of an existing large enterprise. The thesis proposes redesign and migration plan of the legacy network. A partial implementation of Cisco Catalyst SD-WAN part of the redesign is performed in EVE-NG simulation environment, provided by FIT CTU. It serves as proof of concept and can be used for demonstrating Cisco Catalyst SD-WAN functionality in networking laboratories.

Keywords SDN, Cisco SD-Access, Cisco Catalyst SD-WAN, EVE-NG

Abstrakt

Tématem této bakalářské práce je koncept softwarově definovaných sítí (SDN), který je zaměřen na řešení problémů tradičních sítí. Práce začíná teoretickým popisem softwarově definovaných a tradičních sítí, po kterém následuje představení dvou proprietárních řešení SDN společnosti Cisco: Cisco SD-Access a Cisco Catalyst SD-WAN. Jejich použití pro zlepšení funkčnosti reálných sítí je demonstrováno na příkladu existujícího rozsáhlého podniku. Práce navrhuje redesign a plán migrace dané starší sítě. Částečná implementace oblasti pokryté Cisco Catalyst SD-WAN v redesignu dané sítě je provedena v simulačním prostředí EVE-NG, které bylo poskytnuto od FIT ČVUT. Slouží jako proof-of-concept a lze ji využít pro demonstraci funkčnosti Cisco Catalyst SD-WAN v síťových laboratořích.

Klíčová slova SDN, Cisco SD-Access, Cisco Catalyst SD-WAN, EVE-NG

List of abbreviations

5G	5th Generation
AAA	Authentication, Authorization, and Accounting
AAR	Application-Aware Routing
ACL	Access Control List
AP	Access Point
API	Application Programming Interface
ARP	Address Resolution Protocol
ASIC	Application Specific Integrated Circuit
AWS	Amazon Web Services
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
CAPWAP	Control and Provisioning of Wireless Access Points
CE	Customer Edge
CLI	Command Line Interface
CMS	Cisco Metadata Header
CPU	Central Processing Unit
DevOps	Development Operations
DHCP	Dynamic Host Configuration Protocol
DMVPN	Dynamic Multipoint Virtual Private Network
DNA	Digital Network Architecture
DNS	Domain Name System
DTLS	Datagram Transport Layer Security
EID	Endpoint IDentifier
EIGRP	Enhanced Interior Gateway Routing Protocol
ERS	External RESTful Services
ESA	Email Security Appliance
FCAPS	Fault, Configuration, Accounting, Performance and Security
FHRP	First Hop Redundancy Protocol
FMC	Firepower Management Center
FQDN	Fully Qualified Domain Name
FTD	Firepower Threat Defense
GBAC	Group Based Access Control
Gbps	Gigabit per second
GiB	GibiByte
GPO	Group Policy Option
GRE	Generic Routing Encapsulation
GRT	Global Routing Table
GUI	Graphical User Interface
HSRP	Hot Standby Router Protocol
HTDB	Host Tracking DataBase
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IDS	Intrusion Detection System

IoT	Internet of Things
IP	Internet Protocol
IPAM	Internet Protocol Address Management
IPS	Intrusion Prevention System
IPsec	Internet Protocol SECurity
ISE	Identity Services Engine
ISP	Internet Service Provider
ISR	Integrated Services Router
IPvn	IP version n
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network
LISP	Locator ID Separation Protocol
Ln	Layer n
LTE	Long-Term Evolution
MAB	MAC Authentication Bypass
MAC	Media Access Control
MACsec	Media Access Control Security
mGRE	Multipoint Generic Routing Encapsulation
MnT	Monitor and Troubleshooting Node
MPLS	MultiProtocol Label Switching
MR	Map Resolver
MS	Map Server
NETCONF	NETwork CONFiguration protocol
NFV	Network Functions Virtualization
OMP	Overlay Management Protocol
ONF	Open Networking Foundation
ONOS	Open Network Operating System
OS	Operating System
OSI	Open Systems Interconnection
OSPF	Open Shortest Path First
PAN	Policy Administration Node
PC	Personal Computer
PnP	Plug and Play
PSN	Policy Service Node
pxGrid	Platform Exchange Grid
QoS	Quality of Service
RAM	Random-Access Memory
REST	REpresentational State Transfer
RLOC	Routing LOCator
RSA	Rivest–Shamir–Adleman
RSVP	Resource Reservation Protocol
SD	Software-Defined
SDN	Software-Defined Networking
SG	Scalable Group
SGACL	Security Group Access Control List
SGT	Scalable Group Tag
SNMP	Simple Network Management Protocol
SSO	Stateful SwitchOver
STP	Spanning Tree Protocol
SVI	Switch Virtual Interface
SSH	Secure SHell protocol
SSID	Service Set IDentifier

SXP	Scalable Group Tag Exchange Protocol
TCP	Transmission Control Protocol
TLOC	Transport LOCator
TLS	Transport Layer Security
UDP	User Datagram Protocol
vCPU	Virtual Central Processing Unit
VID	Virtual Local Area Network Identifier
VLAN	Virtual Local Area Network
VM	Virtual Machine
VN	Virtual Network
VNI	Virtual Network Identifier
VRF	Virtual Routing and Forwarding
VRRP	Virtual Router Redundancy Protocol
VPCS	Virtual Personal Computer Simulator
VPLS	Virtual Private Local area network Services
VPN	Virtual Private Network
VSS	Virtual Switching System
VXLAN	Virtual Extensible Local Area Network
WAN	Wide Area Network
WLC	Wireless Lane Controller
WSA	Web Security Appliance
ZTP	Zero Touch Provisioning

Chapter 1
Introduction

1.1 Definitions

The thesis is focused on technologies in the area of computer networking. Definitions of several basic concepts in the above mentioned area are presented below for a better understanding of the text.

▶ **Definition 1.1** (Computer network [1]). *Computer network is a wired or wireless connection of two or more computing devices that use networking protocols to exchange, transmit, share data and resources with each other.*

Computing devices being part of the interconnect can be classified into two broad categories:

- End devices: source/destination of transmitted information. Typically, PCs, laptops, mobile or IoT devices are classified as end devices.
- Intermediate devices: devices, responsible for transmitting information between the end devices, by building and selecting optimal paths in the interconnect, forwarding the information units along these paths, filtering data based on various policies. Typically, routers, switches, firewalls [2] are classified as intermediary devices.

It is important to keep in mind, that such categorization is not strict and a device can be classified as an end or intermediate device, depending on the topology/context of communication. For example, a router may be considered an end device, when an administrator communicates to it via SSH.

Computer networks can be classified based on their span and purpose, the categories worth mentioning in the context of the thesis being [1]:

- LAN: a small-scale network, typically interconnecting devices within one physical location. Access networks, connecting user devices to WAN or ISP networks, are usually classified as LANs.
- WAN: a large-scale network, interconnecting multiple, typically geographically distant locations or different networks. The term may be used to refer to interconnection of organization's branches or to Internet, if considered as interconnected international networks.

▶ **Definition 1.2** (OSI and TCP/IP models [3])**.** *OSI and TCP/IP are the two most common system communication models referred to nowadays, which separate the flow of data into logical layers. Each layer has well-defined functions and provides an abstraction of the lower-level functionality for the higher levels. OSI model defines seven layers, TCP/IP – four.*

Both models serve as frameworks for developing standardised system communication protocols, OSI being more generalized and theoretical, TCP/IP – tailored for the use in Internet and similar networks.

Communication of data over the network is possible by cooperation of different protocols, operating at different layers. Lower layer protocols abstract their functionality for higher level protocols by encapsulating the data – adding headers and footers, which carry layer-specific information. Decapsulation is the process opposite to encapsulation [3].

The layers of the theoretical OSI model, important in the context of the thesis, are [4]:

- Layer 2 (Data Link Layer): responsible for data transfer within one network segment. Switches typically operate at this layer. MAC address is an example of L2 physical address.
- Layer 3 (Network Layer): responsible for data transfer between network segments. Routers typically operate at this layer. IP address is an example of L3 logical address.

▶ **Definition 1.3** (Address (networking))**.** *An address in networking is an identifier of a device, more precisely of its physical or virtual network interface [5], that facilitates reaching the device through this interface over the network.*

Network addresses are usually designed to be globally unique, however, the use of virtual networks, or private networks with private to public address translation in play allows them to be duplicated without disrupting the normal network operations.

Devices may have several types of addresses assigned for the use of different protocols.

▶ **Definition 1.4** (Networking protocol [6])**.** *Networking protocol is a set of rules that allows networking devices to communicate. Protocols typically define syntax, semantics, synchronisation of communication and error recovery mechanisms.*

Networking protocols exist for every OSI layer and serve different purposes, from exchanging messages at L2 (Ethernet family protocols), to monitoring the network (SNMP).

▶ **Definition 1.5** (Network policy)**.** *"Network policy is a collection of rules that govern the behaviors of network devices" [7]. Policies are applied by configuring corresponding protocols.*

Network policies can be categorized based on their objective. Access and security policies define which resources on the network are accessible by which users/devices, and can be enforced by firewall rules or ACLs. QoS policies allow to prioritize certain traffic classes [8, p 24], like traffic to/from specific applications with the use of, for example RSVP protocol to reserve bandwidth for each traffic flow.

1.2 Traditional Networks [9, 10, p 87, 8, p 20]

It may be debatable what network is referred to as traditional, because there is a vast number of architectures developed over the years, and it would not be enough to select just one as a basis of definition. Moreover, what is called non-traditional or innovative today may become a new standard in future.

However, it is possible to present a generalized description of approaches, that have been widely used in designing, building and managing networks for the past several decades.

Traditional networks are usually physical networks, built by interconnecting physical devices, which run protocols serving different purposes (communication, policy enforcement, monitoring). Protocols have to be configured by network administrators to provide the desired functionality.

The devices are composed of hardware (ASICs) and software (OSs), both customized to perform specific functions (switching, routing, load-balancing). ASICs are optimized to efficiently process network traffic (forward, filter), OSs control device operation by managing hardware resources and providing CLI for configuring protocols. Such a strong integration of network functions in devices makes operation of traditional networks fast and efficient, but dependent on devices' capabilities and their physical interconnection. For example, a flow of traffic, that requires firewall inspection, depends on the physical location of the firewall. A change to the direction of traffic or enforcing inspection earlier/later in the path would require repositioning the firewall, rewiring, making configuration changes to the affected devices.

Devices in traditional network typically implement all network logical planes (data, control, management) locally. Thus, to affect traffic flows or to configure/monitor devices, each device's control and management plane respectively have to be accessed independently via CLI. Distributed control and management planes offer both advantages, such as lack of a single point of failure and higher scalability, and disadvantages, such as complexity of managing a distributed system consisting of many nodes. The complexity comes from the need of accessing and manually configuring each device, which is time-consuming and error-prone. There are network management protocols, such as SNMP, available to centrally monitor and configure the devices, but they typically suffer from limited functionality, limited scalability, complex configuration and lack of support in devices.

Devices typically run proprietary software, which provides vendor-specific CLI and which can't be rewritten to better suit current needs. Additionally, they lack an identical level of support of various technologies and protocols, for example, monitoring and configuration protocols: SNMP, NETCONF. Moreover, many vendors design proprietary protocols, which are even less likely to be supported by devices from a different vendor. These properties impact management of heterogeneous networks, such networks allowing to use best-of-breed devices, but suffering from issues with devices' interoperability and non-uniform management approaches.

Traditional networks have been existing for decades and were capable of keeping up with demands of the time. For this to be possible, various design guidelines (multi-tiered architecture, network segmentation), technologies (virtualization), protocols have been developed and used.

However, traditional approaches with their limitations are just not suitable to meet current requirements. Nowadays businesses and service providers are expected to keep up with fast evolving technology trends (cloud computing, big data, multimedia), an ever-increasing number of connected devices (mobile, IoT devices), as well as rapidly launch new services. SDN, discussed in the following chapter, is one of the solutions to address these needs.

1.2.1 Challenges

The following list presents limitations of traditional networking, discussed above, as well as describes additional, more specific challenges, typically encountered in network administration.

Management [11, p. 29]
"For more than three decades, network management has been entirely based on the command-line interface (CLI) and legacy protocols such as SNMP" [11, p. 29]. Thus, traditional networks have been managed by manually configuring networking equipment and using limited-capability network monitoring tools.

While in small scale such approaches may be efficient, managing large-scale networks in such a fashion is complicated and time-consuming. Many organizations seek solutions that provide network programmability and automation for easier, faster, less erroneous management and allow to centrally manage the whole infrastructure. However, limitations of traditionally used tools don't allow them to be a basis for programmable and automated networks.

- CLI: Manual device configuration via CLI in large scale is cumbersome, time-consuming and error-prone. Moreover, CLI "is vendor specific, lacks a unified data hierarchy (sometimes even for platforms from the same vendor), and was designed primarily as a human interface" [11, p. 29]. Thus, while CLI scripts may assist in automating management tasks to reduce configuration time and probability of error, they require knowledge of commands, may be complex to write, are vendor-specific, lack cross-platform compatibility and do not present data in a normalized format, which could be used by applications or in machine-to-machine communication.

- SNMP [12, p. 28-29]: SNMP is capable of reading and updating some networking device properties (up/down status, CPU, memory, space utilization), it is supported by nearly each networking device, there are Python libraries written for it – SNMP may be used as a programmatic interface to networking devices. However, SNMP is not highly scalable, is difficult to implement and its configuration management capabilities are not fully supported by all vendors. Thus, SNMP may be suitable for collecting basic information in large infrastructures, but not to configure them programmatically.

Network Visibility [8, p. 26-27]
Organizations often lack an adequate level of visibility into the network. The existing network monitoring tools, like SNMP are either not used, are limited in functionality, or availability of those is limited on some platforms, which makes it complicated to have a comprehensive view of the network, easily troubleshoot it and address issues proactively.

Most often to find a problem network administrators have to analyze extensive devices' configurations. Moreover, without network visibility in play they usually correct the issue itself, rather than analyzing its cause and addressing it instead.

Wireless Segment [8, p. 22]
Traditionally, wired and wireless network segments are configured and managed separately, thus, making it difficult to apply consistent policies over the whole network. Moreover, wireless networks do not support advanced segmentation techniques, like VRF/VLAN on the wired level, which makes it challenging to logically separate traffic flows.

There is a possibility to dedicate separate SSIDs to different user/device groups as an alternative to L2 segmentation. However, these are limited in number and would need to be mapped to VLANs at WLCs, when packets cross the wireless/wired boundary, which introduces additional configuration steps. WLC is unaware of L3 segmentation concept whatsoever, making it difficult to keep consistency in L3 traffic separation between wired and wireless layers.

Network Segmentation [8, p. 22-23]
Most common technologies for L2/L3 network segmentation are VLAN/VRF.

- L2 segmentation (VLAN): To achieve the desired level of user/device separation large organizations end up managing hundreds/thousands VLANs. This complicates IP address planning, decreases visibility of the network and complicates troubleshooting, as it is challenging to keep track of device-VLAN mapping, increases the configuration complexity, because each new VLAN requires adding it to necessary trunks, possibly configuring protocols for loop avoidance and adding new ACLs to control the inter-VLAN traffic flow. Widely-spanned VLANs are vulnerable to L2 loops when loop avoidance protocols are misconfigured. Large L2 designs introduce a high degree of port blocking.

- L3 segmentation (VRF-Lite): VRF-Lite configuration in large networks may become complicated in case of, for example, implementing multiple 802.1q trunks among devices. VRF-Lite deployments are not highly scalable, as more that 8-10 VRFs are cumbersome to manage. Each VRF-Lite process requires a separate routing protocol instance to run on a device, increasing CPU load.

Policies [8, p. 24-25]
Typically, an organization manages several different policies: security and access policies with the help of network segmentation (VLANs, VRFs), ACLs and firewall rules, QoS policies with leveraging, for example, QoS capabilities of protocols (RSVP) and network technologies (MPLS). Use of policies is vital to control and direct traffic flows in the network, as well as manage access to resources and maintain an adequate level of security.

However, especially in a large scale, designing and maintaining policies is complicated. This is caused by the fact that many organizations lack visibility into the network and means to easily identify devices, which inhibits them from effortlessly dividing network resources into logical groups and defining relationship between them. Moreover, in traditional networks applying a policy means manually configuring all policy-aware devices, which is time-consuming and error-prone. It is also a common practice not to remove the existing policy rules when there's a need to apply changes, because typically networks lack a way to provide a human-friendly context into the current policy configurations, making it difficult for administrators to understand the purpose of each rule and to safely remove/alter them without breaking the previously defined logic.

User and Device Identity [8, p. 25-26]
Organizations often struggle with correctly and securely establishing user/device identities and transferring the identity information through the network.

Typically a connected device is assigned a VLAN and an IP address statically or dynamically. Static assignment – hard-coding a VLAN/subnet to a wired port or a wireless SSID – offers little security, because any device connecting to this port/SSID will be granted an equal level of access. Dynamic assignment using various authentication methods, such as 802.1x, comes with limitations such as configuration complexity and lack of support in devices.

In typical deployments IP address and VLAN are ultimately used to determine level of user/device access to the network resources, which comes with disadvantages such as dependency on network architecture changes (IP address space or VLAN segmentation re-planning), lack of mobility (different network segments have different IP address scopes and are aware of different VLANs, making it difficult to grant the same level of access to migrated user), etc.

Devices' Interoperability
Heterogeneous networks, very commonly built to employ functionality of best-of-breed devices from different vendors, suffer from interoperability issues and are complex to manage. Interoperability is influenced by the fact, that vendors may choose not to support some protocols or implement them differently. Complexity in management arises from differences in configuration approaches among vendors and sometimes even among different devices from the same vendor.

1.3 Introduction to SDN

1.3.1 Definitions

The terms presented below are frequently used in SDN context and need to be defined for a better understanding of the text.

▶ **Definition 1.6** (Network Management [11, p. 3]). *Network management encompasses tools, methodologies, standards, aimed at defining and monitoring network behaviour. There is an extensible list of all the network aspects, that can be managed. FCAPS is an exemplar frame-*

work, that groups them in the following areas: fault, configuration, accounting, performance and security management.

▶ **Definition 1.7** (Automation [11, p. 5-6]). *Automation aims at reducing human intervention and increasing software employment in performing tasks. Trivial examples include scripting, or protocols, such as DHCP, which automatically provides hosts with IP addresses. Less trivial is, for example, automated device provisioning (creating configuration files and pushing them to devices), with implementation differing among vendors.*

Automation brings the advantages of reduced human error in network configuration, reduced degree of performing repetitive tasks, agility in implementing network changes, consistency and standardization of system changes.

▶ **Definition 1.8** (Orchestration [11, p. 6-7]). *Orchestration refers to automating workflows, in contrast to automation intended to automate single tasks. In other words, orchestration coordinates automated tasks execution in a specific order across different systems to accomplish a single objective. This, similar to automation, reduces human error and saves time.*

Examples of open-source orchestration tools include Kubernetes [13], a container orchestration system, and Apache Airflow [14], a batch workflow orchestration system.

▶ **Definition 1.9** (Programmability [11, p. 7-8]). *Programmability is the ability of a system to accept a new set of instructions that may alter its behavior.*

A slightly different view on network programmability suggests that it is the ability to manage network devices through a programmable interface, formally known as API. APIs allow communication between software programs and devices.

▶ **Definition 1.10** (Virtualization and abstraction [11, p. 8-9]). *Network abstraction refers to hiding the complexity and unnecessary low-level details of network resources, services and functionality. This allows higher level applications to use them without being aware of implementation details.*

Virtualization is creation of a software device or service, behaving like its physical counterpart. Examples include virtual devices (servers, routers, firewalls) in forms of VMs or containers, virtual L3/L2 isolated networks created using VRFs/VLANs.

Virtualized services are usually built on top of abstractions. VXLAN, a network virtualization technology, can be used as an example to illustrate that. VXLAN extends L2 networks over L3 infrastructures with the help of tunnels, creating virtual L2 topologies. It abstracts L3 transport from end devices to create an illusion they physically belong to the same L2 network. This is achieved by encapsulating Ethernet frames and VXLAN header into UDP packets at the source tunnel endpoint, and decapsulating at the destination one.

1.3.2 SDN Overview

SDN is an approach to networking, intended to address limitations of the traditional one by increasing the degree of software employment in defining network behaviour. This is typically achieved by abstracting the networking resources and allowing applications to control these abstractions programmatically [11, p. 13].

SDN is often associated with OpenFlow-based centralized control plane approach (1.3.3.1), which is considered to be at the roots of SDN revolution. However, other software-defined solutions have emerged over the years, making it difficult to provide a comprehensive SDN definition [12, p. 5]. To address this challenge, selected SDN-related technologies are described in 1.3.3.

1.3.3 SDN-Related Technologies

The technologies presented below are among the ones, that are often associated with SDN. These were selected, because they or their modified versions are used in Cisco SD-Access and Catalyst SD-WAN solutions, which are the main topic of the thesis. The naming convention for Cisco's solution components is discussed in 1.4.2. The solutions are introduced in 1.4.3 and 1.4.4, and described in detail in 2.1 and 2.2.

1.3.3.1 Centralized Control Plane [12, p. 6]

Centralized control plane solutions are based on decoupling network devices' control and data forwarding planes, and consolidating control functions into a centralized location: a software-based controller. Controllers collect reachability information from the underlying infrastructure, centrally make forwarding decisions and download forwarding rules to devices using a specialized protocol (e.g. OpenFlow). Controllers typically expose APIs for applications to program the control plane.

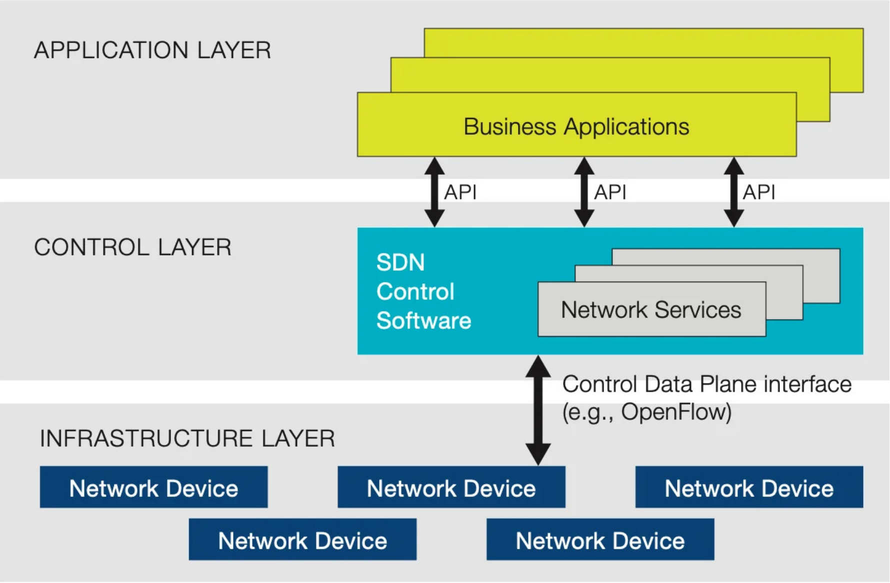

Figure 1.1 Centralized Control Plane Architecture [15]

Centralized control plane approach allows to maintain an inventory of simple cheap networking devices, which perform data forwarding functions only. It additionally comes with benefits of programmable control plane and increased granularity in controlling traffic flows. However, in its pure form it suffers from decreased fault tolerance and increased latency due to central processing. This is typically solved by deploying multiple controllers and employing load-balancing.

Google, for example, adopted a modified centralized control plane approach in its private WAN, B4. Centralized OpenFlow-based control plane functionality is distributed among multiple site-local controllers. Control logic is centrally dictated by Google's traffic engineering algorithm [16].

Cisco leverages centralized control plane approach in SD-WAN overlay. Control functions are performed by SD-WAN Controller, which acts like BGP route reflector. OMP, Cisco's version of control protocol, is used to communicate to devices.

1.3.3.2 Controller Networking [12, p. 15-16]

Controller-based architectures are often associated with SDN. Generally speaking, their main element is a software-based controller, which oversees the whole network and centrally performs a specialized function. Typically, network management or control plane functionality. Centralized control plane architectures are controller-based, for example.

Controllers serve as intermediaries between administrators/applications and underlying infrastructures [17]. They provide GUI/APIs for defining network functionality and maintain connection to devices to program their behavior.

Examples of open-source controllers for centralized network control, configuration and monitoring include ONOS by ONF [18] and OpenDaylight by Linux Foundation [19].

Both Cisco SD-Access and SD-WAN solutions are controller-based. SD-Access relies on Cisco DNA Center for central network management (designing, defining policies, monitoring). SD-WAN relies on SD-WAN Manager, Controller and Validator for centralized management, control plane functionality and device authentication/on-boarding.

1.3.3.3 Device APIs [12, p. 10]

APIs on networking devices are an alternative to user interfaces: CLI or GUI. They allow internal and external applications/scripts to control device's behavior. APIs, specific to a device/vendor, abstract internal implementation of a system, providing a set of calls/subroutines to be used by applications. They enable network programmability and simplify automation tasks.

Examples of APIs on networking devices include NX-API on Cisco Nexus 9000 Series switches [20] and eAPI on Arista EOS (Extensive Operating System) [21].

APIs are common in controller-based architectures, being used as a means of communication between controller software and networking devices. Additionally, controllers expose own APIs, making it possible to program their behaviour via external software.

Management platforms in both Cisco SD-Access (Cisco DNA Center) and SD-WAN (SD-WAN Manager) expose northbound RESTful APIs. These can be used to write custom programs for interacting with the controllers. The APIs support device inventory and provisioning, infrastructure monitoring, event management and other functionality. More on capabilities of DNA Center and SD-WAN Manager APIs can be found in [22] and [23].

1.3.3.4 Network Automation

Network automation aims at using software to perform tasks, that were previously done manually by administrators.

Both Cisco SD-Access and SD-WAN solutions include network automation features. The most common one is automated device provisioning: downloading all the necessary configurations to devices upon joining the network. Underlying technologies are PnP, as well as ZTP for Cisco SD-WAN.

1.3.3.5 Overlay Networking [12, p. 9]

Overlay networking is a network virtualization technique, aimed at creating independent logical networks (overlays) on top of physical infrastructures (underlays).

Underlays provide L2/L3 transport for overlays. Overlays are often built with the help of peer-to-peer tunnels, common tunneling technologies being L2TP, IPsec, GRE, VXLAN. Network services and policies are typically implemented at the overlay level.

Overlay networking is used, for example, in data centers to provide logical L2 connectivity between devices over L3 physical interconnects. VXLAN technology is a common tunneling protocol choice.

Both Cisco SD-Access and SD-WAN solutions are based on overlays, also called fabrics. Tunneling protocols used are VXLAN and IPsec/GRE, respectively.

1.4 Cisco SDN

1.4.1 Motivation

Among the goals of the thesis is proposing a new design of a large enterprise's campus and WAN networks, leveraging SDN principles.

Both campus and WAN segments of the enterprise's network are predominantly composed of Cisco devices. This is the main reason behind selecting Cisco SDN-based solutions for the redesign, mainly Cisco SD-Access for campus segment, and Cisco Catalyst SD-WAN for WAN segment. This will reduce migration costs, since Cisco allows to migrate its devices' OSs to the SDN-enabled versions. Additionally, such an approach will not impact devices' interoperability, since products from a different vendor may not be easily compatible with Cisco devices.

1.4.2 Naming Convention

1.4.2.1 Cisco SD-Access

Cisco SD-Access solution components will be addressed by their full name or in a simplified form, achieved by omitting the word Cisco. E.g. Cisco ISE may be shortened to ISE. Cisco SD-Access Fabric components' names may be shortened by omitting the word Fabric or Node. E.g. Fabric Border Node may be called Border Node or Fabric Border.

Cisco DNA Center has been rebranded to Cisco Catalyst Center [24]. However, since most of the consulted documentation is still using the legacy name, it will also be used in the further text.

1.4.2.2 Cisco Catalyst SD-WAN

Cisco software-defined WAN solution is based on Viptela SD-WAN, Viptela company being acquired by Cisco in 2017 [25]. Since then the solution and its components have been rebranded, causing inconsistencies in documentation.

The solution's name has changed from Cisco SD-WAN to Cisco Catalyst SD-WAN. Controllers, priorly Cisco Controllers, currently Cisco Catalyst SD-WAN Control Components, have experienced brand name changes as well. vManage, vBond and vSmart have become Cisco Catalyst SD-WAN Manager, Cisco Catalyst SD-WAN Validator, Cisco Catalyst SD-WAN Controller [23].

Devices, residing on SD-WAN edge can either be based on Viptela OS (vEdge) or Cisco IOS XE operating system (Cisco IOS XE Catalyst SD-WAN Device or cEdge). Alternative names include WAN Edge router, SD-WAN router [26], SD-WAN Edge router [27], etc.

Further text will address the components by their full brand name or by a shortened name for simplicity. Simplification will be achieved by omitting the following words: Cisco, Catalyst, and sometimes SD-WAN. E.g. Cisco Catalyst SD-WAN Manager may be shortened to SD-WAN Manager or simply Manager. Devices, residing on SD-WAN edge will be referred to as WAN Edge routers/devices.

1.4.3 Cisco SD-Access Overview [8, 28]

Cisco SD-Access is an SDN solution, specialized for campus (access) networks, designed to address limitations of traditional infrastructures.

At the core of the SD-Access architecture is Cisco Digital Network Architecture (DNA) Center – a network controller with GUI-based management dashboard and exposed RESTful APIs, which provides an abstraction layer for managing the whole infrastructure. It allows to program the network via APIs, as well as define, monitor, troubleshoot complex segmented policy-driven access networks, while automating the configuration process. This enhances security and simplifies networking operations.

To enhance security, SD-Access allows to segment the network, define group-based access policies, maintain a high degree of endpoint and whole network visibility using a simple GUI or APIs, with wired and wireless networks managed as one.

To simplify networking operations SD-Access optimizes performance of data, control and policy planes, and leverages Cisco DNA Center to centrally automate and manage their operation. It also separates physical infrastructure from the logical one using overlay technology for better scalability of both. The Cisco DNA Center provides an intuitive GUI interface and APIs to define the desired network behavior, hiding the complexity of manual per-device configurations and logical-to-physical infrastructure mapping.

1.4.3.1 Main Solution Components

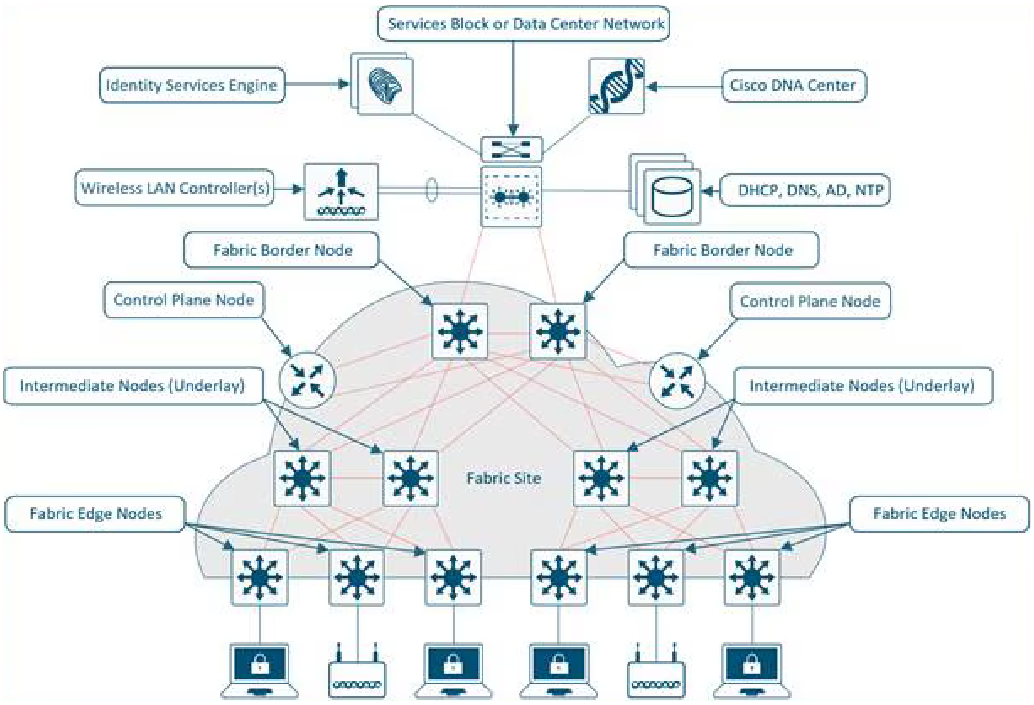

Figure 1.2 SD-Access Components [28]

- Cisco DNA Center: A controller which provides a single dashboard for performing management operations on the network, and automates these operations internally. It allows to design, provision, apply policies, and assure network performance via a GUI, eliminating the need for manual per-device configurations. SD-Access is an application package, part of DNA Center, which uses its capabilities to automate the creation of virtual networks with integrated security and segmentation

- Cisco Identity Services Engine (ISE): Cisco Identity Services Engine helps to implement identity-based network access control. Among other features, it performs device/user authentication, allows to define identity-based policies, and enforces them. Cisco ISE integrates with with Cisco DNA Center for automation of policy configurations.
- Overlay-based interconnect (SD-Access Fabric): SD-Access campus networks are built using Fabric technology, which allows to create multiple programmable isolated logical networks (overlays) on top of a shared physical infrastructure (underlay). The technology also enables software-defined segmentation and policy enforcement based on user identity and group membership within logical networks. Fabric functionality is supported by devices with Fabric roles: Border, Control Plane and Edge Nodes, as well as Fabric WLC and AP for wireless segment.

1.4.4 Cisco Catalyst SD-WAN Overview [23, 26]

Cisco Catalyst SD-WAN is an SDN solution, specialized for WAN, interconnecting multiple locations within an organization, designed to address limitations of traditional infrastructures.

Cisco Catalyst SD-WAN is is an overlay-based WAN architecture, allowing to build transport-independent secure logical connections between locations. It also decouples control and policy, management, orchestration and data planes of the network, and consolidates them in SD-WAN Controller, SD-WAN Manager, SD-WAN Validator and WAN Edge routers, respectively.

Cisco Catalyst SD-WAN provides a single pane of glass for managing the whole network, and integrates features for enhancing network performance (AAR), security (integrated security features), agility (automated device provisioning via PnP/ZTP), etc.

1.4.4.1 Main Solution Components

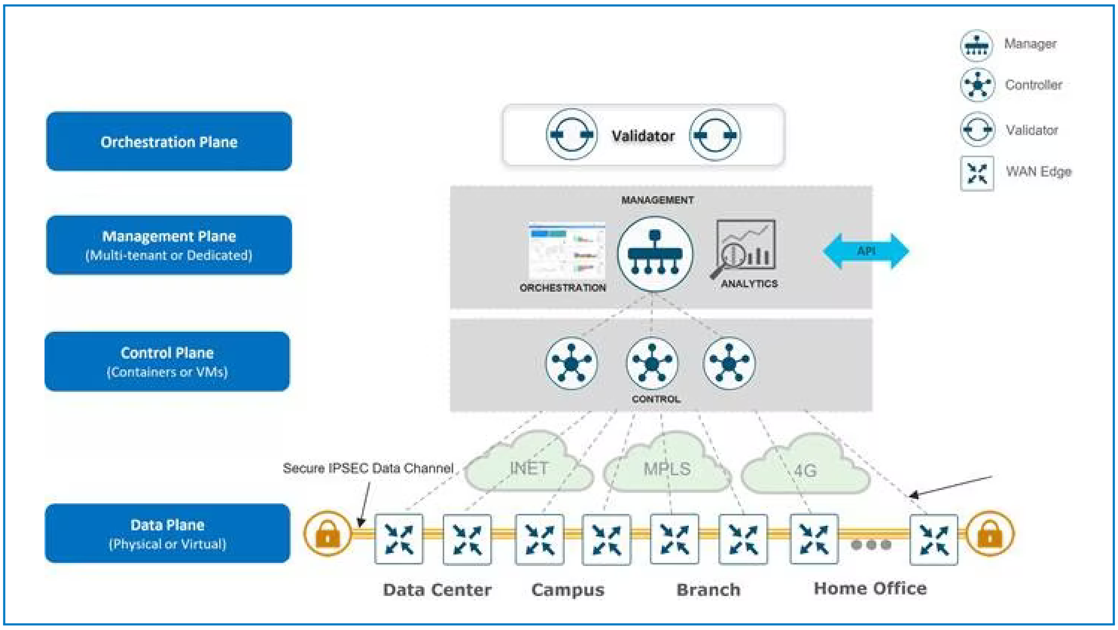

Figure 1.3 SD-WAN Components [26]

- SD-WAN Manager: A software, in which management plane functionality is centralized. It provides a single dashboard (GUI interface) for managing the whole SD-WAN network. This includes monitoring, configuring, maintaining all SD-WAN devices and their connected links in both overlay and underlay networks.

- SD-WAN Controller: A software, in which control and policy plane functionality of the overlay network is centralized. It maintains a centralized routing table for the overlay network, stores and distributes routes and policy information among WAN Edge routers. It also assists SD-WAN Validator in authenticating the routers.
- SD-WAN Validator: A software, in which orchestration plane functionality is centralized. It performs initial authentication of SD-WAN Controllers and WAN Edge routers, and coordinates connectivity between them.
- WAN Edge router: Hardware or software device, located at a physical site or in the cloud, at the border between site-local network and WAN interconnect. Its role is to provide secure data plane connectivity between locations over available WAN transports. WAN Edge routers have local intelligence (control plane) to establish site-local and underlay connectivity. To build overlay networks they make use of control plane centralized in the SD-WAN Controller.
- Overlay-based interconnect: The Cisco Catalyst SD-WAN network consists of underlay (physical) and overlay (logical) networks. Overlays are built on top of underlays with the use of tunnels between WAN Edge routers. Underlay uses distributed control plane (on each WAN Edge router), overlay uses centralized control plane (in SD-WAN Controller) with control plane information being exchanged between SD-WAN Controller and WAN Edge routers.

Chapter 2

Cisco SD-Access and Cisco Catalyst SD-WAN [8, 28]

2.1 Cisco SD-Access Architecture

Cisco SD-Access overview is provided in 1.4.3.

2.1.1 SD-Access Fabric

SD-Access network is overlay-based, consisting of two main components: Underlay and Overlay. Overlay is also called Fabric. Overlay segmentation is provided using SD-Access VNs (macro-segmentation) and SGTs (micro-segmentation).

2.1.1.1 Underlay and Overlay

Underlay

Underlay is the physical network, which performs forwarding functions for Overlays. Although Endpoints typically connect into the Underlay directly or through APs, they belong to the Overlay layer.

The main requirement for Underlay is L3 (IP) connectivity, which implies moving the L2/L3 boundary to Access Layer. This allows to use routing benefits: simplifies troubleshooting, decreases the convergence time on link failure, and eliminates the need for STP, FHRP and similar protocols [29].

When building the Underlay it is possible to configure devices manually or make use of Cisco DNA Center's LAN automation service. It automatically discovers, provisions, and deploys network devices according to Cisco-validated design best practices. As part of this process, Cisco DNA Center assigns IP addresses and pushes routing protocol configuration to devices.

Overlay

Overlays are L2/L3 logical networks built on top of Underlay with the use of tunnels. Underlay provides Overlays with transport services. Endpoints belong to Overlay, the latter providing them with a logical full-mesh connectivity. Overlays are also responsible for software-based policy enforcement and segmentation.

Overlay creation, policy enforcement and segmentation is aided by technologies/protocols, operating in Control, Policy and Data Planes. These technologies lay the foundation of the SD-Access solution and will be described in more detail.

2.1.1.2 Overlay Operational Planes and Associated Protocols

Overlay Data Plane – VXLAN

The fabric data plane is responsible for tunneling L2/L3 data via Underlay using VXLAN encapsulation. This gives raise to L2/L3 overlays.

VXLAN is MAC-in-IP encapsulation method, meaning, it encapsulates the whole L2 frame in a new IP header. Keeping the original MAC header intact allows to create both L2 and L3 logical connections over the Underlay.

Each overlay network is identified using VNI, carried in VXLAN header. VNIs can be L2 or L3, correlated to VLANs or VRFs respectively. This mechanism is used to distinguish between L2 and L3 overlays.

VXLAN-GPO header format allows to specify SGT in addition to VNI. SGT is the Overlay Policy Plane construct, carrying group membership information of users/devices.

Overlay Control Plane – LISP

LISP is an architecture and a set of protocols that allow to separate Endpoint's identity (EID) from its physical location on the network (RLOC). This allows to provide Endpoints with the same level of access independent on their physical location.

In SD-Access Endpoint's MAC or IP is used as EID, IP (Loopback 0) of the tunnel endpoint (Fabric Edge Node) it's directly connected to/through is used as RLOC. "The RLOC address is part of the underlay routing domain, and the EID can be assigned independently of the location" [28].

LISP architecture includes a mapping system to store EID-to-RLOC mappings, and resolve EIDs to RLOCs. This functionality is typically implemented in Control Plane Node.

In a LISP-enabled network when a source RLOC sends data to an EID it "queries the mapping system to identify the destination RLOC for traffic encapsulation" [28].

LISP also removes the need for the tunnel endpoints to store reachability information to every access-layer subnet in their routing tables. Instead, routing tables only store routes to all RLOCs and to remote EIDs communicating through the node.

Overlay Policy Plane – Cisco TrustSec

Fabric Policy Plane operation is powered by Cisco TrustSec security architecture. It replaces IP/VLAN-based access control with Group Based Access Control (GBAC). This is achieved by classifying Endpoints (users/devices) into scalable groups based on identity and defining access policies based on group membership. Groups are identified by Security Group Tags (SGTs), group-based policies are called Security Group Access Control Lists (SGACL).

Operation of Cisco TrustSec solution can be broken down in 3 phases:

- Classification: Assigning SGT to connecting users/devices dynamically (as a result of authentication to ISE PSN) or statically (based on IP, VLAN, port). SD-Access uses dynamic classification.

- Propagation: Carrying SGT information from classification to enforcement point. In SD-Access packets are labeled with SGT at Fabric ingress (source Edge Node), SGT is carried in VXLAN-GPO header.

- Enforcement: Applying a policy (SGACL) based on source and destination SGT by the enforcement device (switch, firewall, router). In SD-Access, enforcement happens at Fabric egress (destination Edge Node) or firewalls. SGACLs are downloaded to enforcement devices by ISE PSN. Edge Nodes only download policies relevant to their directly connected endpoints.

2.1.1.3 Segmentation

Segmentation is hierarchical in SD-Access. Macro-segmentation creates isolated virtual networks, while micro-segmentation dictates communication rules within them.

- Macro-segmentation – VNs: Used to separate groups of Endpoints that don't normally communicate to each other. It is achieved by creating isolated L3 virtual networks (VNs in SD-Access) on top of physical infrastructures. They are identified by L3 VNIs and instantiated as VRFs for isolation.
- Micro-segmentation – SGTs: Controls communication between groups within a single VN. Groups are identified by SGTs, and access policies enforced with the use of SGACLs.

Segmentation in SD-Access is implemented in Overlay. Policy Plane handles Endpoint assignment to VNs/SGs and policies. Data Plane allows to carry VNIs/SGTs in VXLAN-GPO header. Control Plane manages control communication separately for each VN. Finally, Management plane automates these workflows.

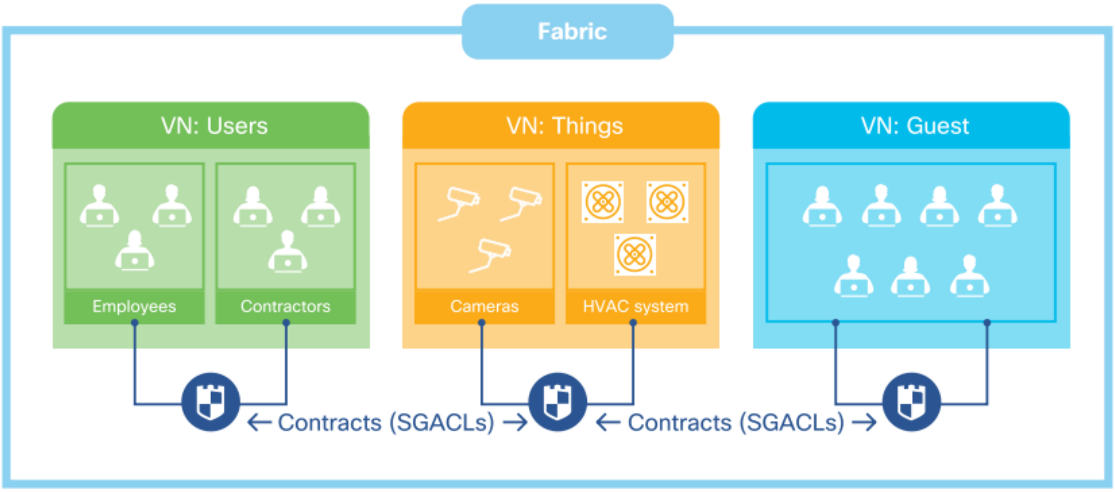

Figure 2.1 SD-Access Macro and Micro Segmentation [23]

2.1.2 Solution Components

Generally, an SD-Access deployment is represented by SD-Access Fabric Sites, interconnected by Transit/Peer Networks, and using Shared Services for operation.

2.1.2.1 Fabric Site

Fabric Site is a portion of the Fabric with its own Control Plane and Edge Node. A Fabric Site may accommodate a WLC for wireless, an ISE PSN for policy implementation, and a Border Node for external network connectivity.

Fabric-enabled sites are composed of interconnected devices. Some of them operate in a Fabric Role (Control Plane, Border, Edge Nodes and WLCs) to support Fabric functionality. Devices are managed by Cisco DNA Center and Cisco ISE, typically residing outside of the Fabric in the Shared Services Block. Fabric Sites are interconnected by Transit/Peer Networks.

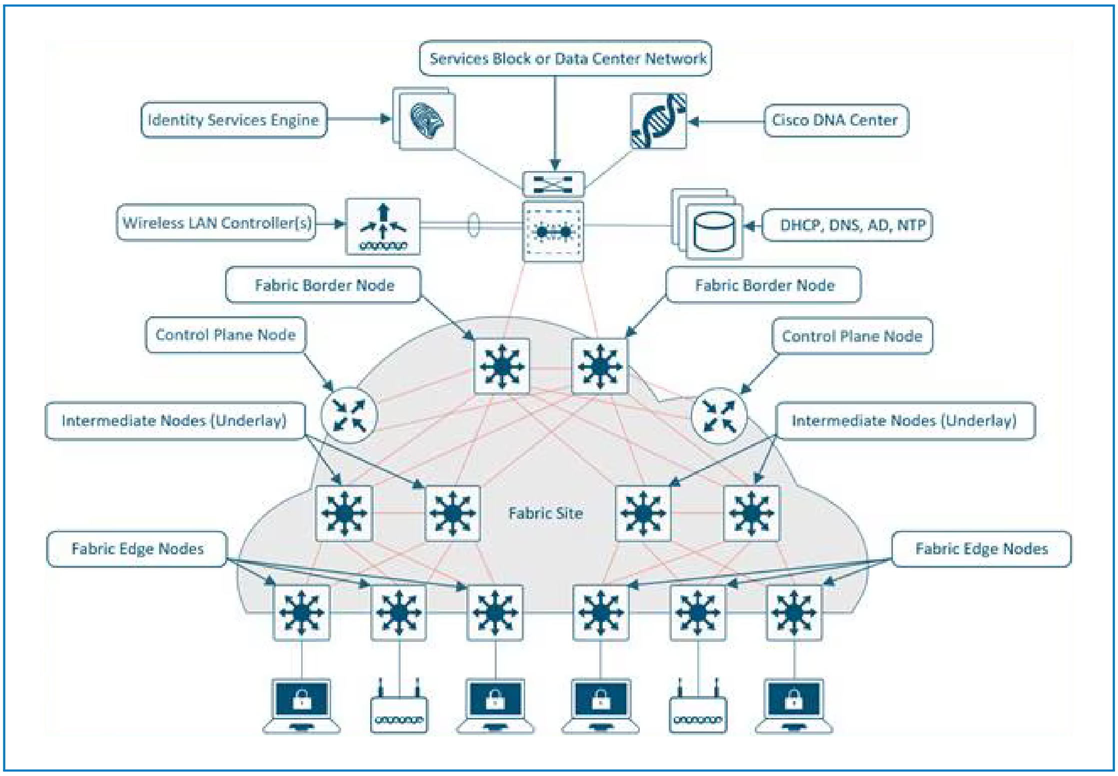

Figure 2.2 SD-Access Architecture and Components [28]

Devices with Fabric Roles

- Fabric Control Plane Node: supports LISP protocol operation by maintaining a Host Tracking Database (HTDB), and acting as LISP Map-Server (MS) and Map-Resolver (MR). HTDB stores EID-to-RLOC mapping. MS receives Endpoint registration and populates the HTDB. MR resolves EID to RLOC in respond to queries from Fabric devices.
- Fabric Border Node: Acts as a VXLAN tunnel endpoint and connects Fabric sites to external networks.

  To enable external connectivity it

  - Advertises EID subnets to the outside network
  - Acts as a gateway of last resort for Edge Nodes
  - Preserves segmentation outside of the Fabric with the help of VRF-Lite and VRF-aware routing protocols
  - Maps SGTs to constructs supported by the outside network
- Fabric Edge Node: Acts as a VXLAN tunnel endpoint and connects Endpoints to the Fabric.

  To enable Endpoint connectivity it

  - Registers Endpoints' EIDs with Control Plane Nodes
  - Maps Endpoints to VLANs statically or dynamically

  - Stretches subnets: it is possible to define the same subnet behind multiple Edge Nodes with the help of the Anycast L3 Gateway
  - Is typically assigned with RLOC

- Fabric WLC: Resides outside of the Fabric and belongs to Shared Services block. It maintains CAPWAP tunnels to Fabric APs, to exchange wireless control traffic. Apart from traditional WLC control functionality it registers wireless Endpoints with Control Plane Node and instructs Fabric APs to form VXLAN tunnels to their adjacent Fabric Edge Node.

Other Fabric Devices

- Fabric Intermediate Node: Fabric devices not assigned with any Fabric Role. They are not required to support Overlay-related technologies (VXLAN, LISP, SGTs), merely providing L3 Underlay transport for Overlays.

- Endpoint: wired and wireless end-devices connected to Edge Nodes and Fabric APs, respectively. Fabric Endpoints are identified by EIDs.

- Fabric AP: Is directly connected to an Edge Node. It is placed in predefined INFRA_VN, which maps to GRT, to allow communication to WLC without route/VRF leaking. Establishes a CAPWAP tunnel to Fabric WLC for control plane communication. Forms VXLAN tunnels to connected Edge Nodes, letting the Fabric handle the data plane communication. Apart from traditional AP functionality, supports subnet stretching.

2.1.2.2 Shared Services. Cisco ISE. Cisco DNA Center

Shared Services include common network services, tools and infrastructure components used by devices in the Overlay. This means, they must be accessible by multiple Fabric VNs. The services typically reside outside of the Fabric, in GRT or dedicated VRF, which requires employment of inter-VRF routing mechanisms.

Shared Services typically include SD-Access infrastructure elements (Cisco DNA Center, Cisco ISE, WLCs), network services (DNS, DHCP, IPAM), monitoring tools and data collectors (SNMP, Netflow), Internet access (common Internet firewalls), etc.

Cisco ISE

Cisco ISE is a platform, performing network access control, policy enforcement and security services. It supports Cisco TrustSec architecture, allowing to define scalable groups, map users/devices to those and create policies/segmentation rules based on group membership.

In SD-Access deployments Cisco DNA Center integrates with ISE using pxGrid and REST APIs to automate access control management.

The distinct services performed by the Cisco ISE platform are called personas. ISE nodes (actual ISE instances) can function as single or multiple personas, these being:

- Policy Administration Node (PAN): Responsible for system-related configurations and administrative operations on Cisco ISE. These include policy, segmentation and security rules definition. PAN functionality can be accessed through a web interface or external systems/applications (Cisco DNA Center) integrated with it via APIs.

- Monitor and Troubleshooting Node (MnT): Provides log collection (from PAN and PSN nodes), monitoring and troubleshooting functionality. It aggregates and correlates information to present it in form of reports.

- Policy Service Node (PSN): Evaluates access requests against configured policies. Performs ”AAA, posture, guest access, client provisioning and profiling” services.

- Platform Exchange Grid (pxGrid): uses pxGrid framework to exchange contextual information with other ISE nodes or external network systems/platforms (Cisco DNA Center).

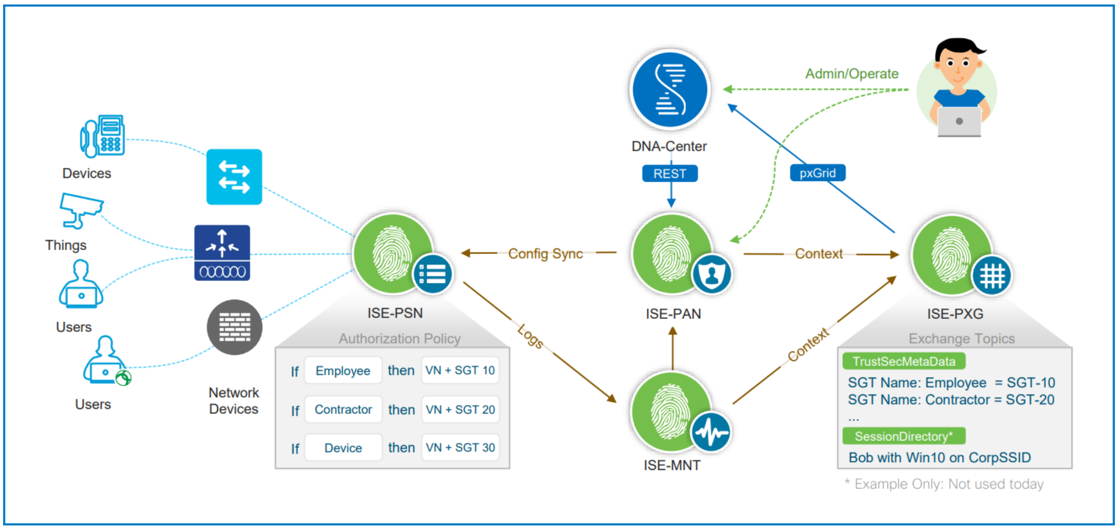

Figure 2.3 ISE Personas in SD-Access [30]

Cisco DNA Center
A centralized GUI-based network management platform with automation and orchestration capabilities. It abstracts physical infrastructures by providing a user-friendly interface for defining the desired network operation, while launching automated workflows in the background.

In addition, Cisco DNA Center continuously collects, correlates and represents information about the network (system health, connectivity status, protocol configuration and associated table contents, etc.).

Moreover, it allows external systems to view ans program the network by exposing northbound APIs.

SD-Access is part of Cisco DNA Center software. It allows to perform solution-specific management operations: deploying new physical networks, creating Fabric Overlays and Transit Networks, managing Cisco TrustSec policies by integrating with Cisco ISE. On the device level this involves automating discovery, inventory, provisioning and configuration.

2.1.2.3 Transit/Peer Network

Transit/Peer Networks connect Fabric Sites to external networks (Fabric and non-Fabric sites, Internet, etc.).

It is possible to configure IP-based, SD-Access or SD-WAN transits. The last two are able to carry VRF/SGT from site to site, while usage of IP-based transits requires remapping of VRFs/SGTs among locations.

2.2 Cisco Catalyst SD-WAN Architecture [23, 26]

Cisco Catalyst SD-WAN overview is provided in 1.4.4.

2.2.1 SD-WAN Fabric

SD-WAN network is overlay-based, consisting of two main components: Underlay and Overlay. Overlay is also called Fabric. Overlay segmentation is provided using SD-WAN VPNs.

2.2.1.1 Underlay and Overlay

Underlay
Underlay is the physical network, running traditional routing protocols, and performing forwarding functions for the Overlay. In SD-WAN it includes WAN Edge routers' connections to WAN transport networks and transport networks themselves.

A WAN transport network is a network infrastructure, used to interconnect an organization's locations together. Underlay may be composed of multiple transport networks: MPLS circuits, public Internet connections, LTE/5G mobile networks, etc.

Transport services are usually provided by ISPs, WAN Edge router connectivity to the Service Provider gateway being established with the help of a static default gateway or a dynamic routing protocol (OSPF, BGP).

WAN Edge routers' ports, connecting to transport networks, are part of VPN 0.

Overlay
Overlay is composed of site-to-site IPsec tunnels between WAN Edge routers, built over WAN transports.

Overlay Control Plane is centralized in SD-WAN Controller and is based on OMP protocol. Control Plane communication is established via DTLS/TLS tunnels between a Controller and each WAN Edge device.

Using OMP, WAN Edge routers share their site-local routes and TLOCs with connected Controller, which applies policies to them and advertises best routes to other WAN Edge devices.

It is important to distinguish between the following Overlay constructs:

- Site ID: a unique 32-bit identifier of a site (data center, branch office, campus, etc.), shared by all WAN components belonging to it. By default, IPsec/GRE tunnels are not formed between WAN Edge routers with the same Site ID.

- TLOC: logical representation of a WAN Edge router's attachment to a WAN transport. It is uniquely identified by a three-tuple: WAN Edge router's System IP, link color, encapsulation method (GRE/IPsec).

  It is said that Overlay tunnels are formed between TLOCs, as TLOCs are key part of the tunnel endpoint identification. They are also used in the Control Plane as next-hop attributes of OMP routes.

- Color: an abstraction, a statically defined keyword, used to identify WAN transports and their TLOCs. A transport may be classified as private or public and assigned a color from pools of private/public colors. “Color dictates the use of either private IP or public IP address when communicating through the control or data plane” [26].

- System IP: persistent, system-level IPv4 address, uniquely identifying a device. It is assigned to system interface, residing in VPN 0. System IP is used when establishing OMP sessions between Controllers and WAN Edge devices. It is also part of a TLOC identifier.

2.2.1.2 Segmentation

L3 segmentation of traffic carried by SD-WAN Overlay is achieved with the help of SD-WAN VPNs, which function in a way similar to VRFs.

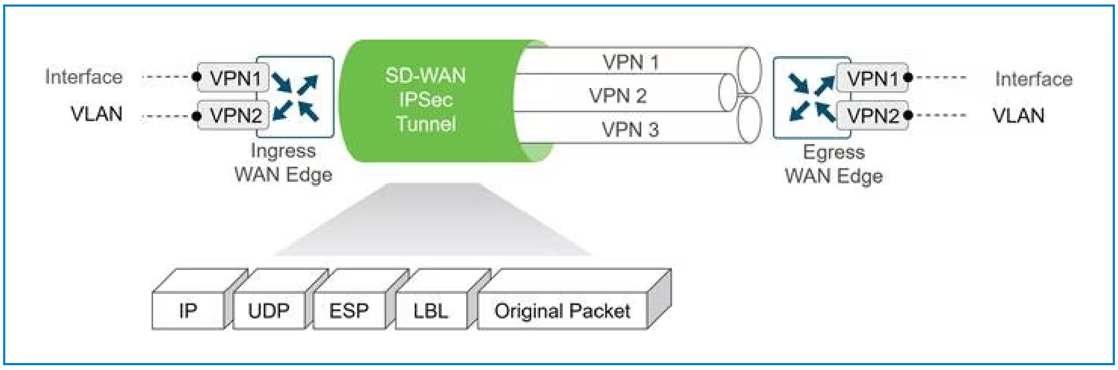

Figure 2.4 End-to-End Segmentation [26]

VPN numbers are used for global VPN naming, while VPN labels are constructs local for each WAN Edge router. Upon VPN creation, WAN Edge routers assign labels to them and share VPN-TO-LABEL mapping with other WAN Edge devices through the Controller. The labels are used to direct traffic to a desired VPN in the Overlay.

Segmentation is enforced at the WAN Edge routers by assigning each interface to a VPN and maintaining a separate routing table instance per VPN. Segmentation information is carried across the Overlay by inserting VPN labels in Data Plane packets. Overlay Control Plane is responsible for maintaining isolation between VPNs. It stores and distributes VPN-to-LABEL mapping among WAN Edge devices, as well as associates OMP routes with specific VPNs using labels.

It is important to distinguish between the following VPNs in the context of SD-WAN:

- Transport VPN (VPN 0): Contains WAN Edge routers' interfaces, connecting to WAN transports. It is used for establishing both Underlay and Overlay connectivity between locations.

  Underlay static routes/routing protocols are configured in VPN 0 to connect WAN Edge routers to Service Provider gateways. Overlay IPsec tunnels also belong to transport VPN.

  Finally, WAN Edge routers' DTLS/TLS connections to control components belong to this VPN.

- Management VPN (VPN 512): carries out-of-band management traffic to and from SD-WAN devices. It is ignored by OMP and not carried across the Overlay.

- Service VPNs: Contain WAN Edge routers' interfaces, facing local-site networks, and carrying user data. Service VPN routes are advertised into the OMP routing protocol. Conversely, OMP routes are redistributed into service VPN routing protocol.

VPNs 0 and 512 are reserved for internal use and are the only VPNs, functional in SD-WAN Controller, Manager and Validator devices. Other numbers can be used to label service VPNs.

2.2.1.3 Overlay Operational Planes and Associated Protocols

Overlay Data Plane – IPsec, BFD

SD-WAN overlays are formed by IPsec or GRE tunnels between locations. TLOCs, exposed by WAN Edge routers, serve as tunnel endpoints. By default, full mesh connectivity between WAN Edge routers is established: between each pair of TLOCs over each transport.

IPsec tunnels are recommended in SD-WAN deployments, because they provide endpoint authentication and data encryption, which GRE lacks. Due to this further text focuses on IPsec.

IPsec

IPsec is a protocol suite, that authenticates and encrypts data to provide secure connectivity between endpoints. It can operate in tunnel mode to build encrypted point-to-point links between devices.

Traditionally IPsec uses the following protocols: AH for data authentication, ESP for data encryption, IKE for exchanging encryption keys between peers.

In SD-WAN authentication algorithm is configurable and is included in TLOC properties, advertised by OMP in Control Plane. Data encryption is provided by ESP. Encryption key exchange is performed by OMP. "Each WAN Edge router generates one AES key per TLOC and transmits this information to the SD-WAN Controller in OMP route packets, which is then distributed to all WAN Edge routers" [26].

BFD

Bidirectional Forwarding Detection (BFD) protocol runs between peers over IPsec tunnels. It starts operating automatically and can't be disabled. It is used to measure tunnel liveliness and path quality (loss, latency, jitter). Tunnel liveliness information allows to detect when a tunnel goes down and redirect traffic to the operational ones. Path quality measurements are useful in AAR, where application traffic can be directed through tunnels matching configured SLA policies.

Overlay Control Plane – SD-WAN Controller – OMP

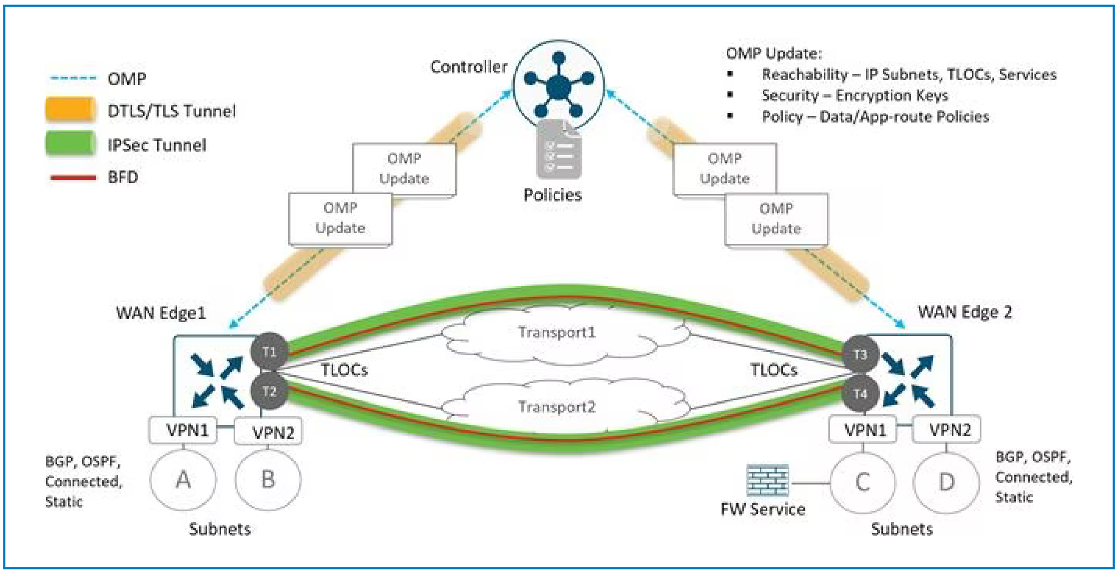

Figure 2.5 OMP Operation [26]

SD-WAN Overlay Control Plane operation is centralized in SD-WAN Controller and assisted by OMP protocol.

OMP runs between SD-WAN Controller and WAN Edge routers over DTLS/TLS connections. OMP peering is automatically established between them upon forming control connections. If WAN Edge router looses connection to Controller it continues forwarding data for the duration of configurable OMP graceful restart timer.

In the context of Control Plane operation SD-WAN Controller acts as a route reflector. It receives routes from WAN Edge devices, processes and applies policies to them, further advertising them to other WAN Edge devices using OMP.

Each WAN Edge router advertises routes, learnt from its local site. The routes are of three types:

- OMP routes: site-local prefixes along with attributes, such as VPN to which they belong and TLOC via which they are reachable. The prefixes are originated as routes: static, connected or learnt via a routing protocol (OSPF, BGP, etc.).
- TLOC routes: TLOCs, exposed by the WAN Edge router, along with attributes (TLOC private and public IPs, associated port number, color, etc.).
- Service routes: routes to shared services (firewall, IPS, IDS, load balancers etc.), available at the site.

OMP is also used in Data Plane to exchange encryption keys for IPsec tunnel establishment. In Policy Plane OMP is used to distribute Centralized Data Policies to WAN Edge routers.

Overlay Policy Plane – SD-WAN Manager
SD-WAN policies influence data flow between WAN Edge routers in the Overlay. They can be Centralized or Localized and apply to Control or Data Plane traffic. Centralized policies are applied to Overlay as a whole, while Localized policies are independently applied to WAN Edge routers. All policies are configured through SD-WAN Manager GUI and pushed to SD-WAN Controllers (Centralized) and WAN Edge routers (Localized) using NETCONF.

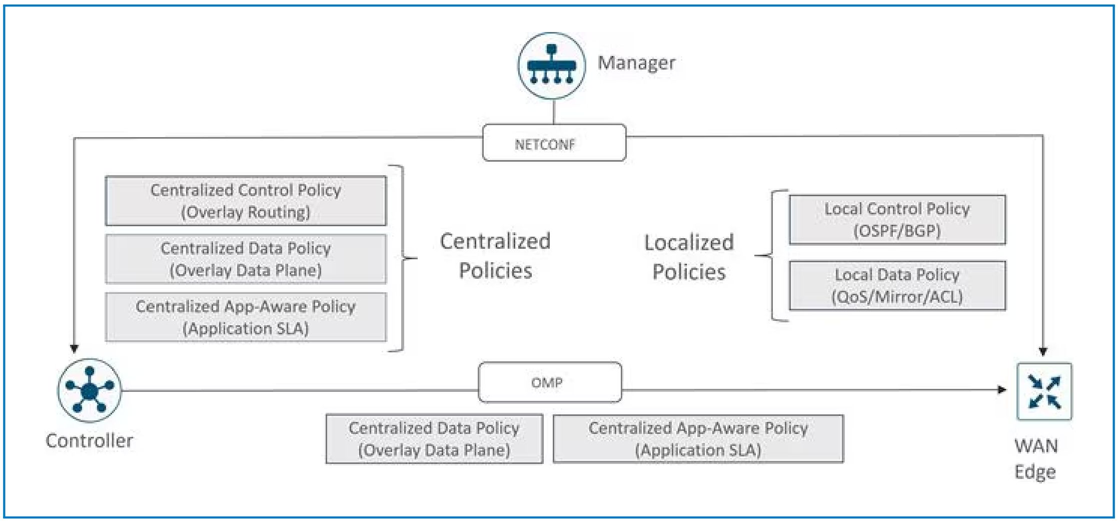
Manager
NETCONF
Centralized Control Policy (Overlay Routing)
Centralized Data Policy (Overlay Data Plane)
Centralized App-Aware Policy (Application SLA)
Centralized Policies
Localized Policies
Local Control Policy (OSPF/BGP)
Local Data Policy (QoS/Mirror/ACL)
OMP
Controller
Centralized Data Policy (Overlay Data Plane)
Centralized App-Aware Policy (Application SLA)
WAN Edge

Figure 2.6 SD-WAN Policy Types and Application [26]

Control Plane policies influence routing decisions and path selection. Centralized Control policies allow to configure per-VPN topologies, WAN Edge routers' VPN membership, etc. based on OMP routing and TLOC information. Localized Control policies affect site-local routing via OSPF/BGP route maps and prefix lists.

Data Plane policies are applied based on packet's IP header and VPN membership. Centralized Data policies include AAR and QoS policies. Since these must be applied on WAN Edge routers, they are communicated to the devices using OMP. Localized Data policies allow to apply ACLs and QoS, as well as perform mirroring and policing of data at the site.

2.2.2 Solution Components

2.2.2.1 Primary Components

SD-WAN deployments are represented by sites' edge devices (WAN Edge routers) connected to WAN transports physically and to each other logically, via tunnels. SD-WAN network operation is governed by SD-WAN Controller, Validator and Manager, implementing Control, Orchestration and Management Plane functionality.

SD-WAN Validator (Orchestration Plane Controller)

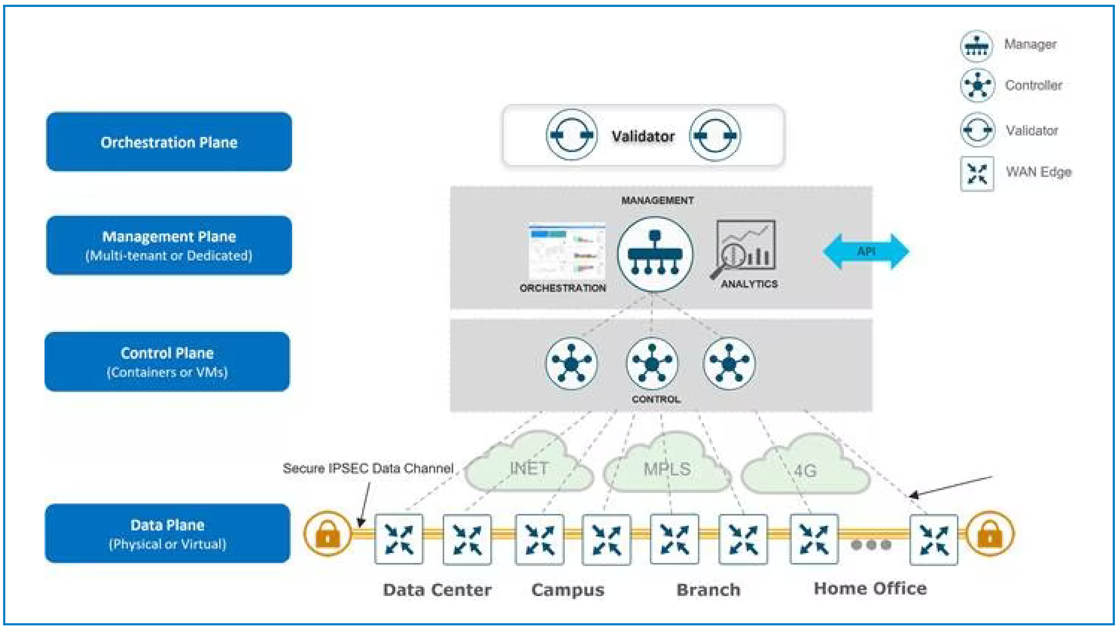

Figure 2.7 SD-WAN Architecture and Components [26]

SD-WAN Validator communicates to other SD-WAN components via DTLS tunnels to perform their initial authentication and assist in forming control connections between them.

Main functions of SD-WAN Validator include:

- Bring-up of SD-WAN devices: automated authentication and validation of WAN Edge routers and controllers, attempting to join the network. Authentication is done using certificates and RSA cryptography.
- Orchestrating connectivity: Validator informs controllers of newly joined WAN Edge routers. It also instructs them to form control connections with each other.
- NAT traversal: Validator is the only SD-WAN device that sits in public address space. Thus, it can communicate to other components sitting behind NAT and solve NAT traversal issues.
- Load balancing: with multiple Controllers in domain Validator evenly pairs newly connected Edge with them.

SD-WAN Controller (Control Plane Controller)

Represents the centralized Control Plane of the Overlay network. It collects and distributes routing and policy information among WAN Edge routers using OMP. It additionally orchestrates

WAN Edge routers' key exchange necessary for IPsec communication encryption.

WAN Edge Routers (Data Plane Devices)
WAN Edge routers are located at sites' perimeters, and provide both Underlay and Overlay connectivity between them. WAN Edge routers are authenticated by Validator, managed by Manager and assisted with Overlay connectivity by Controller.

SD-WAN Manager (Management Plane Controller)
SD-WAN Manager is a centralized management system, that provides a GUI for deploying, monitoring, configuring, troubleshooting the SD-WAN network.

SD-WAN Manager communicates to other WAN components via secure DTLS/TLS channels. Within these it establishes an SSH session with the devices and uses NETCONF for performing management tasks.

One of the Manager's functions is automated provisioning of WAN Edge routers and Controllers. It allows to create and store their configurations to later push them to the devices upon request. It additionally stores and downloads certificate credentials to WAN Edge routers.

Manager also uses other protocols to manage the network: SNMP for data collection, ICMP to detect device liveliness, etc.

2.2.2.2 Control Connections

SD-WAN control connections are built between control components (Controller, Validator, Manager) and WAN Edge routers. They are used to separate control and management from data traffic in SD-WAN network.

Control connections are established in the form of DTLS/TLS tunnels to secure communication independent of the protocol used and its native level of security. DTLS/TLS encrypt data at the transport layer: at UDP and TCP socket level, respectively. TLS relies on TCP for packet delivery, while DTLS implements additional mechanisms for dealing with packet loss (sequence numbers, fragment offsets, retransmission).
Control connections are mainly used for:

- Authentication of WAN Edge routers and controllers to Validator.
- Management Plane Operation: Automated provisioning and monitoring of SD-WAN devices by SD-WAN Manager with the use of NETCONF, SNMP, ICMP, etc.
- Control Plane Operation: Establishing of OMP peering between each WAN Edge and Controller.

2.3 Cisco SD-Access | SD-WAN Pairwise Integration

When an enterprise has multiple SD-Access Fabric Sites it is desirable to preserve segmentation (VNI) and policy (SGT) context when data travels between them.

SD-WAN, used as Transit Network in SD-Access deployments, allows to transfer VNIs and SGTs in IPsec header. To preserve segmentation SD-Access VNs are mapped to SD-WAN Service VPNs. To allow site-to-site connectivity, routes are shared between SD-Access and SD-WAN domains.

However, to use SD-WAN as Transit, it is necessary to interconnect SD-Access and SD-WAN fabrics. Mainly, SD-Access Border Nodes with WAN Edge routers, further referred to as Border Node and WAN Edge router. This interconnect must as well perform route exchange between fabrics, preserve segmentation and carry VNIs and SGTs.

Depending on the way to design this interconnect, SD-Access and SD-WAN can be integrated into Independent or Integrated Domain.

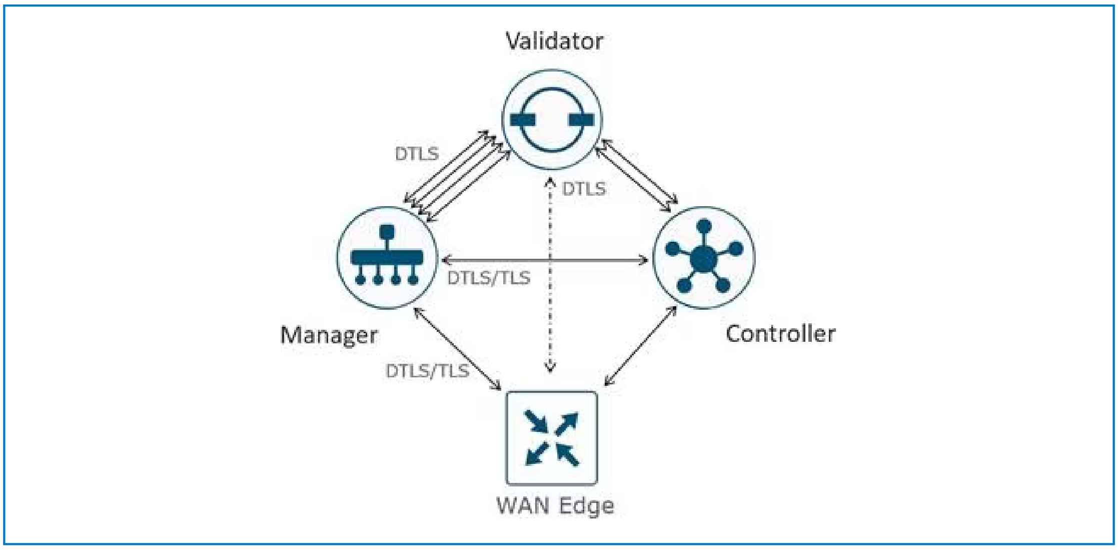

Figure 2.8 SD-WAN Control Connections [26]

2.3.1 Independent Domain [31]

In Independent Domain deployments WAN Edge router functionality and SD-Access Fabric roles are implemented on different devices. The devices are typically directly connected via a L3 link and communicate using IP forwarding. SD-WAN and SD-Access Fabrics are separately managed by SD-WAN Manager and Cisco DNA Center.

Control Plane Integration
SD-Access VNs are mapped to SD-WAN Service VPNs. Traffic from each VN/VPN is placed in a separate VRF-Lite instance while traveling between Border Node and WAN Edge router to preserve segmentation. BGP runs in each VRF to share prefixes withing each virtual segment. WAN Edge router then redistributes them into OMP.

Data Plane Integration
802.1Q VLAN technology facilitates VNI transfer between fabrics. A VLAN is configured per VRF-Lite instance. VLAN ID associated with a VNI is carried between Border Node and WAN Edge router in Ethernet header.

Policy Plane Integration
SGT is carried between Border Node and WAN Edge router in Ethernet header using SGT Inline Tagging. [It requires using a special EtherType of Ethernet frame, which allows to insert Cisco Metadata Header (CMS) in the frame's header. CMS carries SGTs] []. [Inline tagging has to be supported by WAN Edge router's interface, connecting to SD-Access Border Node] [].

Overall Picture
Independent Domain operation is visualized in Fig. 2.9. VNI and SGT placement in VXLAN, Ethernet and IPsec headers is visualized in Fig. 2.10.

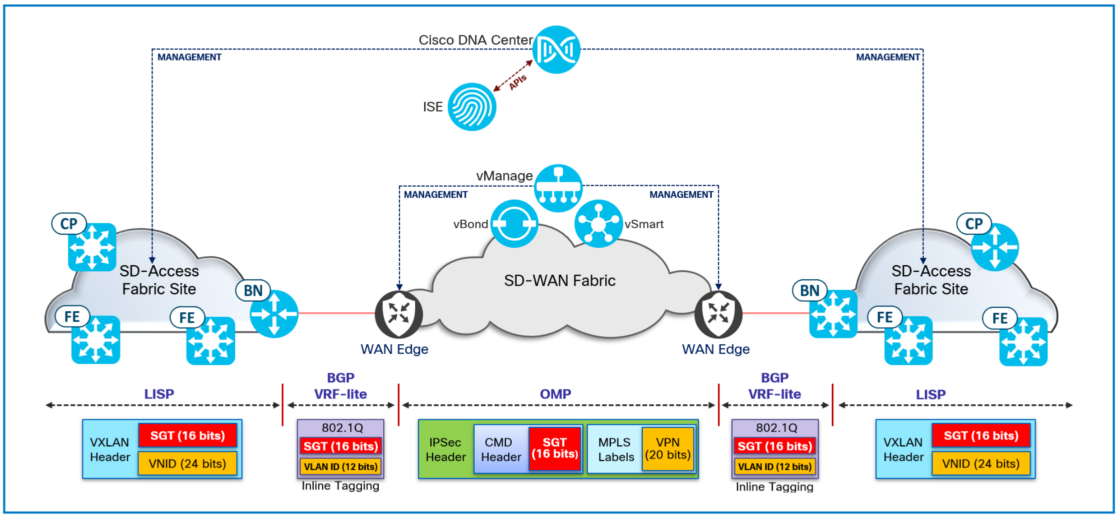

Figure 2.9 Independent Domain Components [31]

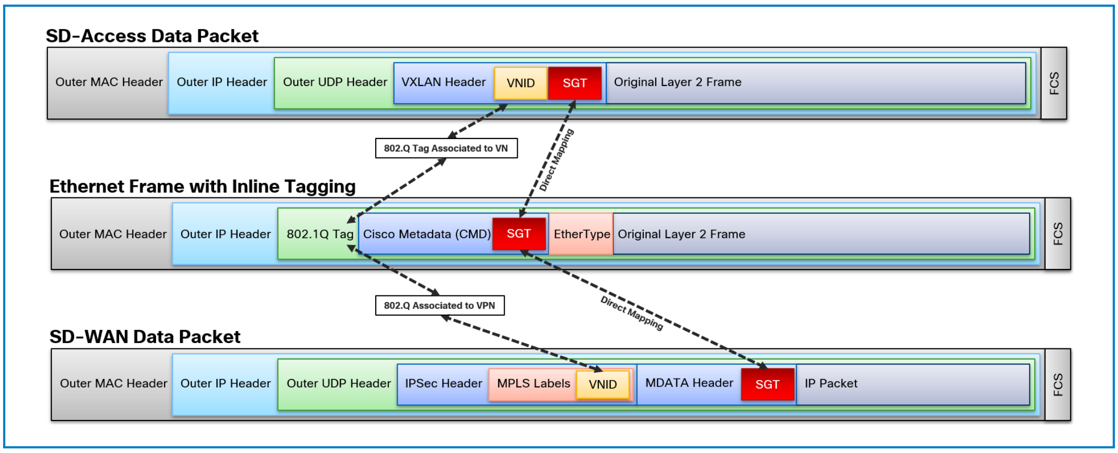

Figure 2.10 VXLAN to IPsec [31]

2.3.2 Integrated Domain [32]

In Integrated Domain deployments SD-WAN Edge functionality is collocated with SD-Access Border and Control Plane Nodes on the same device. With the exception of such devices, SD-Access and SD-WAN components are managed separately. SD-WAN Edges colocated with SD-Access Border/Control Plane Node functionality are provisioned by SD-WAN Manager, which receives configurations for SD-Access components from Cisco DNA Center.

Control Plane Integration

SD-Access VNs are mapped to SD-WAN Service VPNs. Prefixes are shared withing each virtual segment by redistributing LISP into OMP and vice versa.

Data Plane Integration

VNIs are copied from VXLAN header into IPsec header and vice versa when SD-Access - SD-WAN boundary is traversed.

Policy Plane Integration

SGTs are copied from VXLAN header into IPsec header and vice versa when SD-Access - SD-WAN boundary is traversed.

Overall Picture

Integrated Domain operation is visualized in Fig. 2.11. VNI and SGT placement in VXLAN and IPsec headers is visualized in Fig. 2.12

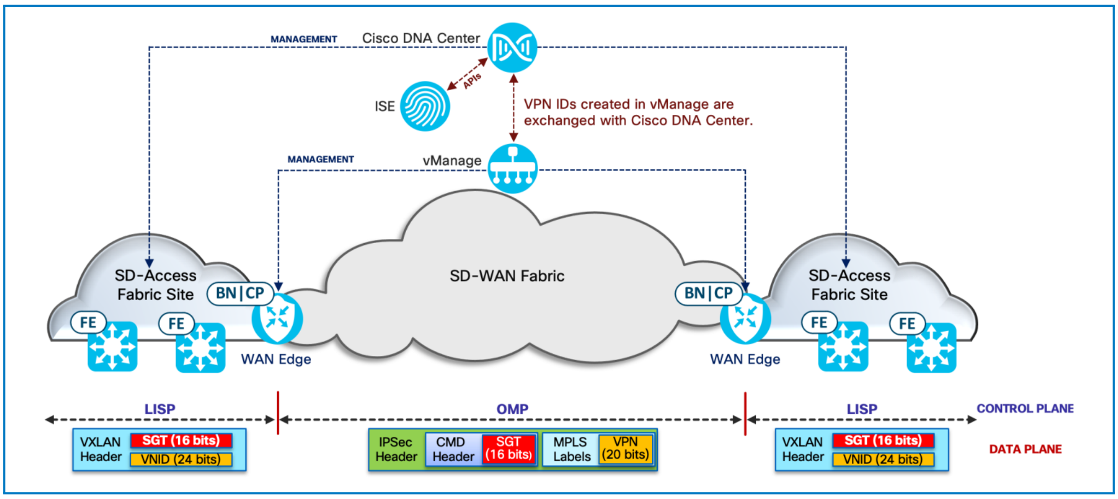

■ **Figure 2.11** Integrated Domain Components [32]

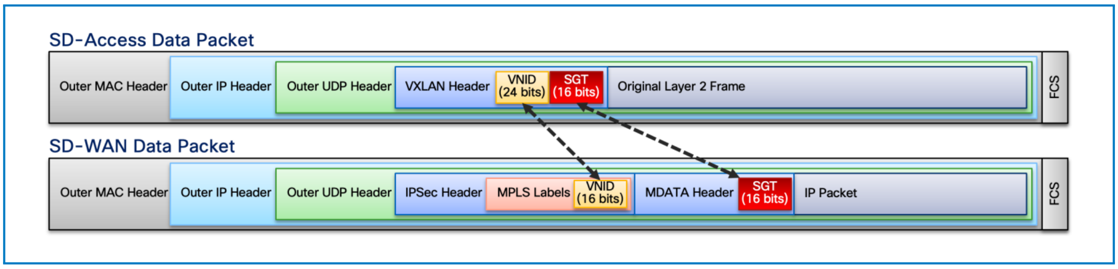

■ **Figure 2.12** VXLAN to Inline Tagging to IPsec [32]

Chapter 3

Analyzing the Traditional Network of an Existing Large Enterprise

3.1 Infrastructure Overview

CNX is network infrastructure of an anonymized existing large classic enterprise. Its logical design is depicted in Fig. 3.1. CNX is represented by Campus Network, interconnected by WAN with more than 100 branches, further referred to as Secondary Locations.

Campus Network is made up of 2 large offices, further referred to as Primary Locations, and 1 Colocation Data Center. They are interconnected by private links. Primary Locations house employees and some data center equipment. Campus Network offers centralized services for Secondary Locations, including Internet access. The enterprise is security-sensitive and does not take advantage of cloud services.

In contrast to Primary Locations, Secondary Locations have a simpler network topology. They house significantly fewer employees, lack data center equipment and locally available services, have a reduced redundancy degree.

CNX is a traditional infrastructure. Its operation heavily relies on physical interconnection of devices and configuration of protocols on each of them. Each device is managed separately through CLI.

The vast majority of the company's network equipment is manufactured by Cisco, Cisco proprietary networking protocols are used (EIGRP, HSRP, DMVPN, etc.).

3.2 Campus Network Description

Campus Network is composed of 2 Primary Locations and 1 Colocation Data Center. It follows a traditional two-tier collapsed core architecture. Core devices of the 3 locations are interconnected by private links.

Primary Locations and Colocation Data Center house data center equipment and offer centralized services for Secondary Locations. Additionally, Primary Locations house most human resources of the company.

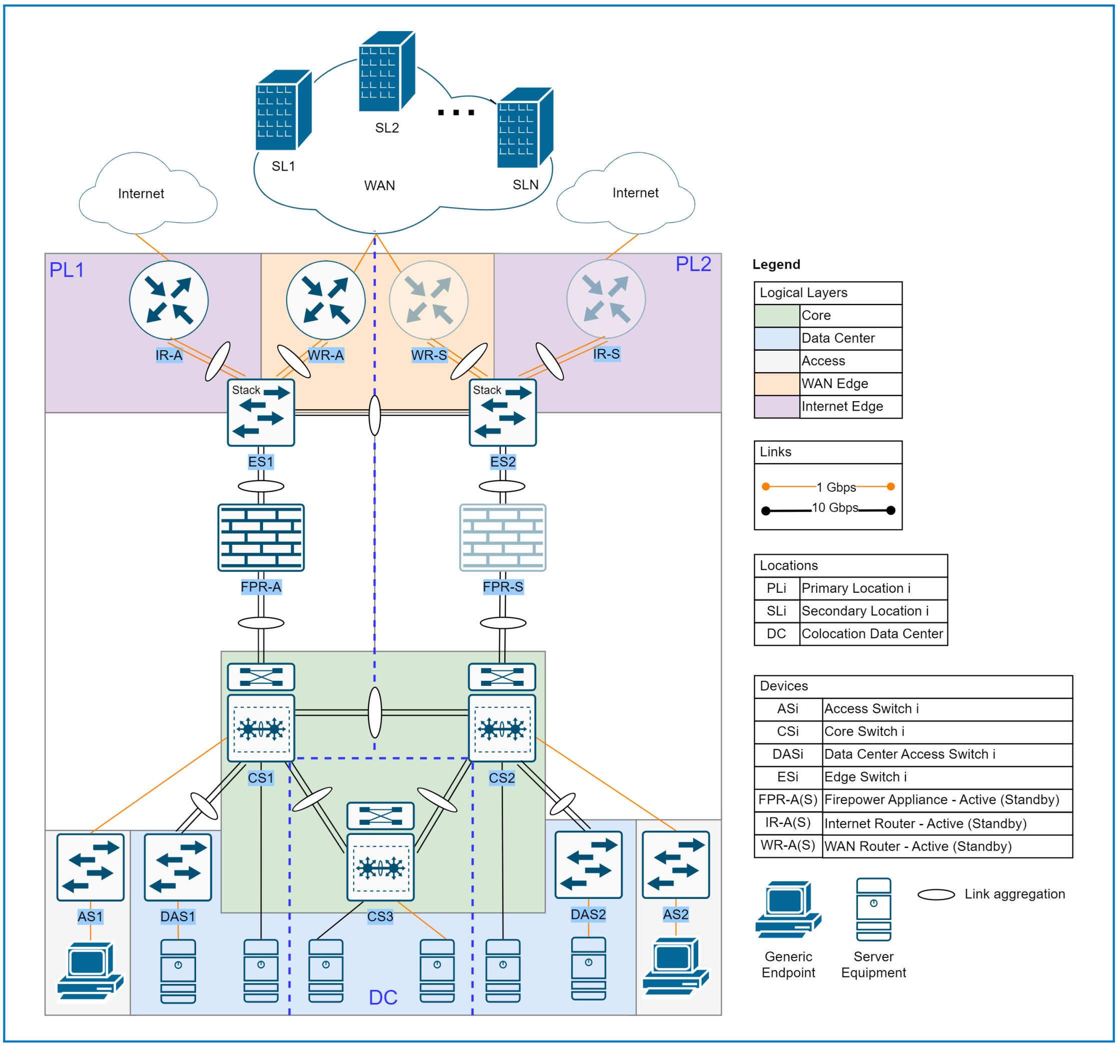

Logical Layers	
🟩	Core
🟦	Data Center
⬜	Access
🟧	WAN Edge
🟪	Internet Edge

Locations	
PLi	Primary Location i
SLi	Secondary Location i
DC	Colocation Data Center

Devices	
ASi	Access Switch i
CSi	Core Switch i
DASi	Data Center Access Switch i
ESi	Edge Switch i
FPR-A(S)	Firepower Appliance - Active (Standby)
IR-A(S)	Internet Router - Active (Standby)
WR-A(S)	WAN Router - Active (Standby)

Figure 3.1 CNX Original Design

3.2.1 Primary Components

Primary Locations' infrastructure can be logically divided into Core, Data Center, WAN Edge, Internet Edge, Access layers. The locations additionally accommodate virtual firewall instances, which perform centralized traffic filtering for the whole network.

Logical layers at Colocation Data Center are Core and Data Center.

Core Layer

Core Layer is made up of a total of three collapsed core devices, further referred to as "core switches", residing at Primary Locations and Colocation Data Center. They are represented by pairs of L3 switches (Cisco Catalyst 6500 Series Switches), each pair forming a VSS – a single logical switch. The core switches are directly interconnected at L3 with 10 Gbps aggregated fiber-optic links for a higher throughput. IS-IS is used as a routing protocol in the interconnect. VPLS is running in the interconnect to provide L2 connectivity between sites. The communication is MACsec encrypted at L2.

Firewalls
Each Primary Location houses a Cisco Secure Firewall 4100 Series appliance, further called ”Firepower appliance”, northbound of core switches. Firepower appliances operate in Active/Standby mode for redundancy. They connect to core switches via 1 Gbps aggregated links. Firepower appliances exchange routes with the rest of the network via OSPF.

Each appliance runs 5 FTD software virtual instances. These are logical firewalls, filtering traffic in WAN Layer (partner VPN), Internet Layer (Internet access, remote access and site-to-site VPN) and Data Center Layer (server access).

FMC, deployed in the form of 2 virtual instances operating in Active/Standby mode, acts as centralized FTD management center. The instances migrate to suitable physical devices in server farms to satisfy requirements for resources.

Firewall rules are IP-based, Firepower appliances are not integrated with AD.

Edge Switches
Northbound of the Firepower appliances are stacks of 2 switches (Cisco Catalyst 3650 Series Switches), 1 at each primary site. They are connected to the appliances via 1 Gbps aggregated links and to each other via 10 Gbps aggregated links. The stacks will be further referred to as “edge switches”.

L2 interconnect between edge switches ensures L2 communication between Firepower appliances, WAN and Internet routers, as well as operation of HSRP protocol.

Data Center Layer
Data Center Layer at the 3 sites is represented by access switches southbound of core switches, server farms and corresponding FTD instances. Server equipment requiring 10 GE ports is connected directly to core switches. Devices requiring 1 GE ports are connected to core switches directly or, in the lack of enough ports, to stacks of 2 access switches (Cisco Catalyst 3750-X Series Switches). The stacks, further referred to as ”data center access switches”, use 10 Gbps aggregated uplinks to core switches.

VLANs of Data Center Layer are terminated at SVIs of core switches. FTD instances and ACLs on SVIs enforce access policies.

Access Layer
Access Layer is represented by single or daisy-chained access switches (Cisco Catalyst 2950/2960/3550 Series Switches) southbound of core switches, connecting users to the corporate network. Switches provide connectivity for around 1000 wired endpoints and 40 APs at each Primary Location. Access switches connect to core switches via 1 Gbps links.

Core switches terminate local VLANs. STP protocol is configured to eliminate loops. ACLs on core switches’ SVIs enforce access policies.

Internet Layer
Internet Layer is represented by 2 redundant Internet routers (Cisco 3900 Series ISR), 1 at each primary site, northbound of edge switches. It also includes corresponding FTD instances, Cisco WSA and Cisco ESA. 1 Gbps aggregated links are used to connect Internet routers to edge switches.

Each Internet router is redundantly connected to 2 ISPs, BGP is used as a routing protocol in provider network. Local network facing interfaces are configured with OSPF for routing.

WAN Layer
WAN Layer is represented by 2 redundant WAN routers (Cisco 4000 Series ISR), 1 at each Primary Location, northbound of edge switches, and corresponding FTD instances. Each WAN router is connected to 5 ISPs, each Secondary Location to 1-2 ISPs. 1 Gbps aggregated links connect WAN routers to edge switches. OSPF is used to exchange routes with local network.

For redundant encrypted WAN connectivity over public Internet links dual hub phase 2 DMVPN is used. Using this technology mGRE+IPsec tunnels are built from each secondary site (spoke) to each primary site (hub) via each ISP network connected to Secondary Locations. Temporary tunnels are built between Secondary Locations on demand to avoid forwarding data through hubs. EIGRP is used in the DMVPN cloud. Firewalls terminate partner VPN connections.

3.3 Secondary Locations Description

Secondary Locations accommodate up to 30 employees. Each location houses a router (Cisco 2900 Series ISR or Cisco 890 Series ISR) and user workstations, connected to the router directly or, in case of shortage of ports, via a switch. Each router is connected to one-two ISPs and maintains redundant DMVPN paths to both Primary Locations via each Internet transport. When two ISP links are available, they are used in Active/Failover manner. Secondary Locations have neither server equipment, nor local services on-site — these are centralized in Primary Locations and Colocation Data Center.

3.3.1 Features, Characteristics

Hierarchical Design

Campus Network follows a two-tier collapsed core architecture.

Access layer is represented by access switches, data center access switches stacks. Stacks are built using StackWise Virtual technology.

Collapsed core and distribution layer functions are performed by Core Switches. Core Switches are VSSs.

Redundancy

CNX includes redundant devices (Active/Standby Internet, WAN routers and Firepower appliances), aggregated links, connections to internet providers. Additionally, Logical Core and data center access switches are composed of redundant physical devices with the help of VSS and StackWise technologies.

Internet and WAN routers operate in Active/Standby mode. When data originates at Primary Site the active router is elected by HSRP protocol. At Secondary Sites data flow to the active WAN router is part of DMVPN configuration. Active ISP link selection is configured using BGP attributes and policies.

L2 interconnectivity between Internet and WAN routers and Firepower appliances, as well as operation of HSRP protocol are ensured by L2 logical connectivity between edge switches.

Network Segmentation

The organization uses a combination of VRF-Lite and 802.1Q VLAN technologies to provide network segmentation.

VRF-Lite technology allows to isolate major data flows (user, server, voice, WAN, Internet, DMZ, etc.). VRF-Lite creates multiple L3 virtual networks on top of a physical infrastructure. Devices maintain separate routing table instances per VRF and have routing protocols configured per VRF. More than 10 VRF-Lite instances are configured in CNX.

802.1Q VLAN technology is used to logically separate company departments, user/device groups within a VRFs. VLANs are separate L2 domains, communication between them is performed at L3. More than 150 VLANs are configured in CNX. Some VLANs span several access switches, some are extended between locations. This requires configuring loop-avoidance protocols (STP, etc.) and enabling L2 communication between locations (VPLS, etc.).

Authentication and User/Device On-Boarding

Users and devices authentication is managed by AD. Cisco ISE is used to manage authentication during remote connections only. Authentication in AD grants access to AD resources.

However, independent on success of authentication a device is given the VLAN/subnet hard-coded to the wired port it connects to.

ACLs and rule-sets on FTDs filter based on IP, rather than identity.

Policies

The organization administers multiple types of policies, but lacks means of unified management of those.

Security policies are implemented using ACLs on VSSs of core switches and rule-sets on FTDs. Access control policies are implemented by placing devices into separate VLANs/subnets based on their role. QoS policies are implemented using queues on network devices to prioritize one application over another.

Network Monitoring

The organization uses legacy monitoring and data collecting tools like SNMP and syslog.

Wireless

CNX has a wireless segment for guest traffic only. It is isolated from corporate network, traffic is redirected to Internet services.

3.4 Analysis of the Current Design

3.4.1 Advantages

Hierarchical Design

Collapsed core architecture at Primary Locations and Colocation Data Center divides roles/-functions between hierarchical layers, offering greater flexibility and simpler management of each layer.

At each Primary Site, distribution layer, colocated with core layer at core switch, is composed of a single logical L3 switch. This is achieved with the use of VSS technology and brings several advantages to the access-distribution block. Use of VSS eliminates the need for FHRP configuration, allows multiple active uplinks per VLAN for access devices and L2 loop avoidance in spanned VLANs. It also leads to faster L3 re-convergence in core-distribution block.

Redundancy

Redundancy raises network uptime due to increased fault tolerance. Redundant aggregated links come with an additional benefit of increased bandwidth.

Network Segmentation Benefits

CNX is segmented using 802.1Q and VRF Lite technologies. Network segmentation simplifies and enhances control of network communication. It allows to easily add/remove logical networks and define communication rules between them without the need to change the physical infrastructure.

VRF Lite provides complete isolation of L3 traffic flows (guest, corporate, etc.) with the help of separate routing table instances. This removes the need to build complex ACLs, since inter-VRF communication is typically configured via route leaking or through a firewall [8, p. 23]. Additionally, VRF Lite offers greater flexibility in designing the logical networks, since VRF instances have separate routing protocol and policies configuration.

802.1Q VLAN provides simple, yet great control of L2 traffic flow between user/device groups. It allows to define inter-VLAN communication rules at L3 boundary (ACLs, Firewall rules) [8,

p. 22]. The technology also limits broadcast domain, reducing bandwidth usage and attack surface.

Firewall Design
Having a separate FTD instance for filtering different types of traffic simplifies overall logical firewall management, including troubleshooting and firewall rules design. Management is furthermore simplified with the help of FMC.

3.4.2 Disadvantages

Network Segmentation Limitations [8, p. 22-23]
VRF-Lite deployments don't scale well beyond 8-10 VRFs. VRF-Lite 802.1Q trunks and route leaking are difficult to implement in large networks. Each additional VRF process increases the CPU load.

VLAN as a segmentation technology is not scalable (number of VLANs is limited by the 12-bit VID field). Extending VLANs across L3 interconnects between locations requires additional configuration of corresponding protocols (VPLS in CNX). Moreover, widely spanned VLANs are vulnerable to L2 loops and require complex configuration of loop-avoidance protocols (e.g STP) per VLAN. Large L2 designs are inefficient (50 % of ports blocking typically), complicate IP address planning and require designing complex ACLs to control traffic flows.

Authentication and User/Device On-boarding Limitation
Hard-coding of VLANs/subnets to wired ports limits wired user mobility. Workplace change may require re-assigning VLAN/subnet to previously and newly connected ports, updating ACLs and firewall rule-sets. Additionally, such an approach offers little security, since any device connecting to some port will get access to the network. Moreover, the level of access will be the same for any device, independent on success of authentication.

Using IP instead of identity for controlling access to the network has its disadvantages: limited visibility into who accesses resources, complexity of creating and troubleshooting access rules. IP addressing is not human-friendly, thus it is difficult to analyze rules based on IP. It is common not to change already existing rules not to cause unintentional misconfigurations. Moreover, if a device/user gets a new IP for any reason, maintaining the same access policy would require changes to ACLs and Firewall rule-sets.

Policies Implementation Complexity [8, p. 24-25]
The organization lacks means of unified management of its policies.

Adding voice to the network would require carving a new set of VLANs and associated subnets, updating ACLs, firewall rule-sets and QoS policies. All the devices affected by the change would need to be configured manually and separately.

Such a task is error-prone and may take several working days. Moreover, complexity in managing, understanding and implementing the policies makes the network more vulnerable to security threats.

Network Monitoring Limitations
Legacy network monitoring tools (SNMP, syslog) don't scale well, lack real-time monitoring, have limited capabilities in monitoring virtualized environments, etc.

Use of Outdated Devices
Many Cisco networking devices in CNX are obsolete. Maintenance releases, bug fixes or software remedy for security vulnerability issues are not released for unsupported devices. The hardware is not serviced/repaired by Cisco. This increases the possibility of device failure and vulnerability exploitation. Cisco policy regarding treatment of obsolete devices, End-of-Life Policy, is

described in more detail in [33].

The most critical devices for which the policy applies include core switches (Cisco Catalyst 6509-E Switch), Internet routers (Cisco 3900 Series Integrated Services Routers), edge switches (Cisco Catalyst 3650 Series Switches) and access switches in the Colocation Data Center (Cisco Catalyst 3750 Series Switches).

Manual configuration via CLI

Network management is solely built on manual, non-centralized device configuration via CLI. In large scale it is time-consuming, error-prone and complicated, considering the number of technologies and protocols used. Some level of network automation via CLI scripts is possible, but complex and vendor-specific.

Chapter 4

Redesigning the Traditional Network into SDN

4.1 Redesigning CNX into SD-Access and SD-WAN

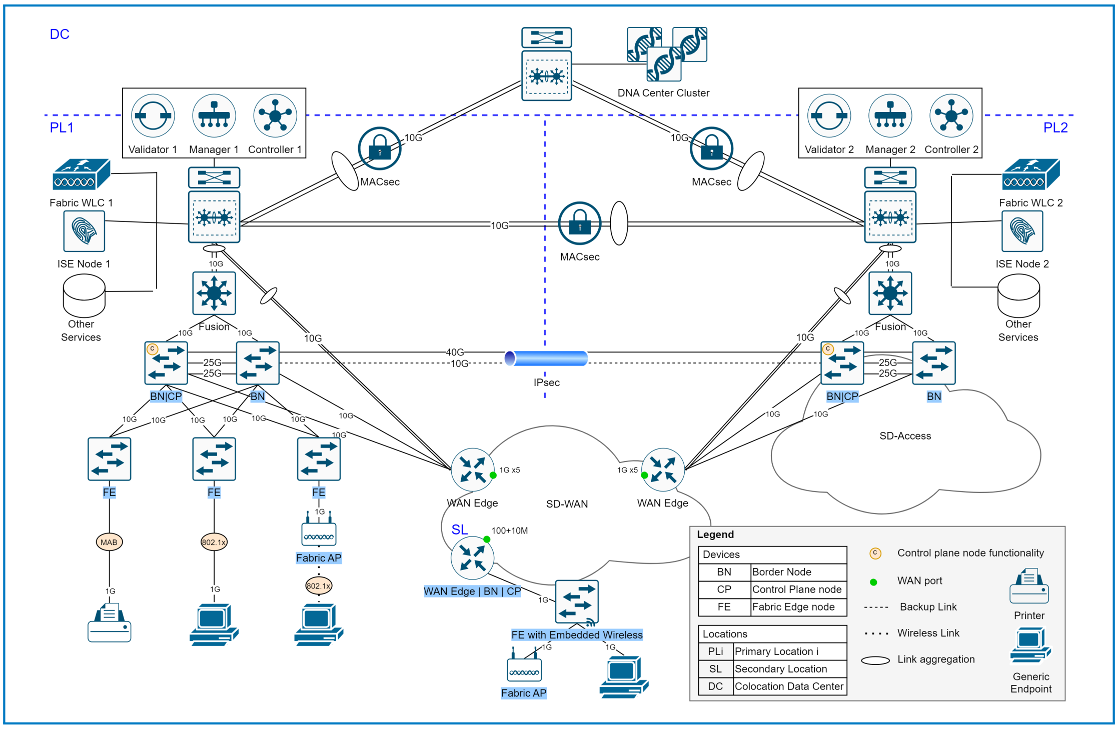

Figure 4.1 CNX Redesign using Cisco SD-Access and Catalyst SD-WAN Solutions

4.1.1 Overview

The proposed redesign of CNX is depicted as a logical schema in Fig. 4.1. Design of physical connectivity is not included in the thesis.

Access and WAN Edge Layers at both Primary and Secondary Locations will be migrated to Cisco SD-Access and Cisco Catalyst SD-WAN, respectively. Wireless segment fully integrated with SD-Access network will also be added at Primary and, optionally, at Secondary Locations. SD-Access and SD-WAN at Primary Locations will be integrated into an Independent Domain, at Secondary Locations – into Integrated Domains. Each standalone solution will bring improvements to corresponding logical layer of CNX, while their integration will allow to preserve segmentation and policy constructs in inter-site communication.

Upgrading company's data center infrastructure at Primary Locations and Colocation Data Center is out of scope of the thesis. Since the legacy infrastructure is not SGT aware, policy context will be lost. Notes on theoretical preservation of SGTs in Data Center Layer can be found in 4.1.6. Endpoints in local and remote SD-Access VNs will be able to access selected resources in the legacy network with the help of Fusion devices and WAN Edge routers (5.4).

Static VLAN assignment to wired hosts will be replaced with dynamic VLAN and SGT assignment based on result of authentication.

All Controllers can be deployed in public cloud, but will be placed on-premises in the data center due to security reasons. Cisco ISE, already present in the network, will be reused.

Considerations regarding SD-Access and SD-WAN device models are presented in 5.1.

4.1.2 SD-Access Design Decisions

4.1.2.1 Control Components

All Cisco SD-Access control components: DNA Center, ISE and WLC are recommended to be placed in Shared Services Block outside the Fabric.

Cisco DNA Center

Cisco DNA Center can be deployed as a single-node or a three-node cluster. The latter provides high availability, but not scaling. When all nodes in a cluster fail, SD-Access network remains operational, even though GUI-based infrastructure management is not possible [28].

Cisco DNA Center can be deployed in a public cloud (AWS) or on-premises. When deployed on-premises, DNA Center nodes can be physical appliances or virtual instances on VMware ESXi. Physical appliances must be deployed in the same physical location. Virtualization offers more flexibility, however, requires additional expenses on VMware licenses. More on deployments in AWS and on ESXi in [34].

Due to the company's size and criticality of its operations, Cisco DNA Center is proposed to be deployed as a three-node cluster. Since the organization is security-sensitive, the cluster will be located on-premises, for example in Colocation Data Center. The company may choose nodes to be physical or virtual.

Cisco ISE

A Cisco ISE deployment is a cluster of one or more ISE nodes. Deployments are classifies as small, medium and large, based on the number of supported active endpoints [35]. Since this number is estimated to be less than 10000, a small deployment is selected for the company. It is represented by 2 ISE nodes, running all personas. Existing ISE nodes will be reused and configured to support integration with Cisco DNA Center. This includes enabling ERS API and pxGrid [36]. The nodes will serve both Primary and, optionally, Secondary Locations, the latter performing authentication over SD-WAN transit.

Fabric WLCs [37]

Fabric WLC can be deployed as a virtual or physical instance. N+1, SSO or hybrid high-availability models are supported for WLCs. SSO HA model is selected, since it can be automated with any version of Cisco DNA Center, unlike N+1 model. Moreover, it allows APs and clients to stay connected upon primary WLC failure. HA SSO model requires deploying a pair of controllers, these will be places in Primary Locations. The company may choose nodes to be physical or virtual.

4.1.2.2 Segmentation

Macro-Segmentation

Since CNX is segmented using VRF-Lite technology legacy VRFs can be easily replaced with SD-Access L3 VNs. For simplicity, CORPORATE_VN will further represent migrated VRFs and newly introduced wireless corporate traffic. Fabric APs, as per Cisco guidelines, are placed in INFRA_VN for connectivity to WLCs without route/VRF leaking. It is also recommended to create a separate VN for guest traffic, e.g. GUEST_VN.

Thus, INFRA_VN, CORPORATE_VN and GUEST_VN compose the minimum recommended set of L3 VNs to be used. These will be created, assigned IP address pools and associated to corresponding Fabric sites via DNA Center GUI.

Macro-segmentation will be preserved between SD-Access and SD-WAN enabled sites by means of carrying VNI in IPsec header. Fusion devices and WAN Edge routers at Primary locations will propagate legacy infrastructure routes into VNs to facilitate communication between non-migrated and migrated parts of the network (5.4).

VLAN Segmentation

Legacy VLANs will also be migrated to SD-Access. However, static VLAN assignment, based on values hard-coded to ports/SSIDs, will be replaced with dynamic one. Moreover, logical separation of endpoint groups and control of communication between them will be achieved with SGs and SGACLs, instead of VLANs and ACLs. Thus, in SD-Access user VLANs will be used solely for limiting broadcast domains.

Micro-Segmentation and Policies

Micro-segmentation using SGs and SGACLs will replace legacy VLANs/ACLs for controlling communication between logical endpoint groups. SGTs and related policies will be managed through Cisco DNA Center. ISE will dynamically assign SGTs to endpoints based on result of authentication and will push SGACLs to Edge nodes.

SGTs will be carried by SD-WAN in IPsec header to preserve policy context in inter-site communication. It would be beneficial to extend micro-segmentation awareness to Data Center Layer. Possibilities are discussed in 4.1.6.

4.1.2.3 Authentication [38]

For host on-boarding, ports of Fabric Edge nodes are configured to perform authentication of connected endpoints according to one of the 4 models. The most secure one, Closed Authentication, forces endpoints to go through 802.1x authentication or MAB if prior fails. If both fail, no access is granted to the network.

This authentication model will be used for wired clients. Authentication of wireless clients will be enforced on APs. 802.1x authentication will be used for corporate traffic, Web Authentication – for guests. Authentication requests will be redirected by Fabric Edge nodes/APs to ISE. Authentication methods: 802.1x, MAB, Web Authentication are described in [39].

4.1.3 Cisco Catalyst SD-WAN Design Decisions

4.1.3.1 Control Components [26, 40]

All SD-WAN controllers are deployed as virtual instances on premises or in cloud. On-premises deployment in Shared Services Block is selected for security reasons. Each of the control components will be made redundant. For each controller type redundancy models are described in [26], and the recommended number of instances in [40]. The latter depends on the number of managed WAN Edge routers. This number will be less than 250.

Cisco Catalyst SD-WAN Validator
SD-WAN Validator redundancy is achieved by deploying multiple independent instances, preferably in different geographic regions. For <250 WAN Edge routers Cisco recommends to deploy 2 Validators. Each Primary Location will accommodate one. To reference the Validators it is recommended to use a single FQDN associated with their IPs. FQDN is configured on all SD-WAN components. WAN Edge routers will try all Validators' IPs in succession until control connections can be successfully formed.

Cisco Catalyst SD-WAN Controller
SD-WAN Controller redundancy is achieved by deploying multiple instances operating in Active/Active fashion, and preferably placed in different geographic regions. For <250 WAN Edge routers Cisco recommends to deploy 2 Controllers. Each Primary Location will accommodate one. To synchronize, Controllers maintain a full mesh of control connections with each other. By default each WAN Edge router maintains control connections to 2 Controllers over each transport to handle the situation when one of them fails. The default settings will not be changed.

Cisco Catalyst SD-WAN Manager
SD-WAN Manager clustering is primarily used for scale, while redundancy is achieved primarily by Active/Standby operation of standalone instances or clusters. Members of a cluster should reside in the same geographic region, while Active/Standby ”nodes” – in different locations. For <250 WAN Edge routers Cisco recommends to deploy 1 small Manager instance. For redundancy, 2 will be deployed and configured in Active/Standby fashion. Each Primary Location will accommodate one. Each WAN Edge router maintains a control connection to one SD-WAN Manager over 1 transport. Preferred transport and TLOC can be configured for each device.

4.1.3.2 Overlay Topology [26]

By default, each WAN Edge forms IPsec tunnels from each local TLOC to each remote TLOC. This results in a so-called full mesh topology between TLOCs.

Since Secondary Locations communicate mostly to Primary Locations, it is inefficient to maintain idle tunnels between them. Hub-to-spoke Overlay topology will be configured, with Primary Locations acting as a single dual-router hub. This will cause infrequent traffic between Secondary Locations to flow through the Primary ones. If not desired, on-demand tunneling can be configured to build tunnels between Secondary Locations when communication is initiated.

Hub-to-spoke topology will be configured with the help of a centralised policy, pushed to SD-WAN Controller.

4.1.4 Primary Location Redesign

Primary Locations will represent a singe SD-Access Fabric and a single SD-WAN site. These will be integrated in SD-Access | SD-WAN Independent Domain for macro/micro-segmentation preservation. Fusion devices and WAN Edge routers will provide integration of SD-Access and SD-WAN networks with legacy part of the infrastructure. The legacy part will include non-migrated Data Center Layer, legacy WAN Layer and Shared Services Block.

4.1.4.1 Shared Services Block

Shared Services block will span both Primary Locations and Colocation Data Center. Apart from legacy network services (DNS, AAA, AD, etc.) it will include existing ISE nodes and newly deployed SD-Access and SD-WAN control components.

Since legacy services are placed in GRT, Shared Services block will reside in GRT as well, rather than in a separate VRF. This is aimed at not disrupting the legacy network operation and at simplifying services management.

Routes to resources in Shared Services Block will be advertised to Fabric VNs with the help of Fusion devices and WAN Edge routers (5.4).

4.1.4.2 Cisco SD-Access Network

Legacy Access Layer in both Primary Locations will be migrated to SD-Access. The Fabric infrastructure described below will be deployed at both Primary Locations. It will have to support more than 1000 connected wired endpoints, as well as wireless users.

Two-tier hierarchical model will be used, where Border nodes and Edge nodes will represent collapsed core and access devices, respectively. Each Primary Location will accommodate 2 redundant Border nodes, 1 of them with colocated Control Plane node functionality. The number of Edge nodes will depend on selected models and on the number of ports required. According to Cisco recommendations, each Edge node will be connected to each site-local Border node via 10 Gbps links. Endpoints and Fabric APs will connect to Edge nodes via 1 Gbps links. Same-site Border nodes will be cross linked to each other using 2 redundant 25 Gbps links [28].

The 2 locations will represent a single Fabric site. This will be achieved by interconnecting pairs of geographically distant Border Nodes with 40 Gbps primary and 10 Gbps backup links. Leased lines secured with IPsec will be used for interconnection.

Communication to legacy part of the infrastructure will be facilitated by Fusion devices and WAN Edge routers (5.4). Fusion devices will be connected via 10 Gbps links to each Border Node and via aggregated 10 Gbps links to Core switch. WAN Edge devices will be connected via 10 Gbps links to each Border Node and via aggregated 10 Gbps links to Edge switch[1].

4.1.4.3 Cisco SD-WAN Network

SD-WAN network will operate alongside the legacy WAN Layer for migration purposes and to allow communication to locations which will not be migrated. Each Primary Locations will accommodate a WAN Edge router, connected to legacy WAN routers, and belonging to a single L2 domain due to a L2 link between Edge switches (5.4.2). This is a measure of redundancy.

WAN Edge routers will operate in the Active/Standby fashion with automatic fail-over in case of primary component failure. On the LAN side they will be accessible via a single virtual IP, acting as a gateway. To make traffic flow symmetric, communication to Primary Locations will be forced through the primary WAN Edge device by controlling TLOC preference. Each time a primary component is elected its TLOC preference is automatically modified to be higher than the one of the secondary component. The described functionality is achieved by configuring VRRP protocol [26].

Border nodes will direct traffic to remote SD-WAN enabled locations to WAN Edge routers, and to Core switches if communicating to non-migrated remote locations. In latter case legacy WAN communication will be used.

[1] Fig. 4.1 depicts WAN Edge as connected to Core switch. This is a logical representation of connection to legacy infrastructure. A more accurate connectivity schema is presented in Fig. 5.4

4.1.4.4 Cisco SD-Access | SD-WAN Integration

Cisco SD-Access and SD-WAN deployments at Primary Locations will be integrated into an Independent domain, where the mentioned deployments are managed separately by Cisco DNA Center and SD-WAN Manager. This will allow preserving macro/micro-segmentation between SD-Access sites.

Links between Border nodes and WAN Edge routers will be used for mapping SD-Access VNs to SD-WAN Service VPNs. Fig. 5.4 depicts mapping of CAMPUS_VN to VPN10 and INFRA_VN to VPN 100. The links will be configured as trunks and will preserve segmentation with the help of 802.1Q tagging.

4.1.5 Secondary Location Redesign

Secondary Locations will be migrated to SD-Access | SD-WAN Integrated Domain, because it is more suitable for small branches. Since they house up to 25 wired endpoints it is proposed to deploy the minimum number of devices. For the selected deployment type the minimum is 2: Colocated SD-Access Border, Control Plane and SD-WAN WAN Edge, as well as an SD-Access Edge node. Selected device platforms will have to support the above mentioned functionality. Colocated device must additionally accommodate 1-2 WAN ports, while Edge node must provide enough LAN ports.

In case wireless infrastructure is decided to be added at Secondary Locations, embedded WLC is proposed to be installed on the Edge node. A Fabric AP will be connected to the Edge node. INFRA_VN will be added to the site for AP to WLC connectivity. For guest access it is recommended to use a dedicated VN: GUEST_VN.

Legacy VRFs will be migrated to SD-Access L3 VNs, VLAN segmentation will be preserved in case of need. Endpoints will be added to security group for policy definition and enforcement.

All required SD-Access and SD-WAN components, except WLC, will be accessed via SD-WAN, because they will reside in centralized Shared Services Block.

4.1.6 Future of Data Center Layer

It is advantageous to extend micro-segmentation awareness to the Data Center Layer, when Access Layer in Primary and Secondary Locations is migrated to SD-Access and uses SGTs for policy enforcement. This would allow to replace VLAN/IP-based ACLs and firewall rules with SGACLs and SGT-based firewall rules respectively. Such an upgrade would eliminate the need to manage VLAN/IP-based ACLs/firewall rule sets and SGT-based policy rules separately. Moreover, this would simplify policy management in general, since designing VLANs/IP-based access rules is complicated.

However, migrating the Data Center Layer to Cisco TrustSec architecture is viable if legacy WAN is migrated as well. This would allow remote users and partners to access data center resources based on SGTs.

The migration would require designing SGs and assigning data center resources and hosts accessing Data Center Layer to these groups statically or dynamically based on result of authentication. Data Center and WAN Layer should be reconfigured to propagate SGTs. In Data Center Layer this could be done with the use of inline tagging or SXP, depending on capabilities of devices [28]. In legacy WAN Layer SGTs can be carried in IPsec header.

However, the steps below require careful redesigning of old undocumented policy rules (ACLs, firewall rule sets) into new SGT-based format to maintain the same access logic. Configuring SGT propagation would require checking compliance of and possibly updating the devices OS or hardware. Reconfiguration will be time-expensive and will require a careful planning of migration not to disrupt the network operation.

Migration of CNX to Cisco TrustSec architecture is a logical step, but the effort and expenses associated with it may not be viable. Legacy policy rules are corresponding to current requirements of the company. Devices in the remaining legacy part of CNX are outdated and unsupported, and are likely to be replaced. Thus, it is difficult to prove that migration to Cisco TrustSec is a vital step before CNX is upgraded, expanded or changes logic of its operation.

Chapter 5

Planning Technical Migration into SDN

When migrating CNX into SD-Access and SD-WAN it is preferable to

1. Reuse as many legacy devices as possible;
2. Properly schedule migration of locations to reduce network downtime;
3. Pick suitable migration strategies for SD-Access and SD-WAN;
4. Configure interoperability between new deployment and remaining legacy infrastructure.

5.1 Reusing Legacy Devices. Model Selection Considerations

At Primary Locations legacy devices in Access Layer will not be reused, since they are outdated and cannot be upgraded to support SD-Access functionality. Legacy WAN routers will remain in place for migration purposes and not to disrupt communication to partner locations. WAN Edge routers running Cisco IOS XE OS will be additionally deployed. At Secondary Locations all legacy networking devices will be replaced due to lack of support of required functionality.

It is worth mentioning, that software versions of control components and other devices in SD-Access/SD-WAN network must be compatible.

Both virtual and physical platform requirements for running Cisco IOS XE software are available at [41]. Compatibility of these platforms with different Control Components software versions can be verified at [42]. Platform requirements for supporting various Fabric roles depending on Cisco DNA software version can be studied at [43].

5.2 Location Migration Order and Number of Maintenance Windows [44]

When migrating to SD-WAN it is recommended to start with hub/data center locations, followed by branches. Moreover, since branches are typically migrated gradually, SD-WAN is configured alongside legacy WAN at hubs/data centers to allow their interoperability [27, p. 25].

For SD-Access | SD-WAN integrated deployments at large sites (hubs/data centers) it is recommended to firstly deploy SD-WAN, followed by SD-Access. Small locations can be migrated in one maintenance window.

In case of CNX, Primary Locations resemble hubs/data centers, Secondary Locations – branches. Thus, the migration process is recommended to consist of the following steps:

1. Migrating standby Primary Location to SD-WAN and half of Secondary Locations to SD-WAN and SD-Access;
2. Migrating active Primary Location to SD-WAN and the remaining half of Secondary Locations to SD-WAN and SD-Access;
3. Migrating Primary Location to SD-Access.

Active and standby Primary Locations are the ones, that accommodate active and standby WAN, Internet routers and Firepower Appliances. Standby location will be migrated first not to interrupt WAN communication in case of failure and to minimize potential downtime.

5.3 Migration Strategies

5.3.1 SD-Access Migration Strategies [28]

5.3.1.1 Parallel Migration

Parallel approach doesn't involve reusing the legacy infrastructure. Instead, a new independent SD-Access network is built alongside. When SD-Access deployment is fully operational endpoints are disconnected from access switches and connected to Fabric Edge nodes. Changes can be easily reverted by connecting users back into old network. When the new deployment functions as desired the legacy network can be decommissioned.

5.3.1.2 Incremental Migration

Incremental approach involves gradually converting the legacy infrastructure into SD-Access network. Portion of the network being migrated at a time is typically an access-distribution block. To migrate a block distribution-access links are reconfigured for L3 communication and access switches are converted to Fabric Edge Nodes. The incremental approach is not possible if the above mentioned operations are not supported by devices.

Incremental migration is assisted by a dedicated Border Node which connects with trunk links to distribution switches of the block being migrated. It performs VLAN translation from Fabric to non-Fabric to allow transparent L2 communication of already migrated to non-migrated endpoints on the same subnet/VLAN. This functionality is referred to as Layer 2 Border Handoff, Fig. 5.1.

5.3.1.3 Strategy Choice for CNX

Since devices at both Primary and Secondary Locations will not be reused, migration to SD-Access will follow the parallel approach.

5.3.2 SD-WAN Migration Strategies

Data centers are recommended to be migrated to SD-WAN prior to branch locations. They serve as transit for legacy WAN and SD-WAN traffic to allow for gradual branch migration [27, p. 25]. This is also required for the cases where legacy WAN cannot be fully replaced with SD-WAN. Branch sites are typically fully migrated to SD-WAN.

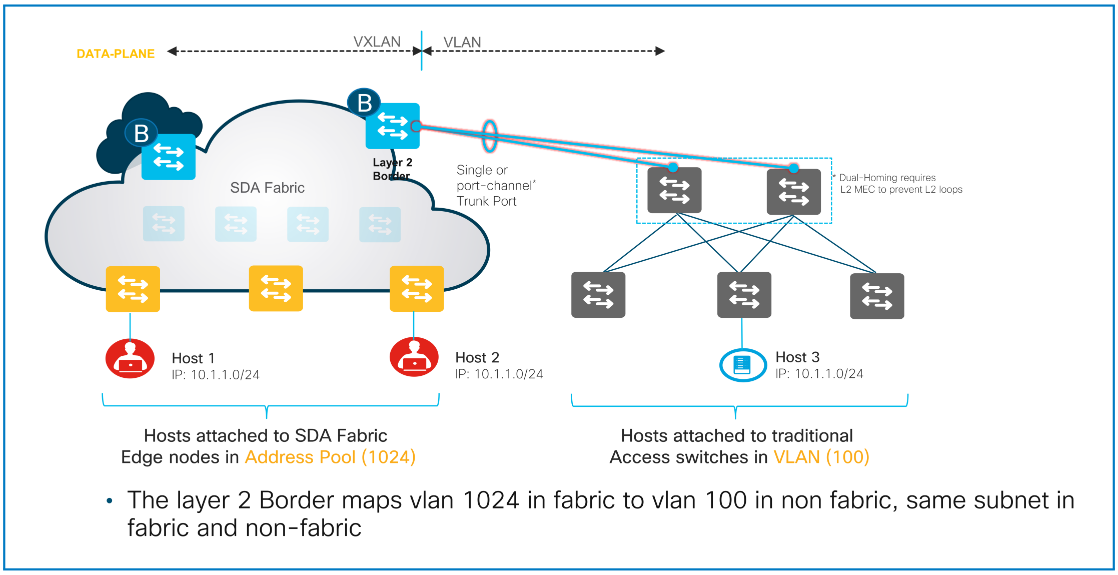
DATA-PLANE
VXLAN
VLAN
B
B
SDA Fabric
Layer 2 Border
Single or port-channel* Trunk Port
Dual-Homing requires L2 MEC to prevent L2 loops
Host 1
IP: 10.1.1.0/24
Host 2
IP: 10.1.1.0/24
Host 3
IP: 10.1.1.0/24
Hosts attached to SDA Fabric Edge nodes in Address Pool (1024)
Hosts attached to traditional Access switches in VLAN (100)

- The layer 2 Border maps vlan 1024 in fabric to vlan 100 in non fabric, same subnet in fabric and non-fabric

Figure 5.1 Layer 2 Border Handoff [45]

5.3.2.1 Data Center Migration [27, p. 25-30]

To integrate SD-WAN components with legacy WAN infrastructure without impacting legacy WAN traffic flow, WAN Edge routers are placed behind existing WAN routers, further referred to as CEs. Thus, WAN Edge routers' connectivity to transports is extended through CEs. WAN Edge routers' service-side interfaces are connected to the same devices CEs are connected to and are configured to run the same routing protocol. WAN Edge routers perform redistribution between OMP and LAN routing protocols. CE routes legacy WAN traffic to the core, and SD-WAN traffic to WAN Edge routers.

A recommended legacy WAN and SD-WAN integration topology, as well as routing protocols operation are illustrated in Fig. 5.2.

For networks with more than one data center location, the locations are migrated sequentially. If there's a direct link between data centers and route advertisement is happening over this link, routing loop avoidance mechanisms might need to be configured.

5.3.2.2 Branch Migration [27, p. 30-37]

When migrating a branch it is recommended to remove legacy WAN routers (CEs) or upgrade them to be SD-WAN capable, unless it is required to preserve them. In the latter case, a parallel migration is performed [27, p. 33-37]. In case a branch has more than one CE, they are migrated sequentially [27, p. 32-33].

When migrating a single router (Fig. 5.3) it is simply replaced with WAN Edge router or upgraded to support SD-WAN. Minimal reconfiguration on the LAN side is required, such as updating default gateway [27, p. 30-32].

5.3.2.3 Strategy Choice for CNX

Since Primary Locations accommodate Data Center Layer and WAN routers which act as hubs in the legacy WAN infrastructure, they should be migrated like data centers. Migration details are described in 5.4.2.

Secondary Locations will be migrated like single-router branches.

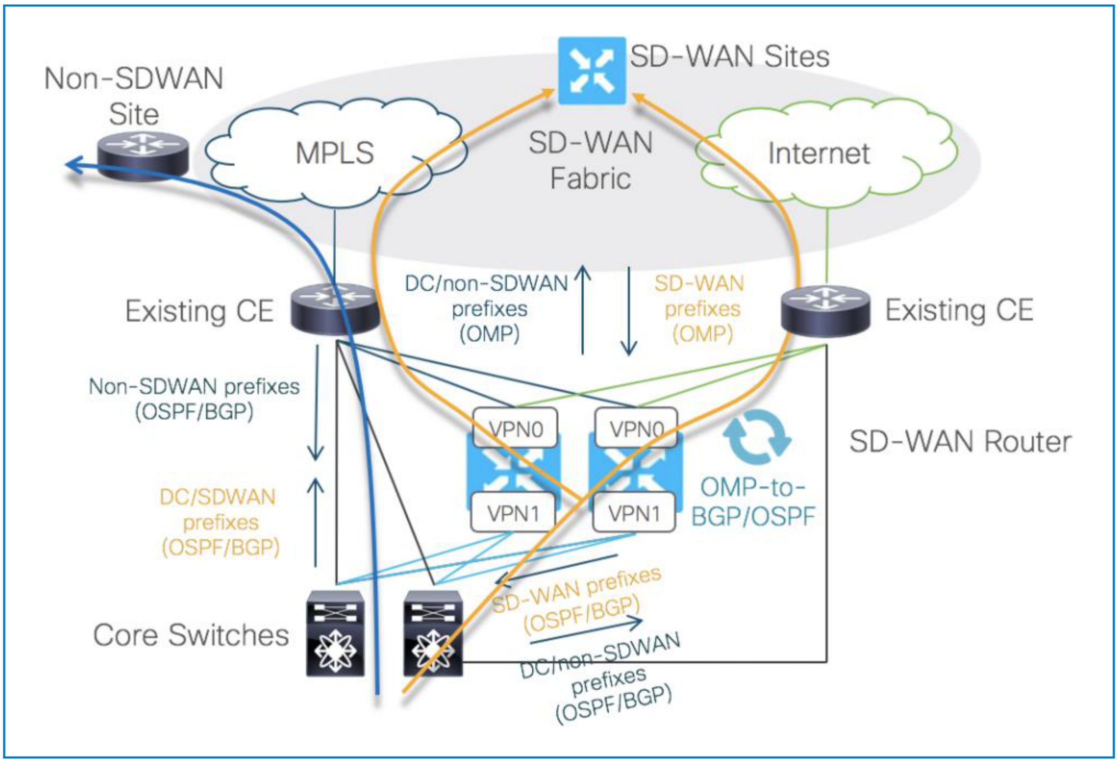

Figure 5.2 Legacy WAN and SD-WAN Integration at Data Center [27]

5.3.3 SD-Access | SD-WAN Integrated Deployment Migration

When a site houses both SD-Access and SD-WAN deployments, it is recommended to firstly enable SD-WAN functionality, then build SD-Access Fabric and, finally, integrate them into Independent or Integrated Domain. See deployment steps overview in [31, p. 30] and [32, p. 30]. Steps are explained and illustrated in [31, p. 29-62] and [32, p. 28-125].

5.4 Interoperability with Legacy Infrastructure

SD-Access, SD-WAN and legacy parts of the company's infrastructure [1] must be integrated for the network to act as a whole and to provide transparent communication for endpoints from different sections. Integration involves sharing routes between SD-Access/SD-WAN VNs/VPNs and legacy network to enable communication, while preserving VN/VPN isolation.

WAN Edge routers at Primary Locations enable communication to legacy infrastructure for the Whole SD-Access domain. Fusion routers at Primary Locations perform a similar functionality only for local SD-Access Fabrics for the purpose of redundancy. WAN Edge routers additionally facilitate route exchange between SD-WAN and legacy WAN infrastructures, that co-exist for the purpose of migration.

The logic of integrating Cisco SDN solutions with legacy network is illustrated at Fig. 5.4.

[1] Legacy infrastructure in the following context is defined as a section of CNX which was not migrated to SD-Access or SD-WAN. It includes Data Center Layer, non-migrated section of WAN Layer, and Shared Services Block. Although the last one accommodates SD-Access and SD-WAN control components, it physically resides in the legacy Data Center Layer.

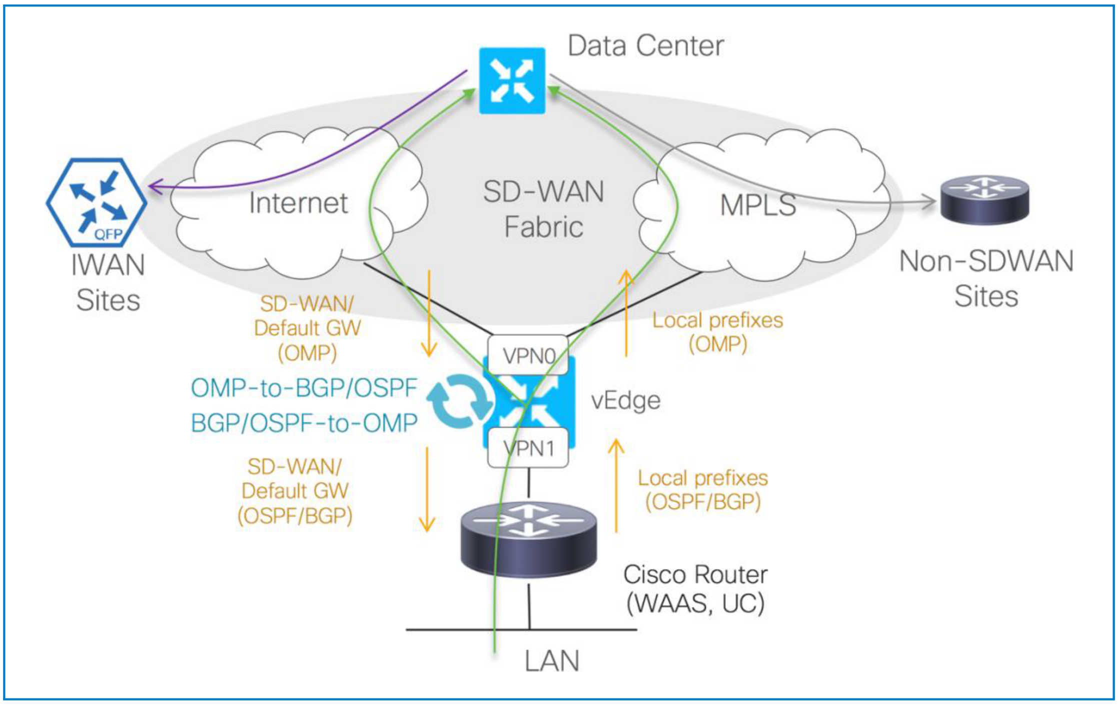

Figure 5.3 Single Router Branch Migration [27, p. 31]

5.4.1 Fusion Device. Legacy Infrastructure and Local SD-Access Network Interoperability

At each Primary Location a Fusion device will be placed between Border Nodes and Core switch. It will perform route leaking to distribute routes from Fabric VNs into legacy-side GRT and vice versa. However, isolation between VNs will be preserved by restricting inter-VN route sharing. Border node to Fusion link will be a trunk to allow mapping Fabric VNs to dedicated VRFs on Fusion device. For example, CAMPUS_VN is mapped to CAMPUS VRF. Fusion device's interface connecting to Core switch will be part of GRT. VRF to GRT route leaking will be configured.

Border nodes configuration will be automated through Cisco DNA Center by assigning Layer 3 Handoff role to the device. Fusion devices are configured manually.

5.4.2 WAN Edge router. WAN to SD-WAN Migration. Legacy Infrastructure and SD-Access Network Interoperability

Following recommended steps for data center migration to SD-WAN, WAN Edge router at each Primary Location will be directly connected to local legacy WAN router and Edge switch.

Link to legacy WAN router will be used for providing Underlay connectivity to WAN Edge device. The number of transport interfaces on WAN Edge router will match the number of Internet transports, used for SD-WAN communication. For simplicity, only 2 ISPs are depicted in Fig. 5.4. Each transport interface in VPN 0 will connect to an interface on legacy WAN router, residing in a VRF dedicated to corresponding ISP[2]. This will allow routes from ISP VRFs to

[2] A switch is required to be placed between devices in case of insufficient number of ports on WAN router. To

propagate to VPN 0, extending WAN Edge router's connectivity to transports via WAN router. BGP will be configured in the interconnect. SD-WAN IPsec tunnels will be terminated at WAN Edge router's transport interfaces.

Connection to Edge switch will be used for exchanging SD-WAN service VPN and legacy infrastructure prefixes. WAN Edge router will connect to Edge switch with 2 aggregated links for redundancy. Its 2 corresponding interfaces will be placed in VPN 20 which will map to legacy WAN VRF on Firepower appliance. This will allow WAN Edge router to be treated as another WAN router at the site, without any reconfiguration of routing logic in the legacy network. OSPF will be configured in VPN 20 to match the routing protocol in WAN VRF.

Route leaking between VPN 20 and other service VPNs will be configured with the help of SD-WAN centralized policy. Legacy network routes will propagate in service VPN routing table, service VPN routes will propagate to GRT in Primary Locations. Care must be taken to restrict inter-VPN route propagation to preserve isolation. Since service VPNs map to SD-Access VNs, the above mentioned configuration provides SD-Access network and legacy infrastructure interoperability.

map transport interfaces to VRFs each will need to map to a VLAN on the switch. Trunk link to WAN router will facilitate mapping VLANs to VRFs.

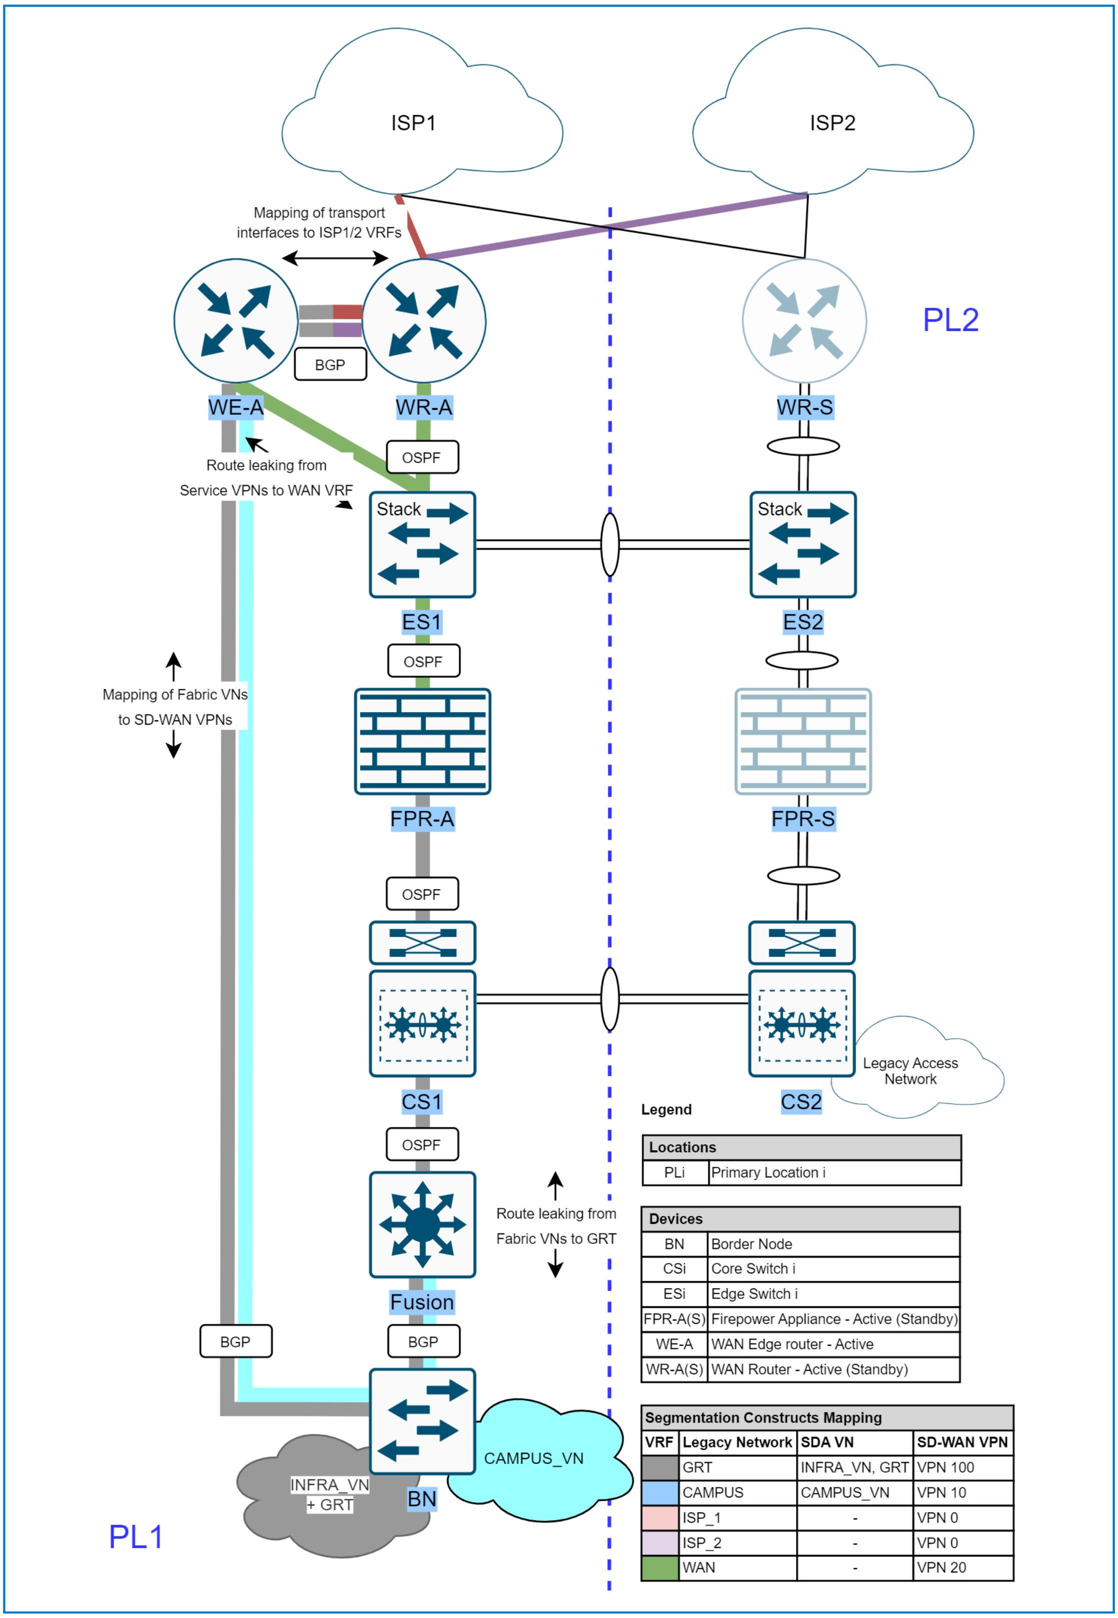

Legend

Locations	
PLi	Primary Location i

Devices	
BN	Border Node
CSi	Core Switch i
ESi	Edge Switch i
FPR-A(S)	Firepower Appliance - Active (Standby)
WE-A	WAN Edge router - Active
WR-A(S)	WAN Router - Active (Standby)

Segmentation Constructs Mapping

VRF	Legacy Network	SDA VN	SD-WAN VPN
⬛	GRT	INFRA_VN, GRT	VPN 100
🟦	CAMPUS	CAMPUS_VN	VPN 10
🟥	ISP_1	-	VPN 0
🟪	ISP_2	-	VPN 0
🟩	WAN	-	VPN 20

Figure 5.4 Primary Location Migration. Macro-Segmentation

Chapter 6
EVE-NG Simulation

This chapter presents a partial implementation of Cisco SD-WAN network, proposed earlier in redesign of CNX (4). EVE-NG simulation environment, as well as pre-launched and cross-authenticated SD-WAN Control Components were provided for this purpose. Although in the redesign SD-WAN was integrated with SD-Access, the latter solution will not be simulated due to inability of EVE-NG to do so [46].

The goal of the chapter is to demonstrate WAN Edge router (vEdge) on-boarding process and configuring SD-WAN network via vManage GUI. Computational and memory requirements of the lab, as well as its suitability for teaching purposes will also be described.

Use of EVE-NG environment to implement Cisco SD-WAN solution bases on a Bachelor thesis [47].

6.1 EVE-NG Description [48]

EVE-NG is one of the leading virtual environment tools for network, security and DevOps simulated scenarios. It provides great versatility and support for systems from various vendors. The complete feature set depends on the deployed license (community or professional). While the community edition focuses mainly on basic capabilities and provides sufficient environment for testing low-end scenarios, the professional edition is geared towards mid- to high-end scenarios where complex functionalities are necessary (such as integrated Wireshark, Docker container support and multi-lab privilege management).

Interaction with the virtual environment is enabled via CLI and web-based GUI. While CLI still remains the feature-first interface and supports all provided functionalities, the web-based interface is quickly catching up and nowadays should be sufficient for most every-day tasks. Apart from said interfaces and self-provided OS images, no other external application should be necessary.

While it is recommended to run EVE-NG as bare-metal hypervisor, launching it within a VM is possible, but may produce nested virtualization problems and degraded performance may be observed.

6.2 Description of the Provided Environment

6.2.1 EVE-NG Environment

Provided EVE-NG environment runs as a VM on the hypervisor, hosted by the supervisor of the thesis. It is allocated 96GiB RAM, 16 virtual CPUs (vCPUs) and 600GiB of storage.

6.2.2 Available Device Images

All device images, desired to be used in a simulation must be uploaded to EVE-NG VM or hypervisor [48]. SD-WAN specific images: Viptela's vManage, vSmart, vBond and vEdge were provided for use. Cisco's vIOS images for a L3 switch and router will be used for traditional switching/routing, and are configured using Cisco CLI. VPCS is a lightweight software with limited network functionality, that allows to configure IP, gateway and use ping/traceroute commands. It uses just 2MB of RAM and doesn't require an image to run [49]. VPCS nodes will act as Endpoints at sites.

All device images used in the laboratory can be retrieved in the attachments [47] or from original sources, also specified in [47].

6.2.3 Initial Provided Laboratory

All components mentioned in the paragraph are present in Fig. 6.1. The initial topology provided for expansion consisted of deployed and cross-authenticated Control Components: vManage, vSmart and vBond connected to the Internet (Cisco vIOS Switch). Internet is an abstraction, since internally it acts as a L2 switch with no additional configurations. The Internet accommodates a Gateway (Cisco vIOS Router), connected to the external real network (Management) via gi0/0 interface. Since the external network has real Internet connectivity, there exists a possibility to provide the laboratory with such. This would require configuring NAT on the Gateway or external device. However, this is not necessary for showcasing the desired SD-WAN functionality and is not implemented. Gateway acts as DHCP server for assigning IP addresses to vEdge's transport interfaces (ge0/0) and as DNS server, primarily to resolve vBond's domain name during vEdge on-boarding. Docker device is used primarily to access vManage GUI interface from web browser.

Additionally, a WAN Edge Whitelist Authorization File [50, p. 75-77] was uploaded to vManage, which distributed it to other Control Components. Cross-authentication of Control Components and of WAN Edge devices to Control Components is an integral part of SD-WAN operation. Authentication process is supported with the use of certificates.

Necessary steps to deploy the initial laboratory in EVE-NG are not covered in the thesis. In case these are desired to be replicated a Bachelor thesis [47] can be consulted. It served as a reference to deploy the laboratory.

6.3 Laboratory Topology Description

Fig. 6.1 demonstrates the final physical topology of SD-WAN laboratory. It will be attached to the thesis in compressed format. The goal of minimizing the computational and memory complexity was not set, thus, the infrastructure accommodates more components than required for a minimal laboratory [47].

6.3.1 Overlay

SD-WAN network consists of 4 sites. Site 1000 is dedicated to Control Components. Site 10 represents Primary Locations in the redesign, with 2 WAN Edge devices operating in Active/-Standby fashion and acting as a single logical hub. VRRP with TLOC priority is configured to

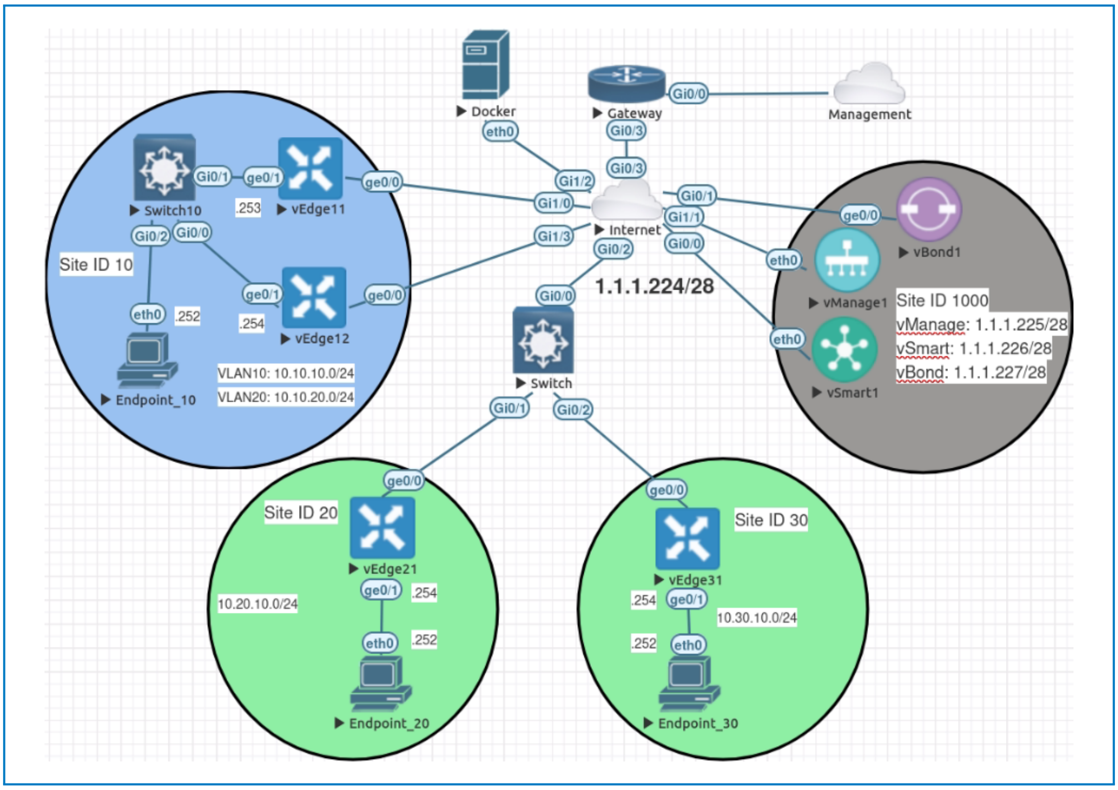

Figure 6.1 SD-WAN Lab Topology

force symmetric traffic through the primary node in VRRP group (6.4.4). Sites 20 and 30 represent Secondary Locations, acting as spokes in the hub-to-spoke Overlay topology. No on-demand tunneling is configured, thus, traffic between sites 20 and 30 traverses site 10. Overlay topology configuration is part of the centralized policy (6.4.5).

ge0/0 interfaces of all vEdges are part of VPN 0 and act as IPsec tunnel endpoints. They are connected to Internet transport of *default* color, the transport representing one of the ISPs in the redesign. Thus, each vEdge has 1 TLOC, which they advertise via OMP.

A total of 2 service VPNs is configured. VPN 10 represents CORPORATE_VN (VPN10) in the redesign. VPN 20 was firstly isolated from VPN 10 and represented GUEST_VN (VPN20). Later, route leaking between VPN 10 and 20 was configured (6.4.6), VPN 20 representing WAN VRF (VPN100) in the redesign. Site 10 is configured with both VPNs, sites 20 and 30 with VPN 10 only.

6.3.2 Underlay

SD-WAN Overlay is built on top of an Underlay, which typically refers to IP connectivity between devices. All vEdges and Control Components are interconnected via Internet, the Underlay transport subnet being 1.1.1.224/28. Control Components have fixed IP addresses as part of their initial configuration. Gateway is configured as DHCP server to assign IP addresses to vEdges' Internet-facing interfaces (gi0/0) during on-boarding process.

LAN-side IP subnets follow the pattern 10.<Site#>.<VPN#>.0/24. Numbers in the last octet of IPs assigned to devices are depicted in Fig. 6.1. On sites 20 and 30 only 1 subnet is used since each site will contain only 1 service VPN. On site 10 two subnets are used, 1 per service VPN. To use only 1 LAN-side interface on each vEdge, ge0/1 physical interface is placed in VPN

0, while its subinterfaces (ge0/1.10, ge0/1.20) in corresponding service VPNs. To distinguish between traffic dedicated to different VPNs each is mapped to a VLAN on Switch10. Links between Switch10 and vEdges are trunk links.

6.4 Laboratory Deployment

6.4.1 vEdge On-Boarding [47]

vEdges can be on-boarded using manual CLI configuration, bootstrap configuration or via ZTP process. ZTP is a preferred method, since it allows automating the on-boarding process. However, this requires deploying a ZTP server and preparing the infrastructure for ZTP process. For simplicity manual configuration is presented.

To on-board a vEdge, it is required to add a respective node to the topology (it will load from provided image) and interconnect it as desired. Then, minimal CLI configuration is entered, root certificate is installed and vEdge is activated. For detailed steps refer to [47]. After that VEdge should form BFD sessions to other vEdges and Control Connections to Control Components. It can be further configured via CLI or vManage templates.

Fig. 6.2 shows the main dashboard of vManage, which apart from other statistics displays vEdges deployed in the Overlay and their connectivity status. BFD sessions and Control Connections can be displayed directly in vEdge's CLI. Fig. 6.3 and Fig. 6.4 show required commands and their output after vEdge31 joined the Overlay. Other vEdges were already present and VRRP configured with vEdge11 as primary component.

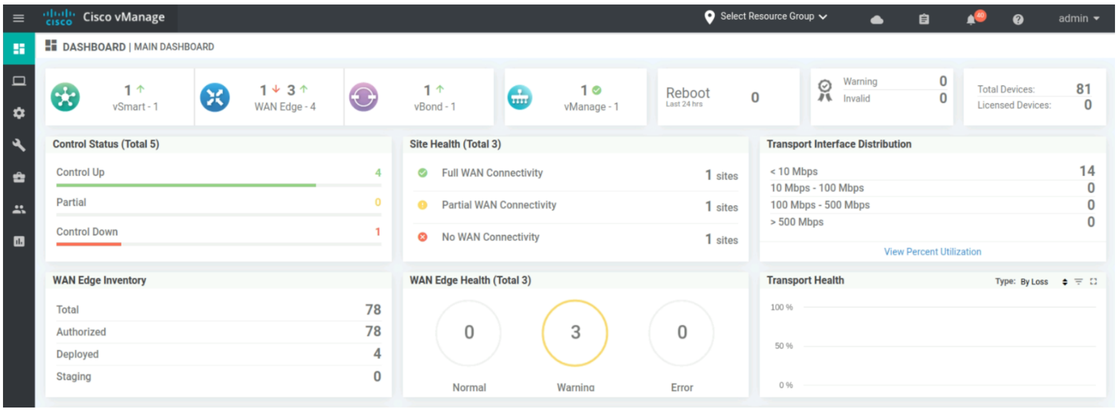

DASHBOARD | MAIN DASHBOARD

1 ↑ vSmart - 1

1 ↓ 3 ↑ WAN Edge - 4

1 ↑ vBond - 1

1 ✓ vManage - 1

Reboot Last 24 hrs 0

Warning 0
Invalid 0

Total Devices: 81
Licensed Devices: 0

Control Status (Total 5)

Control Up	4
Partial	0
Control Down	1

Site Health (Total 3)

✓ Full WAN Connectivity	1 sites
● Partial WAN Connectivity	1 sites
✗ No WAN Connectivity	1 sites

Transport Interface Distribution

< 10 Mbps	14
10 Mbps - 100 Mbps	0
100 Mbps - 500 Mbps	0
> 500 Mbps	0

View Percent Utilization

WAN Edge Inventory

Total	78
Authorized	78
Deployed	4
Staging	0

WAN Edge Health (Total 3)

Normal	Warning	Error
0	3	0

Transport Health

Figure 6.2 vManage Dashboard

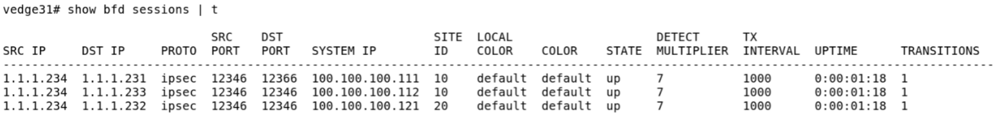

```
vedge31# show bfd sessions | t

                                   SRC    DST                         SITE  LOCAL                   DETECT      TX
SRC IP      DST IP      PROTO  PORT   PORT   SYSTEM IP        ID    COLOR    COLOR    STATE  MULTIPLIER  INTERVAL  UPTIME      TRANSITIONS
-------------------------------------------------------------------------------------------------------------------------------------------
1.1.1.234   1.1.1.231   ipsec  12346  12366  100.100.100.111  10    default  default  up     7           1000      0:00:01:18  1
1.1.1.234   1.1.1.233   ipsec  12346  12346  100.100.100.112  10    default  default  up     7           1000      0:00:01:18  1
1.1.1.234   1.1.1.232   ipsec  12346  12346  100.100.100.121  20    default  default  up     7           1000      0:00:01:18  1
```

Figure 6.3 vEdge31 On-Boarding. BFD Sessions

6.4.2 Device Management [26]

Cisco SD-WAN devices with the exception of vManage are managed via **CLI** or **Device Templates** in vManager GUI. Templates is a means of automating configuration process: they are defined in vManage, converted to CLI and pushed to devices. Templates offer a more secure

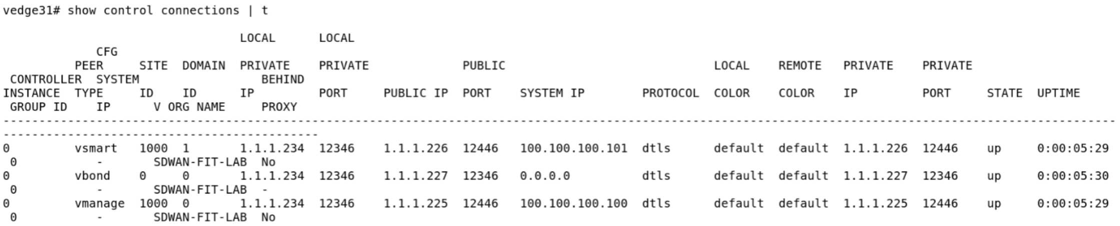

```
vedge31# show control connections | t

                                 LOCAL      LOCAL
             CFG
          PEER     SITE  DOMAIN  PRIVATE    PRIVATE             PUBLIC                            LOCAL    REMOTE   PRIVATE    PRIVATE
 CONTROLLER  SYSTEM                 BEHIND
INSTANCE  TYPE     ID    ID      IP         PORT     PUBLIC IP  PORT    SYSTEM IP        PROTOCOL  COLOR    COLOR    IP         PORT     STATE  UPTIME
 GROUP ID    IP       V ORG NAME     PROXY
------------------------------------------------------------------------------------------------------------------------------------------------------------
--------------------------------------------
0         vsmart   1000  1       1.1.1.234  12346    1.1.1.226  12446   100.100.100.101  dtls      default  default  1.1.1.226  12446    up     0:00:05:29
 0           -       SDWAN-FIT-LAB  No
0         vbond    0     0       1.1.1.234  12346    1.1.1.227  12346   0.0.0.0          dtls      default  default  1.1.1.227  12346    up     0:00:05:30
 0           -       SDWAN-FIT-LAB  -
0         vmanage  1000  0       1.1.1.234  12346    1.1.1.225  12446   100.100.100.100  dtls      default  default  1.1.1.225  12446    up     0:00:05:29
 0           -       SDWAN-FIT-LAB  No
```

Figure 6.4 vEdge31 On-Boarding. Control Connections

way of device management compared to CLI, since they are easier to setup and configuration is checked for validity before being pushed.

Device templates can be **CLI-based** or **Feature-based**. They are specific for each device model, but they can be shared between multiple devices. CLI-based templates resemble CLI configuration, but additionally allow defining variables and are subject to validity checks.

Feature templates are the most preferred configuration method due to their modular structure and GUI-based template creation process. Separate Feature templates exist for configuring interfaces, VPNs, routing protocols, etc. They are created by filling the required variable fields with desired values. Device templates are then built from Feature templates. However, in 6.4.4 it will be demonstrated that with some versions of vManage software Feature templates offer a less extensive set of configuration capabilities compared to CLI. In such cases they have to be converted to CLI templates to configure the required functionality.

Values of fields in templates can be **Global**, **Device-specific** and **Default**. Global values are defined when template is created and applied to every device attached to it. Device-specific values are left empty during template creation and are filled out during attachment process.

6.4.3 vEdge Basic Feature Template

For demonstration purposes a Device Feature template will be used to push full configuration to vEdges after on-boarding process. The initial configuration allows vEdges to join the Overlay, form control connection to Control Components and establish BFD sessions with other vEdges as shown in Fig. 6.3 and Fig. 6.4. Feature template will include the same configuration and additionally set up service-side VPNs.

Feature templates are configured under **Configuration** > **Templates** > **Feature**. Feature templates required to compose a Device template for vEdge11 were of the following categories:

1. **System**: Configures vEdge as a whole system, including **Site ID** and **System IP**. These were set as **Device-Specific** to allow applying the same template to any vEdge;

2. **VPN Interface Ethernet**: Template type used to configure all interfaces and subinterfaces of vEdge11, including **eth0** management interface. A template of this type was created per interface with **Interface Name** specified. **Shutdown** was set to **No** to all interfaces except **eth0**, because it will not be used;

   a. (**ge0/0**): **IPv4** address assignment was set as **Dynamic**. **Tunnel Interface** was set to **On** to enable BFD session forming. **Color** was set to **default**. **Allow Service** > **All** was set to **On** to match the CLI configuration;

   b. (**ge0/1**) **IPv4** address was set as **Static** and not specified since **ge0/1** will be a trunk link interface and will not require an address. Other settings are left to be default;

   c. (**ge0/1.10 and ge0/1.20**) **IPv4 Address** is set to be **Device-specific** to allow to attach the same template to more devices. Moreover, **MTU** was changed to **1496** to accommodate 802.1Q tag;

d. (**eth0**) Doesn't require an IP address.

3. **VPN**: Used to configure VPN 0, 10 and 20 with a template per VPN. These will be attached under corresponding **VPN Interface** templates when building Device policy. It is enough to just specify VPN numbers in template creation.

With all Feature templates prepared, a Device template can be built:

1. Under **Configuration > Templates** click **Create Template > From Feature Template**, specify device model;
2. Under **Basic Information** select created **System** template. Under **Transport & Management VPN** specify created templates in **VPN 0** and **VPN 512** fields. To them add templates for interfaces belonging to these VPNs. Under **Service VPN** add templates created for **VPN 10 and 20** and to them add corresponding subinterfaces templates.

When Device template is created it can be attached to devices by selecting it in the list and choosing **Attach Devices** option. Required device is then selected from the list and procedure of attachment is intuitively followed. Result of configuration validation, as well as success or failure of attachment will be reported by vManage.

The above procedure is repeated for vEdge12. vEdge21 and vEdge31 are configured only with VPN10 and to it interface **ge0/1** is attached, rather than its subinterfaces.

Fig 6.5 shows a list of Device templates and a way to attach devices. It can be seen to which templates devices are currently attached and how many of them. Fig 6.6 shows configuring one of the sections of **VPN Interface Ethernet** template for **ge0/0**. Fig 6.7 shows how to attach VPN templates, and how to assign interfaces to these VPN templates when creating Device templates.

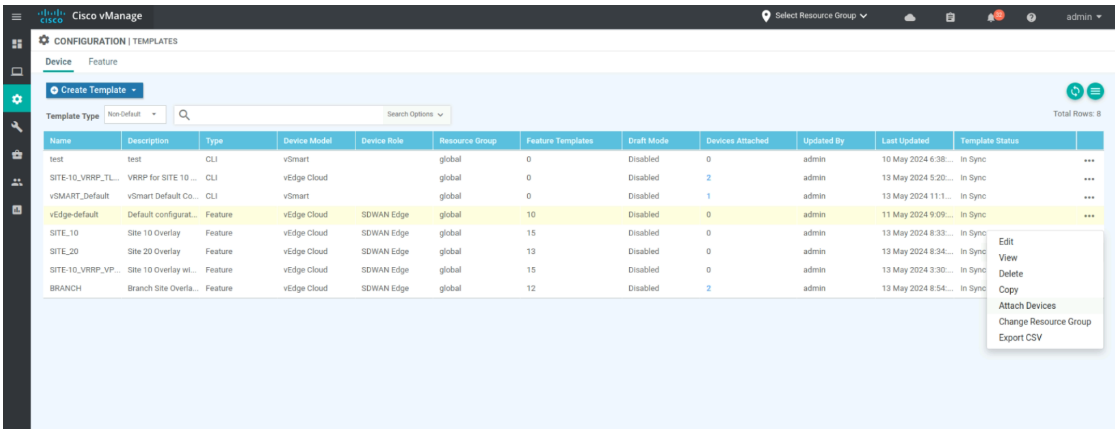

Name	Description	Type	Device Model	Device Role	Resource Group	Feature Templates	Draft Mode	Devices Attached	Updated By	Last Updated	Template Status
test	test	CLI	vSmart		global	0	Disabled	0	admin	10 May 2024 6:38:...	In Sync
SITE-10_VRRP_TL...	VRRP for SITE 10 ...	CLI	vEdge Cloud		global	0	Disabled	2	admin	13 May 2024 5:20:...	In Sync
vSMART_Default	vSmart Default Co...	CLI	vSmart		global	0	Disabled	1	admin	13 May 2024 11:1...	In Sync
vEdge-default	Default configurat...	Feature	vEdge Cloud	SDWAN Edge	global	10	Disabled	0	admin	11 May 2024 9:09:...	In Sync
SITE_10	Site 10 Overlay	Feature	vEdge Cloud	SDWAN Edge	global	15	Disabled	0	admin	13 May 2024 8:33:...	In Sync
SITE_20	Site 20 Overlay	Feature	vEdge Cloud	SDWAN Edge	global	13	Disabled	0	admin	13 May 2024 8:34:...	In Sync
SITE-10_VRRP_VP...	Site 10 Overlay wi...	Feature	vEdge Cloud	SDWAN Edge	global	15	Disabled	0	admin	13 May 2024 3:30:...	In Sync
BRANCH	Branch Site Overla...	Feature	vEdge Cloud	SDWAN Edge	global	12	Disabled	2	admin	13 May 2024 8:54:...	In Sync

Figure 6.5 Device Templates

6.4.4 VRRP

VRRP is configured on vEdges that are part of VRRP group: vEdge11 and vEdge21. VRRP protocol runs per service VPN and has to be enabled on interfaces belonging to these VPNs. For demonstration purposes only feature template for **ge0/1.10** will be modified to include VRRP configuration. Template is edited and under **VRRP** a **New VRRP** is created. **Priority** is set

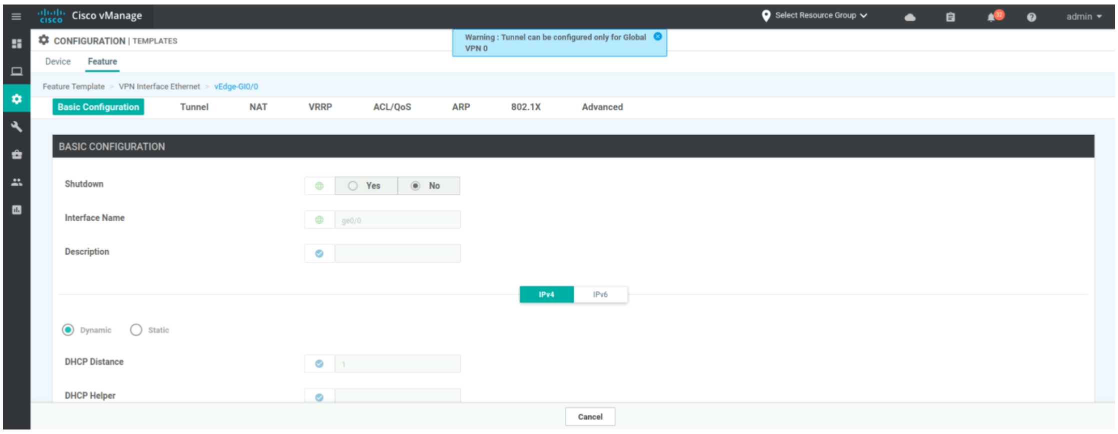

Warning : Tunnel can be configured only for Global VPN 0

BASIC CONFIGURATION

Shutdown ○ Yes ● No

Interface Name ge0/0

Description

● Dynamic ○ Static

DHCP Distance 1

DHCP Helper

Figure 6.6 VPN Interface Ethernet Feature Configuration

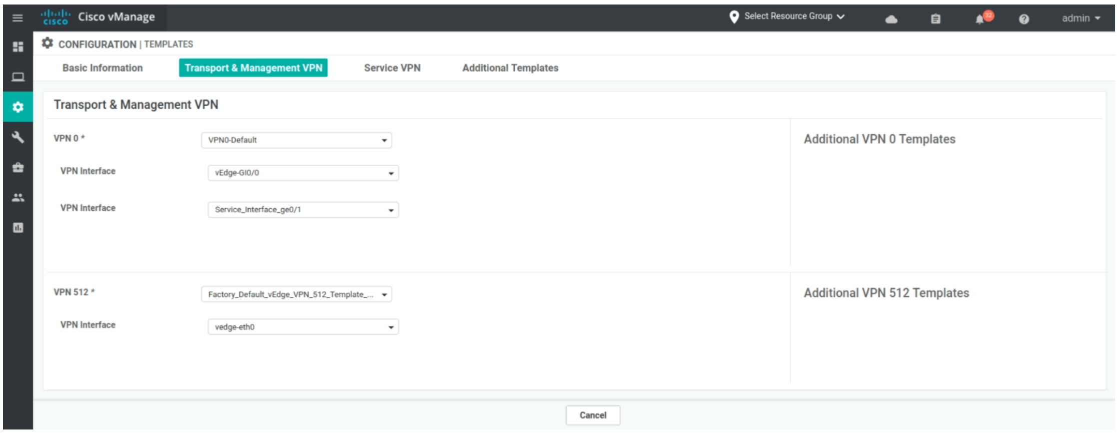
Cisco vManage
Select Resource Group
admin
CONFIGURATION | TEMPLATES
Basic Information
Transport & Management VPN
Service VPN
Additional Templates
Additional VPN 0 Templates
Additional VPN 512 Templates
Cancel

Transport & Management VPN

VPN 0 * VPN0-Default

VPN Interface vEdge-GI0/0

VPN Interface Service_Interface_ge0/1

VPN 512 * Factory_Default_vEdge_VPN_512_Template_...

VPN Interface vedge-eth0

Figure 6.7 Building Device Template

to be **Device specific**, **Global Group ID** is **10** and **IP Address** is **10.10.10.2** (virtual IP of VRRP group).

When the feature template is updated, vManager immediately prompts the user to enter new values into affected Device template. Priority of vEdge11 is set to 200, and of vEdge21 to 100, so that vEdge11 becomes a primary component. Updated Device templates are pushed to devices again. Default gateway of devices on LAN side must point to virtual IP.

The configuration above forced site 10 outbound traffic to be directed to vEdge11. However, for traffic to be symmetric and flow through the same node in the inbound direction TLOC preference automatic change has to be configured. With the current vManage software version this is not possible to do via feature templates. Device template for vEdges at site 10 has to be converted to CLI template, where a line *tloc-change-pref* is added under VRRP configuration. This causes primary components to always have TLOC preference 1 and backup components 0, even after fail-over. TLOC preferences are advertised via OMP and force other vEdges to prefer current primary VRRP group member.

Fig 6.8 shows the result of template application: vEdge12 is a Backup node. Fig 6.9 shows OMP routes on vEdge21 after VRRP was configured. There are 2 routes to the same prefix in site 10 VPN 10. It can be seen that the first route to vEdge11 is preferred, judging from the flag **C (Chosen)**. If vEdge11 fails, which can be caused by *reboot* command, vEdge12 becomes a Primary node (Fig 6.10). Observe, how **TLOC Real Preference value** changes from 0 to 1.

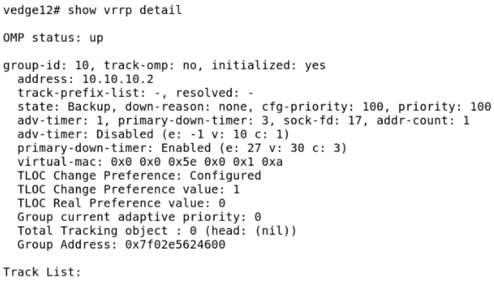

```
vedge12# show vrrp detail

OMP status: up

group-id: 10, track-omp: no, initialized: yes
  address: 10.10.10.2
  track-prefix-list: -, resolved: -
  state: Backup, down-reason: none, cfg-priority: 100, priority: 100
  adv-timer: 1, primary-down-timer: 3, sock-fd: 17, addr-count: 1
  adv-timer: Disabled (e: -1 v: 10 c: 1)
  primary-down-timer: Enabled (e: 27 v: 30 c: 3)
  virtual-mac: 0x0 0x0 0x5e 0x0 0x1 0xa
  TLOC Change Preference: Configured
  TLOC Change Preference value: 1
  TLOC Real Preference value: 0
  Group current adaptive priority: 0
  Total Tracking object : 0 (head: (nil))
  Group Address: 0x7f02e5624600

Track List:
```

Figure 6.8 vEdge12 as a Backup VRRP Node

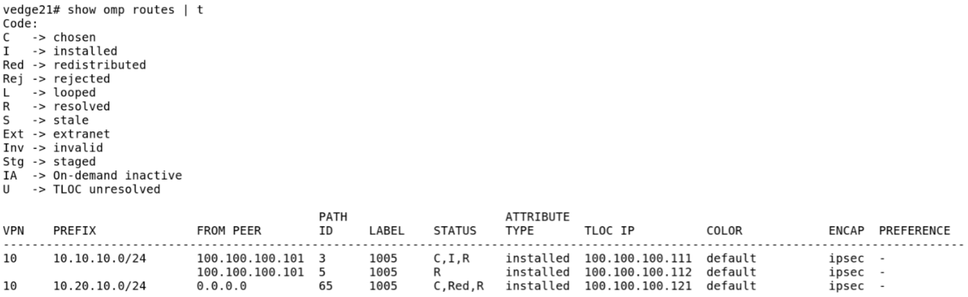

```
vedge21# show omp routes | t
Code:
C   -> chosen
I   -> installed
Red -> redistributed
Rej -> rejected
L   -> looped
R   -> resolved
S   -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA  -> On-demand inactive
U   -> TLOC unresolved

                                         PATH                      ATTRIBUTE
VPN    PREFIX              FROM PEER     ID     LABEL    STATUS    TYPE       TLOC IP          COLOR            ENCAP  PREFERENCE
---------------------------------------------------------------------------------------------------------------------------------
10     10.10.10.0/24       100.100.100.101 3    1005     C,I,R     installed  100.100.100.111  default          ipsec  -
                           100.100.100.101 5    1005     R         installed  100.100.100.112  default          ipsec  -
10     10.20.10.0/24       0.0.0.0       65     1005     C,Red,R   installed  100.100.100.121  default          ipsec  -
```

Figure 6.9 OMP Routes on vEdge21 After Final VRRP Configuration

6.4.5 Hub-to-Spoke Topology

Overlay topology is determined by what OMP routes and TLOCs are advertised to sites. Hub-to-Spoke topology will be enforced by applying 2 rules:

1. A spoke's TLOCs will not be shared to other spoke sites for BFD sessions not to form.
2. Each spoke site will only receive OMP routes with hub site's TLOCs as next hop. This way routes to other spoke sites will also be shared, but communication through hub will be enforced.

In Cisco SD-WAN Overlay topology policies are part of centralized policy. Centralized policies are configured under **Configuration > Policies > Centralized Policy** (Fig. 6.11). Since required topology is determined by OMP TLOC and route updates it will be configured as **Custom Control (Route & TLOC)** under **Custom Options > Topology > Custom Topology**.

Centralized policies typically instruct vSmart to perform certain actions. Different policies, which are part of centralized policies, are applied in either **Outbound** or **Inbound** direction from the perspective of vSmart. Direction determines whether actions are taken on the incoming or outgoing packets. The topology policy will be applied to sites in **SPOKE** site list (20, 30) in the **Outbound** direction. Thus, vSmart will be instructed to filter out certain OMP advertisements sent to sites 20 and 30.

It is important to note that vSmart must be configured from templates if centralized control policy is desired to be pushed to it.

Hub-to-Spoke Topology configuration was performed in the following steps:

1. A **Custom Control (Route & TLOC)** policy was created, followed by creating 2 rule-sets of **Sequence Type Route** and **TLOC**. These influence OMP route and TLOC advertisements;

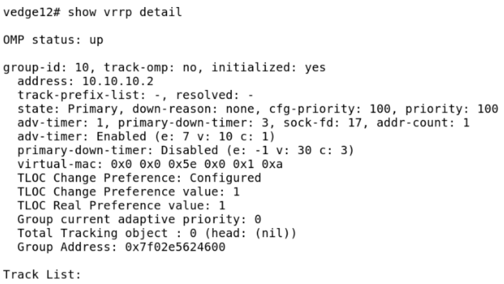

```
vedge12# show vrrp detail

OMP status: up

group-id: 10, track-omp: no, initialized: yes
  address: 10.10.10.2
  track-prefix-list: -, resolved: -
  state: Primary, down-reason: none, cfg-priority: 100, priority: 100
  adv-timer: 1, primary-down-timer: 3, sock-fd: 17, addr-count: 1
  adv-timer: Enabled (e: 7 v: 10 c: 1)
  primary-down-timer: Disabled (e: -1 v: 30 c: 3)
  virtual-mac: 0x0 0x0 0x5e 0x0 0x1 0xa
  TLOC Change Preference: Configured
  TLOC Change Preference value: 1
  TLOC Real Preference value: 1
  Group current adaptive priority: 0
  Total Tracking object : 0 (head: (nil))
  Group Address: 0x7f02e5624600

Track List:
```

Figure 6.10 vEdge12 After Fail-Over

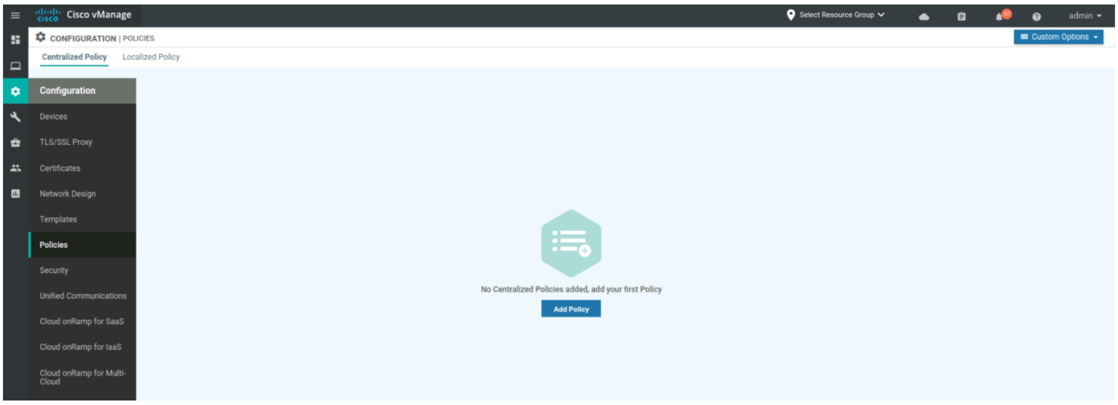

No Centralized Policies added, add your first Policy

Figure 6.11 Centralized Policy Creation

2. A **Sequence Rule** is added under each of them;

   a. **TLOC** policy rule instructs vSmart to accept TLOC advertisements only from sites in **HUB** site list (10), see Fig 6.12. All other TLOC advertisements are dropped by default;

   b. **Route** policy rule instructs vSmart to only accept OMP routes coming from **HUB_TLOCs** list, containing TLOCs of vEdge11 and vEdge12. All incoming routes are checked for specified TLOCs.

3. A new centralized policy is created

   a. Under **Topology** the newly created topology is added;

   b. Under **Policy Application** the newly created topology policy is applied in the **Outbound** direction to **SPOKE** site list.

It is now possible to see that at spoke vEdges routes to other spoke vEdges are pointing to TLOCs of HUB vEdges. Correctness of configuration is verified with the **traceroute** command. Fig 6.13 and Fig 6.14 illustrate the tests performed in vEdge31 CLI.

6.4.6 VPN Route Leaking

As a last step route leaking between VPN 10 and 20 is enabled. It is configured in a way similar to Hub-to-spoke topology. It will be applied in the **Inbound** direction to both **HUB** and **SPOKE** site lists. Thus, vSmart will apply rules to OMP packets received from vEdges in specified sites.

VPN Route Leaking configuration was performed in the following steps:

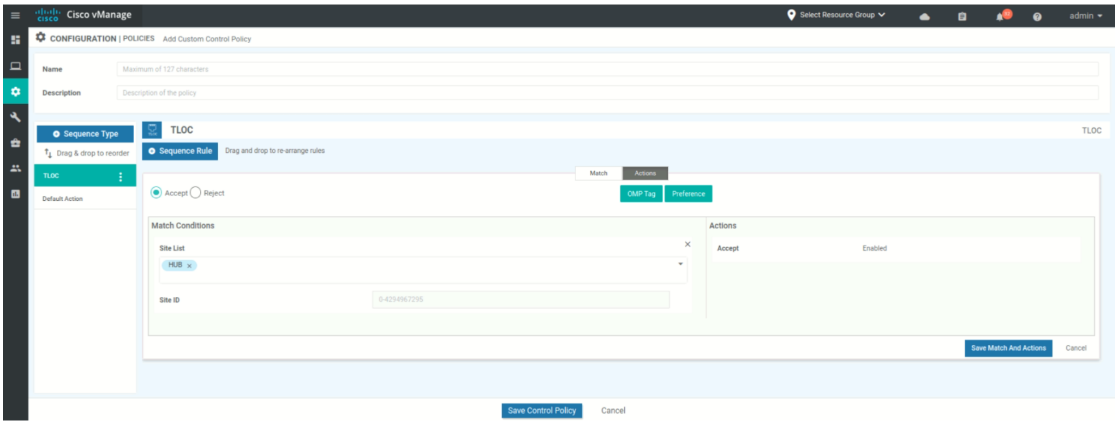

● Accept ○ Reject

Site List: HUB

Accept Enabled

Figure 6.12 Setting Up a TLOC Rule

```
vedge31# show omp routes | t
Code:
C   -> chosen
I   -> installed
Red -> redistributed
Rej -> rejected
L   -> looped
R   -> resolved
S   -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA  -> On-demand inactive
U   -> TLOC unresolved

                                           PATH                      ATTRIBUTE
VPN    PREFIX              FROM PEER       ID     LABEL    STATUS    TYPE       TLOC IP          COLOR            ENCAP  PREFERENCE
-----------------------------------------------------------------------------------------------------------------------------------
10     10.10.10.0/24       100.100.100.101 6      1005     C,I,R     installed  100.100.100.111  default          ipsec  -
                           100.100.100.101 7      1005     R         installed  100.100.100.112  default          ipsec  -
10     10.20.10.0/24       100.100.100.101 8      1005     C,I,R     installed  100.100.100.111  default          ipsec  -
                           100.100.100.101 9      1005     R         installed  100.100.100.112  default          ipsec  -
10     10.30.10.0/24       0.0.0.0         65     1005     C,Red,R   installed  100.100.100.131  default          ipsec  -
```

Figure 6.13 OMP Routes on Site 30

1. Two **Sequence Rules** of type **Route** are created under newly created **Custom Control (Route & TLOC)** policy and create a **Route** sequence. Rules instruct vSmart to export routes from VPN 10 in VPN 20 and vice versa. It is important to set the **Default Action** to **Accept** not to drop advertisements from other VPNs. See final configuration in Fig. 6.15;

2. Under **Configuration > Policies** the currently active centralized control policy is copied and edited. This is due to the fact that a single policy of this type can be active at a time and the currently active policy can't be modified;

   a. Under **Topology** the newly created topology is added alongside the Hub-to-Spoke one;

   b. Under **Policy Application** the newly created topology policy is applied in the **Inbound** direction to both **HUB** and **SPOKE** site lists (Fig 6.16).

   c. The new centralized policy is applied.

It is now possible to see that among VPN 10 routes there are VPN 20 prefixes. Correctness of configuration is verified with the ping command. Fig 6.17 illustrates the tests performed in vEdge31 CLI.

6.5 Computational and Memory Requirements

Table 6.1 provides resource requirements per device image, as well as the number of images used in the simulation. It can be concluded that theoretically the laboratory requires 29/96 GiB of

```
vedge31# traceroute vpn 10 10.20.10.254
Traceroute  10.20.10.254 in VPN 10
traceroute to 10.20.10.254 (10.20.10.254), 30 hops max, 60 byte packets
 1  10.10.10.253 (10.10.10.253)  44.773 ms  65.878 ms  65.912 ms
 2  10.20.10.254 (10.20.10.254)  141.709 ms  141.489 ms  158.970 ms
```

Figure 6.14 Traceroute From Site 30 to Site 20

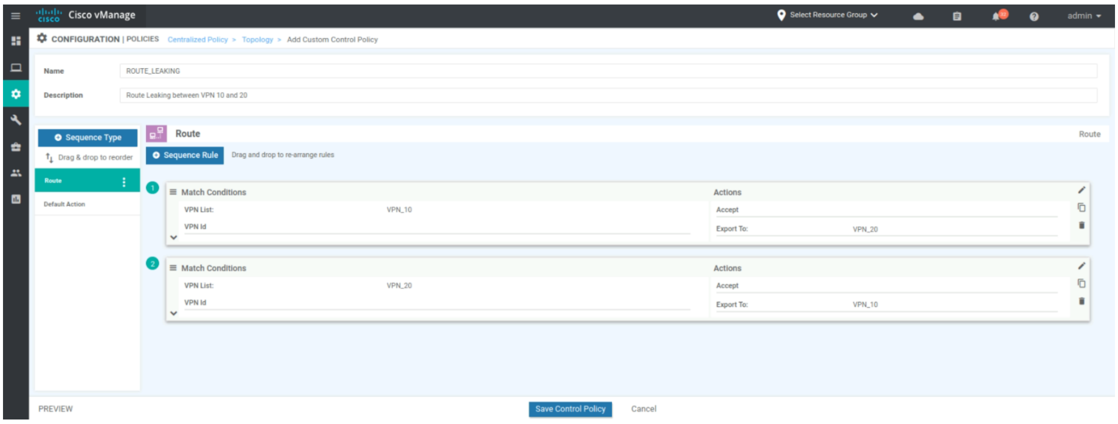

Name: ROUTE_LEAKING

Description: Route Leaking between VPN 10 and 20

1. Match Conditions — VPN List: VPN_10; VPN Id. Actions — Accept; Export To: VPN_20
2. Match Conditions — VPN List: VPN_20; VPN Id. Actions — Accept; Export To: VPN_10

Figure 6.15 Route Leaking Rules

RAM and 17/16 vCPUs. Although vCPU usage is pushed above the limit, this only affects user experience by slowing down the environment. Thus, the laboratory can be replicated and used for demonstration purposes, since the disadvantage of degraded experience is offset by possibility to test more advanced SD-WAN capabilities.

6.6 Suitability for Teaching [48]

From the perspective of simulation capabilities EVE-NG is one of the few tools that allow to implement and test Cisco SD-WAN deployments. It can be used to demonstrate both SD-WAN components on-boarding, as well as perform basic and advanced Overlay configuration in either vManager's GUI or CLI. Thus, it is suitable for demonstrating SDN concepts in networking laboratories.

On the other hand, EVE-NG licensing has to be considered. Both community and professional licences are suitable for simulating Cisco SD-WAN labs, with community edition being free while professional edition offering a more extensive feature set. However, community license offers simultaneous access to the environment to one user only, basic professional – to 2 users. For

Table 6.1 SD-WAN Lab Resource Requirements

Device	Count	EVE-NG Image Name	RAM (GiB)	vCPU	Ethernet
Viptela vManage	1	vtmgmt-20.5.1	16	4	2
Viptela vSmart	1	vtsmart-20.5.1	2	2	2
Viptela vBond	1	vtbond-20.5.1	1	1	2
Viptela vEdge	4	vtedge-20.5.1	1	1	4
Cisco vIOS Switch	3	viosl2-adverterprisek9-m	1	1	8
Cisco vIOS Router	1	vios-adverterprisek9-m	1	1	4
Docker.io	1	-	2	2	1
Virtual PC (VPCS)	3	-	-	-	-
Total	15		29	17	

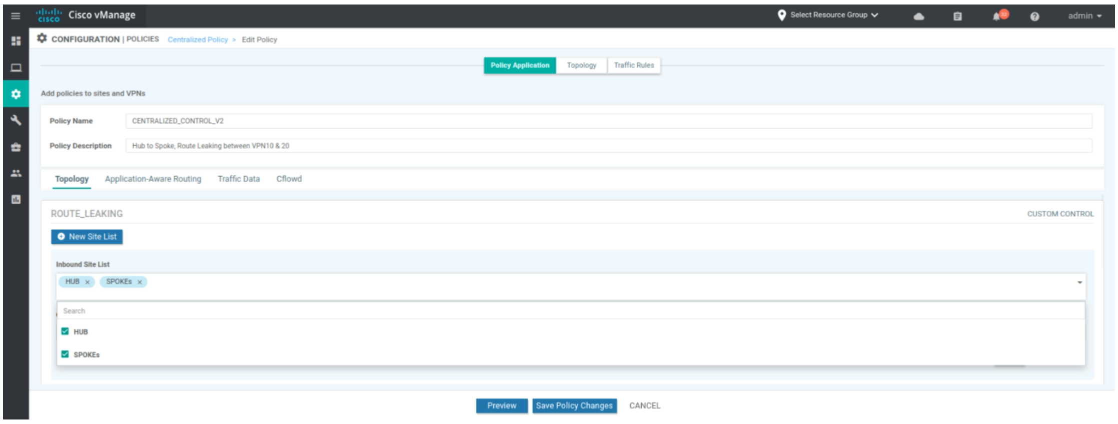

Cisco vManage

Select Resource Group

admin

CONFIGURATION | POLICIES Centralized Policy > Edit Policy

Add policies to sites and VPNs

Policy Name: CENTRALIZED_CONTROL_V2

Policy Description: Hub to Spoke, Route Leaking between VPN10 & 20

ROUTE_LEAKING

Inbound Site List: HUB × SPOKEs ×

☑ HUB

☑ SPOKEs

Figure 6.16 Applying Policy in Inbound Direction

```
vedge31# show omp routes | t
Code:
C   -> chosen
I   -> installed
Red -> redistributed
Rej -> rejected
L   -> looped
R   -> resolved
S   -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA  -> On-demand inactive
U   -> TLOC unresolved

                                            PATH                      ATTRIBUTE
VPN    PREFIX              FROM PEER        ID     LABEL    STATUS    TYPE       TLOC IP          COLOR            ENCAP  PREFERENCE
--------------------------------------------------------------------------------------------------------------------------------------
10     10.10.10.0/24       100.100.100.101  26     1005     C,I,R     installed  100.100.100.111  default          ipsec  -
                           100.100.100.101  27     1005     R         installed  100.100.100.112  default          ipsec  -
10     10.10.20.0/24       100.100.100.101  22     1005     C,I,R     installed  100.100.100.111  default          ipsec  -
                           100.100.100.101  23     1005     R         installed  100.100.100.112  default          ipsec  -
10     10.20.10.0/24       100.100.100.101  19     1005     C,I,R     installed  100.100.100.111  default          ipsec  -
                           100.100.100.101  20     1005     R         installed  100.100.100.112  default          ipsec  -
10     10.30.10.0/24       0.0.0.0          65     1005     C,Red,R   installed  100.100.100.131  default          ipsec  -

vedge31# ping vpn 10 10.10.20.254
Ping in VPN 10
PING 10.10.20.254 (10.10.20.254) 56(84) bytes of data.
64 bytes from 10.10.20.254: icmp seq=1 ttl=64 time=78.4 ms
64 bytes from 10.10.20.254: icmp seq=2 ttl=64 time=72.3 ms
64 bytes from 10.10.20.254: icmp seq=3 ttl=64 time=67.9 ms
64 bytes from 10.10.20.254: icmp seq=4 ttl=64 time=69.1 ms
```

Figure 6.17 Pinging VPN 20 Prefixes From VPN 10

more users to access the simulator, a licence per user must be purchased in addition to basic professional edition. Thus, it is possible to host a free EVE-NG deployment if single simultaneous user access and limited feature set is acceptable. Otherwise, paid licences must be acquired.

6.7 Laboratory Replication

Attached to the thesis there is a compressed EVE-NG laboratory file, as well as configuration of selected devices. VPCS configurations are not presented since they are trivial. If a device configuration is missing in the attachment, default configuration can be used. A way to retrieve images for the devices is described in [47].

Laboratory file can be exported to EVE-NG environment. Configurations of all devices except the SD-WAN specific ones can be exported as well. Viptela devices must be configured from scratch. Guidelines in [47] can be followed.

Chapter 7
Conclusion

This Bachelor thesis is a highly theoretical work, mostly focused on how Cisco solutions and SDN solutions in general can improve the world of traditional networking. The thesis doesn't include any mention of SDN implementation in real scenarios, its effectiveness, cost and complexity. Thus, no conclusions of SDN advantage over traditional networking must be made without consulting additional resources, dedicated to practical aspects of it.

This work dived deeply into Cisco SD-Access and Catalyst SD-WAN solutions, aiming at demonstrating how one of the vendors tried to address challenges of traditional networking and simplify networking operations. It is obvious, that the presented solutions have a high potential in the future of networking, because they make automation, programmability and security integral parts of networking. However, it is important not to overlook solutions presented by other vendors.

An example of an existing large enterprise was taken to propose migration to Cisco solutions to showcase how real networks can benefit from SDN. While it is possible to create large, complex theoretical Software-Defined deployments, this is usually cost-effective and requires network specialists to have the knowledge of implemented concepts.

Simulation of Cisco SD-WAN in EVE-NG was presented to demonstrate how simple it can be to configure a network. Although the deployed laboratory is small and cannot be compared to real-life networks, performed steps in configuring networking operations should not differ much. Nevertheless, the provided simulation can be replicated in teaching environment to showcase SDN concepts.

Thus, it can be concluded that SDN has a potential to replace traditional networking, however, costs and complexity of migration, as well as frequent lack of solution testing in real-life scenarios have to be considered.

Bibliography

1. *What is Computer Networking? | Cisco* [https://www.cisco.com/c/en/us/solutions/enterprise-networks/what-is-computer-networking.html]. 2024. Accessed: 2024-05-11.
2. *Network Firewall vs Network Switch vs Network Router | FS Community* [https://community.fs.com/article/network-switch-router-firewall-why-need-all-three.html]. 2020. Accessed: 2024-05-02.
3. *What is the OSI Model | AWS* [https://aws.amazon.com/what-is/osi-model/]. 2024. Accessed: 2024-05-02.
4. *What is the OSI Model? | CloudFlare* [https://www.cloudflare.com/learning/ddos/glossary/open-systems-interconnection-model-osi/]. 2024. Accessed: 2024-05-11.
5. *Network interface controller | Wikipedia* [https://en.wikipedia.org/wiki/Network_interface_controller]. 2024. Accessed: 2024-05-02.
6. *Communication protocol | Wikipedia* [https://en.wikipedia.org/wiki/Communication_protocol]. 2024. Accessed: 2024-05-02.
7. *What Is Network Policy? | Cisco* [https://www.cisco.com/c/en/us/solutions/enterprise-networks/what-is-network-policy.html]. 2024. Accessed: 2024-05-02.
8. HILL, Craig et al. *Cisco Software-Defined Access. Enabling intent-based networking.* 2nd ed. San Jose, CA: Cisco Systems, Inc., 2019.
9. *SDN vs Traditional Networking: Which Leads the Way?* [https://www.chinacablesbuy.com/sdn-vs-traditional-networking-which-leads-the-way.html]. 2018. Accessed: 2024-05-11.
10. FEAMSTER, Nick; REXFORD, Jennifer; ZEGURA, Ellen. The Road to SDN: An Intellectual History of Programmable Networks. *ACM SIGCOMM Computer Communication Review.* 2014, vol. 44, pp. 87–98. Available from DOI: 10.1145/2602204.2602219.
11. ABUELENAIN, Khaled; KARNELIUK, Anton; DOYLE, Jeff; JAIN, Vinit. *Network Programmability and Automation Fundamentals.* 1st ed. Hoboken, NJ: Cisco Press, 2021. Networking Technology.
12. EDELMAN, Jason; LOWE, Scott; OSWALT, Matt. *Network Programmability and Automation.* 1st ed. Sebastopol, CA: O'Reilly Media, Inc., 2018.
13. *Kubernetes* [https://kubernetes.io/]. 2024. Accessed: 2024-05-02.
14. *Apache Airflow, Release 2.9.0.* The Apache Software Foundation, 2024. Available at https://airflow.apache.org/docs/apache-airflow/2.9.0/index.html. Accessed: 2024-05-02.

15. *Software-Defined Networking: The New Norm for Networks.* ONF, 2012. Available at https://opennetworking.org/wp-content/uploads/2011/09/wp-sdn-newnorm.pdf. Accessed: 2024-05-07.

16. JAIN, Sushant et al. B4: Experience with a Globally-Deployed Software Defined WAN. In: 2013, vol. 43, pp. 3–14. Available from DOI: 10.1145/2486001.2486019.

17. *What Is a Network Controller? | Cisco* [https://www.cisco.com/c/en/us/solutions/enterprise-networks/what-is-a-network-controller.html]. 2024. Accessed: 2024-05-02.

18. *ONOS* [https://opennetworking.org/onos/]. 2024. Accessed: 2024-05-02.

19. *OpenDaylight* [https://www.opendaylight.org/]. 2024. Accessed: 2024-05-02.

20. *Cisco Nexus 9000 Series NX-OS Programmability Guide, Release 6.x.* NX-API. San Jose, CA: Cisco Systems, Inc., 2019. Available at https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/programmability/guide/b_Cisco_Nexus_9000_Series_NX-OS_Programmability_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Programmability_Configuration_Guide_chapter_0101.pdf. Accessed: 2024-05-02.

21. *Arista eAPI.* Santa Clara, CA: Arista Networks, Inc., 2023. Available at https://www.arista.com/assets/data/pdf/Whitepapers/Arista7500RSwitchArchitectureWP.pdf. Accessed: 2024-05-02.

22. *Cisco DNA Center 2.3.5 Data Sheet.* System and platform capabilities. San Jose, CA: Cisco Systems, Inc., 2023. Available at https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/dna-center/nb-06-dna-center-data-sheet-cte-en.html. Accessed: 2024-05-02.

23. *Cisco Catalyst SD-WAN Getting Started Guide.* San Jose, CA: Cisco Systems, Inc., 2024. Available at https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-gs-book/appendix-vmanage-how-tos.html. Accessed: 2024-05-02.

24. ILIESIU, Adrian. Cisco DNA Center Has a New Name and New Features [https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html]. 2023.

25. *Cisco Completes Acquisition of Viptela* [https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2017/m08/cisco-completes-acquisition-of-viptela.html]. 2017. Accessed: 2024-05-11.

26. *Design Zone for Branch/WAN - Cisco SD-WAN Design Guide | Cisco* [https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-design-guide.html]. 2023. Accessed: 2024-05-04.

27. *Cisco SD-WAN Migration Guide.* Cisco Public, 2019. Available at https://www.cisco.com/c/dam/en/us/td/docs/routers/sdwan/migration-guide/cisco-sd-wan-migration-guide.pdf. Accessed: 2024-05-07.

28. *Cisco SD-Access Solution Design Guide (CVD) | Cisco* [https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html]. 2024. Accessed: 2024-05-04.

29. AL-SHAWI, Marwan. Enterprise Campus Architecture Design. In: *CCDE Study Guide.* 1st ed. Hoboken, NJ: Cisco Press, 2015, pp. 123–124.

30. *SD-Access Design. Important Design & Migration Principles.* Cisco Public, 2021. Available at https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/R6BGArNQ/TECCRS-2812.pdf. Accessed: 2024-05-10.

31. *SD-Access | SD-WAN Independent Domain Guide Deployment Guide.* Cisco Public, 2021. Available at https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Campus/Cisco-SD-Access-SD-WAN-Independent-Domain-Guide.pdf. Accessed: 2024-05-04.

32. *SD-Access | SD-WAN Integrated Domain Guide Deployment Guide.* Cisco Public, 2021. Available at https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Campus/Cisco-SD-Access-SD-WAN-Integrated-Domain-Guide.pdf. Accessed: 2024-05-04.

33. *Products - End-of-Life Policy | Cisco* [https://www.cisco.com/c/en/us/products/eos-eol-policy.html]. 2024. Accessed: 2024-05-04.

34. *Cisco Catalyst Center - Install and Upgrade Guides | Cisco* [https://www.cisco.com/c/en/us/support/cloud-systems-management/dna-center/products-installation-guides-list.html]. 2024. Accessed: 2024-05-04.

35. *Performance and Scalability Guide for Cisco Identity Services Engine.* Cisco Systems, Inc., 2024. Available at https://www.cisco.com/c/en/us/td/docs/security/ise/performance_and_scalability/b_ise_perf_and_scale.html. Accessed: 2024-05-04.

36. *Cisco DNA Center & ISE Management Infrastructure Deployment Guide.* Process 2: Installing Cisco Identity Services Engine nodes. Cisco Systems, Inc., 2020. Available at https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-dnac-ise-deploy-guide.html. Accessed: 2024-05-04.

37. *SD-Access Wireless Design and Deployment Guide, Cisco DNA Center 2.1.1.* High availability in SD-Access Wireless. Cisco Systems, Inc., 2023. Available at https://www.cisco.com/c/dam/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/deploy-guide/cisco-dna-center-sd-access-wl-dg.pdf. Accessed: 2024-05-04.

38. *Cisco ISE configuration for onboarding hosts in Cisco SD-Access | Cisco Community* [https://community.cisco.com/t5/networking-knowledge-base/cisco-ise-configuration-for-onboarding-hosts-in-cisco-sd-access/ta-p/4106696]. 2020. Accessed: 2024-05-15.

39. *ISE Secure Wired Access Prescriptive Deployment Guide | Cisco Community* [https://community.cisco.com/t5/security-knowledge-base/ise-secure-wired-access-prescriptive-deployment-guide/ta-p/3641515]. 2023. Accessed: 2024-05-15.

40. *Cisco Catalyst SD-WAN Control Components Compatibility Matrix and Recommended Computing Resources.* Recommended Computing Resources for Cisco Catalyst SD-WAN Control Components Release 20.13.x. Cisco Systems, Inc., 2024. Available at https://www.cisco.com/c/en/us/td/docs/routers/sdwan/release/notes/compatibility-and-server-recommendations/ch-server-recs-20-13-combined.html. Accessed: 2024-04-25.

41. *Cisco Catalyst SD-WAN Device Compatibility - Cisco* [https://www.cisco.com/c/en/us/td/docs/routers/sdwan/release/compatibility/sdwan-device-compatibility.html]. 2024. Accessed: 2024-05-10.

42. *Cisco Catalyst SD-WAN Control Components Compatibility Matrix and Recommended Computing Resources - Compatibility Matrix [Cisco SD-WAN] - Cisco* [https://www.cisco.com/c/en/us/td/docs/routers/sdwan/release/notes/compatibility-and-server-recommendations/comp-matrix-intro-chapter-map.html]. 2024. Accessed: 2024-05-10.

43. *Cisco Software-Defined Access Compatibility Matrix* [https://www.cisco.com/c/dam/en/us/td/docs/Website/enterprise/sda_compatibility_matrix/index.html]. 2024. Accessed: 2024-05-10.

44. BOWMAN, Jeremy. *1 to 100 Master All Steps of Deployment, Integration, and Migration of Large SDA and SD-WAN Networks.* Cisco Public, 2024. Available at https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/BRKENS-3834.pdf. Accessed: 2024-05-14.

45. SANDEEP, Joseph. *Cisco SD-Access. Connecting the Fabric to External Networks.* Cisco Public, 2020. Available at https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/BRKCRS-2811.pdf. Accessed: 2024-05-10.

46. *Cisco SD-Access home lab* [https://notes.networklessons.com/cisco-sd-access-home-lab]. 2024. Accessed: 2024-05-15.

47. LANČA, Matěj. *Analysis and implementation of simulated environment for software defined networks*. Czech Technical University in Prague, Faculty of Information Technology, 2022. Available also from: http://hdl.handle.net/10467/101781. Bachelor's thesis. Accessed: 2024-05-16.

48. *EVE-NG Documentation* [https://www.eve-ng.net/index.php/documentation/]. 2024. Accessed: 2024-05-15.

49. *VPCS* [https://docs.gns3.com/docs/emulators/vpcs/]. 2024. Accessed: 2024-05-16.

50. *Cisco SD-WAN: WAN Edge Onboarding Prescriptive Deployment Guide*. Cisco Systems, Inc., 2022. Available at https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/SDWAN/sdwan-wan-edge-onboarding-deploy-guide-2020nov.pdf. Accessed: 2024-05-16.

Contents of the attachment

```
readme.txt ......................................... brief overview of attachment contents
src
  eve-ng-laboratory.zip ................................. exported EVE-NG laboratory
  thesis-src.zip ................................... source code of the thesis in LaTeX
  running-config.zip ...................... running configuration of laboratory devices
text
  thesis.pdf ........................................... text of the thesis in PDF format
```