Cisco SD-WAN Migration Guide

July 23, 2019

Version 1

Cisco Public. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Contents

CONTENTS **2**

LIST OF FIGURES AND TABLES **4**

ABOUT THIS GUIDE **7**

HISTORY 7

REVIEW 7

1 INTRODUCTION **8**

1.1 AUDIENCE 8

1.2 DOCUMENT SCOPE 8

1.3 ASSUMPTIONS AND CONSIDERATIONS 8

1.4 RELATED DOCUMENTS 8

2 PROJECT OVERVIEW **9**

2.1 CISCO SOLUTION OVERVIEW 9

2.1.1 SD-WAN Layered Design 10

2.2 CISCO SD-WAN CONTROLLERS 11

2.2.1 vManage Overview 11

2.2.2 vSmart Overview 12

2.2.3 vBond Overview 12

2.3 CISCO SD-WAN ROUTERS 12

2.3.1 Upgrading IOS-XE Routers to SD-WAN 14

3 MIGRATION PLANNING **15**

3.1 MIGRATION CONSIDERATIONS 15

3.1.1 Controller Considerations 15

3.1.2 Datacenter Considerations 15

3.1.3 Region/Branch Considerations 15

3.2 DEPLOYMENT STAGES 16

3.3 SEQUENCE OF MIGRATIONS 16

3.4 SMART ACCOUNT AND VIRTUAL ACCOUNT 17

3.5 DEPLOYING CONTROLLERS 19

3.5.1 Firewall Traffic Requirements 19

3.5.2 High Availability and Scalability of Controllers 19

3.6 SD-WAN CONFIGURATION TEMPLATES AND POLICIES 20
3.6.1 System IP 20
3.6.2 Site ID 20
3.6.3 Port Numbering and Mapping 21
3.6.4 Color 22
3.6.5 Segmentation 22
3.6.6 Migrating IOS-XE Configuration to SD-WAN 23
3.6.7 Configurations Template 23
4 MIGRATION SCENARIOS 25
4.1 DATACENTER MIGRATION 25
4.1.1 Legacy DC Migration 25
4.1.2 IWAN DC Migration to SD-WAN 28
4.2 BRANCH MIGRATION 30
4.2.1 Single Router Branch Migration 30
4.2.2 Inline Branch Migration 32
4.2.3 Parallel Branch Migration 33
5 MIGRATION CASE STUDY 38
5.1 GETTING READY 38
5.2 TOPOLOGY 39
5.3 MIGRATION STEPS 40
5.3.1 Smart Account 40
5.3.2 Controller Deployment 40
5.3.3 Datacenters Migration 40
5.3.4 Branch 1 Migration 45
5.3.5 Branch 2 Migration 48
5.3.6 Legacy WAN Branch 3 Migration 51
5.3.7 Branch 4 Migration 54
5.3.8 Phase out DC IWAN BRs 55

List of Figures and Tables

FIGURE 1: CISCO SD-WAN ARCHITECTURE 9
FIGURE 2: CISCO SD-WAN LAYERED ARCHITECTURE 10
FIGURE 3: SD-WAN CONTROLLERS HOSTING MODELS 11
FIGURE 4: THE CISCO SD-WAN ROUTER FAMILY 13
FIGURE 5: IOS-XE UPGRADE TO SD-WAN 14
FIGURE 6: DEPLOYMENT/MIGRATION STAGES 16
FIGURE 7: MIGRATION SEQUENCE 16
FIGURE 8: BOOTSTRAP PROCESS 18
FIGURE 9: CONTROL PLANE SCALABILITY 20
FIGURE 10: CISCO SD-WAN ARCHITECTURE 23
FIGURE 11: CONFIGURATION CONVERSION TOOL 23
FIGURE 12: DATACENTER SD-WAN EDGES BEHIND CE ROUTERS 25
FIGURE 13: SD-WAN BEHIND CES WITH IWAN BRS 28
FIGURE 14: SINGLE ROUTER BRANCH MIGRATION – ROUTING 31
FIGURE 15: INLINE BRANCH DEPLOYMENT 32
FIGURE 16: PARALLEL BRANCH DEPLOYMENT – UNDERLAY/OVERLAY 34
FIGURE 17: PARALLEL BRANCH DEPLOYMENT – WITH SERVICES 36
FIGURE 18: LEGACY WAN TOPOLOGY 39
FIGURE 19: IWAN TOPOLOGY 39
FIGURE 20: MIGRATING LEGACY DCS - SD-WAN EDGE ROUTERS PLACEMENT 40
FIGURE 21: MIGRATING IWAN DC1 - SD-WAN EDGE ROUTERS PLACEMENT 42
FIGURE 22: MIGRATING IWAN DC1 - ROUTING DURING MIGRATION 43
FIGURE 23: MIGRATING IWAN DC2 - ROUTING DURING MIGRATION 45
FIGURE 24: MIGRATING LEGACY BRANCH 1 TO SD-WAN 46
FIGURE 25: MIGRATING IWAN BRANCH 1 TO SD-WAN 46
FIGURE 26: MIGRATING BRANCH 1 – ROUTING 47
FIGURE 27: MIGRATING LEGACY WAN BRANCH 2 48
FIGURE 28: MIGRATING IWAN BRANCH 2 49
FIGURE 29: MIGRATING BRANCH 2 – ROUTING 49
FIGURE 30: MIGRATING LEGACY BRANCH 3 51

FIGURE 31: PARALLEL BRANCH MIGRATION – ROUTING 51
FIGURE 32: MIGRATING IWAN BRANCH 3 53
FIGURE 33: MIGRATING LEGACY WAN BRANCH 4 54
FIGURE 34: MIGRATING IWAN BRANCH 4 54
FIGURE 35: IWAN DC BRS AND CONTROLLERS REMOVED 55

TABLE 1: DOCUMENT HISTORY 7
TABLE 2 : DOCUMENT REVIEWERS 7
TABLE 3: SUMMARY OF THE CISCO SD-WAN COMPONENTS 10
TABLE 4: CISCO SD-WAN EDGE ROUTER FAMILY 13
TABLE 5: SD-WAN FIREWALL TRAFFIC REQUIREMENT 19
TABLE 6: SD-WAN SITE-ID PLANNING 21
TABLE 7: SD-WAN EDGE VPN TYPES 22
TABLE 8: SD-WAN EDGE ROUTERS CONNECTIVITY BEHIND CE ROUTERS 26
TABLE 9: CE ROUTERS CONNECTIVITY IN FRONT OF SD-WAN ROUTERS 27
TABLE 10: DATACENTER CORE ROUTERS CONNECTIVITY 27
TABLE 11: SD-WAN EDGE ROUTERS CONNECTIVITY BEHIND CE ROUTERS (IWAN DC) 29
TABLE 12: CE ROUTERS CONNECTIVITY IN FRONT OF SD-WAN ROUTERS (IWAN DC) 29
TABLE 13: IWAN: DC1 IWAN ROUTERS CONNECTIVITY AND ROUTING 29
TABLE 14: DATACENTER CORE ROUTERS CONNECTIVITY (IWAN DC) 30
TABLE 15: BRANCH 1 SD-WAN ROUTERS CONNECTIVITY AND ROUTING 32
TABLE 16: BRANCH INLINE DEPLOYMENT 33
TABLE 17: BRANCH SD-WAN EDGE ROUTER IN PARALLEL DEPLOYMENT – UNDERLAY/OVERLAY 35
TABLE 18: BRANCH CE ROUTER CONNECTIVITY IN PARALLEL DEPLOYMENT – UNDERLAY/OVERLAY 35
TABLE 19: BRANCH LAN SWITCH CONNECTIVITY IN PARALLEL DEPLOYMENT – UNDERLAY/OVERLAY 36
TABLE 20: BRANCH SD-WAN EDGE ROUTER IN PARALLEL DEPLOYMENT – WITH SERVICES 37
TABLE 21: BRANCH CE ROUTER CONNECTIVITY IN PARALLEL DEPLOYMENT – WITH SERVICES 37
TABLE 22: IWAN & SD-WAN ARCHITECTURES 38
TABLE 23: SD-WAN EDGE ROUTERS CONNECTIVITY BEHIND CE ROUTERS 41
TABLE 24: CE ROUTERS CONNECTIVITY IN FRONT OF SD-WAN EDGE ROUTERS 41
TABLE 25: DATACENTER CORE ROUTERS CONNECTIVITY 42
TABLE 26: IWAN: DC1 SD-WAN EDGE ROUTERS CONNECTIVITY AND ROUTING 43
TABLE 27: IWAN: DC1 CE ROUTERS CONNECTIVITY AND ROUTING 44

TABLE 28: IWAN: DC1 IWAN ROUTERS CONNECTIVITY AND ROUTING 44
TABLE 29: IWAN: DATACENTER CORE ROUTERS CONNECTIVITY 45
TABLE 30: BRANCH 1 SD-WAN ROUTERS CONNECTIVITY AND ROUTING 47
TABLE 31: BRANCH 2 SD-WAN ROUTERS CONNECTIVITY AND ROUTING 50
TABLE 32: BRANCH SD-WAN EDGE ROUTER IN PARALLEL DEPLOYMENT 52
TABLE 33: BRANCH CE ROUTER CONNECTIVITY IN PARALLEL DEPLOYMENT 52
TABLE 34: BRANCH CORE/DISTRIBUTION SWITCH CONNECTIVITY IN PARALLEL DEPLOYMENT 53

About this Guide

History

Table 1: Document History

Version No.	Issue Date	Status	Reason for Change
1	April 24th, 2019	Complete	Version 1

Review

Table 2 : Document Reviewers

Reviewer's Details	Version No.	Date
Cisco SD-WAN TMEs, Technical Leads	1	April 24th, 2019

1 Introduction

1.1 Audience

This document is intended for use by network engineers engaged in the architecture, planning, design, and implementation of migrating WAN to Cisco SD-WAN (powered by Viptela). The recommendations in this document should be used as a foundation for migrating any existing WAN to Cisco SD-WAN architecture.

1.2 Document Scope

The scope of this document includes:

- Overview of the SD-WAN (powered by Viptela) solution and its architecture
- Differences between SD-WAN and legacy WAN
- Legacy WAN migration scenarios including IWAN migration
- Preparing for migration
- Migration workflow
- Interworking with external domains
- Case study

WAN migrations are each very unique to the customer environment. This document gives general design options and guidelines, but does not specify configurations for connectivity and SD-WAN use cases, as these are particular to your environment. This document attempts to give you the tools, best practices, and designs to implement a migration customized to your existing environment.

The scope of this document is limited to the explanation of the migration strategy. For a detailed deployment guide, see Cisco Validated Deployment Guide.

1.3 Assumptions and Considerations

The following assumptions have been made in creating this document:

- Engineer(s) performing the migrations have technical knowledge of legacy WAN, IWAN, and SD-WAN solutions.
- Engineer(s) can configure and verify complex routing, IWAN and SD-WAN solutions individually.

1.4 Related Documents

- SD-WAN Product Documentation & Release Notes
- Plug and Play Guide
- SD-WAN CVDs
- Migration to Next-Gen SD-WAN – BRKCRS-2111

2 Project Overview

2.1 Cisco Solution Overview

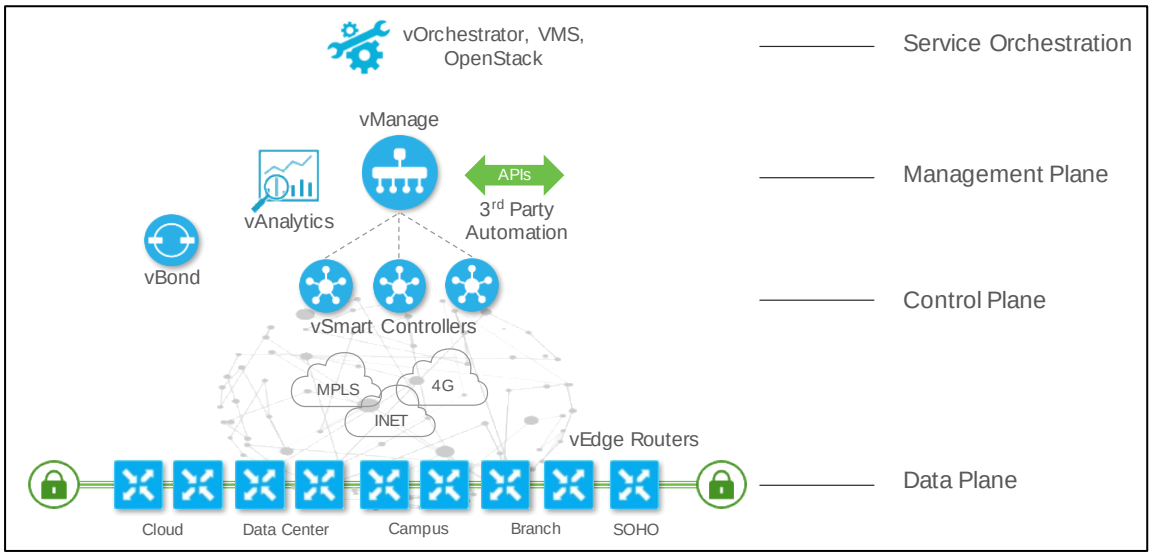


Figure 1: Cisco SD-WAN Architecture

Cisco SD-WAN architecture applies the principles of Software Defined Network (SDN) to the wide area network environment. By clearly separating control plane, data plane, and management plane functions, Cisco SD-WAN fabric achieves high degree of modularity.

Common SD-WAN use cases include:

- Hybrid WAN (MPLS, Internet, 4G) for bandwidth augmentation
- Application Aware Routing and SLA protection
- Direct Cloud Access (IaaS and SaaS)
- Cloud provisioning and management

The Cisco SD-WAN fabric is Cisco's next generation, Cisco cloud-based SD-WAN solution, providing customers with a turnkey solution for a virtual IP fabric that is secure, automatically deployed and provides any-to-any connectivity for next generation software services. This architecture is made up of four fundamental components:

Component	Description
Cisco vManage	The vManage NMS is a centralized network management system that lets you configure and manage the entire overlay network from a simple graphical dashboard.
Cisco vSmart Controller	The vSmart controller is the centralized routing and policy engine of the SD-WAN solution, controlling the flow of data traffic throughout the network.

Component	Description
	The vSmart controller works with vBond orchestrator to authenticate SD-WAN devices as they join the network and to orchestrate connectivity among edge routers.
Cisco SD-WAN Edge Routers 	Cisco SD-WAN edge routers are full-featured IP routers that perform standard functions such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), ACLs, QoS, and various routing policies in addition to the overlay communication. The edge routers sit at the perimeter of a site (such as remote offices, branches, campuses, data centers) and provide connectivity among the sites. They are either hardware devices or software, such as vEdge Cloud routers, which run as virtual machines. Edge routers handle the transmission of data traffic. vEdge cloud router can run as a Virtual Network Function (VNF) on Cisco Enterprise Network Compute System (ENCS) platforms. IOS-XE integration with Viptela vEdge capabilities is on the roadmap for ISR4k, ASR1k and CSR.
Cisco vBond Orchestrator 	The vBond orchestrator automatically orchestrates connectivity between edge routers and vSmart controllers. To allow an edge router or a vSmart controller to sit behind NAT, the vBond orchestrator also serves as an initial STUN server.

Table 3: Summary of the Cisco SD-WAN Components

2.1.1 SD-WAN Layered Design

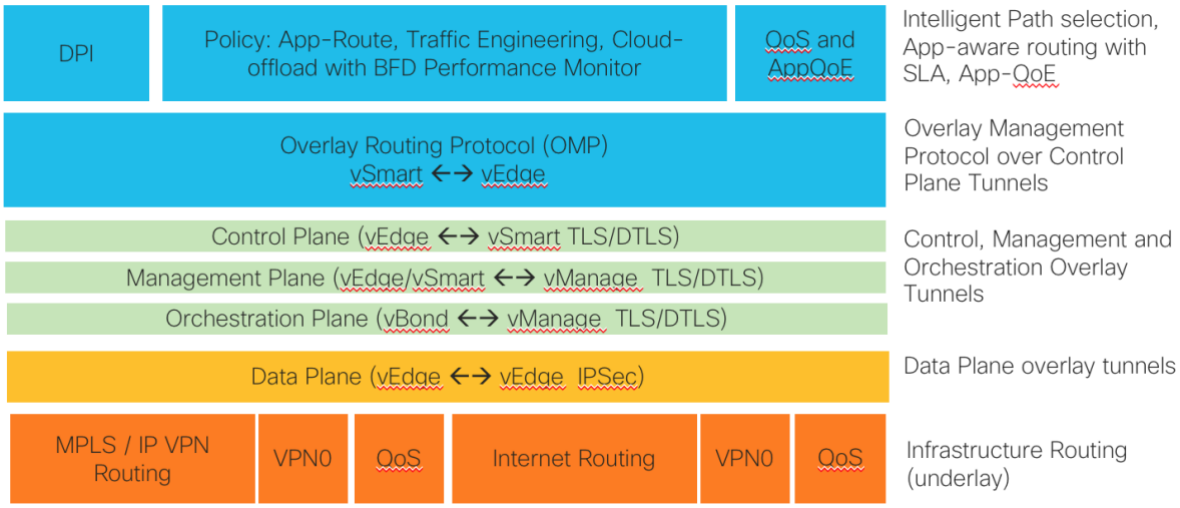


Figure 2: Cisco SD-WAN Layered Architecture

The software defined network deployment proposed uses the Cisco SD-WAN solution. The solution is described as a series of functional layers, using **Figure 2** as reference.

2.2 Cisco SD-WAN Controllers

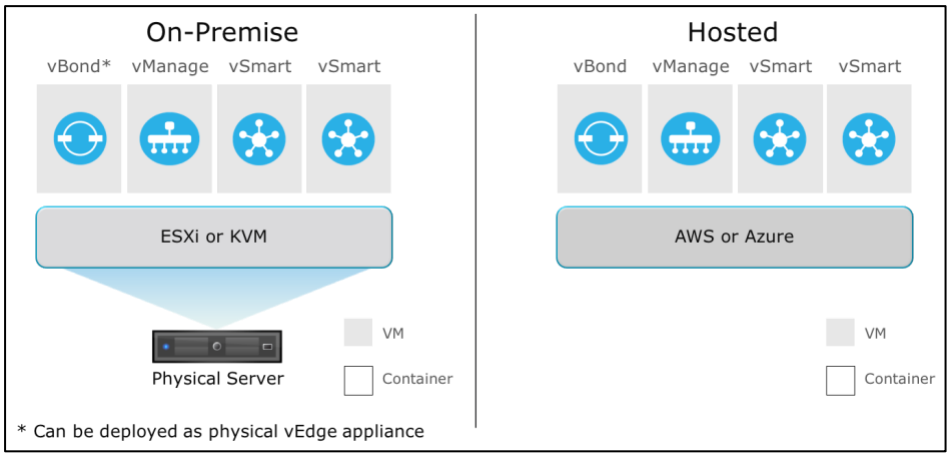


Figure 3: SD-WAN Controllers Hosting Models

Cisco vManage, vBond and vSmart controllers are software-only components. They are delivered in a virtual machine (VM) format supporting popular hypervisors (ESXi, KVM, AWS, Azure).

They can be deployed on customer premises, private cloud, or public cloud. They can be hosted by Cisco on AWS or Azure VPC. A combination of options is also possible, but in practice, this usually translates to Cisco-hosted control plane (vBond, vSmart) and on-premise management plane (vManage).

The clear advantage of the hosted model is outsourcing of compute, storage, configuration, and lifecycle management of the control and management planes. While this is secure, the enterprise must be prepared to give up private control and storage of data to a multi-tenant, virtual-private cloud service provider.

2.2.1 vManage Overview

The vManage Network Management System (NMS) is a centralized network management system that provides a GUI to easily monitor, configure, and maintain all Cisco SD-WAN devices and links in the overlay network to provide a dashboard for the network. The vManage NMS runs as a virtual machine (VM) on a network server, typically situated in a centralized location, such as a data center. The vSmart controller, discussed below, is also software that runs on a centralized server, and it is possible for both the vManage and vSmart software to run on the same physical server.

vManage sits on top of the hierarchy in the SD-WAN architecture and provides the following functions and services:

- Authentication Management – Device authentication database management
- Certificate Management – Certificate signing and lifetime management
- Device Configuration Management using Templates
- Element Management – Network element inventory, operational status, performance data, software information, element statistics
- Policy Management – Policy creation, management, and application
- Software Management and Upgrades
- Traffic and Performance Statistics

2.2.2 vSmart Overview

The vSmart controller is the brain of the overlay network, establishing, adjusting, and maintaining the connections that form the fabric of the overlay network. In these functions, it oversees the control plane of the Cisco SD-WAN overlay network. The vSmart controller participates only in the overlay network and has no direct peering relationships with any of the devices that an edge router is connected to on the host-facing side.

The vSmart-controllers are managed by vManage and are crucial components in the network. They are implemented using a redundant approach assisting both, in control-plane redundancy and scaling. Every edge router must be in session with a vSmart controller at all times, although there are redundancy and recovery measures built into the architecture.

vSmart is the centralized controller platform in the Cisco SD-WAN architecture and provides the following functions and services in the network:

- Dynamic control plane (OMP) peering other controllers and endpoints
- Network-wide policy enforcement
- Dynamic distribution of routing information, encryption keys and policies

2.2.3 vBond Overview

In the Cisco SD-WAN overlay network, the vBond orchestrator automatically facilitates the initial bring-up of vSmart controllers and edge routers. It also facilities the connectivity between vSmart controllers and edge routers. During the bring-up process, the vBond orchestrator authenticates and validates the devices that are ready to join the overlay network. This automatic orchestration process prevents having to bring up the devices manually, which can be a tedious and error-prone process.

The vBond orchestrator also functions as a Session Traversal of User Datagram Protocol Through Network Address Translator (STUN) server to determine the public-private IP mappings for all Edge, vSmart and vManage devices. During the bring-up and onboarding, the vBond orchestrator monitors and gets these mappings from network transactions, builds the mappings, and then sends them back to the Edge/vSmart/vBond for use during overlay routing.

The vBond orchestrator is the only Cisco SD-WAN device that must be located in a public address space for automated zero-touch provisioning. This design allows the vBond orchestrator to communicate with vSmart controllers and Edge routers that are located behind NAT devices. The design also allows the vBond orchestrator to solve any NAT-traversal issues of these Cisco SD-WAN devices.

2.3 Cisco SD-WAN Routers

The Cisco SD-WAN Edge routers are available as hardware Customer Premises Equipment (CPE) or as VNFs for vCPE solutions. They are responsible for the data traffic sent across the network. When you place an edge router into an existing network, it also acts as a standard router, as it can run OSPF, BGP, and static routing as well as standard router functions such as VLAN tagging, QoS, ACLs, and route policies.

The hardware SD-WAN Edge routers include a tamper-proof module – a trusted ID chip - which is a secure crypto-processor that contains the private key and public key for the router, along with a signed certificate. All this information is used for device authentication.

Cisco SD-WAN provides multiple form factors of SD-WAN Edge routers.

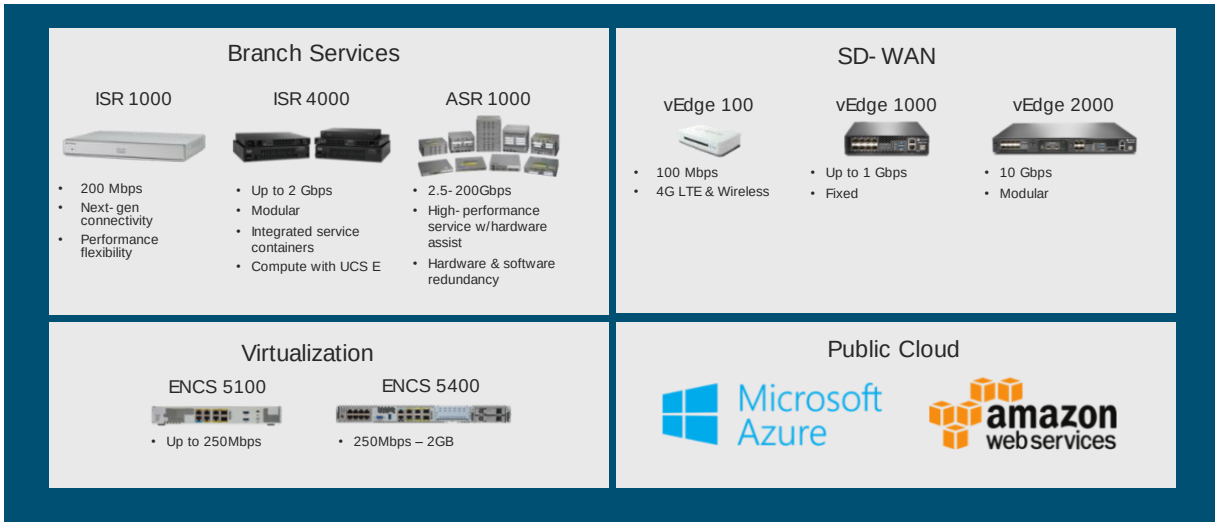


Figure 4: The Cisco SD-WAN Router Family

The table below provides more details.

Cisco SD-WAN Router Family
Original Viptela Products
- Positioned for transport-optimized and lower bandwidth applications where rich services such as WAAS, UC and embedded full-stack security (UTM) are not required. - Run the original Viptela software - Referred to as ‘vEdge’ routers, for example: vEdge 100, vEdge 1000, vEdge 2000
IOS-XE SD-WAN Routers
- Cisco’s integrated software with advanced multi-core hardware scale and services - Runs on Cisco ISR4000, ISR1100, and ASR1000 products - Referred to as ‘cEdge’ in the SD-WAN component family - ISR platforms are positioned for branch locations and ASRs for hub/datacenters
Cloud SD-WAN Routers
- Cisco SD-WAN offers virtual edge devices in the form of vEdge Cloud and CSRv - Instances are certified for and can be deployed in Amazon AWS and Microsoft Azure. - Cloud Edge instances can be implemented in any cloud environment that supports an ESXi, KVM, or HyperV hypervisor.

Table 4: Cisco SD-WAN Edge Router Family

2.3.1 Upgrading IOS-XE Routers to SD-WAN

In some migration scenarios, customers prefer to use their existing IOS-XE routers as cEdge routers. The Migration Scenarios section in this document explains the scenarios in which IOS-XE conversion to SD-WAN can be utilized. Upgrading an IOS-XE device involves the following steps.

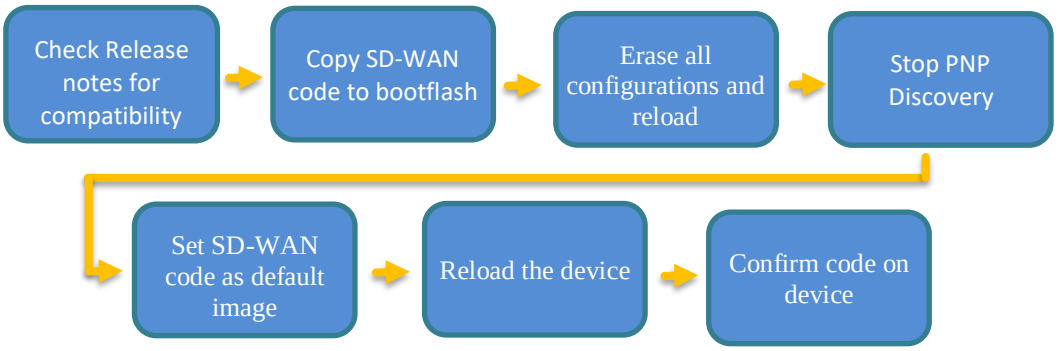


Figure 5: IOS-XE Upgrade to SD-WAN

The link below provides the detailed process on converting an IOS-XE router to an SD-WAN cEdge router.

https://sdwan-docs.cisco.com/Product_Documentation/Getting_Started/Hardware_and_Software_Installation/Software_Installation_and_Upgrade_for_Cisco_IOS_XE_Routers

The same process can be followed to upgrade an IOS-XE router running IWAN to an SD-WAN router.

3 Migration Planning

Deployment of SD-WAN provides tremendous advantages in terms of deployment speed, network automation, security, policy deployment, and cloud integration. Additionally, the consumption model is vastly simplified and has elastic scalability.

3.1 Migration Considerations

Migration planning is critical because moving from legacy WAN to SD-WAN requires changes to control plane and data plane architecture, design, as well as functional partitioning of the network. Below are some of the areas that must be considered to plan a successful migration and deployment of the required SD-WAN design.

3.1.1 Controller Considerations

When you are planning to deploy controllers, consider the following areas.

- Deployment Model
  - On-Prem
  - Cloud-hosted
- Redundancy
- Scale
- Firewall Ports

3.1.2 Datacenter Considerations

When you are planning the migration of your datacenter, consider the following.

- New/existing circuits
- Design
  - Dual MPLS/Hybrid/LTE
  - Full mesh/Hub-spoke/Arbitrary
- Routing
- Traffic path during migration
- Policy requirements
  - App-route/SLA
  - Traffic Engineering
  - Service Advertisement
  - QoS

3.1.3 Region/Branch Considerations

When you are planning site migration at a regional or branch level, consider the following:

- New/existing circuits
- Working applications/services
- Design
  - Dual MPLS/Hybrid/LTE
  - Full mesh/Spoke-hub/Arbitrary
- Policy requirements
  - DIA/Backhaul to DC
  - Regional breakout for services
  - Access to SaaS/IaaS

3.2 Deployment Stages

When deploying Cisco SD-WAN, the initial state of your network could be Green Field, Brown Field or IWAN deployed networks. Regardless of the initial stage, at a high level, the deployment/migration stages are the same.

The diagram below depicts the different stages of the Cisco SD-WAN deployment, starting from the different initial stage.

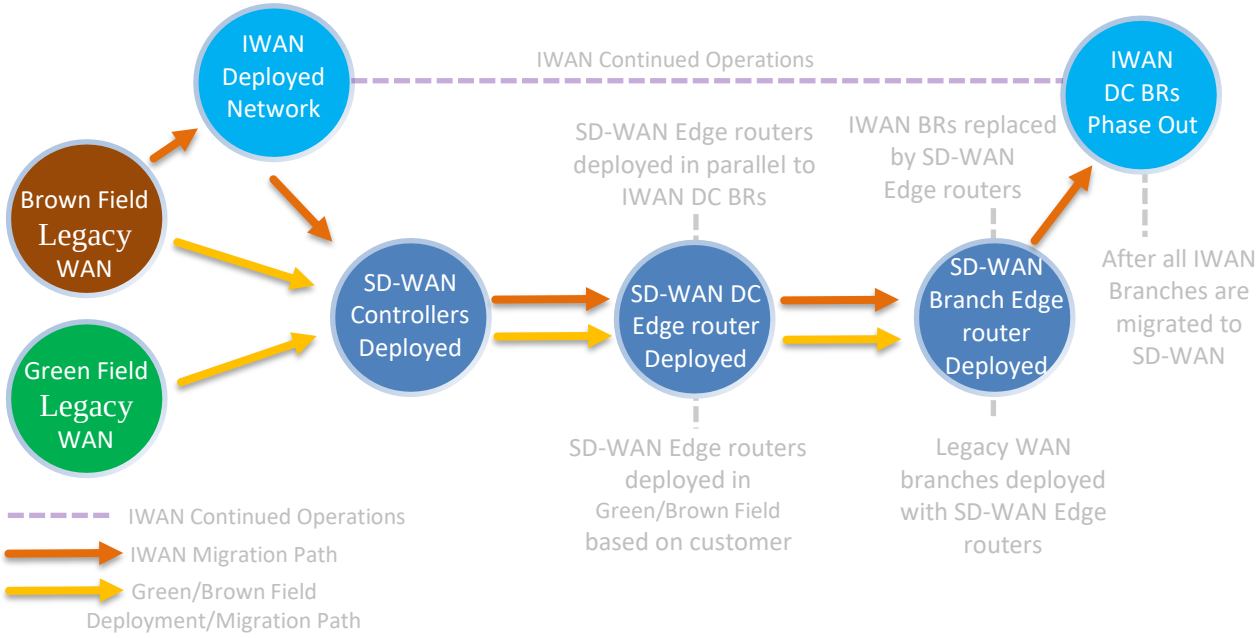


Figure 6: Deployment/Migration Stages

3.3 Sequence of Migrations

For successful migration from legacy models to Cisco SD-WAN, follow the sequence recommended below.

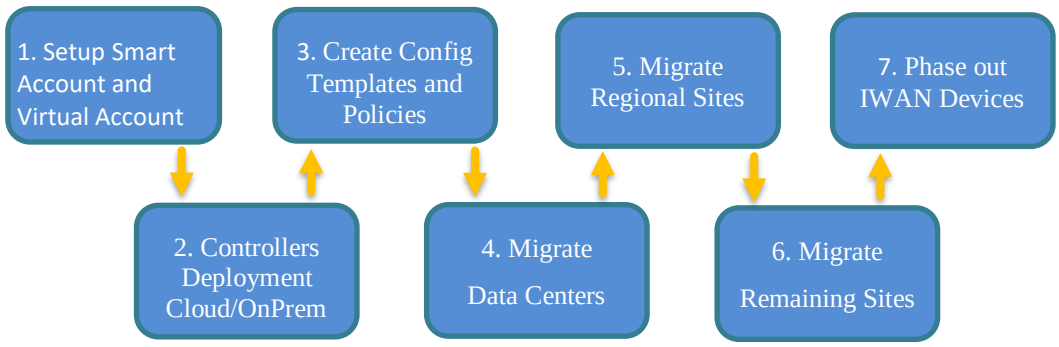


Figure 7: Migration Sequence

The same steps are followed for migration from IWAN to SD-WAN. However, during the migration, the configurations must support routing between IWAN and SD-WAN sites. After every step, it is imperative to verify that the existing and new routing flows are working as required.

1. The Smart Account provides the white-list file of the serial numbers for all SD-WAN Edge routers. This file is secure and signed to verify authenticity while bringing up SD-WAN. This file is uploaded on vManage, which then shares the white-list with other controllers. Only the edge routers listed in the white-list are allowed to join your SD-WAN fabric. Customers can use their existing Smart Accounts for SD-WAN.
2. Deploy controllers on Cloud or on premise. The controllers must be accessible over Internet and/or MPLS transport.
3. On vManage, create Edge routers configurations and define policies before the migration of a site. Test these configurations and policies in the lab environment before deployment.
4. Cisco recommends migrating data center sites first and using them for communication between the legacy and SD-WAN migrated sites until whole migration is complete. During IWAN migration to SD-WAN, ensure that there is routing between the legacy, IWAN, and SD-WAN sites.
5. Next migrate Regional hub or large branch sites in specific regions that act as regional exit points to the public cloud, host services for security, provide WAN optimization, etc.
6. Migrate the smaller branch sites for each region.
7. Finally, remove the data center legacy/IWAN routers followed by the IWAN controllers.

3.4 Smart Account and Virtual Account

Each customer needs to have a Smart Account. After creating a Smart Account, customers can create virtual accounts, that reflect their organizational departments, then associate licenses and devices with those departments. Smart Accounts and virtual accounts are essential in successful on-boarding of an SD-WAN edge router to its corresponding network.

The virtual account within the Smart Account is linked to a single SD-WAN overlay. All SD-WAN devices that are ordered by the customer are listed under the specific overlay virtual account, to be the part of the same SD-WAN overlay. Customers can also manually add their existing devices to their virtual account. Within the virtual account, create a controller profile, add devices, and capture the serial file in preparation for device redirection using the PnP portal. The serial file is uploaded on vManage, which then shares the white-list with other controllers. To access the Smart Account and Virtual Account, log in to https://software.cisco.com with your CEC credentials.

For any additional details on Plug and Play process, visit the following support guide. https://sdwan-docs.cisco.com/Product_Documentation/Getting_Started/Plug_and_Play_Support_Guide_for_Cisco_SD-WAN_Products

In some scenarios, zero touch provisioning is not possible, for example, when the DHCP service in unavailable. In such cases, cEdge can be booted with a bootstrap configuration. From vManage, generate bootstrap config file for the device. A config file (which includes basic interface configuration, Root CA, Organization Name, vBond information, etc.) is fed into the PnP process. Upon bootup, SD-WAN XE router searches bootflash: or usbflash: for the filename ciscosdwan.cfg. The router then continues the normal ZTP process.

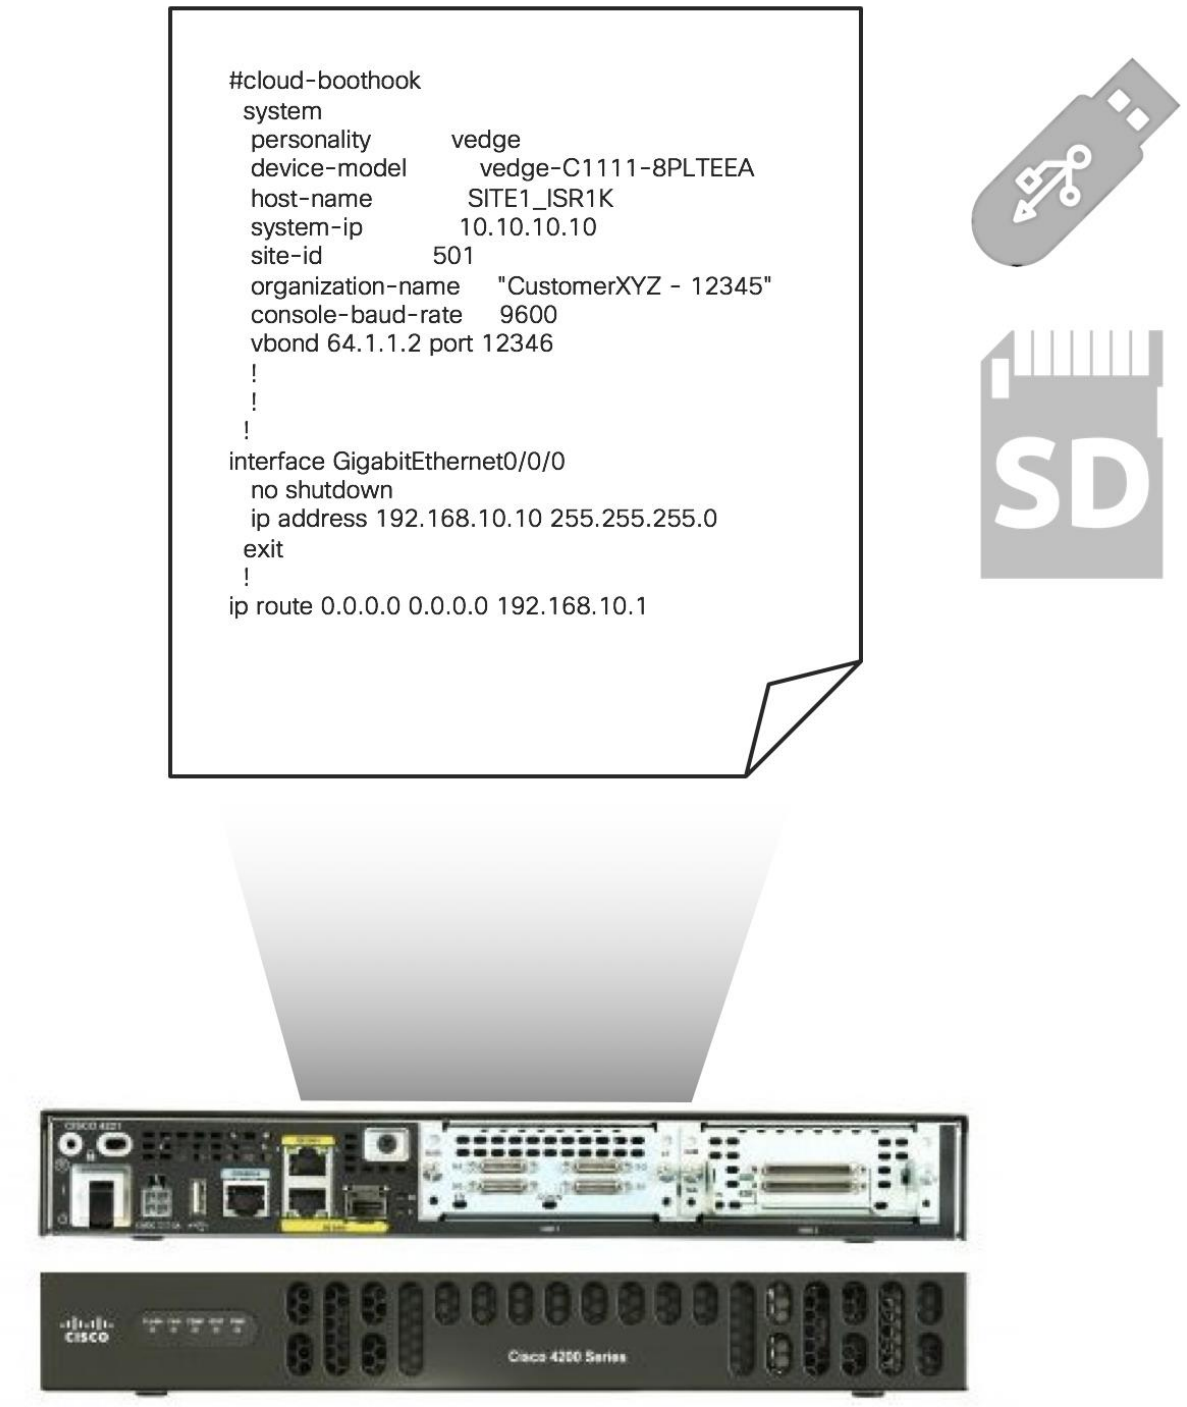


Figure 8: Bootstrap Process

3.5 Deploying Controllers

Controllers can be deployed in hosted cloud or on-premise environments. Refer to the Overlay Bringup Guide for more details on how to deploy controllers on-premise. For more information on how to deploy hosted cloud, visit the PNP Guide.

3.5.1 Firewall Traffic Requirements

The Cisco SD-WAN architecture separates control-plane and data-plane traffic. Control plane traffic requires communication using specific TCP/UDP ports. Ensure that any firewalls in the network allow to-and-from traffic between the SD-WAN devices. Here is a summary of the ports used, assuming the controllers are configured to not use port-hopping:

Source Device	Source Port	Destination Device	Destination Port
vSmart/vManage(DTLS)	UDP/12346-13065	vBond	UDP 12346
vManage	UDP/12346-13065 TCP random port number > 1024	vSmart	UDP/12346-12445 TCP/23456-23555
Edge (DTLS/TLS)	UDP/12346-12445 TCP random port number > 1024	vManage	UDP/12346-13065 TCP/23456-24175
Edge (DTLS)	UDP/12346-12445	vBond	UDP 12346
Edge (DTLS/TLS)	UDP/12346-12445 TCP random port number > 1024	vSmart	UDP/12346-13065 TCP/23456-24175
Edge (IPSec)	UDP 12346+n, 12366+n, 12386+n, 12406+n, and 12426+n, where n=0-19 and represents the configured offset	Edge	UDP 12346+n, 12366+n, 12386+n, 12406+n, and 12426+n, where n=0-19 and represents the configured offset

Table 5: SD-WAN Firewall Traffic Requirement

For more detailed information on firewall ports, visit Firewall Ports for Viptela Deployments.

3.5.2 High Availability and Scalability of Controllers

The Cisco SD-WAN solution is designed to scale horizontally as needed to meet WAN capacity. To increase capacity and have redundancy/high availability, add additional controllers horizontally. At a minimum, the Cisco SD-WAN solution needs one component of each controller. Edge routers establish a temporary connection with the vBond orchestrator at the time of bring-up, and permanent connections with vManage and vSmart. The following image shows the scalability numbers for each of the controllers. It also shows how many components from each controller can be deployed in a single overlay.

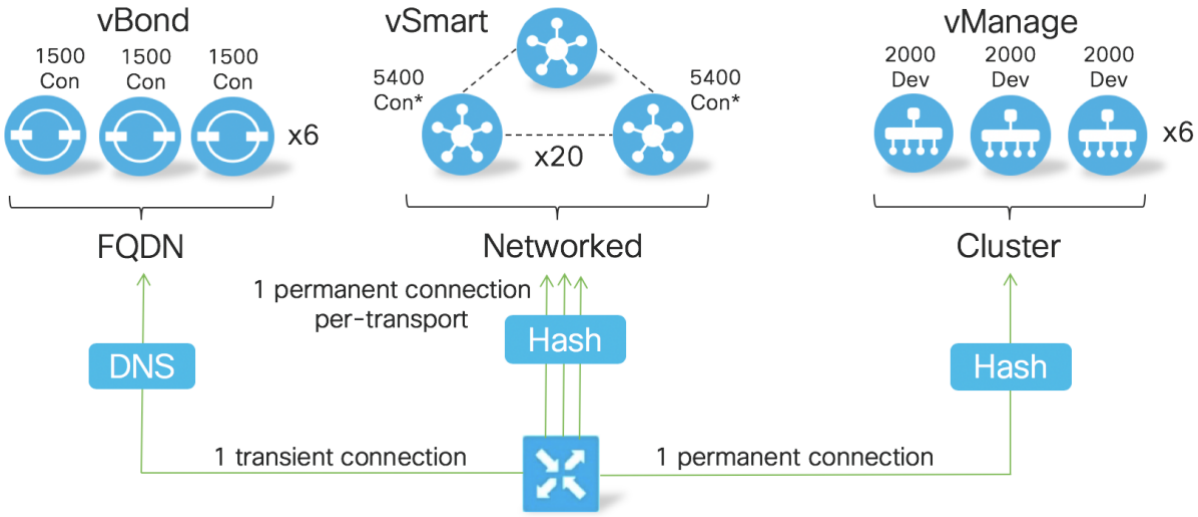


Figure 9: Control Plane Scalability

For more details on High Availability, visit High Availability and Scaling.

3.6 SD-WAN Configuration Templates and Policies

On vManage, create edge router configurations using templates and define policies before the migration of a site. The number of policies defined can vary by the customers' specific use cases. For details on policies refer to the product documentation on cisco.com.

Here are the basic guidelines to help build configurations for SD-WAN deployment.

3.6.1 System IP

System IP is a persistent, system-level IPv4 address that uniquely identifies the device independent of any interface addresses. It acts much like a router id, so it doesn't need to be advertised or known by the underlay. A best practice, however, is to advertise this system IP address in the service VPN and use it as a source IP address for SNMP and logging, which makes it easier to correlate network events with vManage information. You need to configure a system IP address for your controllers to authenticate an Edge router and bring it onto the overlay network.

A logical scheme for your system IP addresses is recommended to make sites more easily recognizable.

3.6.2 Site ID

A site ID is a unique identifier of a site in the SD-WAN overlay network with a numeric value starting from 1 through 4294967295. This ID must be the same for all of the edge devices that reside on the same site. A site could be a data center, a branch office, a campus, or something similar. A site ID is required to be configured in order for an edge router to be authenticated by the controllers and brought into the overlay network. By default, IPSec tunnels are not formed between edge routers within the same site.

A site ID scheme should be chosen carefully, as this makes it easier to apply a policy. You must apply a policy to a list or range of site IDs (example 100,200-299). Note that there is no wildcard support.

Although there are several different ways to organize a site-id scheme, here is an example of a scheme that uses 9 digits:

Nine-digit site ID example

Digit	Representation	Examples
1	Country/Continent	1=North America, 2=Europe, 3=APAC
2	Region	1=US West, 2=US East, 3=Canada West, 4=Canada East
3-6	Site Type	0000-0099=Hub locations, 1000-1999=Type 1 sites, 2000-2999=Type 2 sites, 3000-3999 = Type 3 sites, 4000-4999=Type 4 sites, 5000-9999 = future use
7-9	Store/site/branch number, or any other ID specifier	001, 002, 003

Table 6: SD-WAN Site-ID Planning

Grouping according to geography is helpful in cases where you prefer a regional datacenter over another for centralized Internet access, or for connectivity to hubs in other countries/regions.

Site types should be created according to types of policies applied to make policy application easier. When a new site is created, just creating a site ID that falls into the matching range of a policy will automatically cause the policy to be applied to it. Below are some examples of how you can group branches according to type:

- Branches that use a centrally-located firewall or another centrally-located service
- Branches that use direct Internet access
- Lower versus higher bandwidth sites – in case you want different topologies for each. Low bandwidth sites could use a hub and spoke topology to save bandwidth, while higher bandwidth sites use a full-mesh topology.
- Sites with different SLA/transport requirements, such as using MPLS for critical traffic, voice, and video while everything else traverses the Internet circuit, and perhaps some sites using MPLS for voice only, while everything else traverses the Internet circuit.

You can also have overlapping types, but the idea is to put them in categories that makes applying policy configurations easier. It helps to consider the requirements and policies before assigning site IDs,

3.6.3 Port Numbering and Mapping

Cisco recommends having a port-numbering scheme that is consistent throughout the network. Consistency assists in easier configuration and troubleshooting. The consistent port

mapping simplifies policy management and enhances the usability of configuration templates. For example: Gig0/0 always has Internet and Gig0/1 has MPLS transport.

In addition, the default factory configuration of an edge router specifies certain ports in VPN 0 for DHCP so that the Edge router can automatically obtain a DHCP address, resolve DNS, and communicate with the ZTP/PNP server. So, if you utilize ZTP/PNP, ensure that this port can reach the DHCP and DNS servers by connecting them to the most appropriate place in the network.

3.6.4 Color

On Edge routers, the color attribute helps to identify an individual WAN transport tunnel and is one of the Transport location (TLOC) parameters associated with a tunnel. You cannot use the same color twice on a single Edge.

Colors by themselves have significance. The colors metro-ethernet, MPLS, and private1, private2, private3, private4, private5, and private6 are considered private colors. They are intended to be used for private networks or in places where you will have no NAT addressing of the transport IP endpoints. When an Edge router uses a private color, it attempts to build IPSec tunnels to other Edge routers using the native, private, underlay IP. With public colors, Edge routers try to build tunnels to the post NAT IP address (if there is NAT involved).

If you are using a private color and need the NAT to communicate to another private color, the carrier setting in the configuration dictates whether you use private or public IP. Use this setting and two private colors to establish a session when one or both are using NAT. The public colors are 3g, biz, internet, blue, bronze, custom1, custom2, custom3, default, gold, green, lte, public-internet, red, and silver.

3.6.5 Segmentation

SD-WAN uses VPN technology. Edge routers maintain a separate FIB for each VPN and the VPN label is carried within the data tunnels to maintain end-to-end segmentation. The WAN interfaces on an edge router, connected to WAN transport, are always VPN 0. The management interface is in VPN 512. For LAN side of the network, VPN 1-511 can be used.

VPN Type	Use
VPN 0	Transport (WAN) type connectivity
VPN 512	Out-of-Band Management
VPN 1-511	Service (LAN) side connectivity

Table 7: SD-WAN Edge VPN Types

In a legacy network, routing segmentation is achieved by using VRFs. When migrating to SD-WAN, configure VPNs on the LAN side to maintain segmentation similar to VRF.

The interfaces and sub-interfaces can be assigned to the VPNs. Interfaces are configured for either layer 3 routing protocol or layer 2 dot1q trunk support.

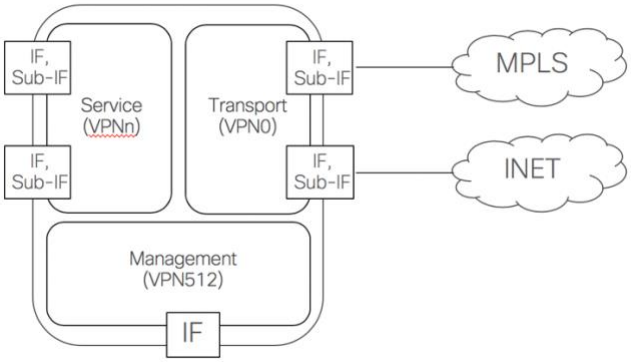


Figure 10: Cisco SD-WAN Architecture

3.6.6 Migrating IOS-XE Configuration to SD-WAN

For cEdge deployments, Cisco provides a configuration conversion tool to migrate the IOS-XE configurations to SD-WAN. The tool identifies the feature parity between the IOS-XE to SD-WAN code and lets you know of the configurations that should be removed before migration. You can add SD-WAN specific information, if required, in the configuration. Once the configuration is finalized, the tool makes an API call and creates templates in the vManage.

Note: The tool is currently available internally to Cisco employees only. Customers can contact Cisco Sales Account Team for assistance.

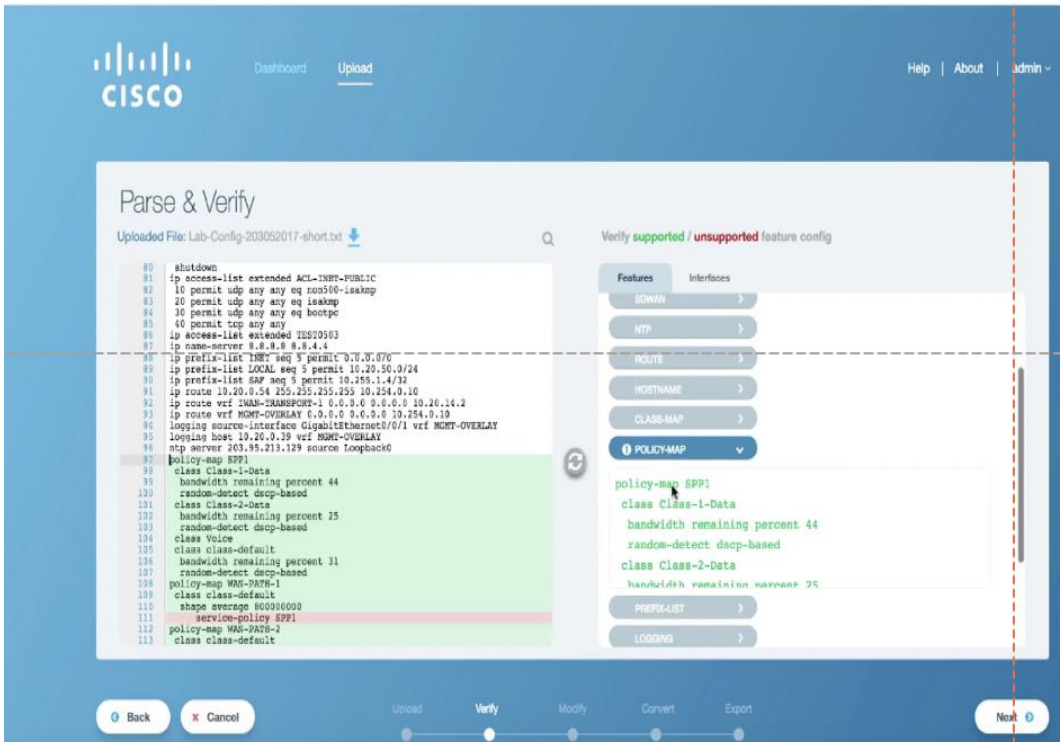


Figure 11: Configuration Conversion Tool

3.6.7 Configurations Template

All devices in a Cisco SD-WAN overlay network that are managed by the vManage NMS must be exclusively configured from the NMS. The configuration procedure is as follows:

1. **Create feature templates**: Feature templates are the fundamental building blocks of device configuration. For each feature that you can enable on a device, the vManage NMS provides a factory default template form that you customize for your deployment. The form allows you to set global values for all devices, or variables that can be customized during site specific provisioning.
2. **Create device templates**: Device templates contain the complete operational configuration of a device. You create device templates for different device types (Data Center, small branch, large branch...etc.) by consolidating multiple feature templates. For each device type, if multiple devices have the same configuration, you can use the same device template for them. For example, many of the edge routers in the overlay network might have the same basic configuration, so you can configure them with the same templates. If the configuration for the same type of devices is different, you create separate device templates.
3. **Attach devices to device templates**. To configure a device on the overlay network, you attach a device template to the device.
4. **Input site specific values into template variables**. Populate templates with site specific configuration by providing values to variables and deploy to the device.

If the device being configured is present and operational on the network, the configuration is sent to the device and takes effect immediately. If the device has not yet joined the network, the configuration to the device is scheduled to be pushed by vManage NMS as soon as the device joins the network.

4 Migration Scenarios

This section provides different scenarios of migration to SD-WAN for datacenter and branch sites.

4.1 Datacenter Migration

Datacenter is the first site that is migrated to SD-WAN. This is because the migration of the branch sites is typically gradual and during the migration the datacenter serves as the transit site for the traffic between non-SDWAN and SD-WAN sites. Since data centers become transit sites, plan for adequate bandwidth utilization that may be required at the datacenter.

In very large networks, where applications can experience latency issues if traffic needs to transit to DC during the migration, \ designate one of the regional sites to be the transit site between SD-WAN and non SD-WAN sites.

This document only explains the recommended migration method, but the actual migration of the datacenter site may vary based on setup of the customers.

4.1.1 Legacy DC Migration

The recommended method to migrate from Legacy WAN to SD-WAN is to connect SD-WAN Edge routers behind the CEs, without much effect to the existing routing. The internet and MPLS connectivity is extended to Edge routers via CEs. Edge routers are also connected to core switches on LAN side as shown in figure 12.

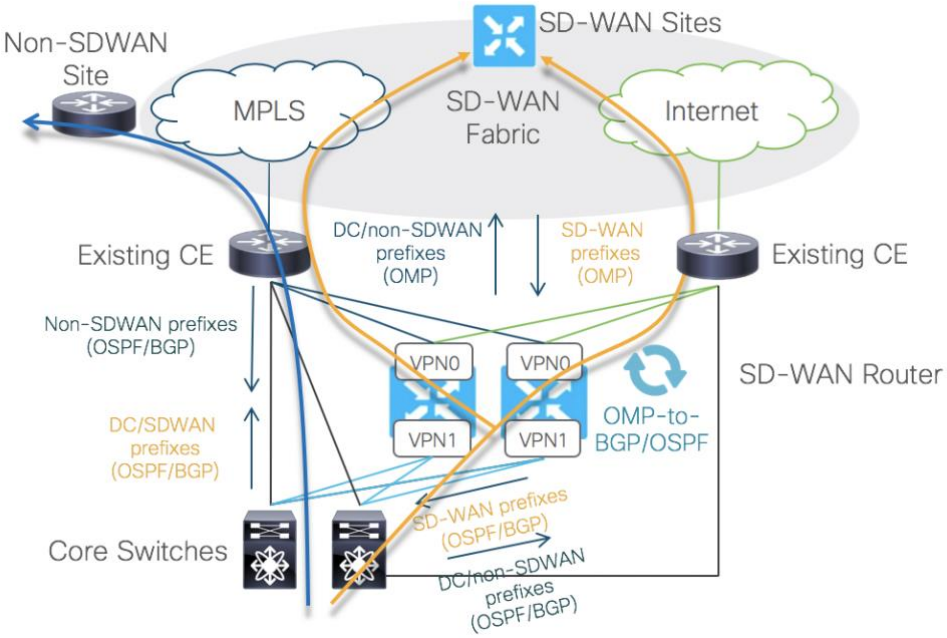


Figure 12*: *Datacenter SD-WAN Edges Behind CE Routers

The table below explains the routing design at the datacenter.

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	- Connect VPN0 WAN side to both CE routers - Use /30 on MPLS CE-to-SDWAN router link - Internet link gets IP through DHCP
	LAN Service VPNs	- Connect to L3 LAN Core switches to Service-side VPN interfaces.
WAN Advertisements	IN	- SD-WAN Prefixes from SD-WAN sites over Internet and MPLS connections through OMP
	OUT	- Through OMP advertise Datacenter LAN prefixes, default GW, aggregate routes for non-SD-WAN prefixes to SD-WAN sites
LAN Advertisements	IN	- Local LAN prefixes, default GW and aggregate routes for Non SD-WAN site prefixes from L3 LAN Core Switch.
	OUT	- SD-WAN sites prefixes to L3 LAN Core Switch

Table 8: SD-WAN Edge Routers Connectivity Behind CE Routers

Note that in certain branch deployments, a static route is used on SD-WAN edge as a default gateway for local internet breakout. If a static route is used at the branch, then the default gateway advertised from the datacenter won't be used and may cause traffic to black hole if there is no other better match for the prefixes. In such scenarios, either use data policy at the branch to perform the local internet break out, or advertise specific prefixes (aggregated routes) from the datacenter.

Since there are typically two edge routers at the datacenter and both devices perform redistribution between OMP and LAN routing protocol, there can be a routing loop. Make sure that prefixes learned from an SD-WAN site are not redistributed into OMP again at the datacenter, which can allow loops. If BGP is the datacenter LAN protocol, then configure both edge routers in the same autonomous system (AS) and create eBGP neighborhood between the core routers and edge routers. Because of the same BGP AS-PATH, the second edge router will not install any of the routes that were originally redistributed by the other edge router from OMP.

When LAN uses OSPF/EIGRP, use tags to mark the prefixes when redistributing from OMP to OSPF/EIGRP on both SD-WAN edge routers. Use these tags to filter the prefixes when redistributing from LAN to OMP.

		CE Routers
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS and Internet - Extend MPLS to SD-WAN Edge over /30 physical link - Extend Internet to SD-WAN Edge over L2/L3 DHCP relay
	LAN	- Connect to L3 LAN Core switch

WAN Advertisements	IN	- Non SD-WAN prefixes from internet and MPLS WAN links
	OUT	- DC, Non-SD-WAN, and SD-WAN site prefixes to Non-SD-WAN sites
LAN Advertisements	IN	- DC and SD-WAN site prefixes from L3 LAN Core Switch
	OUT	- Non SD-WAN site prefixes to L3 LAN Core Switch

Table 9: CE Routers Connectivity in front of SD-WAN Routers

		Datacenter Core Routers
Physical/L3 Connectivity	WAN	- Connect to CE routers - Connect to SD-WAN Edge routers
	LAN	- Connect to L3/L2 Distribution/Access switches as per DC design
WAN Advertisements	IN	- Non SD-WAN prefixes from CE routers (MPLS/Internet) - SD-WAN prefixes from SD-WAN Edge routers
	OUT	- DC, Non-SD-WAN, and SD-WAN site prefixes to CE routers - DC prefixes, default GW, aggregate routes of Non SD-WAN sites to SD-WAN Edge routers
LAN Advertisements	IN	- DC Prefixes from LAN network
	OUT	- Non SD-WAN and SD-WAN site prefixes to LAN as per DC LAN design requirement

Table 10: Datacenter Core Routers Connectivity

Typically, there are more than one datacenters for HA/redundancy requirements. After successful migration of the first datacenter, migrate the second datacenter in a similar method as explained in this section. Note that a routing loop can occur if there is a backdoor link between the datacenter sites and route advertisement is configured between the two datacenters. To avoid the loop, any of the three methods explained below can be used:

1. Use the same Autonomous System Numbers (ASN) on edge routers of the two datacenters. Because of the same ASN, the AS-PATH attribute will avoid learning the same prefixes on the edge routers that are advertised by the other datacenter towards the LAN side.

2. Use overlay-AS to insert Overlay Management Protocol (OMP) AS number when redistributing the routes from OMP into LAN side towards DC LAN. Configure all DC SD-WAN edge routers with the same overlay-as. This allows the edges to filter the routes advertised by the other DCs edge devices towards the LAN side and prevents redistributing the same routes back into OMP.

3. Use tags or communities to mark the prefixes at one datacenter when redistributing to DC LAN and filter on the edge of the other datacenter when learning advertisements from the LAN side.

4.1.2 IWAN DC Migration to SD-WAN

The migration methodology described in section 4.1.1 is also recommended for migrating from IWAN to SD-WAN. The IWAN Border Routers (BRs) are already connected to CEs. SD-WAN routers are added to the topology behind CEs as well, as shown in figure 13. The core routers advertise IWAN prefixes, gateways for non SD-WAN routers, and DC prefixes to SD-WAN routers. The core routers also advertise SD-WAN, non SD-WAN and DC prefixes to IWAN BRs.

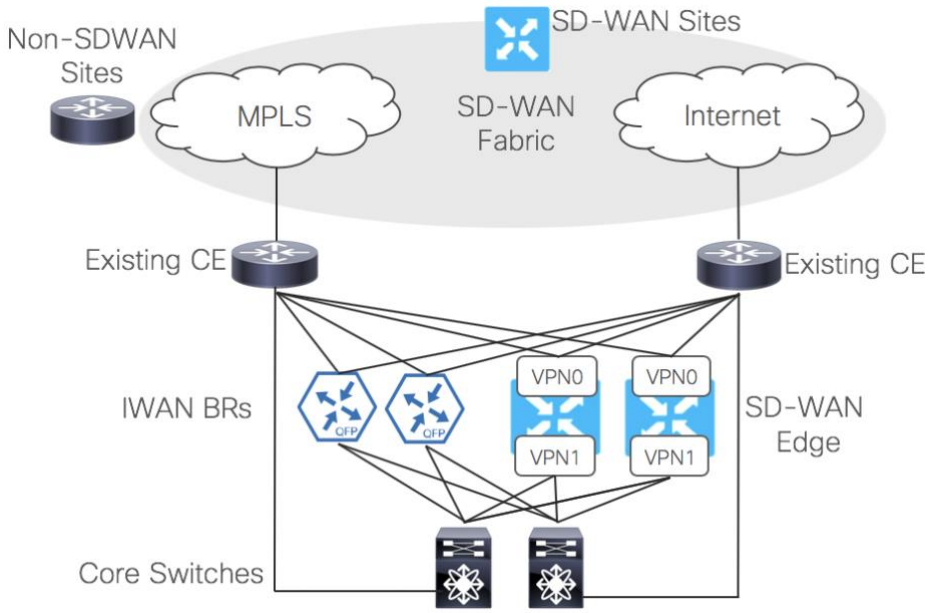


Figure 13: SD-WAN Behind CEs With IWAN BRs

The table below explains the routing design for IWAN migration at the datacenter.

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	- Connect VPN0 WAN side to both CE routers - Use /30 on MPLS CE-to-SDWAN router link - Internet link gets IP through DHCP
	LAN Service VPNs	- Connect to L3 LAN Core switches to service-side VPN interfaces
WAN Advertisements	IN	- SD-WAN prefixes from SD-WAN sites over Internet and MPLS connections through OMP
	OUT	- Through OMP advertise datacenter LAN prefixes, default GW, aggregate routes for non-SD-WAN and IWAN prefixes to SD-WAN sites

LAN Advertisements	IN	- Local LAN prefixes, default GW and aggregate routes for Non SD-WAN and IWAN prefixes from L3 LAN Core Switch.
	OUT	- SD-WAN sites prefixes to L3 LAN core switch

Table 11: SD-WAN Edge Routers Connectivity Behind CE Routers (IWAN DC)

		CE Routers
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS and Internet - Extend MPLS to SD-WAN edge over /30 physical link - Extend Internet to SD-WAN Edge over L2/L3 DHCP relay - No changes to connections to IWAN BRs
	LAN	- Connect to L3 LAN Core switch
WAN Advertisements	IN	- Non SD-WAN prefixes from internet and MPLS WAN links
	OUT	- DC, Non-SD-WAN, IWAN and SD-WAN site prefixes to Non SD-WAN sites
LAN Advertisements	IN	- DC, IWAN and SD-WAN site prefixes from L3 LAN Core Switch
	OUT	- Non SD-WAN site prefixes to L3 LAN Core Switch

Table 12: CE Routers Connectivity in Front of SD-WAN Routers (IWAN DC)

		IWAN BRs
Physical/L3 Connectivity	WAN	- No change in connections to CEs
	LAN	- No change in connections to core routers
WAN Advertisements	IN	- IWAN site prefixes from Internet and MPLS connections
	OUT	- Advertise Datacenter LAN, SD-WAN and non-SD-WAN prefixes to IWAN sites
LAN Advertisements	IN	- Local LAN, SD-WAN and Non SD-WAN site prefixes from L3 LAN core Switch
	OUT	- IWAN site prefixes to L3 LAN Core Switch

Table 13: IWAN: DC1 IWAN Routers Connectivity and Routing

		Datacenter Core Routers
Physical/L3 Connectivity	WAN	- No change to connections to CE routers - No change to connections to IWAN routers - Connect to SD-WAN Edge routers
	LAN	- Connect to L3/L2 Distribution/Access switches as per DC design
WAN Advertisements	IN	- Non SD-WAN prefixes from CE routers (MPLS/Internet) - IWAN prefixes from IWAN BRs - SD-WAN prefixes from SD-WAN edge routers
	OUT	- DC, Non-SD-WAN, and SD-WAN site prefixes to CE routers - DC, Non-SD-WAN, and SD-WAN site prefixes to IWAN routers - Local LAN prefixes, default GW and aggregate routes for Non SD-WAN and IWAN prefixes to SD-WAN routers
LAN Advertisements	IN	- DC prefixes from LAN network
	OUT	- Non SD-WAN, IWAN and SD-WAN site prefixes to LAN as per DC LAN design requirement

Table 14: Datacenter Core Routers Connectivity (IWAN DC)

This method maintains the router level redundancy for both IWAN and SD-WAN fabric. In addition, it provides internet and MPLS connectivity to both fabrics. This allows more flexibility in the migration of the branch sites. The IWAN BRs are removed only after all IWAN branches are migrated.

4.2 Branch Migration

Once a datacenter site is migrated to SD-WAN, the legacy WAN branches can be migrated too. The branch sites can have different topologies depending on the type and number of WAN circuits and HA design. Migration of the branches is done in a single cutover at each branch.

4.2.1 Single Router Branch Migration

Migration of a single router branch requires downtime.

To minimize the branch migration time, take the following steps.

1. Pre-stage the router with the minimum configuration that is required to establish the control connections and data tunnels. If PNP process is not used, then the device must be configured with System-IP, site ID, vBond address, and WAN interfaces under VPN0. The service side VPNs and service interfaces should also be configured well according to the segmentation requirements of that particular branch.

2. Verify on the SD-WAN controllers that the serial is listed in the list of devices with the state Valid.
3. Attach required configuration templates to the device on the vManage NMS.
4. If the branch router is SD-WAN capable, upgrade the router from IOS-XE code to SD-WAN. If the branch router cannot be upgraded to SD-WAN, connect a separate SD-WAN router to the LAN network and move the internet circuit from the existing router to SD-WAN Edge. Make sure that the existing router is the gateway for the router until SD-WAN edge is completely onboarded.
5. Once the SD-WAN edge router establishes connection with the controllers, vManage pushes the configuration to the device.
6. Then make SD-WAN edge router the LAN gateway and move the MPLS circuit to SD-WAN edge.

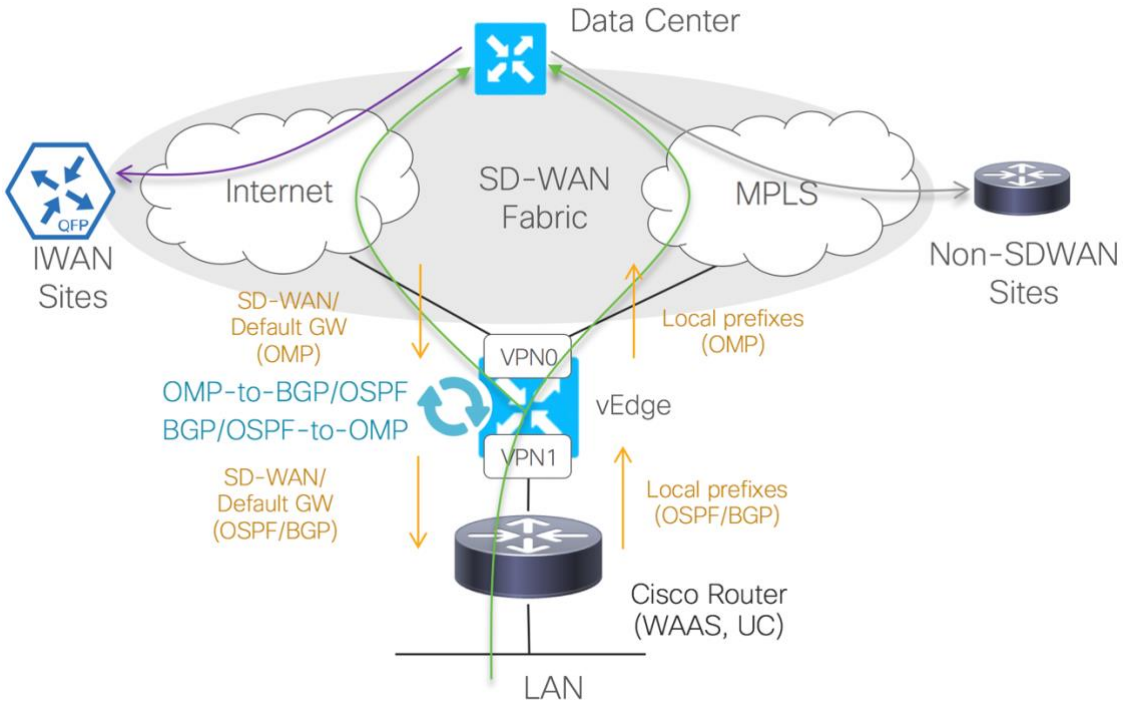


Figure 14: Single Router Branch Migration – Routing

The table below explains the routing design at the branch.

		SD-WAN Edge Router
Physical/L3 Connectivity	WAN VPN 0	- MPLS and Internet connections terminate on SD-WAN Edge router on interfaces under VPN 0
	LAN Service VPNs	- Connect to LAN switches in service VPNs. LAN design will dictate if subinterfaces are needed.

WAN Advertisements VPN 0	IN	- SD-WAN prefixes, aggregate routes and default GW from datacenter from OMP session over Internet and MPLS connections
	OUT	- Redistribute local LAN prefixes into OMP
LAN Advertisements Service VPNs	IN	- Local LAN prefixes – SD-WAN Edge router typically is the GW
	OUT	- With L3 connection on the LAN side – advertise prefixes learned through OMP to LAN - With L2 connection on LAN side – no advertisements are needed as SD-WAN Edge router is the GW for the VLAN/VPN segments

Table 15: Branch 1 SD-WAN Routers Connectivity and Routing

4.2.2 Inline Branch Migration

In scenarios where a branch has two WAN edge routers, downtime can be minimized during the migration. In this type of migration, one router is migrated to SD-WAN first and then the second one. Make sure the gateways for LAN are pointed to the second router before upgrading the router to SD-WAN. Commonly, the SD-WAN controllers are accessible over Internet transport, so first migrate the router connected to Internet to SD-WAN by either upgrading the router to the SD-WAN image or by replacing the router with an SD-WAN router. Keep the SD-WAN router in the staging state through vManage, where it establishes control connections with the controllers and learns the prefixes but does not create data tunnels. Once the Internet router is successfully migrated, mark the device on vManage as Valid, point the LAN gateway to the SD-WAN router and then replace or upgrade the MPLS router to SD-WAN.

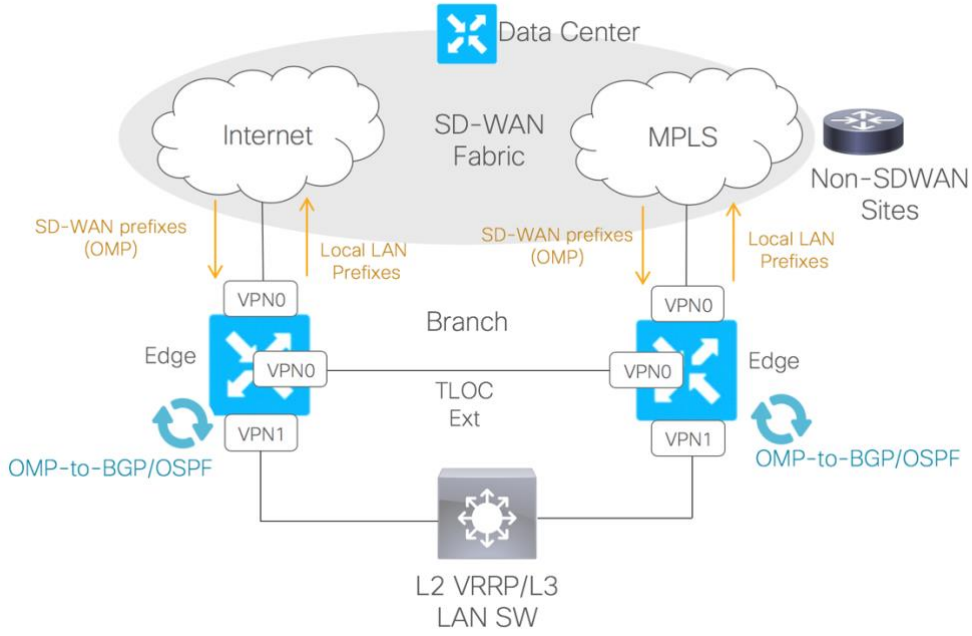


Figure 15: Inline Branch Deployment

If each router is terminating one transport as shown in Figure 15, then the TLOC-Extension feature is configured between the SD-WAN Edge routers to extend the WAN links connectivity. After a branch has been migrated, it communicates with non-SD-WAN sites by using the aggregate routes or default router learned from the datacenter. For more information about TLOC-Extension, visit the Extend the WAN Transport VPN guide.

The table below explains the routing design at the branch.

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	- One Edge router connects to the Internet - One Edge router connects to MPLS - Using TLOC-Extension MPLS connectivity is extended to internet Edge router - Using TLOC-Extension Internet connectivity is extended to MPLS Edge router
	LAN Service VPNs	- Connect to L2/L3 LAN Core switches in Service VPNs
WAN Advertisements	IN	- SD-WAN prefixes, aggregate routes and default GW from SD-WAN sites from Internet and MPLS connections over OMP
	OUT	- Local LAN prefixes over OMP to SD-WAN fabric
LAN Advertisements	IN	- With L3 LAN side, LAN prefixes from LAN switch - With L2 VRRP, no advertisements
	OUT	- With L3 LAN advertise SD-WAN prefixes to LAN switch - With L2 LAN, no advertisement needed. VRRP routers are the GW

Table 16: Branch Inline Deployment

The same methodology can be used to migrate IWAN BRs. First migrate the IWAN BR that is not an IWAN Master Controller (MC). Once the first router is migrated to SD-WAN, migrate the MC/BR router. Note that if LAN connectivity is L2, that uses Hot Standby Router Protocol (HSRP) on IWAN, also consider the migration from HSRP to Virtual Router Redundancy Protocol (VRRP) before migrating BRs, because SD-WAN edge routers only supports VRRP.

In some IWAN deployments, static prefixes are used to advertise to Performance Routing version 3 (PfRv3). Once the site is migrated to SD-WAN, remove the related static prefixes from the datacenter site to maintain the Performance Routing (PfR) operation only on the prefixes remaining on IWAN.

4.2.3 Parallel Branch Migration

In certain scenarios at the branch location, a CE router needs to be maintained with only MPLS connectivity along with the SD-WAN Edge router. SD-WAN Edge router provides direct connectivity to the internet. This parallel migration is typically used when either MPLS connectivity is needed to support applications that need to use MPLS circuits directly

between sites; or when there are services like WAAS or UC connected to the MPLS CE. Note that after parallel migration, there is no redundancy for transport connectivity.

Parallel migration is not recommended for IWAN migration as it adds to the complexity. The recommended method for branch IWAN migration is explained in section 4.2.2.

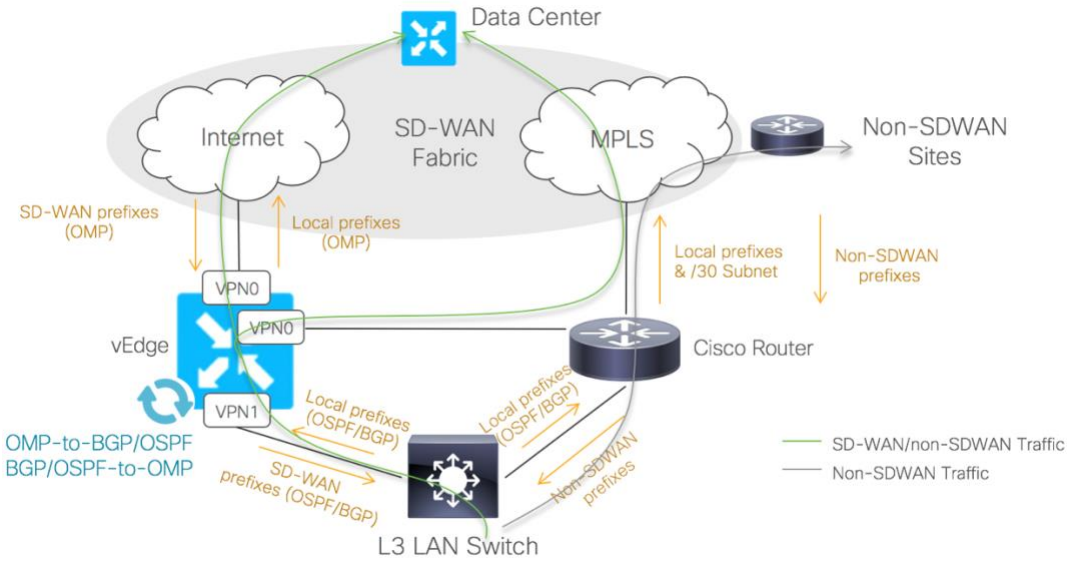


Figure 16: Parallel Branch Deployment – Underlay/Overlay

In the first scenario, where an MPLS CE needs to be maintained for direct application connectivity over MPLS to non-SDWAN sites, you can perform parallel migration and establish underlay/overlay connectivity at the branch. The CE connected to Internet is migrated to SD-WAN edge and the MPLS connectivity is extended from MPLS CE to the SD-WAN edge router.

The CE router learns the non-SDWAN site prefixes and the edge router learns SD-WAN site prefixes. If core/distribution layer is layer 3 on LAN, then configure the route filters on the L3 device to stop re-advertisements from SD-WAN sites to non SD-WAN sites and vice versa.

The table below explains the routing design at the branch

		SD-WAN Edge Router
Physical/L3 Connectivity	WAN VPN0	- Direct Internet connectivity - MPLS connectivity via /30 physical link to MPLS CE router
	LAN Service VPNs	- Connect to L3/L2 LAN switch
WAN Advertisements	IN	- SD-WAN prefixes, aggregate routes and default GW over Internet and MPLS connections via OMP

	OUT	- Local LAN prefixes into OMP to SD-WAN fabric
LAN Advertisements	IN	- With L3 LAN side, LAN prefixes from LAN switch. - With L2 VRRP, no advertisements.
	OUT	- With L3 LAN, advertise SD-WAN prefixes, aggregate routes and default GW from OMP to LAN - With L2 LAN, no advertisement needed. VRRP routers are the GW.

Table 17: Branch SD-WAN Edge Router in Parallel Deployment – Underlay/Overlay

		Branch CE Router
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS - Extending MPLS to SD-WAN Edge over /30 physical link
	LAN	- Connect to LAN switch
WAN Advertisements	IN	- Non SD-WAN site advertisements
	OUT	- Local LAN Branch prefixes and /30 subnet that is between CE router and Edge router towards MPLS
LAN Advertisements	IN	- With L3 LAN side, LAN prefixes from LAN - With L2 VRRP, no advertisements from LAN
	OUT	- With L3 LAN, advertise Non-SDWAN prefixes that need direct MPLS access to LAN switch - With L2 LAN, no advertisement needed. VRRP routers are the GW

Table 18: Branch CE Router Connectivity in Parallel Deployment – Underlay/Overlay

		Branch LAN Switch
Physical/L3 Connectivity	WAN	- Connect to SD-WAN Edge router and CE router
	LAN	- Connect to Distribution/Access switches
WAN Advertisements	IN	- From SD-WAN Edge router, receive SD-WAN prefixes, aggregate routes and default GW route of DC - From CE Router, receive prefixes of non SD-WAN sites that require MPLS transport
	OUT	- LAN prefixes to SD-WAN Edge router and CE router
LAN Advertisements	IN	- LAN prefixes from LAN
	OUT	- Default GW towards LAN side/or prefixes as per Branch design requirements

Table 19: Branch LAN Switch Connectivity in Parallel Deployment – Underlay/Overlay

In the second scenario, where parallel migration can be performed, an MPLS CE is connected to services, like WAAS or UC.

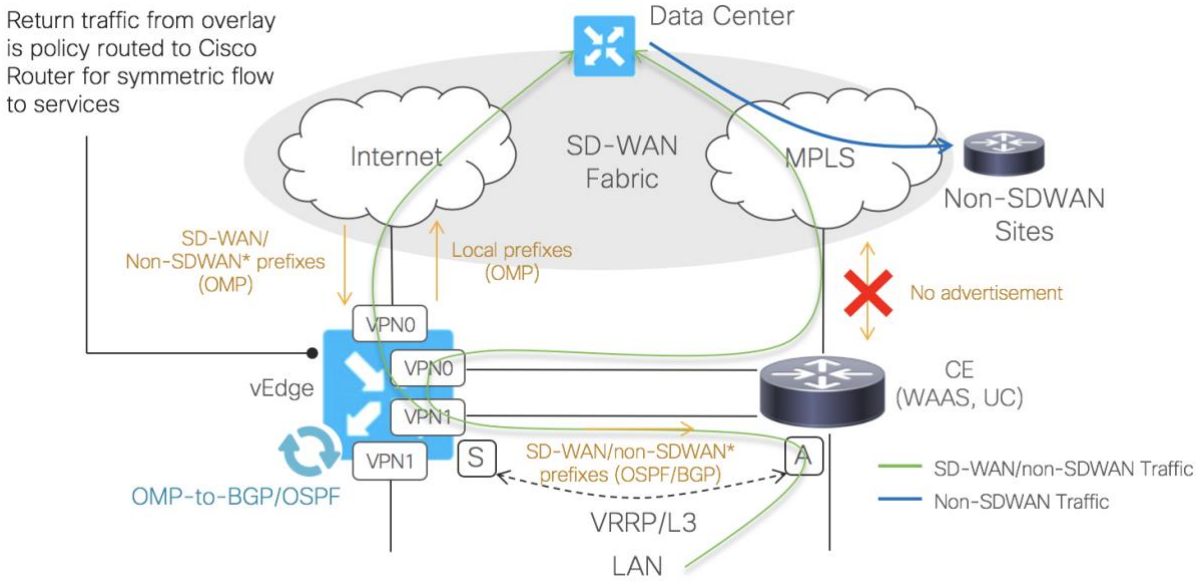


Figure 17: Parallel Branch Deployment – with Services

The CE connected to Internet transport is migrated to SD-WAN edge and the MPLS connectivity is extended from MPLS CE to the SD-WAN edge router. Additionally, the LAN side service VPN connectivity is also established from the SD-WAN edge router to MPLS CE router. Both, MPLS and service LAN side connectivity can be established on an edge router using either physical interfaces or two subinterfaces if the ports are limited.

With L2 LAN, CE router is active for VRRP for traffic that requires services. To maintain the symmetry of routes to the services connected to CE, traffic from SD-WAN overlay is routed to CE router using policies.

To prevent making this branch site a transient site, no prefixes are advertised or learned on MPLS transport from CE. The traffic destined to non SD-WAN site transits through the DC.

The table below explains the routing design at the branch.

		SD-WAN Edge Router
Physical/L3 Connectivity	WAN VPN0	- Direct Internet connectivity - MPLS connectivity via /30 physical link to MPLS CE router
	LAN Service VPNs	- Service Side VPN connectivity with MPLS CE - Connect to L2 LAN switch

WAN Advertisements	IN	- SD-WAN prefixes, aggregate route and default GW over Internet and MPLS connections via OMP
	OUT	- Local LAN prefixes into OMP
LAN Advertisements	IN	- With L2 VRRP, no advertisements.
	OUT	- With L2 LAN, no advertisement needed. VRRP routers are the GW

Table 20: Branch SD-WAN Edge Router in Parallel Deployment – with Services

		Branch CE Router
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS - Extending MPLS to SD-WAN Edge router over /30 physical link
	LAN	- Connected to Service side of SD-WAN Edge router - Connect to LAN switch
WAN Advertisements	IN	- No advertisements
	OUT	- /30 subnet that is between CE router and Edge router towards MPLS
LAN Advertisements	IN	- With L2 VRRP, no advertisements from LAN
	OUT	- Extending L2 LAN connectivity to Edge Service side. - With L2 LAN, no advertisement needed. VRRP is the GW

Table 21: Branch CE Router Connectivity in Parallel Deployment – with Services

5 Migration Case Study

This section describes the migration procedures from legacy WAN and IWAN topology to Cisco SD-WAN using a series of diagrams. Every customer may have specific and unique configurations, however, this guide allows users to view the migration end-to-end and create their own migration plan that best suits their environment.

For detailed migration configurations, refer to the Cisco SD-WAN Migration Configuration Guide.

5.1 Getting Ready

The SD-WAN solution separates control and data planes, which makes the migration from WAN and IWAN to SD-WAN simple. In addition, the Cisco SD-WAN solution is highly scalable and controllers can deployed in cloud hosting or on-prem. This allows the solution to be deployed in parallel to the existing environment, especially IWAN, without disrupting the current network and traffic in any way. Furthermore, all the configurations and policy aspects of the solution are configured on the controller elements (vManage and vSmart specifically).

Before migration, review the IOS-XE configuration for supportability with the SD-WAN image. See section 3.6.6 in this document for details. The table below shows direct comparison of some features of IWAN and SD-WAN.

	IWAN	SD-WAN
Control Elements	APIC-EM	vManage, vSmart, vBond
Edge Routers	PfR Border routers (ISRs, ASRs, CSRv)	vEdge Cloud, vEdge Appliance, SD-WAN capable Cisco Routers (ISRs, ASRs, CSRv)
Control Protocol	Performance Routing (PfR)	Overlay Management Protocol (OMP)
Data Tunnel Security	DMVPN - IPSEC	IPSEC/GRE
Data Plane Monitoring	PfR Smart Probes	BFD
Software	IOS-XE 3.16S and higher	IOS-XE SD-WAN 16.9.1 and higher

Table 22: IWAN & SD-WAN Architectures

5.2 Topology

The topology below shows two datacenters and four branches.

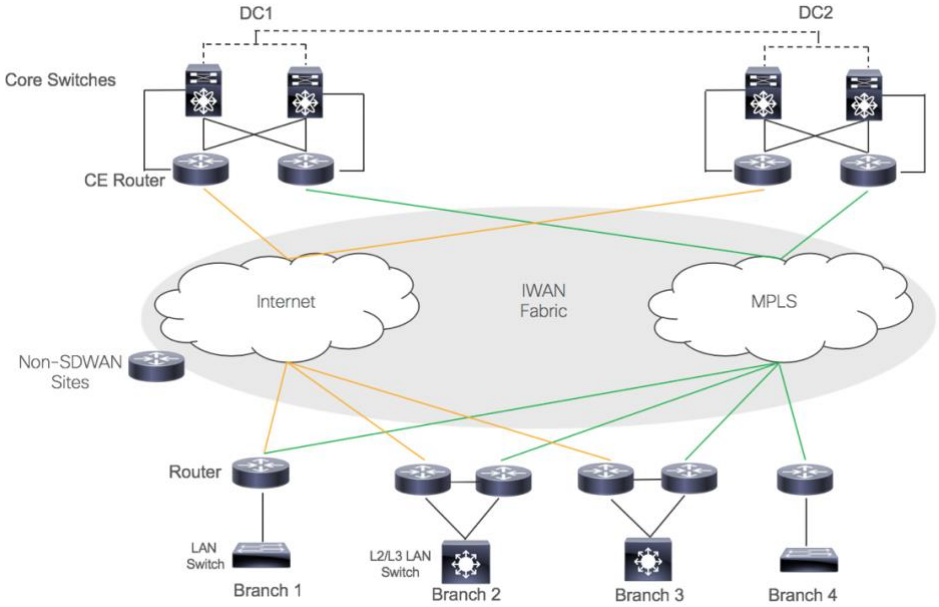


Figure 18: Legacy WAN Topology

In case of IWAN deployments, the PfR BRs are connected behind the CE routers in the datacenters, where MPLS and Internet circuits terminate on the CE routers. The domain controller and master controller is connected to the core routers as shown in figure 19. The BRs at branches are all directly connected to the transports.

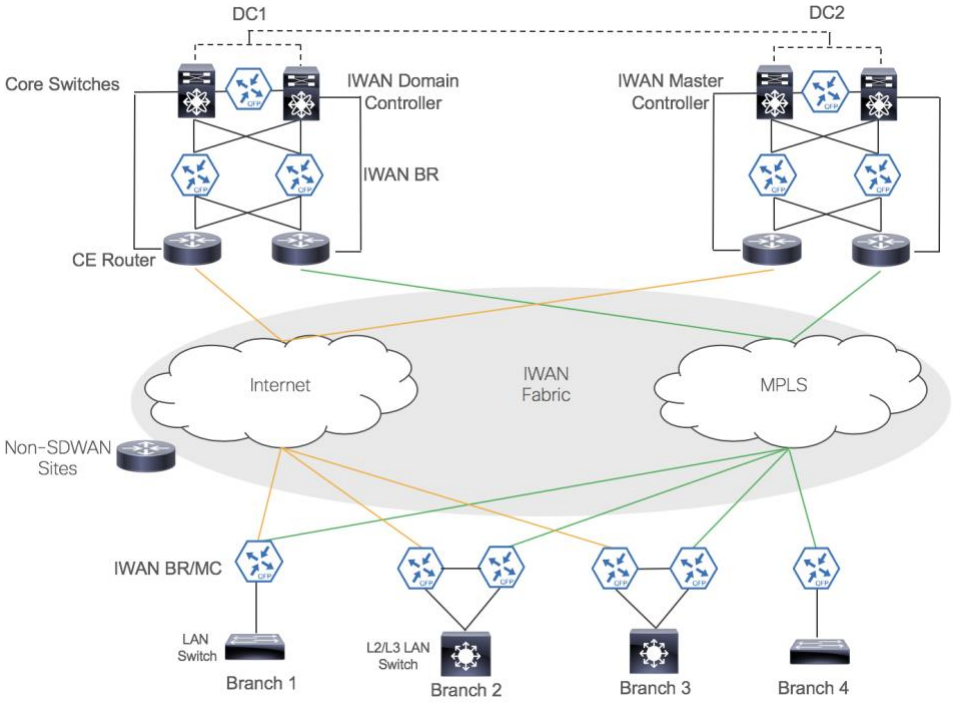


Figure 19: IWAN Topology

5.3 Migration Steps

Follow these steps for migration.

1. Setup Smart account
2. Deploy the controller
3. Migrate the Datacenter
4. Migrate branches
5. Phase out IWAN DC BRs.

5.3.1 Smart Account

Follow the steps explained in section 3.1.1 in this document to setup the Smart Account.

5.3.2 Controller Deployment

Follow the steps in section 3.1.2 to deploy the controllers.

5.3.3 Datacenters Migration

For migrating from legacy WAN to SD-WAN, Cisco recommends deploying the SD-WAN routers between the CE and core routers as explained in section 4.1.1.

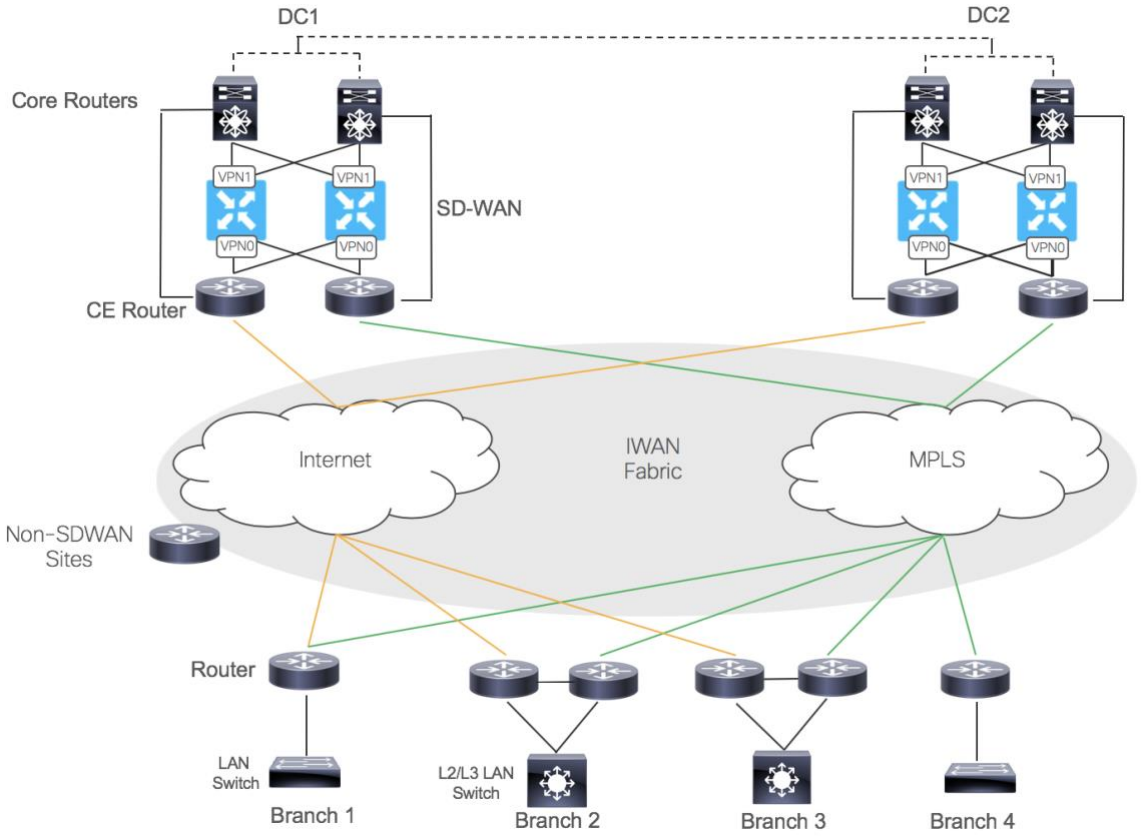


Figure 20: Migrating Legacy DCs - SD-WAN Edge Routers Placement

The table below explains the routing design at the datacenter.

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	- Connect VPN0 WAN side to both CE routers - Use /30 on MPLS CE-to-SDWAN router link - Internet link gets IP through DHCP
	LAN Service VPNs	- Connect to L3 LAN Core switches to Service-side VPN interfaces.
WAN Advertisements	IN	- SD-WAN Prefixes from SD-WAN sites over Internet and MPLS connections through OMP
	OUT	- Through OMP advertise Datacenter LAN prefixes, default GW, aggregate routes for non-SD-WAN prefixes to SD-WAN sites
LAN Advertisements	IN	- Local LAN prefixes, default GW and aggregate routes for Non SD-WAN site prefixes from L3 LAN Core Switch.
	OUT	- SD-WAN sites prefixes to L3 LAN Core Switch

Table 23: SD-WAN Edge Routers Connectivity Behind CE Routers

		CE Routers
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS and Internet - Extend MPLS to SD-WAN Edge over /30 physical link - Extend Internet to SD-WAN Edge over L2/L3 DHCP relay
	LAN	- Connect to L3 LAN Core switch
WAN Advertisements	IN	- Non SD-WAN prefixes from internet and MPLS WAN links
	OUT	- DC, Non-SD-WAN, and SD-WAN site prefixes to Non-SD-WAN sites
LAN Advertisements	IN	- DC and SD-WAN site prefixes from L3 LAN Core Switch
	OUT	- Non SD-WAN site prefixes to L3 LAN Core Switch

Table 24: CE Routers Connectivity in front of SD-WAN Edge Routers

		Datacenter Core Routers
Physical/L3 Connectivity	WAN	- Connect to CE routers - Connect to SD-WAN Edge routers
	LAN	- Connect to L3/L2 Distribution/Access switches as per DC design

WAN Advertisements	IN	- Non SD-WAN prefixes from CE routers (MPLS/Internet) - SD-WAN prefixes from SD-WAN Edge routers
	OUT	- DC, Non-SD-WAN, and SD-WAN site prefixes to CE routers - DC prefixes, default GW, aggregate routes of non-SDWAN sites to SD-WAN Edge routers
LAN Advertisements	IN	- DC Prefixes from LAN network
	OUT	- Non SD-WAN and SD-WAN site prefixes to LAN as per DC LAN design requirement

Table 25: Datacenter Core Routers Connectivity

Typically, the datacenter sites are connected through a backbone link and prefixes of DC site are advertised to the other DC site. To avoid loops and non-optimal routing, the core routers should be configured to explicitly filter prefixes learned over backbone from getting advertised to the SD-WAN edge routers.

In IWAN migrations, the recommended migration strategy for datacenters is to deploy SD-WAN edge routers in parallel to IWAN BRs behind CEs as explained in section 4.1.2 of the document.

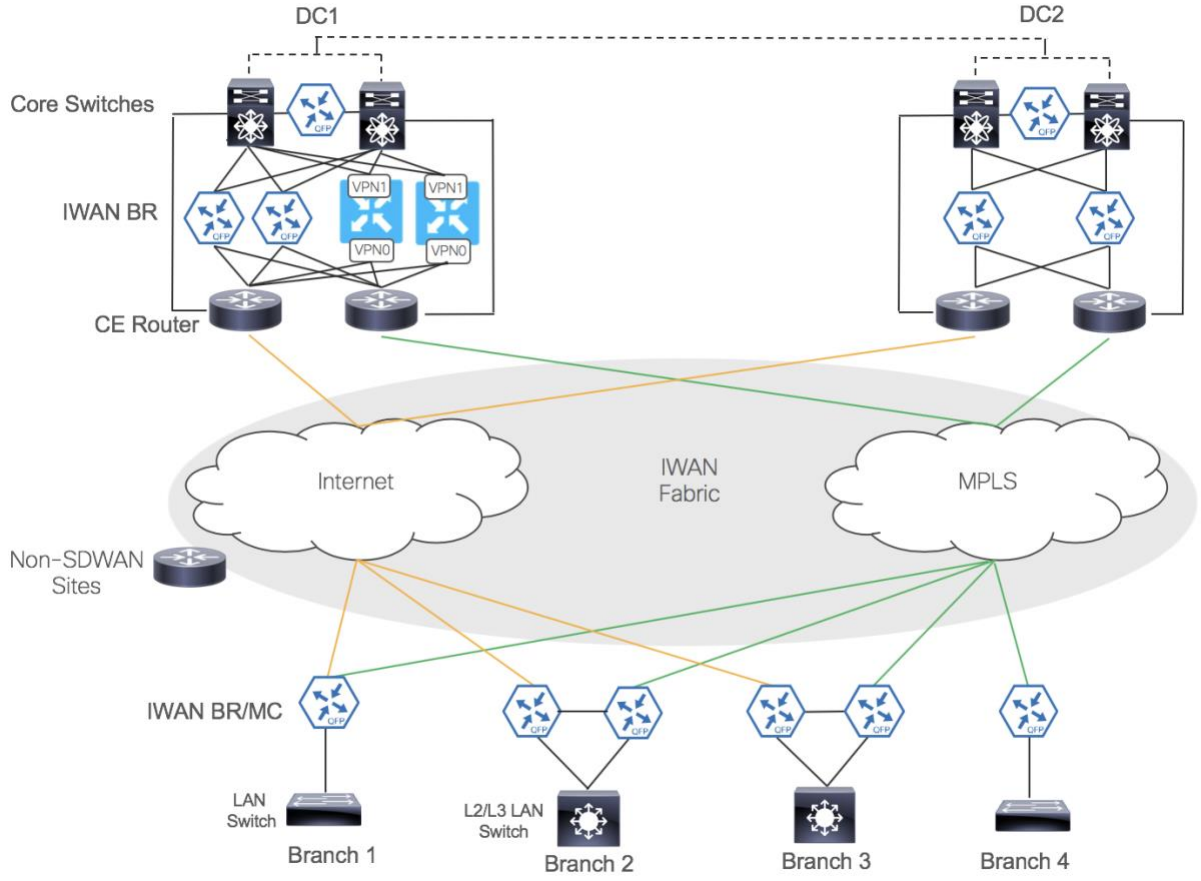


Figure 21: Migrating IWAN DC1 - SD-WAN Edge Routers Placement

During the migration, the goal is to maintain the routing between legacy, IWAN and SD-WAN branches.

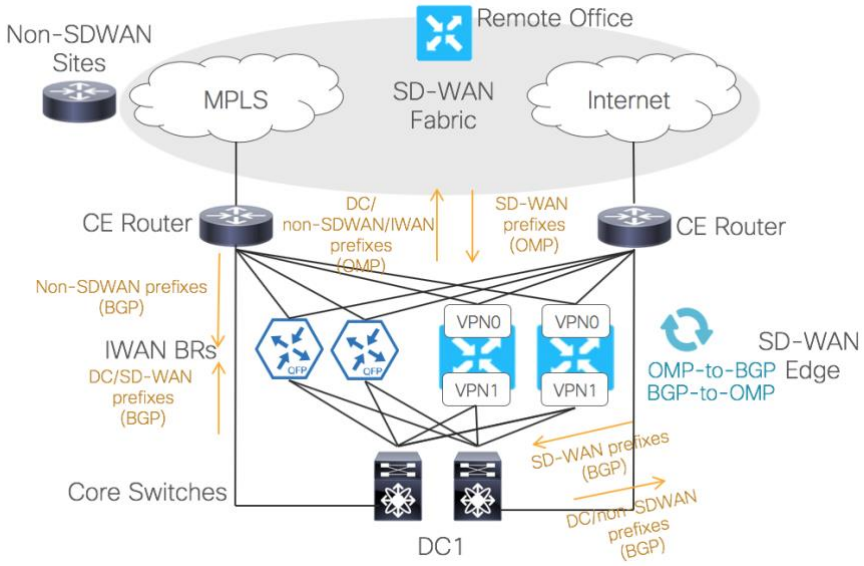


Figure 22: Migrating IWAN DC1 - Routing During Migration

The table below explains the routing design at the datacenter for IWAN migration.

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	- Connect VPN0 WAN side to both CE routers - Use /30 on MPLS CE-to-SDWAN router link - Internet link gets IP through DHCP
	LAN Service VPNs	- Connect to L3 LAN Core switches to Service-side VPN interfaces
WAN Advertisements	IN	- SD-WAN Prefixes from SD-WAN sites over Internet and MPLS connections through OMP
	OUT	- Through OMP advertise Datacenter LAN prefixes, default GW, aggregate routes for non-SD-WAN and IWAN prefixes to SD-WAN sites
LAN Advertisements	IN	- Local LAN prefixes, default GW and aggregate routes for Non SD-WAN and IWAN prefixes from L3 LAN Core Switch.
	OUT	- SD-WAN sites prefixes to L3 LAN Core Switch

Table 26: IWAN: DC1 SD-WAN Edge Routers Connectivity and Routing

		CE Routers
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS and Internet - Extend MPLS to SD-WAN Edge over /30 physical link - Extend Internet to SD-WAN Edge over L2/L3 DHCP relay - No changes to connections to IWAN BRs
	LAN	- Connect to L3 LAN Core switch
WAN Advertisements	IN	- Non SD-WAN prefixes from internet and MPLS WAN links
	OUT	- DC, Non SD-WAN, IWAN and SD-WAN site prefixes to Non-SD-WAN sites
LAN Advertisements	IN	- DC, IWAN and SD-WAN site prefixes from L3 LAN Core Switch
	OUT	- Non SD-WAN site prefixes to L3 LAN Core Switch

Table 27: IWAN: DC1 CE Routers Connectivity and Routing

		IWAN BRs
Physical/L3 Connectivity	WAN	- No change in connections to CEs
	LAN	- No change in connections to Core routers
WAN Advertisements	IN	- IWAN site prefixes from Internet and MPLS connections
	OUT	- Advertise Datacenter LAN, SD-WAN and non-SD-WAN prefixes to IWAN sites
LAN Advertisements	IN	- Local LAN, SD-WAN and Non SD-WAN site prefixes from L3 LAN Core Switch
	OUT	- IWAN site prefixes to L3 LAN Core Switch

Table 28: IWAN: DC1 IWAN Routers Connectivity and Routing

		Datacenter Core Routers
Physical/L3 Connectivity	WAN	- No change to connections to CE routers - No change to connections to IWAN routers - Connect to SD-WAN Edge routers
	LAN	- Connect to L3/L2 Distribution/Access switches as per DC design
WAN Advertisements	IN	- Non SD-WAN prefixes from CE routers (MPLS/Internet) - IWAN prefixes from IWAN BRs - SD-WAN prefixes from SD-WAN Edge routers

WAN Advertisements	OUT	- DC, Non-SD-WAN, and SD-WAN site prefixes to CE routers - DC, Non-SD-WAN, and SD-WAN site prefixes to IWAN routers - Local LAN prefixes, default GW and aggregate routes for Non SD-WAN and IWAN prefixes to SD-WAN routers
LAN Advertisements	IN	- DC Prefixes from LAN network
	OUT	- Non SD-WAN, IWAN and SD-WAN site prefixes to LAN as per DC LAN design requirement

Table 29: IWAN: Datacenter Core Routers Connectivity

Similarly, migrate the DC2 site ensuring that the routing between IWAN, SD-WAN, and non SD-WAN sites is still maintained as required.

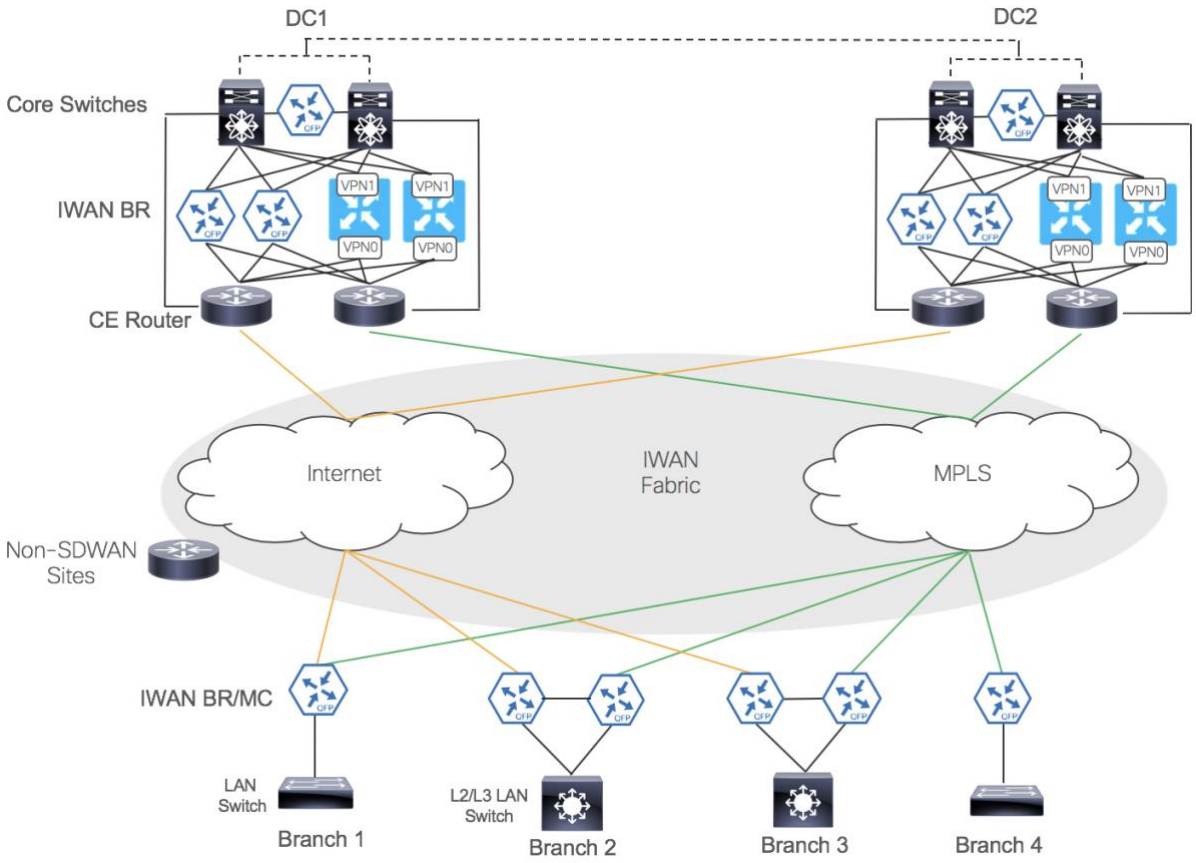


Figure 23: Migrating IWAN DC2 - Routing During Migration

Modify the routing on core routers for the routes learned from the backbone link between the two datacenters.

5.3.4 Branch 1 Migration

Single router branch requires a downtime window for the migration. Follow the detailed steps of migration as explained in section 4.2.1.

The diagrams below show that legacy WAN/IWAN BR is replaced at Branch1 with SD-WAN edge router.

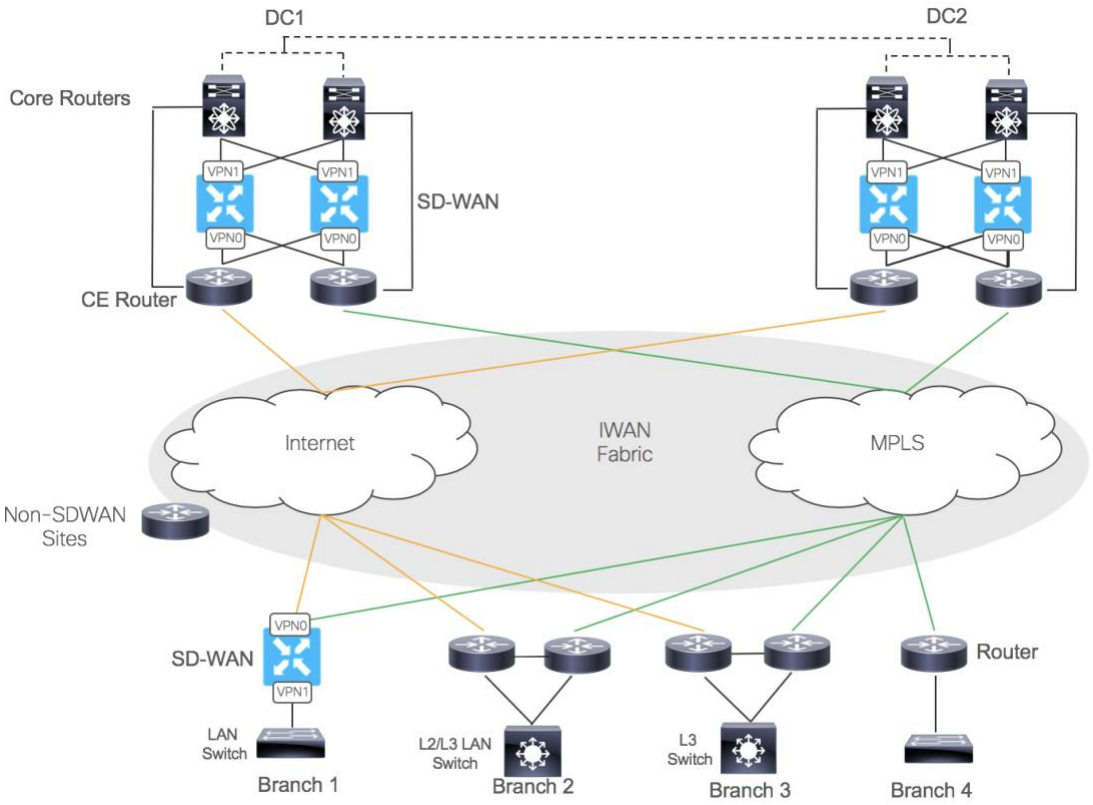


Figure 24: Migrating Legacy Branch 1 to SD-WAN

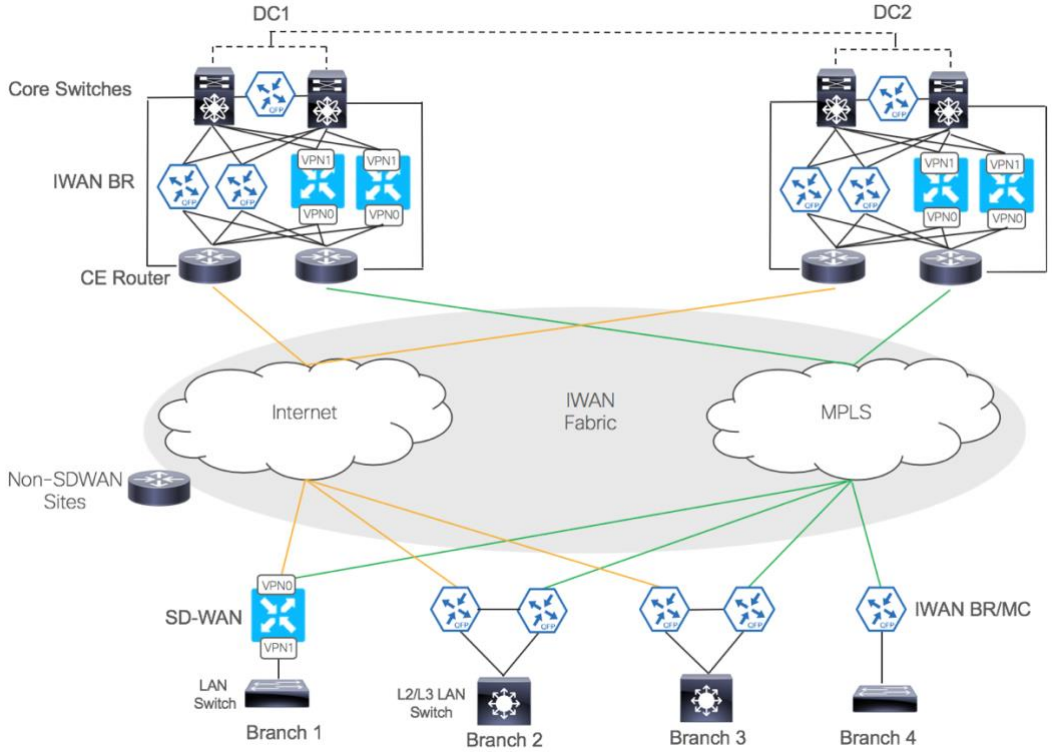


Figure 25: Migrating IWAN Branch 1 to SD-WAN

Once the branch is migrated to SD-WAN, the exchange of prefixes at the site is for the local prefixes over OMP with the SD-WAN control plane. The branch communicates with non SD-WAN sites and IWAN sites through the DCs.

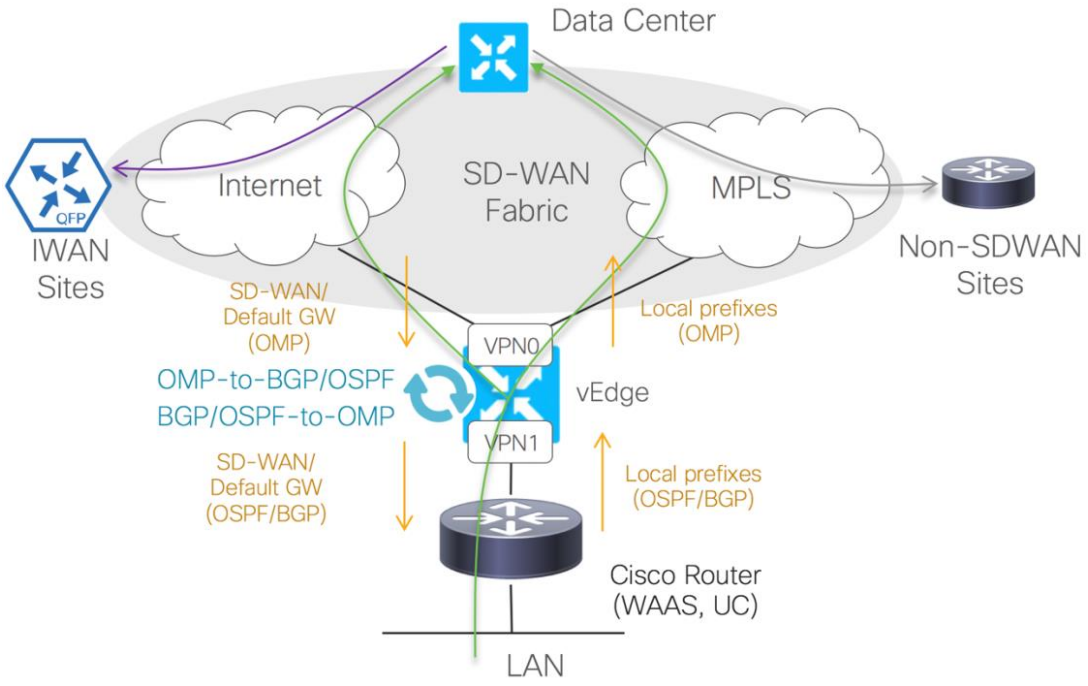


Figure 26: Migrating Branch 1 – Routing

The table below explains the routing design at the branch.

		SD-WAN Edge Router
Physical/L3 Connectivity	WAN VPN 0	- MPLS and Internet connections terminate on SD-WAN Edge router on interfaces under VPN 0
	LAN Service VPNs	- Connect to LAN switches in Service VPNs. LAN design will dictate if subinterfaces are needed.
WAN Advertisements VPN 0	IN	- SD-WAN prefixes and default GW from Datacenter from OMP session over Internet and MPLS connections
	OUT	- Redistribute local LAN prefixes into OMP
LAN Advertisements Service VPNs	IN	- Local LAN prefixes – SD-WAN Edge typically is the GW
	OUT	- With L3 connection on LAN side – advertise prefixes learned through OMP to LAN - With L2 connection on LAN Side – no advertisements are needed as SD-WAN Edge is the GW for the VLAN/VPN Segments

Table 30: Branch 1 SD-WAN Routers Connectivity and Routing

5.3.5 Branch 2 Migration

It is recommended that you migrate the sites with redundant WAN routers in the same cutover. Follow the steps discussed in section 4.2.2. to reduce the downtime of the maintenance window.

In the diagram below, branch 2 has redundant routers and each router connects to each of the transport circuits on the WAN side. The LAN side of the SD-WAN edge routers can be configured for Layer 2 VRRP redundancy or can be connected to an L3 switch on LAN side. The IWAN migration for Branch 2 should be carried out in a similar way.

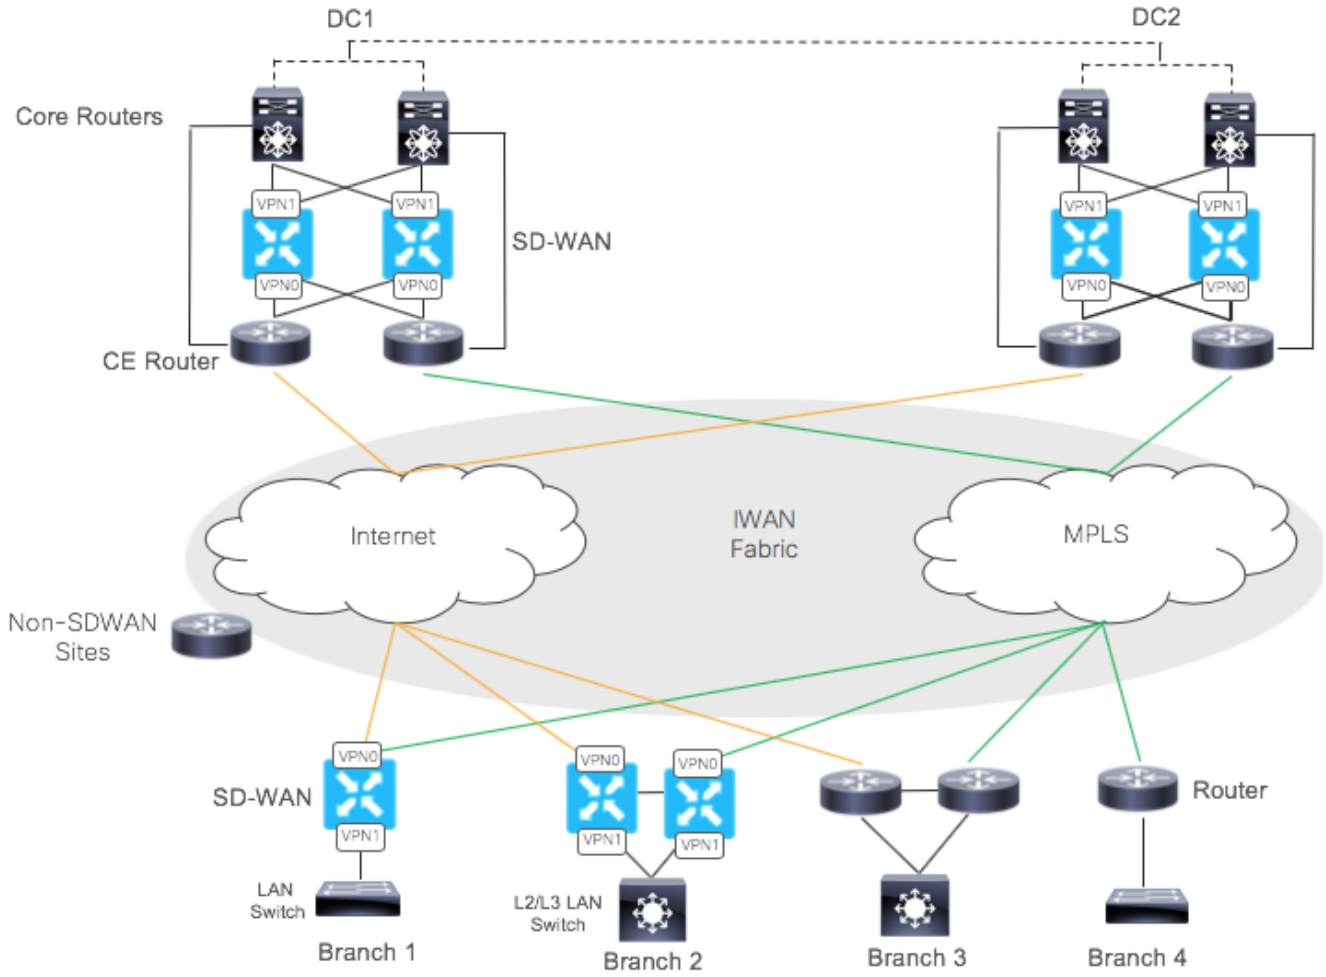


Figure 27: Migrating Legacy WAN Branch 2

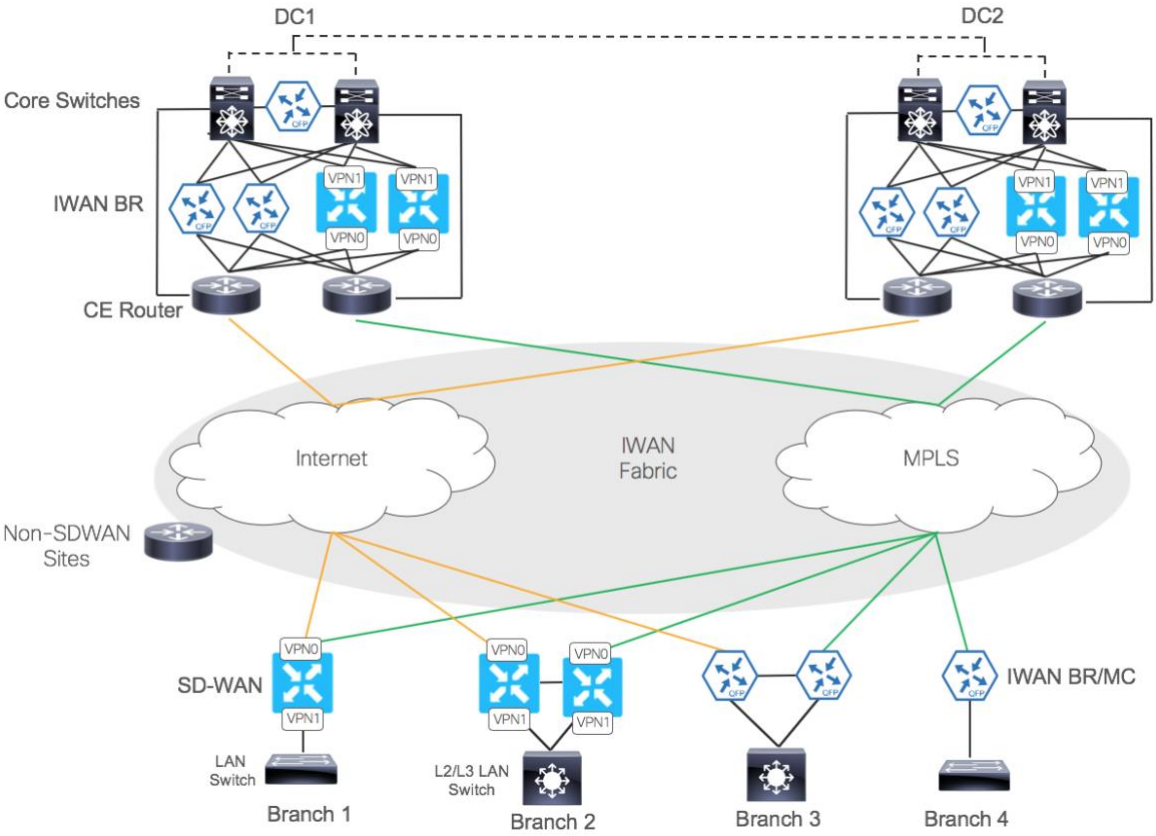


Figure 28: Migrating IWAN Branch 2

Since each SD-WAN edge router terminates either MPLS or Internet transport link, you can extend the WAN connectivity using the TLOC-Extension feature between the edge routers, which extends the WAN connectivity to the other edge device.

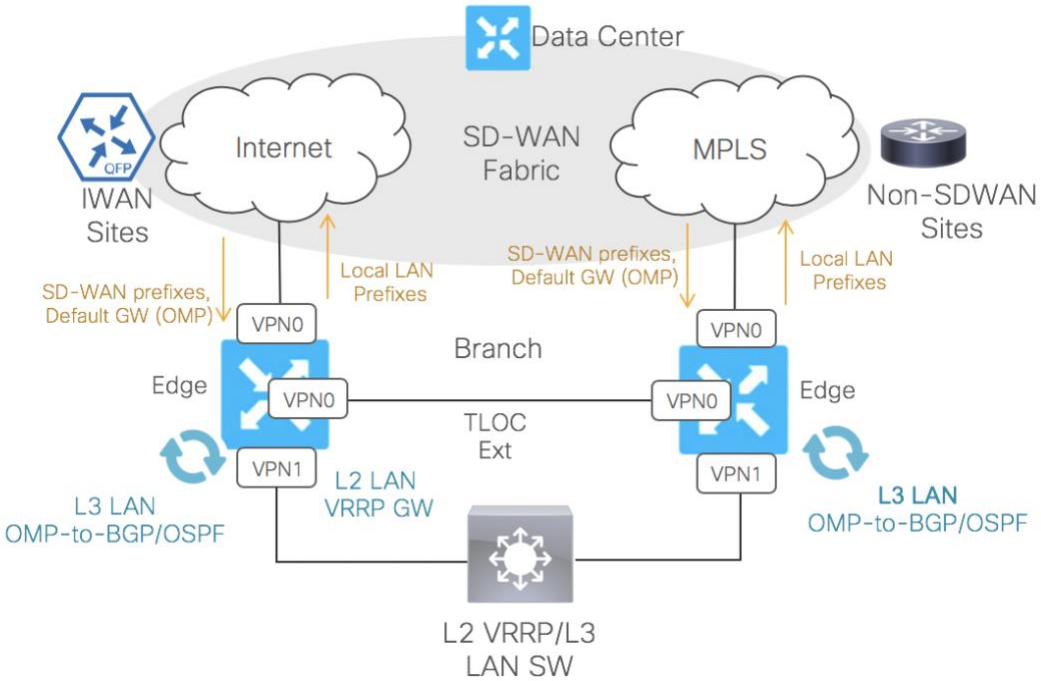


Figure 29: Migrating Branch 2 – Routing

The table below explains the routing design at the branch.

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	- One Edge router connects to the Internet - One Edge router connects to MPLS - Using TLOC-Extension MPLS connectivity is extended to Edge1 - Using TLOC-Extension Internet connectivity is extended to Edge2
	LAN Service VPNs	- Connect to L2/L3 LAN Core switches in Service VPNs
WAN Advertisements	IN	- SD-WAN Prefixes, default GW and aggregate routers from SD-WAN sites from Internet and MPLS connections over OMP
	OUT	- Local LAN prefixes over OMP
LAN Advertisements	IN	- With L3 LAN side, LAN prefixes from LAN switch - With L2 VRRP, no advertisements
	OUT	- With L3 LAN advertise SD-WAN prefixes to LAN switch - With L2 LAN, no advertisement needed. VRRP routers are the GW

Table 31: Branch 2 SD-WAN Routers Connectivity and Routing

5.3.6 Legacy WAN Branch 3 Migration

In the diagram below, the requirement for branch 3 is to deploy SD-WAN edge routers in parallel to the CE router connected to MPLS. The section 4.2.3 explains the parallel migration in detail. Primarily, the underlay/overlay routing and symmetry of the flows is maintained after the migration.

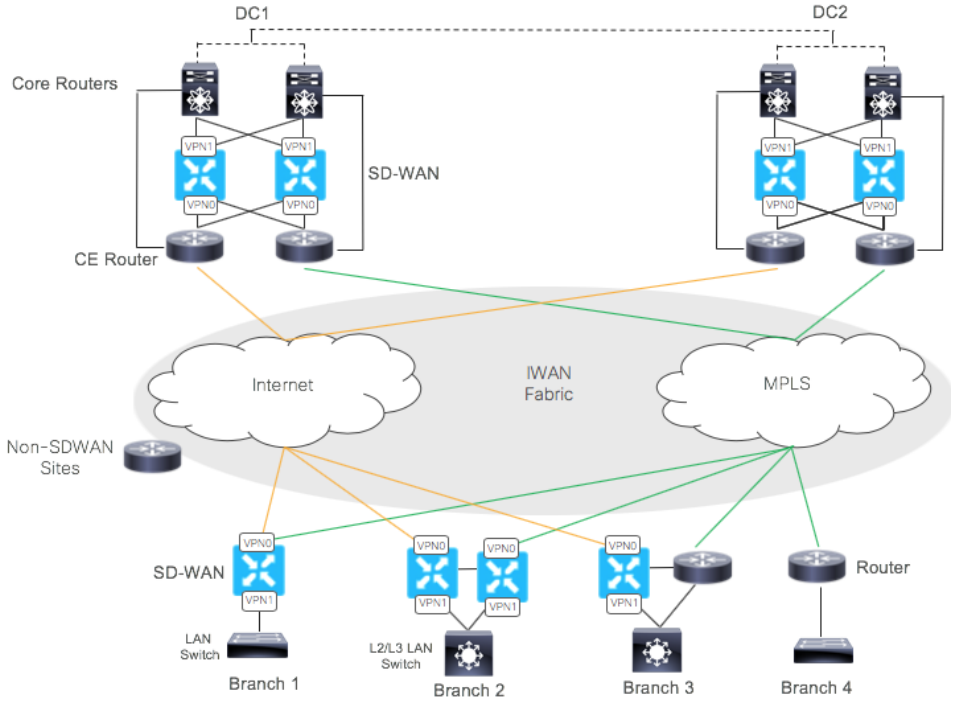


Figure 30: Migrating Legacy Branch 3

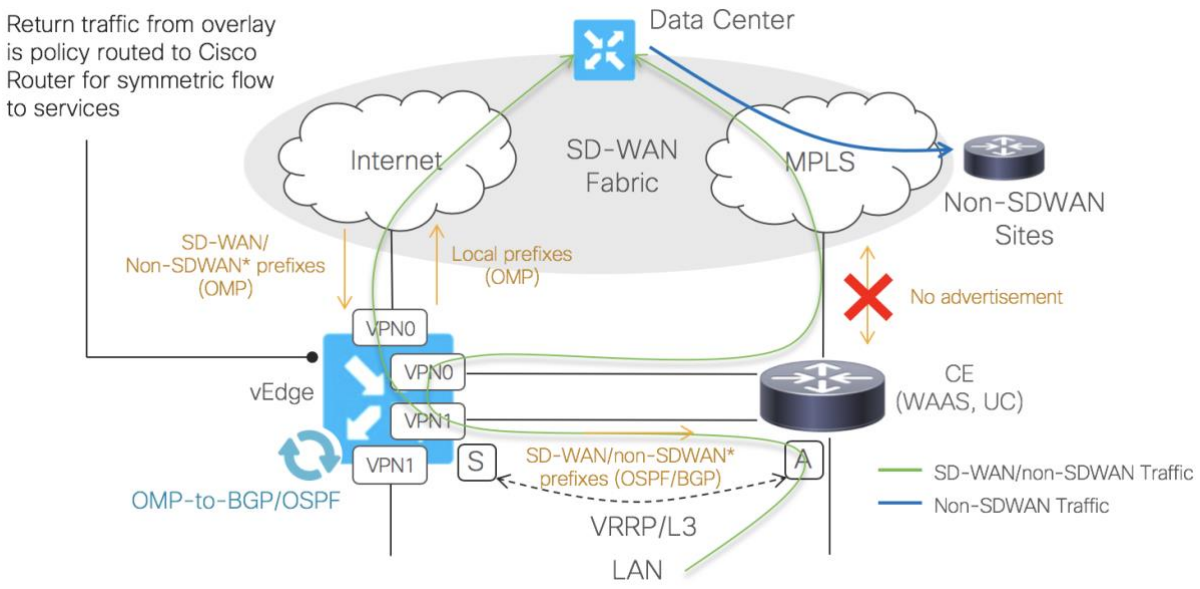


Figure 31: Parallel Branch Migration – Routing

The table below explains the routing design at the branch.

		SD-WAN Edge Router
Physical/L3 Connectivity	WAN VPN0	- Direct Internet connectivity - MPLS connectivity via /30 physical link to MPLS CE router
	LAN Service VPNs	- Service Side VPN connectivity with MPLS CE - Connect to L2 LAN switch
WAN Advertisements	IN	- SD-WAN prefixes, aggregate route and default GW over Internet and MPLS connections via OMP
	OUT	- Local LAN prefixes into OMP
LAN Advertisements	IN	- With L2 VRRP, no advertisements.
	OUT	- With L2 LAN, no advertisement needed. VRRP routers are the GW.

Table 32: Branch SD-WAN Edge Router in Parallel Deployment

		Branch CE Router
Physical/L3 Connectivity	WAN	- Directly connecting to MPLS - Extending MPLS to SD-WAN Edge router over /30 physical link
	LAN	- Connected to Service side of SD-WAN Edge router - Connect to LAN switch
WAN Advertisements	IN	- No advertisements
	OUT	- /30 subnet that is between CE router and Edge router towards MPLS
LAN Advertisements	IN	- With L2 VRRP, no advertisements from LAN
	OUT	- Extending L2 LAN connectivity to Edge Service side. - With L2 LAN, no advertisement needed. VRRP is the GW

Table 33: Branch CE Router Connectivity in Parallel Deployment

		Branch LAN Switch
Physical/L3 Connectivity	WAN	- Connect to SD-WAN Edge and CE router
	LAN	- Connect to Distribution/Access switches

WAN Advertisements	IN	- From SD-WAN Edge, receive SD-WAN prefixes,. Aggregate routes and default GW route of DC - From CE Router, receive prefixes of non-SDWAN sites that require MPLS transport
	OUT	- LAN prefixes to SD-WAN Edge and CE
LAN Advertisements	IN	- LAN prefixes from LAN
	OUT	- Default GW towards LAN side/or prefixes as per Branch design requirements

Table 34: Branch Core/Distribution Switch Connectivity in Parallel Deployment

As explained in section 4.2.3, the parallel migration is not recommended for IWAN migration. The IWAN branch 3 should be migrated similar to branch 2 as explained in section 5.3.5.

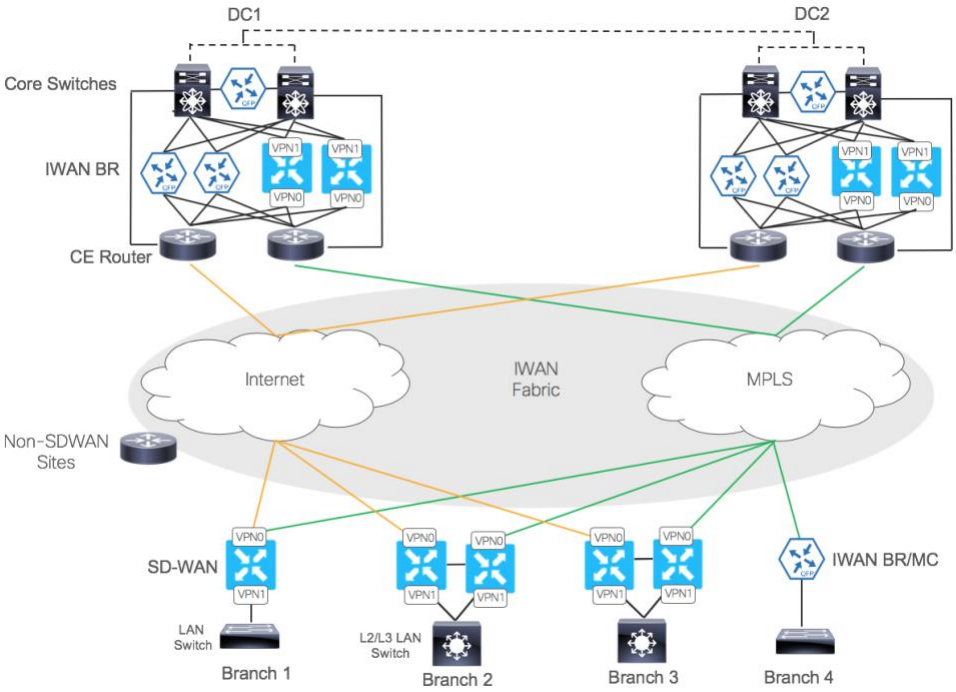


Figure 32: Migrating IWAN Branch 3

5.3.7 Branch 4 Migration

The steps to migrate branch 4 are the same as branch 1. See section 5.3.4 for more details. The only difference is that the edge router is terminating a single WAN link of Internet instead of MPLS.

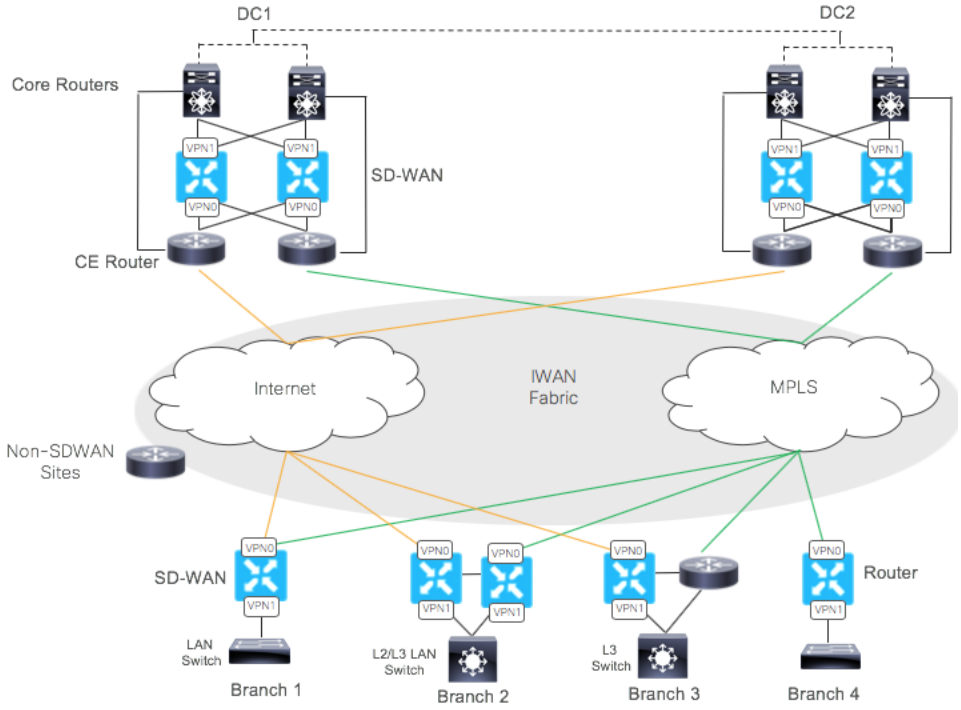


Figure 33: Migrating Legacy WAN Branch 4

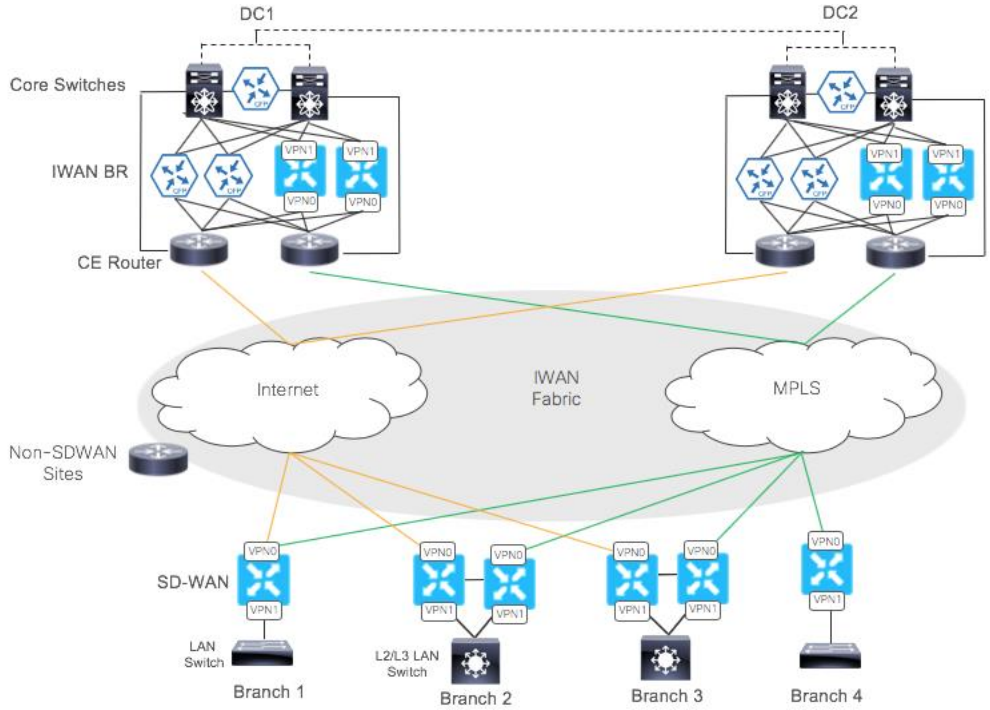


Figure 34: Migrating IWAN Branch 4

This completes the migration from legacy WAN to SD-WAN.

5.3.8 Phase out DC IWAN BRs

In an IWAN migration, once all IWAN branches are successfully migrated to SD-WAN edge routers, you can remove the DC IWAN BRs, Domain Controller and Master Controller from the network.

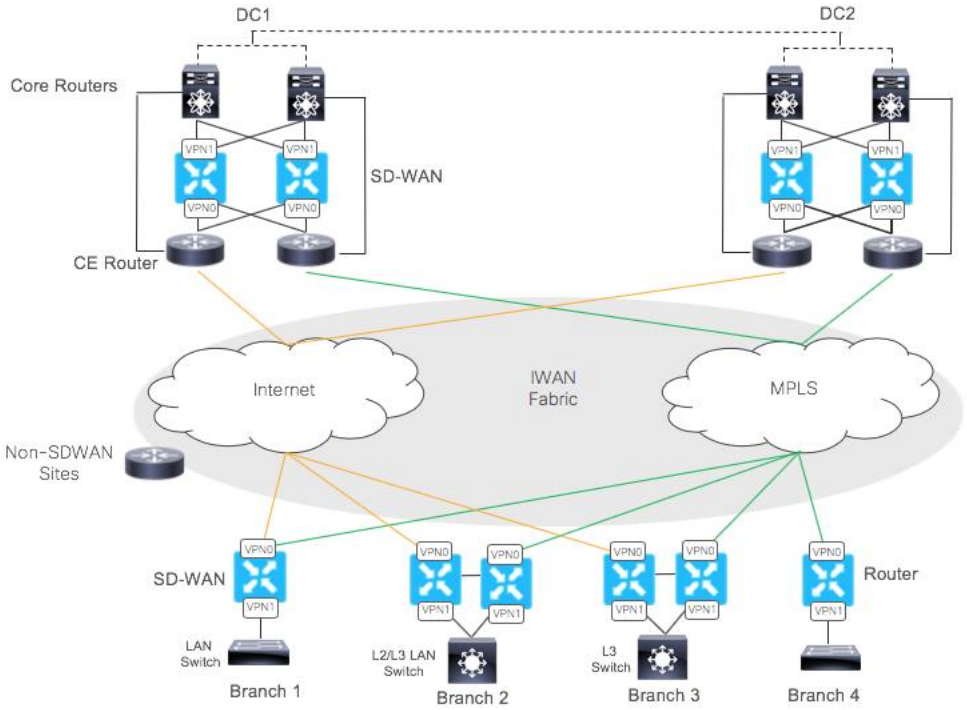


Figure 35: IWAN DC BRs and Controllers Removed