



# IWAN to Cisco SD-WAN

## Migration Guide

**A Customer Journey**

November 25, 2020

**Version 1.0**

---

# Contents

Contents .....	2
List of Figures and Tables .....	5
About this Guide.....	6
<b>History</b> .....	6
<b>Review</b> .....	6
1 Introduction .....	7
<b>1.1 Audience</b> .....	7
<b>1.2 Document Scope</b> .....	7
<b>1.3 Assumptions and Considerations</b> .....	7
<b>1.4 Related Documents</b> .....	7
2 IWAN and Cisco SD-WAN Overview .....	9
<b>2.1 Intelligent WAN (IWAN) Overview</b> .....	9
2.1.1 Transport-independent design .....	9
2.1.2 Intelligent path control .....	9
2.1.3 Application Optimization.....	9
2.1.4 Secure connectivity.....	10
<b>2.2 Cisco SD-WAN Solution Overview</b> .....	10
3 Migration Planning .....	13
<b>3.1 Identify sites/regions for migration</b> .....	14
<b>3.2 Identify Use Cases and Feature/Configuration Analysis</b> .....	15
3.2.1 Identify Use Cases.....	15
3.2.2 WAAS Use Case with Cisco SD-WAN .....	16
3.2.3 Feature and Configuration Analysis .....	17
<b>3.3 Current Inventory, Bandwidth &amp; Platform requirement</b> .....	19
3.3.1 Current Inventory .....	19
3.3.2 Identify Bandwidth and Platforms .....	19
3.3.3 Identify Licensing.....	20
<b>3.4 Identify Controller Deployment Model and Requirements</b> .....	20
<b>3.5 Detailed Site Migration Planning</b> .....	21
4 Migration Deployment Steps .....	22
<b>4.1 Cisco SD-WAN Onboarding</b> .....	22
<b>4.2 Deploying Controllers</b> .....	23
4.2.1 Cloud Hosted Controllers.....	23
4.2.2 OnPrem Hosted Controllers .....	24

4.2.3 Firewall Traffic Requirements .....	24
4.2.4 High Availability and Scalability of Controllers .....	25
<b>4.3 IWAN Configuration and SD-WAN Configuration Templates and Policies.....</b>	<b>25</b>
4.3.1 Configurations Template .....	26
4.3.2 Configuration Policies .....	26
4.3.3 Migrating Data Centers to SD-WAN .....	27
4.3.4 Data Center Migration Prerequisites .....	31
4.3.5 Data Center Migration Steps .....	31
4.3.6 Single IWAN Branch Router Inline Migration .....	33
4.3.7 Summary of bootstrap migration Steps .....	34
4.3.8 Dual IWAN Branch Router Inline Migration .....	36
<b>5 Customer IWAN to Cisco SD-WAN Migration Case Study .....</b>	<b>39</b>
5.1.1 Legacy IWAN Deployment Overview .....	39
5.1.2 IWAN Use Cases Deployed.....	40
5.1.3 Planned Cisco SD-WAN Design.....	40
5.1.4 SD-WAN Use Cases .....	41
5.1.5 Migration State – Parallel IWAN and SD-WAN Infrastructure .....	41
<b>5.2 IWAN Deployment Deep Dive.....</b>	<b>43</b>
5.2.1 DMVPN Design and Business Intent.....	43
5.2.2 The DMVPN design can be summarized as: .....	43
5.2.3 DMVPN Verification.....	44
5.2.4 DMVPN Hub Router Configuration walkthrough (DC1-SJC-BR1) .....	48
5.2.5 Table: Hub DMVPN Tunnel and Encryption Template Walkthrough .....	50
5.2.6 The BR3-LAX-MCBR branch DMVPN configuration (Crypto, tunnels, QoS and routing) is shown below: .....	55
5.2.7 Performance Routing (PfRv3) Deployment .....	59
5.2.8 Direct Internet Access (DIA) Deployment .....	68
<b>5.3 Migration Planning of Case Study .....</b>	<b>76</b>
5.3.1 Identify Sites/Regions.....	78
5.3.2 Identify Use Cases and Feature/Configuration Analysis .....	79
5.3.3 Current Inventory, Bandwidth and Platform Requirement .....	81
5.3.4 Identify Controller Deployment Model and Requirements .....	81
5.3.5 Detailed Site Migration Planning .....	82
<b>5.4 Migration Deployment Steps for Case Study .....</b>	<b>83</b>
5.4.1 Cisco SD-WAN Onboarding .....	83
5.4.2 Deploying Controllers .....	83
5.4.3 Migrating Data Centers to SD-WAN .....	84
5.4.4 Basic Configurations .....	85
5.4.5 Branch Routers Migration basic config of LA branch missing .....	95
5.4.6 Verification .....	100

---

5.4.7 Mapping IWAN to SD-WAN Configuration Policies and Advanced Use Cases .....	103
5.4.8 Day2 Monitoring and Serviceability .....	130
5.4.9 API .....	132
6 Cisco SD-WAN Advance Use-cases .....	133
6.1 AppQoE Features .....	133
6.2 Security Use Cases .....	134
6.3 Unified Communications .....	136
Appendix.....	137
Configurations and Policies.....	137

---

# List of Figures and Tables

FIGURE 1: CISCO SD-WAN ARCHITECTURE .....	10
FIGURE 2: DEPLOYMENT/MIGRATION STAGES.....	13
FIGURE 3: MIGRATION PLANNING.....	14
FIGURE 4: MIGRATION SEQUENCE.....	22
FIGURE 5: CONTROL PLANE SCALABILITY.....	25
FIGURE 6: SD-WAN BEHIND CEs WITH IWAN BRs .....	28
FIGURE 7: SINGLE ROUTER BRANCH MIGRATION – ROUTING .....	35
FIGURE 8: INLINE BRANCH DEPLOYMENT.....	37
FIGURE 11: DEPLOYMENT/MIGRATION STAGES.....	76
FIGURE 12: MIGRATION PLANNING.....	77
FIGURE 13: MIGRATION SEQUENCE.....	75
TABLE 1: DOCUMENT HISTORY .....	6
TABLE 2 : DOCUMENT REVIEWERS.....	6
TABLE 3: SUMMARY OF THE CISCO SD-WAN COMPONENTS.....	12
TABLE 4: SD-WAN EDGE ROUTERS CONNECTIVITY BEHIND CE ROUTERS (IWAN DC).....	28
TABLE 5: CE ROUTERS CONNECTIVITY IN FRONT OF SD-WAN ROUTERS (IWAN DC) .....	29
TABLE 6: IWAN: DC1 IWAN ROUTERS CONNECTIVITY AND ROUTING .....	29
TABLE 7: DATA CENTER CORE ROUTERS CONNECTIVITY (IWAN DC) .....	30
TABLE 8: BRANCH 1 SD-WAN ROUTERS CONNECTIVITY AND ROUTING .....	36
TABLE 9: BRANCH INLINE DEPLOYMENT .....	38

---

# About this Guide

## History

**Table 1: Document History**

Version No.	Issue Date	Status	Reason for Change
1			

## Review

**Table 2: Document Reviewers**

Reviewer's Details	Version No.	Date
Cisco SD-WAN TMEs, Technical Leads	1	
Cisco SD-WAN TSA Team	1	

---

# 1 Introduction

## 1.1 Audience

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

## 1.2 Document Scope

IWAN migrations are each very unique to the customer environment. This document gives general migration steps and guidelines, with configuration example of a case study, for migration from a specific IWAN environment to Cisco SD-WAN.

The scope of this document includes:

- Section 2: Provides high level overview of IWAN and SD-WAN architectures
- Section 3: The migration planning section provides the guidelines about the information needed to plan a successful IWAN to SD-WAN migration.
- Section 4: After the requirements are identified in Section 3, this section identifies the common migration steps to follow during migration.
- Section 5: A case study of migrating an IWAN lab setup to SD-WAN. This section provides a walkthrough of migrating an IWAN deployment to SD-WAN by using the guidelines provided in Section 3 and 4.
- Section 6: An overview of advanced use cases for Cisco SD-WAN is presented.

Note that customers must validate the migration scenarios, for their specific environment in the lab environment, before migrating production sites.

For tools, best practices, and different designs to implement a migration customized to your existing environment, see general [Cisco SD-WAN Migration Guide](#). Also, for a detailed SD-WAN 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 IWAN and Cisco 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](#)

- 
- [Cisco SD-WAN WAAS Migration Guide](#)
  - [Cisco GitHub](#)



---

## 2 IWAN and Cisco SD-WAN Overview

This section provides overview of IWAN and Cisco SD-WAN.

### 2.1 Intelligent WAN (IWAN) Overview

IWAN is Cisco's first Software defined wide area network architecture. IWAN completed 5 years of innovation since its inception, with its final Release IWAN 2.3.2.

**The four components of Cisco Intelligent WAN (IWAN) are:**

#### 2.1.1 Transport-independent design

Transport-independent design: Using Dynamic Multipoint VPN (DMVPN) IWAN provides capabilities for building an IPsec/GRE tunnel overlay on top of any carrier service offering, including MPLS, broadband Internet, leased lines or cellular 3G/4G/LTE. An overlay design simplifies the WAN design as it presents a single routing control plane and minimal peering to providers, making it easy for organizations to mix and match and change providers and transport options.

#### 2.1.2 Intelligent path control

Cisco Performance Routing (PfR) improves application delivery and WAN efficiency by dynamic traffic steering, load balancing and performance based routing. PfR dynamically controls data packet forwarding decisions by looking at application type, performance, policies, and path status. PfR protects business applications from fluctuating WAN performance while intelligently load-balancing traffic over the best performing path based on the application policy. PfR monitors the network performance - jitter, packet loss, delay - and makes decisions to forward critical applications over the best performing path based on the application policy. Cisco PfR consists of border routers that connect to the broadband service, and a master controller application supported by Cisco IOS® Software on a router. The border routers collect traffic and path information and send it to the master controller, which detects and enforces the service policies to match the application requirement. Cisco PfR can select an egress WAN path to intelligently load-balance traffic based on circuit costs, to reduce a company's overall communications expenses. IWAN intelligent path control is the key to providing a business-class WAN over Internet transport.

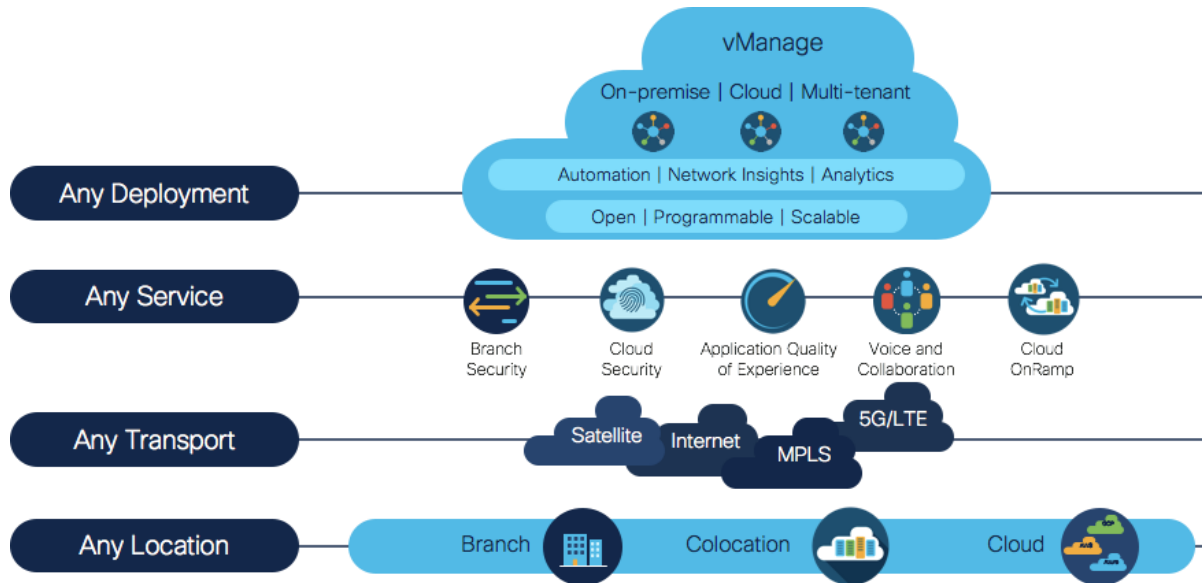
#### 2.1.3 Application Optimization

Cisco Application Visibility and Control (AVC) and Cisco Wide Area Application Services (WAAS) provide application performance visibility and optimization over the WAN. With applications becoming increasingly opaque due to increase reuse of well-known ports such as HTTP (port 80), static port classification of application is no longer sufficient. Cisco AVC provides application awareness with deep packet inspection of traffic to identify and monitor applications' performance. Visibility and control at the application level (layer 7) is provided through AVC technologies such as Network-Based Application Recognition 2 (NBAR2), NetFlow, quality of service (QoS), Performance Monitoring, Medianet, and more. Cisco AVC allows IT to determine what traffic is running across the network, tune the network for business- critical services, and resolve network problems. With increased visibility into the applications on the network, better QoS and PfR policies can be enabled to help ensure that critical applications are properly prioritized across the network. Cisco WAAS provides application-specific acceleration capabilities that improve response times while reducing WAN bandwidth requirements.

## 2.1.4 Secure connectivity

IWAN Protects the WAN and offloads user traffic directly to the Internet. Strong IPsec encryption, zone-based firewalls, and strict access lists are used to protect the WAN over the public Internet. Routing branch users directly to the Internet improves public cloud application performance while reducing traffic over the WAN. Cisco Cloud Web Security (CWS) service provides a cloud-based web proxy to centrally manage and secure user traffic accessing the Internet.

## 2.2 Cisco SD-WAN Solution Overview



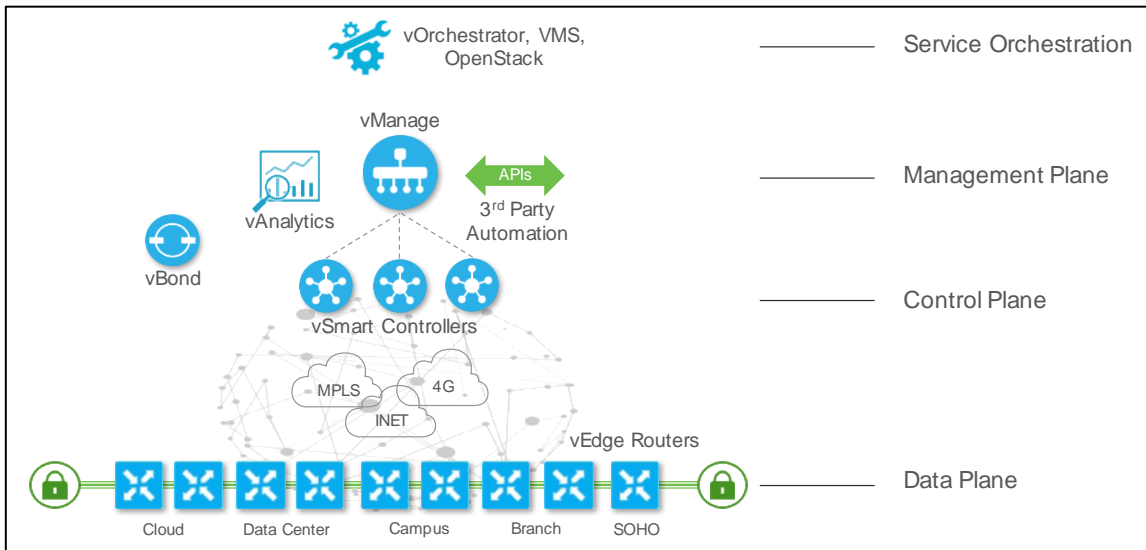
**Figure 1: Cisco SD-WAN Architecture**





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
<b>Cisco vManage</b> 	<p>The vManage NMS is a centralized network management system that lets you configure and manage the entire overlay network from a simple graphical dashboard.</p>
<b>Cisco vSmart Controller</b> 	<p>The vSmart controller is the centralized routing and policy engine of the SD-WAN solution, controlling the flow of data traffic throughout the network.</p> <p>The vSmart controller works with vBond orchestrator to authenticate SD-WAN devices as they join the network and to orchestrate connectivity among edge routers.</p>
<b>Cisco SD-WAN Edge Routers</b>  vEdge  cEdge	<p>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.</p> <p>vEdge cloud router can run as a Virtual Network Function (VNF) on Cisco Enterprise Network Compute System (ENCS) platforms.</p> <p>Cisco SD-WAN image is supported on IOS XE devices ISR1100s, ISR4k, ASR1k and CSR 1000v.</p>

---

Component	Description
<p data-bbox="224 296 480 323">Cisco vBond Orchestrator</p> 	<p data-bbox="544 342 1390 422">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.</p>

**Table 3: Summary of the Cisco SD-WAN Components**

For detailed introduction to Cisco SD-WAN Design, please refer to [Cisco SD-WAN Design Guide CVD](#)

# 3 Migration Planning

When deploying and migrating to Cisco SD-WAN, the initial state of the production “brown field” WAN may be a traditional/legacy WAN (MPLS-VPN, IP VPN, L2 VPN/Metro Ethernet) or a Cisco Intelligent WAN (IWAN) built according to the prescriptive guidance of the architecture. For general SD-WAN deployment and migration guidance from a traditional/Legacy WAN, refer to the Cisco SD-WAN Migration Guide. This guide focuses on specific design, implementation migration considerations for migrating from an existing Cisco IWAN to Cisco SD-WAN.

During the migration, two parallel overlay networks (IWAN and SD-WAN) will exist side-by-side until all sites are migrated. This allows for an incremental migration from branch sites, minimizing the chance of disruption. As IWAN branches are migrated, the size of the IWAN network will shrink, and ultimately decommissioned. The first steps to deploying the new Cisco SD-WAN network will be the SD-WAN controllers and WAN edge aggregation routers, typically installed in customer data centers or other central sites where breakout access to private and public cloud services. Network to Network interconnect between the two overlay networks is achieved at the DC core layer, where traditional routing protocols such as BGP, OSPF or EIGRP exchange the prefixes associated with each domain.

The diagram below depicts the different stages of the Cisco SD-WAN deployment for both IWAN migration and traditional WAN (Brownfield Legacy WAN migration to Cisco SD-WAN). The recommended flow of migrating to Cisco SD-WAN is to first deploy Controllers, then deploy SD-WAN Edge routers in DC and then migrated branch sites. Once all IWAN branch sites are migrated to Cisco SD-WAN, the IWAN Master Controller(s) (MC) and Border Routers (BRs) can be removed from DCs.

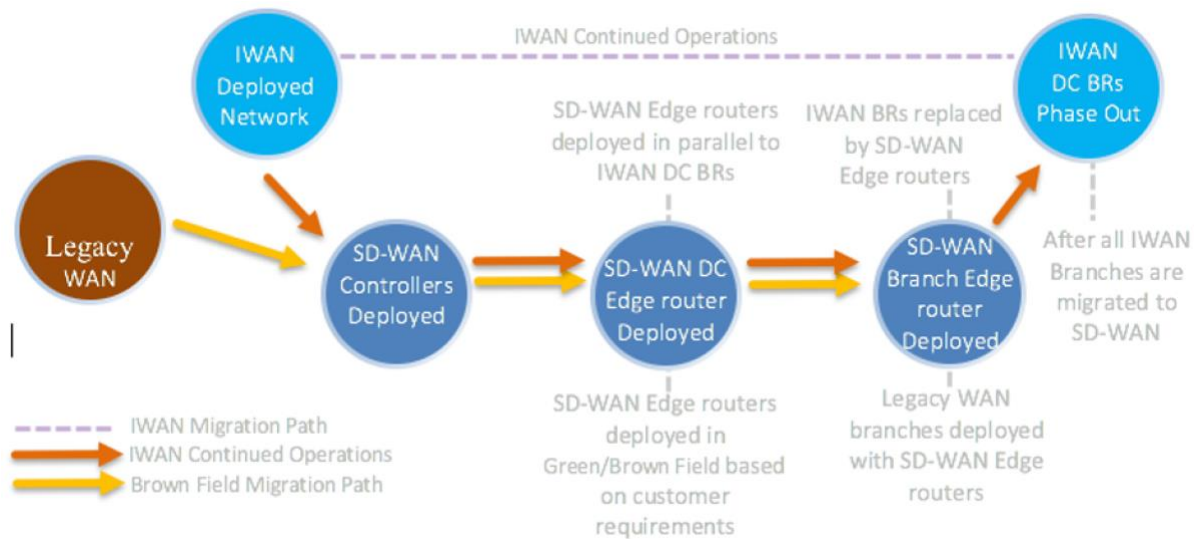
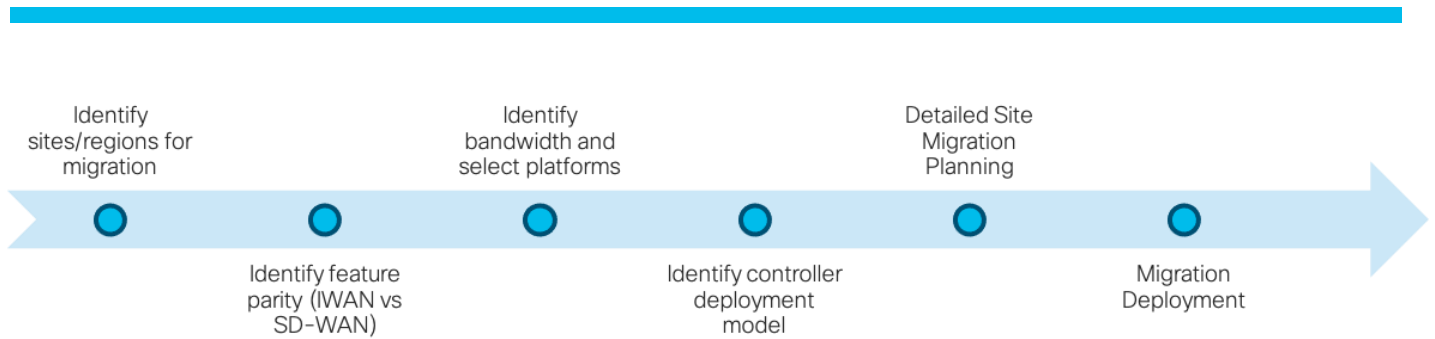


Figure 2: Deployment/Migration Stages

Migration planning is critical because moving from IWAN 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.



**Figure 3: Migration Planning**

### 3.1 Identify sites/regions for migration

The first step is to identify where to start the migration. If the existing IWAN deployment is distributed across multiple regions interconnected by data centers, it is recommended that the first migrations be contained to a common region for simplicity. When determining the order of migration within a region, it is always recommended to start in a DC and then move to the branches, one at a time. Things to consider before planning a particular regional migration include:

- Hardware/software inventory of router platforms across all regions to determine which are SD-WAN capable and may be candidates of in-line migration. (See section 3.4)
- Configuration analysis to understand the IWAN use cases and features deployed
- Details of the WAN transport providers, to include access methods, QoS and circuit speeds
- Application traffic patterns in each region. Strict branch to hub, or also branch to branch?
- Site Profiles/Site groupings in each region
  - Small, Medium, Large sites based on bandwidth or number of branch users?
  - Functional structure? Corporate offices, retail offices, sales offices, ATM machines?
- End state SD-WAN topology
  - Single region with full mesh between all sites?
  - Single region with hub and spoke only?
  - Multiple regions interconnected with traditional IP core?
  - Multiple regions interconnected with SD-WAN core? (hierarchical)
- Data Center LAN connectivity, IP addressing, VLAN schema, QoS and routing details
- Traffic path between legacy, IWAN and SD-WAN branches during migration
- Site ID and SD-WAN Policy planning
  - App-route policies for SLA based routing preferences
  - Control or Data policy for administrative based routing preferences
  - Service Advertisements/Service chaining
  - QoS architecture to include application requirements and marking
- Branch specific policy requirements

- DIA/Backhaul to DC
- Regional breakout for services
- Access to SaaS/IaaS cloud services (backhauled to DC, DIA, SD-WAN)
- Branch security requirements (on-prem or cloud security)

### 3.2 Identify Use Cases and Feature/Configuration Analysis

In this step identify the IWAN use cases that are deployed in the IWAN network. When migrating to Cisco SD-WAN, the business intent of the IWAN use cases should be deployed with Cisco SD-WAN. Also, perform feature and configuration analysis of both IWAN and Cisco SD-WAN solution that will assist in the configuration development and the time of deployment of SD-WAN Edge devices.

#### 3.2.1 Identify Use Cases

- Identify the currently deployed IWAN use cases and business problems they are solving
- Identify any IWAN use cases that were planned to be deployed prior to SD-WAN and business problems they intended to solve
- Translate existing IWAN use cases to the equivalent SD-WAN use cases
  - Understand that SD-WAN approach is different than IWAN and same business intent may get deployed differently with Cisco SD-WAN.
  - if gap exist in features achieving the business intent or use case, identify whether a workaround better solution exists.
- Identify new Cisco SD-WAN use-cases for the deployment optimization

Refer to below table for comparison of the use cases, terminology and features between IWAN and Cisco SD-WAN:

Use Case Comparison	
IWAN	Cisco SD-WAN
<p><b>Transport Independent Design (TID)</b></p> <ul style="list-style-type: none"> <li>• Secure site-to-site VPN communications</li> <li>• DMVPN IPsec-encrypted mGRE (DMVPN) tunnel overlay on top of any IP transport (MPLS, Internet, 4G LTE, etc.....)</li> <li>• Hub and spoke or full mesh with BGP or EIGRP routing over the top</li> <li>• Up to 2000 remote sites in a single IWAN domain</li> <li>• 2-5 WAN transports per branch, more than 3 transports require dual routers</li> </ul>	<p><b>Secure Automated WAN</b></p> <ul style="list-style-type: none"> <li>• Secure site-to-site VPN communications</li> <li>• IPsec/GRE-encrypted P2P tunnels on top of any IP transport</li> <li>• Hub and spoke, full mesh, or custom topologies with controller-based policies and overlay routing that simplify WAN routing using OMP routing protocol</li> <li>• Up to 2000 remote sites with a single vManage Controller. No limit on number of devices with horizontally scalable design.</li> <li>• No architectural limit to WAN transports per branch with horizontal scalability</li> </ul>
<b>Intelligent Path Control</b>	<b>Application Performance Optimization</b>

<ul style="list-style-type: none"> <li>• Performance Routing (PFRv3) for application aware/SLA based routing</li> <li>• Proprietary smart probes between sites for path quality measurements and monitoring (loss, delay, jitter)</li> <li>• Every site requires local master controller, typically collocated with one Border router where enforcement occurs</li> <li>• IWAN policies learned by branch MCs from central domain controller</li> <li>• Deterministic path selection with primary and fallback policies</li> <li>• Load balancing of non-performance traffic-classes</li> </ul>	<ul style="list-style-type: none"> <li>• App-Aware routing data policies for SLA based tunnel selection</li> <li>• BFD probes between WAN edge routers for path quality measurements and monitoring (loss, delay, jitter)</li> <li>• Policies distributed to WAN edge by centralized vSmart controller.</li> <li>• Deterministic path selection with intelligent app-aware routing policies or traffic engineering data policies, full flexibility of path selection supporting active/standby, or primary/secondary, tertiary routing designs.</li> <li>• Per-Tunnel QoS</li> <li>• Forward Error Correction (FEC)</li> <li>• Packet Duplication</li> <li>• Application visibility with NBAR (cEdge) Qosmos (vEdge)</li> <li>• Flexible NetFlow (cEdge) and Cflowd (vEdge) for flow visibility</li> <li>• TCP optimization on the WAN edge</li> <li>• App-NAV XE redirection to WAAS appliance (refer to the next section for details)</li> </ul>
<p><b>Application Optimization</b></p> <ul style="list-style-type: none"> <li>• Application visibility with NBAR deep packet inspection and NetFlow</li> <li>• Appliance based Wide Area Application Services (WAAS) for TCP optimization, compression, data redundancy elimination, caching</li> </ul>	
<p><b>Secure Direct Internet Access (DIA)</b></p> <ul style="list-style-type: none"> <li>• Local Internet access with split tunneling through static default routing on IWAN router terminating ISP circuit</li> <li>• IWAN border router with NAT on ISP facing interface</li> <li>• High availability with backup path over DMVPN to central Internet gateway site</li> <li>• Security with Zone Based Firewall on IWAN Border router, provisioned through cli</li> </ul>	<p><b>Secure Direct Internet Access (DIA)</b></p> <ul style="list-style-type: none"> <li>• Local Internet access with split tunneling through static default routing on Cisco SD-WAN edge router or through Cisco SD-WAN data policies for specific traffic DIA requirements.</li> <li>• SD-WAN edge router with NAT on ISP facing interface</li> <li>• High availability with backup path over SD-WAN tunnel to central Internet gateway site</li> </ul> <p><b>DIA Security options include:</b></p> <ul style="list-style-type: none"> <li>• Embedded router features including ZBFW, IPS/IDS, AMP, DNS Security, URL filtering, orchestrated through vManage</li> <li>• Redirect to Cisco Umbrella Secure Internet gateway (SIG) or to Third part SIG providers</li> <li>• Redirect to Cloud OnRamp for Colocation security service chain</li> </ul>
<p><b>Direct Cloud Access (DCA) for SaaS</b></p> <ul style="list-style-type: none"> <li>• Restricts local Internet breakout for known/trusted SaaS applications and/or domains, all others are backhauled over DMVPN to Internet GW</li> <li>• NAT on local Border Router designated for DCA</li> <li>• Path optimization with IPSLA probes to SaaS servers to determine best performing Internet path</li> <li>• Complex solution leveraging Umbrella DNS, NBAR, PFRv3 and IPSLA for SaaS probes, all configured through cli</li> </ul>	<p><b>MultiCloud Connectivity</b></p> <ul style="list-style-type: none"> <li>• Cloud OnRamp for SaaS</li> <li>• Trusted SaaS applications redirected to local Internet breakout or designated Internet gateway site across SD-WAN fabric.</li> <li>• Path selection with SaaS probe monitoring <ul style="list-style-type: none"> <li>• Cloud OnRamp for IaaS</li> </ul> </li> <li>• SD-WAN tunnels extended from cloud branches to IaaS providers (AWS, Azure)</li> <li>• Fully orchestrated through vManage</li> <li>• Cloud OnRamp for Colocation</li> <li>• Internet and SaaS traffic backhauled to Colocation for Service chaining and Cloud access</li> </ul>

### 3.2.2 WAAS Use Case with Cisco SD-WAN

Many customers have integrated WAAS with IWAN solution for optimization of the application traffic. Cisco SD-WAN allows the WAAS deployment using AppNav capability to support the WAAS migration from IWAN to Cisco SD-WAN network designs. In addition, as mentioned in previous table, Cisco SD-WAN supports traffic optimization features such as Forward Error Correction, Packet Duplication, HTTPS Proxy, App Aware SLA etc.



Refer to the [Cisco SD-WAN WAAS Deployment and Migration Guide](#) for complete details on how WAAS can be migrated to Cisco SD-WAN.

### 3.2.3 Feature and Configuration Analysis

Perform analysis of the IWAN features that are deployed in the network and related with the Cisco SD-WAN features.

Review [Release Notes](#) on cisco.com to ensure existing features are supported with IOS XE SD-WAN software.

Perform a comparison of IWAN and SD-WAN configurations.

- Perform configuration audit to identify deployed features, for example, routing, QoS, and features outside of IWAN such as Voice.
- Identify the target SD-WAN code version based on required features and platforms to be migrated  
 Perform bug scrubbing from release notes and forums  
 Lab test the deployment scenario (SVS/partners/customer labs/Cisco dCloud/Cisco Modeling Labs)  
 Involve Cisco Teams (Account team, Customer Experience) for additional support if required

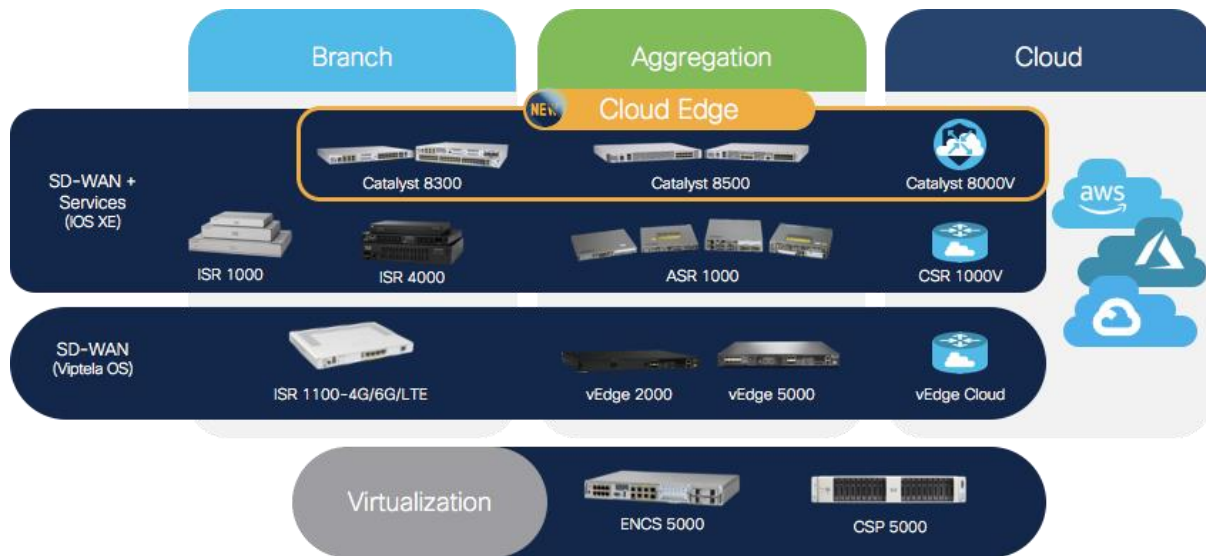
Feature Descriptions for Each Solution	
Cisco IWAN	Cisco SD-WAN
<b>Control Plane:</b> Master Controller (MC) on every site (software feature on router)	<b>Control Plane:</b> vSmart centralized controller
<b>Data Plane:</b> Border Routers (BR) monitor path quality and enforce path for egress traffic based on policy and SLA	<b>Data Plane:</b> WAN edge monitor path quality and enforce path for egress traffic based on policy and SLA
<b>Path quality monitoring:</b> Proprietary smart probes between PfR Border routers	<b>Path quality monitoring:</b> BFD probes between WAN edge routers
Custom routing with BGP or EIGRP configured on top of DMVPN (ECMP or active/backup with route policy)	Custom routing with overlay management protocol (OMP) policies on vSmart controller. OMP peering is automatically establishes between WAN edge and vSmart controllers and defaults to ECMP load sharing across multiple tunnels to same destination.
PfRv3 for intelligent path selection for applications matched by DSCP or through NBAR deep packet inspection	Application-Aware Routing policies for intelligent path selection matched by DSCP or through NBAR or Qosmos (vEdge) deep packet inspection
Application optimization through Wide Area Application Services (WAAS). Application visibility with NBAR and NetFlow/IPFIX.	FEC, TCP opt, packet duplication, QoS, per tunnel QoS, adaptive QoS, intelligent path steering using app-aware, AppNav only (no DRE, caching), Cloud OnRamp for SaaS, SD-AVC
Secure encrypted communications with IPsec over mGRE tunnels. Pre-shared keys or PKI integration for authentication with a variety of strong encryption protocols.	Default SD-WAN IKE-less IPsec and GRE, secure control/data plane using combination of PKI and secure symmetric key exchange, optional pair wise keys, legacy IPsec/GRE, Greatwall UTM (ZBFW, IPS, AMP, URL filtering, Cisco Umbrella), SIG tunnels, segmentation, secure control/data plane, ACL, SSL proxy

cloud-based applications	Cloud OnRamp for IaaS, SaaS and colocation, Cisco Umbrella DNS/SIG, Third party SIG
Secure bootstrap, Cisco Plug and Play (PnP), SUDI	Secure PnP based ZTP, ZTP other methods (OTP with cloud_init, usb boot, on-prem ZTP), ZTP without DHCP (automatic IP), SUDI, TPM support, allow-list for trusted devices serial with Smart account Sync
Cisco IOS (classic) and Cisco IOS XE	Cisco IOS XE and Viptela OS
PKI	Cisco CA, Enterprise CA, vManage as CA
Automation	Full automation support by vManage using templates with REST API support
Manageability APIC-EM with IWAN APP	vManage for complete day0,1,2 configs, monitoring and troubleshooting, vAnalytics for advance analytics
Routing protocols support for (EIGRP, BGP, OSPF)	OMP for Overlay, BGP, OSPF, Static for WAN, BGP, EIGRP, OSPF, Static for LAN.
IPv6	IPv6 (see specific feature details)
Master Controller	vSmart (centralized controller with redundancy)
Scalability 2000 sites	IKE less IPsec provides scalability of IPsec tunnels, no limit on number of sites with horizontal scaling of controller and routers
Multicast Support	Multicast Support (see Release Notes for details)
NAT	NAT DIA, NAT POOL, Service Side NAT
Segmentation 20 hub, 7 remote VRFs	Multi VPN (VRF) up to 300, inter VRF route leak
Trackers	DIA tracker, SLA monitoring with BFD, Static Route tracker, zScaler L7 health check, VRRP tracker of OMP peering and prefixes
NBAR (N/W Based App Recognition)	NBAR and Customer App Recognition, (Qosmos DPI engine for vEdges running Viptela OS)
Direct Cloud Access	Available as Direct Internet Access both from policy and route within Service VPN. Also provide Cloud OnRamp services for optimized SaaS and auto-integration with IaaS
NetFlow v9	Available
VRRP/HSRP	VRRP Supported. vEdge supports up to five VRRP groups per physical/sub-interfaces for Primary and Secondary IP addresses. cEdge support for multiple groups is planned for 17.4.  No HSRP support.
Port Channel interfaces for additional bandwidth capacity and redundancy	SD-WAN routers currently do not support port link aggregation technology. Alternatively, L3 ECMP can be used using standard routing protocols. Note: Link aggregation is part of the roadmap.

### 3.3 Current Inventory, Bandwidth & Platform requirement

Cisco SD-WAN can be deployed with vEdge Viptela devices or on Cisco IOS XE platforms that are Cisco SD-WAN capable. Customers should identify the platforms that can have in-place migration without replacing the hardware. Devices running version lower than IOS XE 17.2 image requires code upgrade to run SD-WAN. However, versions from 17.2 and higher are universal images, which can run SD-WAN image with cli config change to Controller Mode and a reboot. Detailed upgrade process is described in Section 5 of Case Study.

Base on the requirements, identify the appropriate platform to support SD-WAN. At the time of writing of this document, Cisco SD-WAN portfolio includes below platforms:



#### 3.3.1 Current Inventory

Conduct hardware inventory to include PID, CPU, memory, SUDI certificate of devices at the sites to assist in identifying devices for in-place migration.

#### 3.3.2 Identify Bandwidth and Platforms

For scalability, licensing and feature compatibility identify below information:

- Conduct traffic analysis to determine PPS rates across IWAN tunnels at branches and hubs, to identify the bandwidth requirements per site.
- Map identified features in previous section to platforms and site types (some features might not be supported on all the platforms/some sites might not be using all the features).
- Identification of the required bandwidth and use cases needed at each site will assist in determining the license type as well.
- Identify platforms for Data Center and branches based on scale of design (expected PPS and number of IPsec tunnels).
- Check if horizontal scaling is required for large sites, primarily based on number of tunnels and bandwidth.

- 
- Obtain physical topology of sites to identify LAN and WAN transport connections and map to port density required in platforms.

### 3.3.3 Identify Licensing

Customer may already have Cisco devices in their network that can be upgraded to Cisco SD-WAN from IWAN. You should identify sites with those devices and licenses associated with those devices that will be upgraded to support Cisco SD-WAN. This will allow you to plan the transition of the licensing when migrating to SD-WAN.

Information to identify:

- Current Licensing with IWAN
- Licensing with SD-WAN
- Licensing differences between IWAN and SD-WAN
- License and device migration to Plug and Play (PnP) Virtual Account (VA)

Refer [Cisco DNA Software SD-WAN and Routing Matrices](#) for details about SD-WAN License. Please contact your Cisco Account/Sales team representative for latest information about migration process of license.

## 3.4 Identify Controller Deployment Model and Requirements

Cisco SD-WAN Controllers can be deployed in two models, On-Prem and Cisco Cloud hosted.

Cisco recommends Cisco Cloud hosted deployments for Controllers because of below benefits:

- Easily Scalable
- Monitored SLA
- Geo-redundant
- Cisco Ops Support
- ZTP with automation

On-Prem Customer hosted controllers are installed and managed by customer:

- Hardware and software maintenance and monitoring
- Opening Firewall ports for SD-WAN overlay communication
- Scalability may require additional hardware

For On-Prem controller deployment, things to consider:

- Identify scalability and affinity requirements for Controllers.
  - Is vManage Cluster needed?
  - How many vBonds and vSmarts are needed?
- Determine physical hardware host requirements for virtual Controllers
- What are the backup and storage requirements?

- 
- Determine placement of Controllers in the network (typically DMZ.)
  - Open Firewall ports for SD-WAN overlay communication
  - How will branches reach Controllers?
  - Identify allowed IPs for mgmt.
  - [Hardware recommendations for On-Premise deployments](#)

## 3.5 Detailed Site Migration Planning

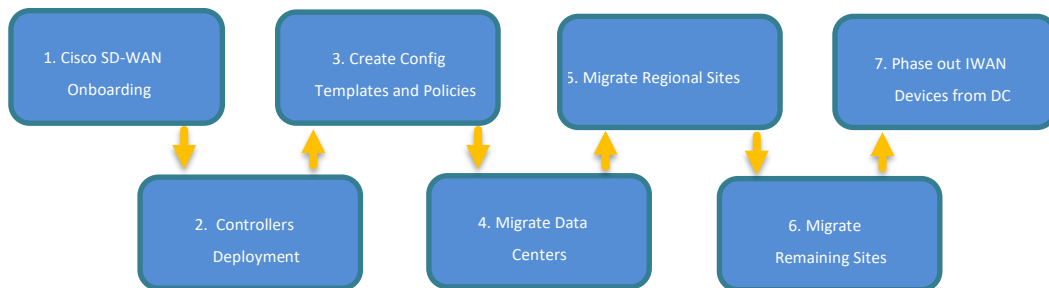
Identify the individual site-level details to develop a migration plan. Some of the important considerations are:

- Determine port availability
- Build IP addressing and port/VLAN schema
- Code/Rommon upgrades
- Controllers- redundancy, affinity, clustering, horizontal scaling, compute, backup and storage requirements, firewall ports, allowed IP addresses for mgmt.
- Generate and document expected template/cli configurations from lab setup
- Develop a test plan to validate migration
- Determine automation requirements (API, Netconf)
- Build SD-WAN policies based on traffic flow requirements (refer to [Cisco SD-WAN policy framework guide](#))
- Document connectivity physical and logical topologies during and after migration
- Determine OSS integration requirements (Cflowd, SNMP, email, webhooks, syslog)
- Determine a fallback procedure in case of issues

# 4 Migration Deployment Steps

The Section 3 provides the guidelines for gathering the information on the existing IWAN network and guidelines to provide the Cisco SD-WAN migration deployment. This section provides the guidelines about the steps that should be taken to perform the migration.

The migration from IWAN to Cisco SD-WAN should follow below recommended sequence of steps for a successful migration:



**Figure 4: Migration Sequence**

During the migration, the configurations must support routing between IWAN and SD-WAN sites. After every step, it is imperative to verify the existing and new routing flows are working as required.

1. The onboarding process involves steps of choosing the hardware, preferred management, license subscription and end customer’s Smart Account and Virtual Account on PnP Portal. The PnP portal is used to manage the devices, controller profile and licensing.
2. Deploy Controllers on Cloud or On-Premise. The Controllers must be accessible over Internet and/or MPLS transport.
3. Perform IWAN Configuration Analysis in comparison to SD-WAN. 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. It is recommended to migrate Data Center sites first and use them for communication between the legacy and SD-WAN migrated sites. During IWAN migration to SD-WAN, assure there is routing between 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. In the end migrate the smaller branch sites for each region.
7. Remove Data Center Legacy/IWAN routers and then Master Controllers.

## 4.1 Cisco SD-WAN Onboarding

Order submission is the first and most important step during the onboarding process. All required information in this step is used throughout the process.

1. Choose the hardware platform or migrate existing IWAN devices to Virtual Account

- 
2. Choose the preferred management.
  3. Choose the subscription tier and term.
  4. Choose the requisite subscription term.
  5. Choose the requisite bandwidth tier for the offer.
  6. Enter the end customer's Smart Account and Virtual Account.
  7. Choose optional services. These are highly recommended.
  8. For any problem regarding Smart Account and Virtual Account setup, please reach out to Cisco TAC.

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 (SA) and Virtual Accounts (VA) are essential in a successful on-boarding of a SD-WAN Edge router to its corresponding network. Migrate existing IWAN devices to VA that will be upgraded to Cisco SD-WAN, along with the licenses.

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 Overlay. Customer can also manually add their existing devices to their Virtual. 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 Login at [software.cisco.com](https://software.cisco.com) with your CEC credentials.

The onboarding process for Cloud controllers, On-prem controllers and software devices is available in detail in [Cisco SD-WAN Onboarding Guide](#) For any additional details on Plug and Play process, visit [support guide](#).

In some scenarios, Zero Touch Provisioning is not possible, for example, due to unavailability of DHCP service. In such cases, cEdge can be booted with a bootstrap configuration. From vManage generate bootstrap Config file for the device. Config file (which includes basic interface configuration, Root CA, Organization Name, vBond information, etc.) is fed into the PnP process through. Upon bootup, SD-WAN XE router will search bootflash: or usbflash: for filename ciscosdwan.cfg. After that Router continues normal ZTP process. The upgrade process using bootstrap is described in detail in Case Study section 5 later in this document.

## 4.2 Deploying Controllers

Controllers can be deployed in hosted Cloud or On-Premise. Refer to the [Overlay Bringup Guide](#) for more details on how to deploy Controllers On-Premise.

### 4.2.1 Cloud Hosted Controllers

The next step is to check the information on the Cisco Plug and Play (PnP) portal. For the Cisco hosted cloud controller deployment, if all required information during the procurement was provided accurately, no further action is needed on the PnP portal.

1. After order submission, the hardware serial number is pushed into the [PnP portal](#) automatically.

- 
2. An email with information required for accessing to vManage is sent to the email address associated with the order that has to be replied with information.
    - No further action is needed for first-time greenfield deployment (first overlay for the SA). Aforementioned email includes vManage link.
  3. To migrate existing hardware:
    - Purchase a cloud subscription from Cisco Commerce Workspace (CCW) (**L-Lic-DNA-ADD**).
    - Add hardware manually to the PnP portal under the corresponding VA. (refer to page 81 in Plug and Play Connect Capability)
  4. For more details, please refer to the [Plug and Play Support Guide for Cisco SD-WAN](#) and [Cisco Network Plug and Play Connect Capability Overview](#).
  5. In case of failure at any step on PnP portal, contact the Cisco TAC for further assistance.

The last step is to sync up the vManage and PnP portal information. The vManage information is sent to the overlay administrator via the email address used during procurement.

1. After the Cisco team has spun up the controllers, an email with vManage information is sent to the email associated with the order.
2. To add devices to the overlay:
  - Log in to vManage (the default is admin/admin).
  - Sync up vManage with the Smart Account/Virtual Account. The Cisco.com credentials of the VA administrator role are required (Configuration à Devices à Sync Smart Account).
  - You must re-sync vManage with the Smart Account/Virtual Account for any new devices added to the PnP portal.
3. After you've transferred device information to vManage, the overlay is ready to be set up.

This process is described in [Cisco SD-WAN Onboarding Guide](#).

## 4.2.2 OnPrem Hosted Controllers

Similar to Cloud Hosted Controllers, the OnPrem Controllers must be onboarded to PnP as described in the Onboarding guide. The process to deploy OnPrem controllers is described in the Cisco Live presentation [BRKRST-2559](#). At the time of writing this document the CVD document for OnPrem controller deployment was in progress.

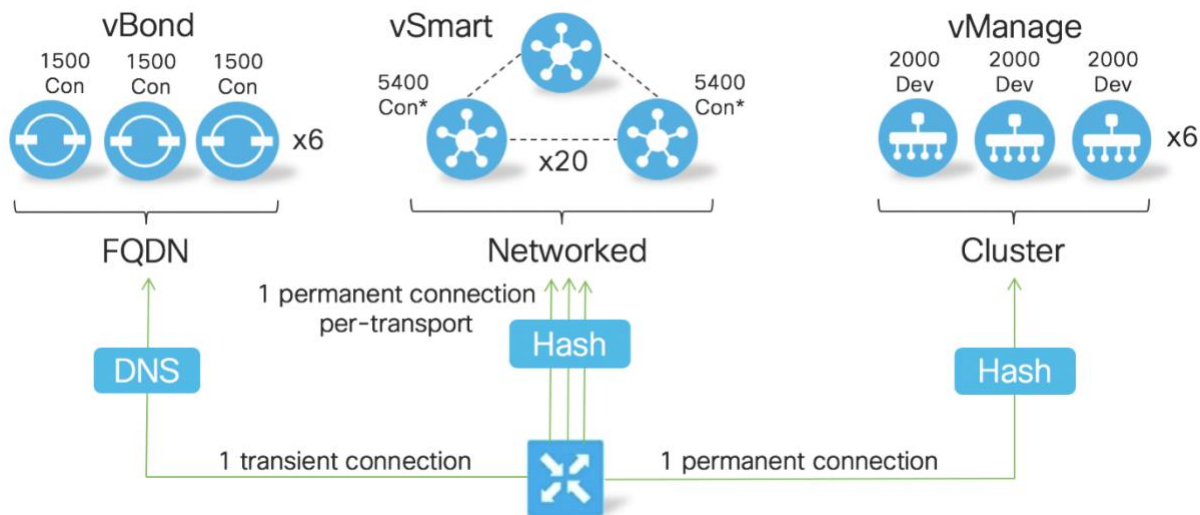
## 4.2.3 Firewall Traffic Requirements

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. Refer to Table 2. in [End to End Deployment CVD](#).



## 4.2.4 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 5: Control Plane Scalability**

Visit [high availability](#) guide for more detail. Also visit [vManage Cluster deployment guide](#) for cluster creation and troubleshooting.

Verify all Controllers are up, vManage lists vBond and vSmart in Up states, and vManage has the device list of the SD-WAN Edge devices that will join the overlay in the migration process.

## 4.3 IWAN Configuration and SD-WAN Configuration Templates and Policies

At this stage of the migration, the IWAN deployed use cases are captured in section 3. It is also captured in section 3 that what use cases should be deployed with Cisco SD-WAN to achieve same business intent after migration. Using configuration analysis from Migration planning section 3, create SD-WAN configurations templates for each site and policies for the SD-WAN network. 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 customer's specific use cases. Section 5 shows the configuration templates and policy creation for this Case Study.

- The Cisco SD-WAN router's specific attributes are configured like site-id, system-ip, vBond etc. (see documentation for more details)
- The transports are named as colors, for example biz-internet and MPLS. This helps in policy making for traffic control in the Cisco SD-WAN overlay.
- SD-WAN routers connect to the legacy LAN infrastructure with a traditional routing protocol such as BGP (OSPF/EIGRP/Static routing also supported) on LAN side under VRF 1. (segmentation use-case)

- 
- Two default routes will be used in the underlay pointing to INET FW and MPLS CE to form control plane connectivity (transport independence use case)
  - Once the control plane connectivity is done, OMP peering will be automatically established using the underlay. Unlike in IWAN where BGP/EIGRP is used and manual configuration is required (Overlay routing)
  - OMP distributes overlay information including security keys for IPsec, tunnels endpoints, routes and other Cisco SD-WAN attributes for the overlay topology etc.
  - BFD tunnels are created using both colors (biz-internet and MPLS) using IPsec encapsulation automatically (see more details below)
  - Service side routes from BGP gets redistributed in OMP and vice versa and the routes are received and distributed to other sites by vSmart.
  - Once the control plane, bfd tunnels and route propagation are complete the sites can ping each other and basic Cisco SD-WAN use cases of transport independence, secure overlay are already deployed.

### 4.3.1 Configurations Template

Below is high level detail of Configuration Templates and Policies on Cisco SD-WAN.

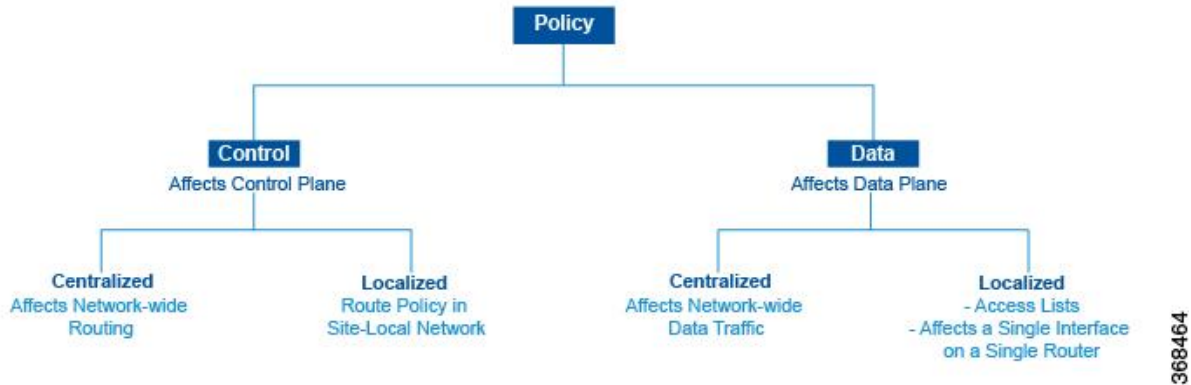
All devices in a Cisco SD-WAN overlay network that are managed by the vManage NMS must be exclusively configured from the NMS. The configuration procedure is as follows:

4. **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.
5. **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.
6. **Attach devices to device templates.** To configure a device on the overlay network, you attach a device template to the device.
7. **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.

### 4.3.2 Configuration Policies

Policies will be defined on per customer use case. Cisco SD-WAN policies includes Control Policies and Data Policies and provides granular control on how SD-WAN network operates.



368464

Refer to [Cisco SD-WAN Configuration Guides](#) for the release that is going to be deployed for detailed steps for configuration templates and policies.

Apart from centralize and localized policies, security policies can be separately configured under configuration tab of vManage which provides one window workflow to configure Cisco SD-WAN Greatwall Security features. This includes (Zone based Firewall, IDS/IPS (Intrusion Detection/Prevention Systems), URL filtering, DNS based security with Cisco Umbrella security and Malware protection using Cisco AMP/Threatgrid.


### 4.3.3 Migrating Data Centers to SD-WAN

A data center 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 data center serves as the transit site for the traffic between non-SD-WAN and SD-WAN sites. Since data centers become transit sites, plan for adequate bandwidth utilization that may be required at the data center.

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 and the methodology used to migrate the use case topology, but the actual migration of the data center site may vary based on setup of the customers.

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 6: SD-WAN Behind CEs With IWAN BRs**

The table below explains the routing design for IWAN migration at the data center:

		SD-WAN Edge Routers
Physical/L3 Connectivity	WAN VPN0	<p>Connect VPN0 WAN side to both CE devices (router or Firewall) that currently provide transport access to the IWAN BR's</p> <p>Recommend /30 addressing on CE-to-SD-WAN router links</p> <p>The /30 link prefixes connecting to the MPLS CE router must be advertised into the MPLS core for remote site reachability.</p> <p>The /30 link prefixes connecting to the Internet may be allocated from Internet routable address space or RFC 1918 private space. In the case of the latter, it is assumed the Internet CE performs NAT</p>
	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 data center 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 4: 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 5: 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 data center 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 6: IWAN: DC1 IWAN Routers Connectivity and Routing**

Data Center 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
	<b>OUT</b>	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
<b>LAN Advertisements</b>	<b>IN</b>	DC prefixes from LAN network
	<b>OUT</b>	Non SD-WAN, IWAN and SD-WAN site prefixes to LAN as per DC LAN design requirement

**Table 7: Data Center 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.

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 data center will not 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 data center.

Since there are typically two edge routers at the data center 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 data center, which can allow loops. If BGP is the data center 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.

Typically, there are more than one data centers for HA/redundancy requirements. After successful migration of the first data center, migrate the second data center in a similar method as explained in this section. Note that a routing loop can occur if there is a backdoor link between the data center sites and route advertisement is configured between the two data centers. 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 data centers. 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 data center 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 data center when redistributing to DC LAN and filter on the edge of the other data center when learning advertisements from the LAN side.

---

#### 4.3.4 Data Center Migration Prerequisites

1. SD-WAN overlay established with vManage, vBond and vSmart controllers active and onboarded
2. DC WAN edge router details (PID, chassis serial#, device certificate) associated with the customer's SD-WAN virtual account in the [Cisco Software Connect - Plug and Play \(PnP\) Connect Portal](#)
3. WAN edge list including DC WAN edge routers uploaded to customer vManage
4. DC WAN edge Device templates and policies created on vManage
5. vSmart Central policy created on vManage

#### 4.3.5 Data Center Migration Steps

The data center migration steps involved are:

- Step 1.** Baseline current network before any changes
- Step 2.** Pre-stage WAN edge
- Step 3.** Activate central policy
- Step 4.** Attach device templates
- Step 5.** Validate device certificates
- Step 6.** Onboard to vManage
- Step 7.** Validate DC routing
- Step 8.** Verify NetOps

##### Step 1: Baseline current network before any changes

1. Application Performance (NMS tools, or application performance tools)
  - a. Enterprise DC hosted Apps: Critical business applications, Real-time voice/video
  - b. Cloud apps (aaS applications)
2. Traffic Path Performance (traceroute and/or other synthetic traffic tests)
3. end-to-end application path
4. loss, latency, jitter
5. WAN Performance
6. utilization, errors, QoS drops
7. loss, latency, jitter
8. Platform (router) performance and health checks
9. CPU, MEM, I/O and Storage
10. Firewall/IPS, other Security appliance performance
11. CPU, MEM, I/O session counts

---

## Step 2: Pre-stage WAN edge

1. Upload target Software image
2. Identify the target image based on use case requirements, platform and release notes
3. secure copy (SCP) the target SD-WAN image to IWAN router bootflash.
4. Update rommon if necessary
5. Upload bootstrap configuration (ciscosdwan.cfg) file
6. vManage creates bootflash, must include
7. vBond/FQDN/IP, org name, system IP, site ID, VPN0 interfaces and routes.
8. Can load on bootable USB or router bootflash

## Step 3: Activate central policy

Centralized policy must be enabled prior to cutting over remote sites. Full mesh is the deployment unless otherwise configured.

1. vManage > configuration > policies
2. select centralized policy > preview (to review), activate
3. verify policy activates and pushed to vSmart

## Step 4: Attach device templates

1. Attach device templates to DC WAN edge routers and supply device specific values for variables
2. Upload completed config and verify routing

## Step 5: Validate device certificates

1. Ensure device certificates for DC WAN edge router are active state in vManage
2. If not, move to active and "send to controllers" from vManage

## Step 6: Onboard to vManage

1. Power connected devices
2. Open DC Internet edge Firewall ports
3. Ensure control plane connections establish
4. Ensure devices managed by vManage inventory

## Step 7: Validate DC routing

1. Ensure reachability from vManage
2. Ensure LAN routing to DC
3. Ensure default route to MPLS and INET aggregation routers in DC
4. Ensure no routing loops or instability with DC WAN edge routers activated



---

## Step 8: Verify NetOps

1. Verify normal NetOps checkouts
2. Integration with NMS - (AAA, Logging, SNMP, Cflowd)

Once a data center 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.

### 4.3.6 Single IWAN Branch Router Inline Migration

**Before you deploy an IOS XE router in the overlay network, ensure the following:**

1. The controller devices—vBond orchestrators, vManage NMSs, and vSmart controllers—are running Cisco SD-WAN Software Release 18.3.
2. If you deploy both IOS XE and vEdge routers in the overlay network, the vEdge routers are running Release 17.2.1 or higher of the Cisco SD-WAN software. With these software versions, the vEdge and IOS XE software can interoperate, allowing BFD tunnels to be established between vEdge routers and IOS XE routers.
3. If you deploy both IOS XE and vEdge routers in the same site, the vEdge routers are running Cisco SD-WAN Software Release 18.3.
4. The ISR 4000 series router has at least 4 gigabytes (GB) of DRAM installed. It is recommended that the router have 8 GB of DRAM.
5. The ASR 1000 series router has at least 8 GB of DRAM installed. The ASR 1002-HX router has at least 16 GB of DRAM installed.
6. The router's bootflash has a minimum of 1.5 GB space available for the XE SD-WAN image.
7. If using your enterprise root certificate to authenticate the router, the certificate is copied to the router's bootflash before installing the XE SD-WAN software.
8. All unsupported modules are removed from the router before installing the XE SD-WAN software. For a list of supported modules, see Supported Interface Modules and Supported Crypto Modules.
9. The updated device list is uploaded to the vManage NMS and sent to the vBond orchestrator. To do so:
10. Obtain the router's chassis and board ID serial number by issuing the `show crypto pki certificates CISCO_IDEVID_SUDI` command at the system prompt. If running Release 16.6.1 or earlier on an ASR series router, issue the `show sdwan certificate serial` command.
11. Add the router's serial number to Plug and Play (PnP) Connect portal. See Add the IOS XE Router to the PnP Portal.
12. In the vManage NMS Configuration ► Devices screen, click the Sync Smart Account button to download the updated device list to vManage NMS and send it to the vBond orchestrator.
13. Device configuration templates are created and attached to the router using the vManage NMS Configuration ► Templates screen. This ensures that the router can obtain a configuration and establish full control connections when it comes up.

14. If the router exceeds the unidirectional encrypted bandwidth of 250 Mbps and if the HSECK9 license is not already installed, the license file is copied to the router's bootflash and license installed on the router license install file path
15. The ASR 1000 series, ISR 1000 series, and ISR 4000 series router is running the required version of the ROM monitor software (ROMMON), as shown in the table below. To verify the ROMMON version running on the router, issue the show rom-monitor or show platform command at the system prompt.

Required ROM Monitor Software Version	
ASR 1000 series	16.3 (2r)
ISR 1000 series	16.9 (1r)
ISR 4000 series	16.7 (3r)

### 4.3.7 Summary of bootstrap migration Steps

#### Prior to Migration:

1. vManage NMS: Verify that the WAN edge device information (product ID, chassis number and device serial number) is present in the devices pane. This may require a manual upload of the latest device serial file from the PnP portal, or if enabled, a "Sync smart account" action on vManage
2. vManage NMS: From the device certificates pane, move the WAN Edge device from "valid" to "staging" mode. (This mode allows control plane tunnels to be built, but prevents forwarding plane tunnels to be built until full verification can be completed during a maintenance window)
3. vManage NMS: Attach device template to WAN edge device and supply site-specific values for template variables to generate full device configuration
4. vManage NMS: Generate bootstrap configuration file from devices pane
5. Branch Router: Upgrade router rommon to required version (if necessary)
6. Branch Router: Upload the new IOS XE SD-WAN image onto the router bootflash
7. Branch Router: Copy the bootstrap configuration (ciscosdwan.cfg) generated by vManage to the router bootflash (bootflash:/ciscosdwan.cfg)
8. Branch Router: Remove existing boot statements and add boot variable that points to the new IOS XE SD-WAN image
9. Branch Router: Save the existing configuration as a named file in the router bootflash


#### During Migration

1. Branch Router: Reboot the router
2. Branch Router: Verify the router comes up on the target image in controller mode
3. Branch Router: If router comes up in autonomous mode, the previous IWAN config will be present - Change to controller mode which will trigger a reboot and reformat

4. Branch Router: Verify control connections formed and WAN edge router receives its full configuration file from vManage, and is placed into vManage mode
5. vManage NMS: Place WAN Edge device in valid mode in the certificates pane
6. vManage NMS: Verify WAN Edge control plane and forwarding plane (BFD sessions) formed.

**Backout Procedure (if required)**

1. Change branch router to autonomous mode (configure terminal, “controller-mode disable”) which will trigger a reboot, erase filesystem and SD-WAN configuration
2. Escape from configuration dialog, change the boot statement to original IOS XE image running IWAN and reboot a second time
3. Copy saved IWAN configuration from the router bootflash to running configuration
4. Verify IWAN control and forwarding planes



**Figure 7: 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

	<b>LAN Service VPNs</b>	Connect to LAN switches in service VPNs. LAN design will dictate if sub-interfaces are needed.
<b>WAN Advertisements VPN 0</b>	<b>IN</b>	SD-WAN prefixes, aggregate routes and default GW from data center from OMP session over Internet and MPLS connections
	<b>OUT</b>	Redistribute local LAN prefixes into OMP
<b>LAN Advertisements Service VPNs</b>	<b>IN</b>	Local LAN prefixes – SD-WAN Edge router typically is the GW
	<b>OUT</b>	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

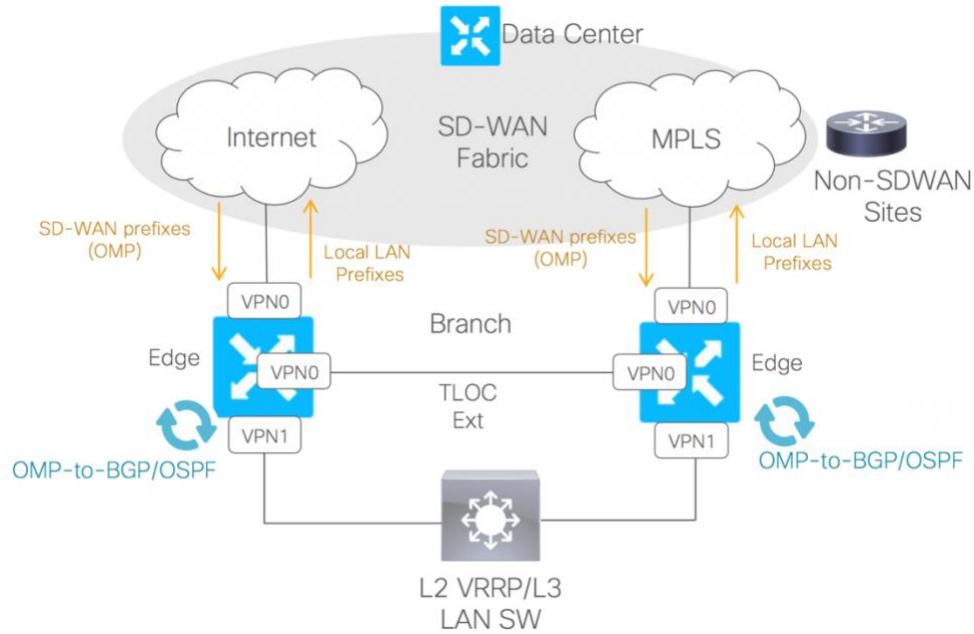
**Table 8: Branch 1 SD-WAN Routers Connectivity and Routing**

#### 4.3.8 Dual IWAN Branch Router Inline Migration

In scenarios where a branch has two IWAN routers, downtime can be minimized during the migration. 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. Also consider, that 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.

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 data center site to maintain the Performance Routing (PfR) operation only on the prefixes remaining on IWAN.



**Figure 8: Dual IWAN Router 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 data center. 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 9: Branch Inline Deployment**

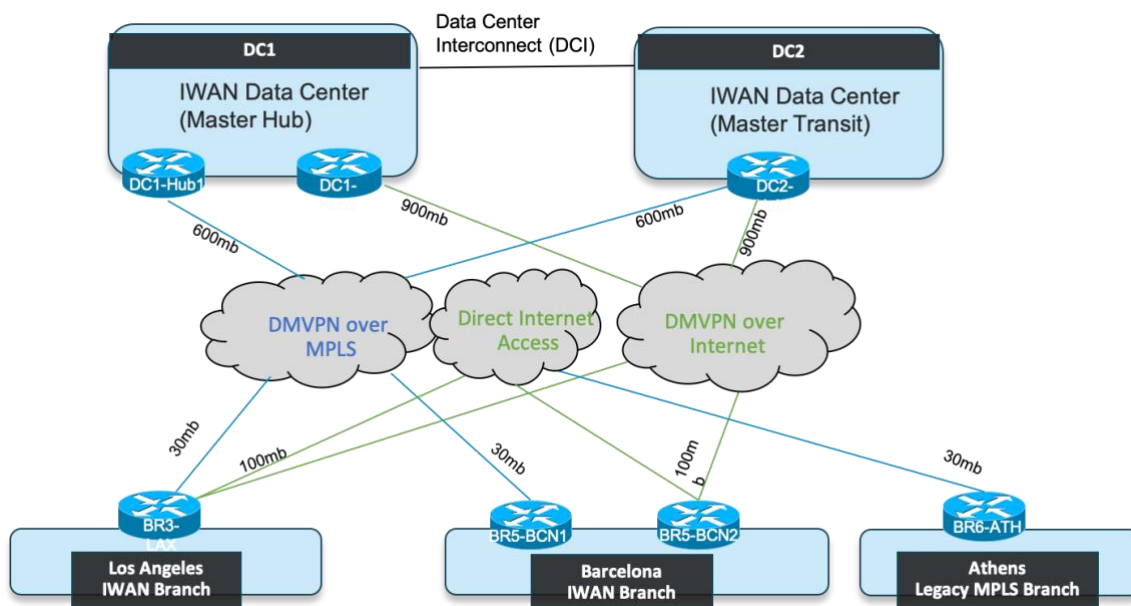
For complex migration scenarios like parallel migration where the existing CE is retained to maintain usage of existing customer services offered by native CE refer to [Cisco SD-WAN migration guide section 4.2.3](#)

# 5 Customer IWAN to Cisco SD-WAN Migration Case Study

This section describes a customer case study for IWAN to Cisco SD-WAN migration. For the purposes of this migration guide, the output for this case study was generated from a simulated customer environment in a Cisco lab. The goal of the case study is to provide guidelines on how information from Section 2, 3 and 4 can be used in migration.

## 5.1.1 Legacy IWAN Deployment Overview

An enterprise customer deployed a Cisco Intelligent WAN (IWAN) as part of a branch refresh project that replaced aging ISRG2 routers that were reaching end of life with ISR4K XE routers. As part of this refresh, broadband Internet circuits were brought to each branch to increase branch bandwidth as existing MPLS circuits were becoming congested at most locations. The business intent for IWAN specifies critical applications should continue to use MPLS services when available and within performance SLAs, switching over to Internet during impairments or disruptions of the MPLS service. Noncritical applications are forwarded on to Internet transports, falling back to MPLS during complete ISP circuit failures. Direct Internet access (DIA) has been enabled at the IWAN branches to improve performance for Internet browsing and Cloud/SaaS applications. By enabling DIA, performance is optimized by removing additional latency incurred by backhauling through one of the centralized Internet gateways at the Data Centers.



**Figure 1: Enterprise IWAN Topology**

---

## 5.1.2 IWAN Use Cases Deployed

The primary IWAN use cases for this Case Study are:

- Secure site-to-site VPN between remote sites and Data Centers with DMVPN over hybrid MPLS and Internet transports
- Application-aware routing with PfRv3 with performance monitoring, SLA protection and preferred path selection
- Secure Direct Internet Access at the branch - Local Internet exit with Zone Based Firewall protection and fallback to MPLS overlay path to DC

## 5.1.3 Planned Cisco SD-WAN Design

The Cisco SD-WAN architecture building blocks include WAN edge routers in the forwarding plane and centralized controllers that handle management, control and orchestration plane tasks. These controllers can be deployed and managed as a service by Cisco Cloud Operations in the cloud (AWS, Azure or GC), or deployed and managed by the customer on premise.

After careful evaluation of each option, this customer decided to self-deploy SD-WAN controllers in a carrier neutral facility (CNF) that was currently being used to host other enterprise services. With this decision, the customer assumed all responsibilities of an on-premise controller deployment and ongoing operations. The deployment tasks included:

- Installation of the compute and storage appliances
- Deployment of the SD-WAN controller VMs
- Initial setup and onboarding of the Controllers to the overlay
- Internet edge security between the CNF and enterprise network.

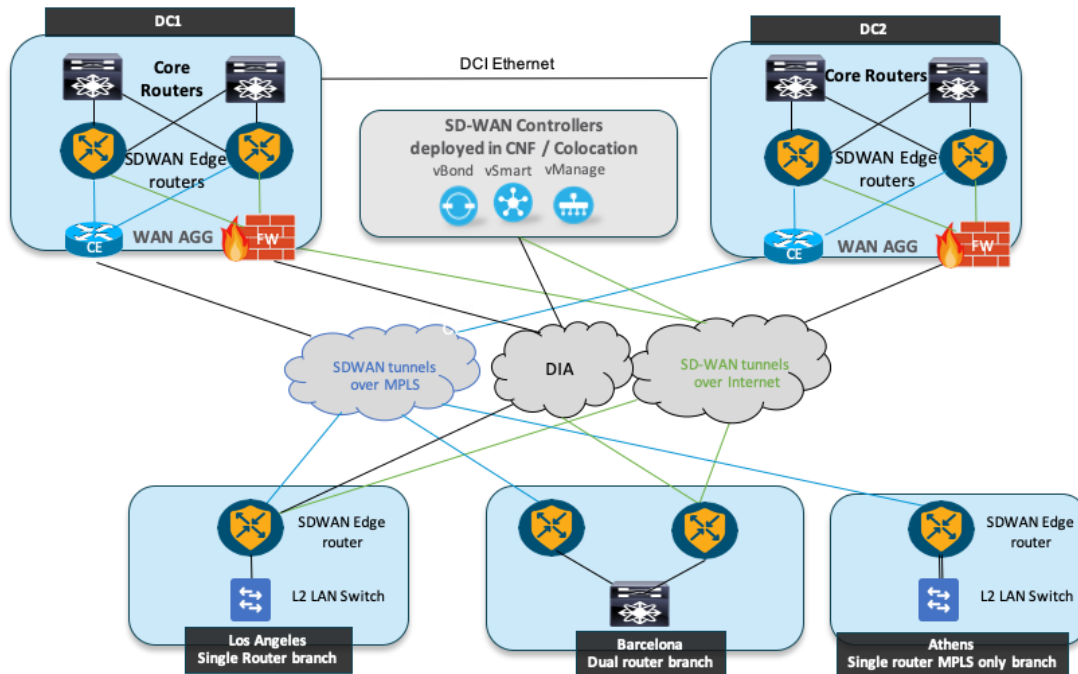
The customer also owns all of the routing operational tasks such as certificate management, monitoring, troubleshooting, capacity planning, backup and restore.

IP connectivity between the WAN edge routers and controllers is gained directly from the branch Internet transports, and indirectly via the MPLS backhaul to the DC Internet breakout points in the DMZ.

In order to prepare for the branch migrations, new SD-WAN edge routers were installed in each DC, connecting to the DC LAN core and WAN aggregation devices (MPLS CE routers and Internet edge Firewalls) No new hardware was needed at the remote sites as the ISR4K IWAN routers would be migrated to SD-WAN WAN edge routers through code upgrades. The same WAN transports that were used for IWAN DMVPN tunnels would be used for SD-WAN IPsec tunnels.

**The following diagram captures the end state SD-WAN design, once all sites are migrated:**





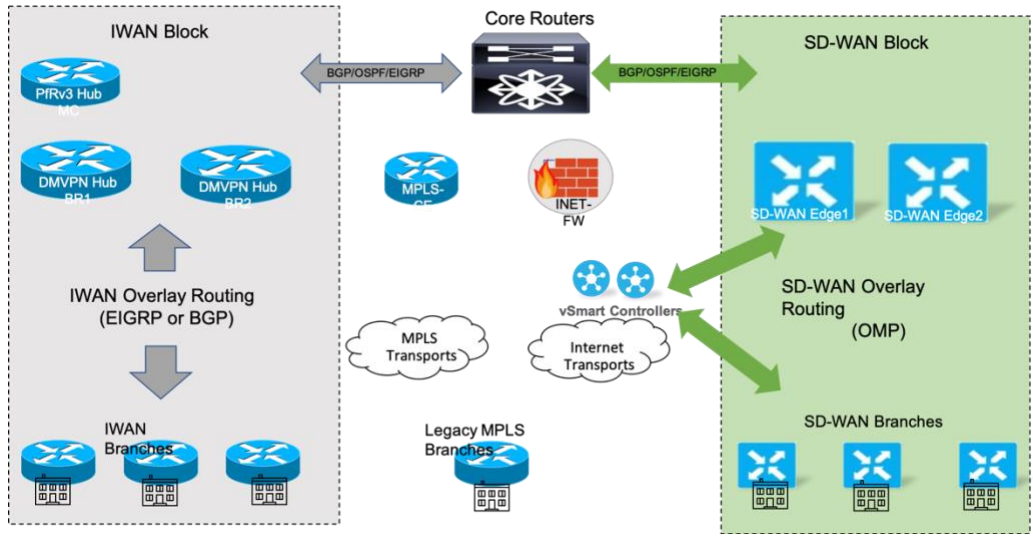
**Figure 1: Planned SD-WAN design**

#### 5.1.4 SD-WAN Use Cases

- Secure Automated WAN: Foundational use case of SD-WAN including controllers and WAN edge routers that interconnect over an IPsec tunnel fabric utilizing hybrid MPLS and Internet transports. The Secure Automated WAN policy for this customer defines a hub and spoke mesh of persistent tunnels between DC and branch routers, with dynamic on-demand tunnels enabled to create temporary site to site tunnels when required. The policy also utilizes a primary and secondary DC preference that prefers DC1 for all upstream or Internet backhaul connectivity.
- Application Performance Optimization: Application aware routing policies that offer SLA protection and preferred path selection for specific applications. (App-Route SLA classes derived from the IWAN Pfrv3 domain policy)
- Secure Direct Internet Access at the branch: Local Internet exit with Zone Based Firewall protection and fallback to MPLS overlay path via DC Internet breakout

#### 5.1.5 Migration State – Parallel IWAN and SD-WAN Infrastructure

During the transition, some sites remain on the IWAN as others are incrementally migrated to SD-WAN. This requires a transition period where the IWAN and SD-WAN planes run independently, side by side. The Data Center is where the IWAN and SD-WAN overlays come together from an aggregation and inter-routing standpoint. An SD-WAN block consisting of two new WAN Edge aggregation routers is installed beside the existing IWAN Block as shown in the diagram below:



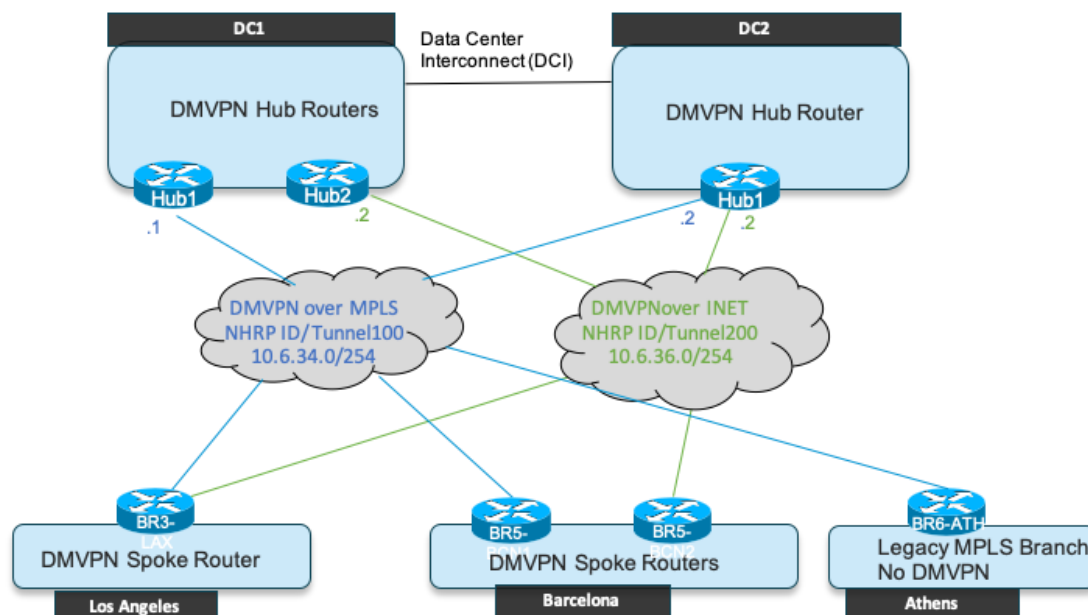
**Figure 1: IWAN and SD-WAN parallel deployments**

## 5.2 IWAN Deployment Deep Dive

It is important to understand the legacy IWAN deployment in order to extract the use cases and features that provide the business intent. This will ensure that appropriate use cases and features are included in the SD-WAN design. This section describes the IWAN design based on a deep dive of the configurations and command line output.

### 5.2.1 DMVPN Design and Business Intent

Enterprise business-critical applications are hosted out of DC1, with DC2 serving as a backup DC for production apps and primary for development. Direct branch to branch traffic is allowed over DMVPN to support peer-to-peer video collaboration apps and file sharing. Transit routing between DC1 and DC2 is facilitated by a Data Center Interconnect leveraging BGP for dynamic routing. The primary DC1 site has separate DMVPN Hubs for resilience, where the backup DC2 has a single Hub router terminating both MPLS and Internet.



**Figure 2: DMVPN deployment**

### 5.2.2 The DMVPN design can be summarized as:

- Hub and Spoke DMVPN design with DC1 (primary) and DC2 (secondary) serving as hubs, Los Angeles and Barcelona serving as spokes.
- Dual transport DMVPN with all sites connecting to common MPLS carrier in addition to business class Internet from multiple carriers.
- IPsec Encryption with Pre-shared keys for IKE authentication
- DMVPN Phase3 shortcuts enabled for spoke-to-spoke dynamic tunnels
- BGP routing on top of the DMVPN overlay with local pref policies that enforce the business intent for primary/secondary DC and transport selection

- 8 class Enterprise QoS model, with sub-rate shaping on the WAN interfaces, queuing/scheduling, and remarking to 5 SP classes
- Per-tunnel DMVPN QoS enabled to ensure that low-speed branch circuits are not overrun by high volumes of data coming from the high bandwidth Hubs.

### 5.2.3 DMVPN Verification

The following command line interface output is shown to better understand the DMVPN overlay design.

- show dmvpn
- show dmvpn detail
- show ip bgp summary
- show ip route
- show policy-map interface <>

#### DMVPN Hub Verification - DC1

Show DMVPN detail on the hub router shows which spokes have registered and additionally the per-tunnel QoS NHRP group and dynamic shaper applied to the output QoS Policy

```

DC1-SJC-HUBBR-1#show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
N - NATed, L - Local, X - No Socket
T1 - Route Installed, T2 - Nexthop-override
C - CTS Capable, I2 - Temporary
# Ent --> Number of NHRP entries with same NBMA peer
NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
UpDn Time --> Up or Down Time for a Tunnel
=====
Interface Tunnel100 is up/up, Addr. is 10.6.34.1, VRF ""
Tunnel Src./Dest. addr: 10.1.101.21/Multipoint, Tunnel VRF "IWAN-TRANSPORT-2"
Protocol/Transport: "multi-GRE/IP", Protect "DMVPN-IPSEC-PROFILE"
Interface State Control: Disabled
nhrp event-publisher : Disabled
Type:Hub, Total NBMA Peers (v4/v6): 2
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb Target Network
-----
1 10.30.0.2 10.6.34.41 UP 00:05:06 D 10.6.34.41/32
NHRP group: RS-GROUP-30MBPS-80
Output QoS service-policy applied: RS-GROUP-30MBPS-80-POLICY
1 10.50.0.2 10.6.34.45 UP 00:05:53 D 10.6.34.45/32
NHRP group: RS-GROUP-30MBPS-80
Output QoS service-policy applied: RS-GROUP-30MBPS-80-POLICY

Crypto Session Details:
-----
Interface: Tunnel100
Session: [0x7FBC9DC76E38]
Session ID: 11
IKEv2 SA: local 10.1.101.21/500 remote 10.30.0.2/500 Active
Output QoS service-policy applied: RS-GROUP-30MBPS-80-POLICY
Capabilities:U connid:1 lifetime:23:54:53
Crypto Session Status: UP-ACTIVE
fvrf: IWAN-TRANSPORT-2, Phase1 id: 172.16.3.1
IPSEC FLOW: permit 47 host 10.1.101.21 host 10.30.0.2
Active SAs: 2, origin: crypto map
Inbound: #pkts dec'ed 368 drop 0 life (KB/Sec) 4607941/3293
Outbound: #pkts enc'ed 375 drop 0 life (KB/Sec) 4607958/3293

```

```

Outbound SPI : 0xA6663EC2, transform : esp-gcm 256
Socket State: Open
Interface: Tunnel100
Session: [0x7FBC9DC76FB8]
Session ID: 10
IKEv2 SA: local 10.1.101.21/500 remote 10.50.0.2/500 Active
Capabilities:U connid:2 lifetime:23:54:06
Crypto Session Status: UP-ACTIVE
fvrf: IWAN-TRANSPORT-2, Phase1_id: 172.16.5.2
IPSEC FLOW: permit 47 host 10.1.101.21 host 10.50.0.2
Active SAs: 2, origin: crypto map
Inbound: #pkts dec'ed 410 drop 0 life (KB/Sec) 4607939/3246
Outbound: #pkts enc'ed 435 drop 0 life (KB/Sec) 4607952/3246
Outbound SPI : 0xECBACFD9, transform : esp-gcm 256
Socket State: Open
Pending DMVPN Sessions:

```

### Branch DMVPN Spoke Verification - Los Angeles Branch Router

Check to see which DMVPN next hop server (Hub) peers are configured and their state. An UP state indicates peers are reachable, and configured properly, with IPsec encryption. In this case there are two peers for each tunnel. Tunnel 100 and Tunnel 200 have peers working to Hub1 in DC1 and Hub1 in DC2

```

BR3-LAX-MCBR#show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
N - NATed, L - Local, X - No Socket
T1 - Route Installed, T2 - Nexthop-override
C - CTS Capable, I2 - Temporary
# Ent --> Number of NHRP entries with same NBMA peer
NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
UpDn Time --> Up or Down Time for a Tunnel
=====
Interface: Tunnel100, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
1 10.1.101.21 10.6.34.1 UP 00:06:18 S
1 10.1.102.5 10.6.34.2 UP 00:06:18 S
Interface: Tunnel200, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
1 192.0.2.10 10.6.36.1 UP 00:03:12 S
1 192.0.2.20 10.6.36.2 UP 00:06:38 S

```

Show DMVPN detail provides more statistics, particularly about IPsec encryption

```

BR3-LAX-MCBR#show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
N - NATed, L - Local, X - No Socket
T1 - Route Installed, T2 - Nexthop-override
C - CTS Capable, I2 - Temporary
# Ent --> Number of NHRP entries with same NBMA peer
NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
UpDn Time --> Up or Down Time for a Tunnel
=====
Capabilities:DNU connid:5 lifetime:23:58:10
Interface Tunnel100 is up/up, Addr. is 10.6.34.41, VRF ""

```

```

Tunnel Src./Dest. addr: 10.30.0.2/Multipoint, Tunnel VRF "IWAN-TRANSPORT-1"
Protocol/Transport: "multi-GRE/IP", Protect "DMVPN-IPSEC-PROFILE"
Interface State Control: Enabled
nhrp event-publisher : Disabled
IPv4 NHS:
10.6.34.1 RE NBMA Address: 10.1.101.21 priority = 0 cluster = 0
10.6.34.2 RE NBMA Address: 10.1.102.5 priority = 0 cluster = 0
Type:Spoke, Total NBMA Peers (v4/v6): 2
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb Target Network
-----
1 10.1.101.21 10.6.34.1 UP 00:02:25 S 10.6.34.1/32
1 10.1.102.5 10.6.34.2 UP 00:02:25 S 10.6.34.2/32
Interface Tunnel200 is up/up, Addr. is 10.6.36.41, VRF ""
Tunnel Src./Dest. addr: 192.0.2.40/Multipoint, Tunnel VRF "IWAN-TRANSPORT-2"
Protocol/Transport: "multi-GRE/IP", Protect "DMVPN-IPSEC-PROFILE"
Interface State Control: Enabled
nhrp event-publisher : Disabled
IPv4 NHS:
10.6.36.1 RE NBMA Address: 192.0.2.10 priority = 0 cluster = 0
10.6.36.2 RE NBMA Address: 192.0.2.20 priority = 0 cluster = 0
Type:Spoke, Total NBMA Peers (v4/v6): 2
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb Target Network
-----
1 192.0.2.10 10.6.36.1 UP 00:01:47 S 10.6.36.1/32
1 192.0.2.20 10.6.36.2 UP 00:01:58 S 10.6.36.2/32

```

Crypto Session Details:

```

-----
Interface: Tunnel100
Session: [0x7F9F8895E590]
Session ID: 10
IKEv2 SA: local 10.30.0.2/500 remote 10.1.101.21/500 Active
Capabilities:DU connid:2 lifetime:23:57:34
Crypto Session Status: UP-ACTIVE
fvrf: IWAN-TRANSPORT-1, Phasel_id: 172.16.1.2
IPSEC FLOW: permit 47 host 10.30.0.2 host 10.1.101.21
Active SAs: 2, origin: crypto map
Inbound: #pkts dec'ed 185 drop 0 life (KB/Sec) 4607972/3453
Outbound: #pkts enc'ed 165 drop 0 life (KB/Sec) 4607980/3453
Outbound SPI : 0xB3E3DA8F, transform : esp-gcm 256
Socket State: Open
Interface: Tunnel100
Session: [0x7F9F8895E710]
Session ID: 9
IKEv2 SA: local 10.30.0.2/500 remote 10.1.102.5/500 Active
Capabilities:DU connid:1 lifetime:23:57:34
Crypto Session Status: UP-ACTIVE
fvrf: IWAN-TRANSPORT-1, Phasel_id: 172.16.3.4
IPSEC FLOW: permit 47 host 10.30.0.2 host 10.1.102.5
Active SAs: 2, origin: crypto map
Inbound: #pkts dec'ed 185 drop 0 life (KB/Sec) 4607972/3454
Outbound: #pkts enc'ed 140 drop 0 life (KB/Sec) 4607984/3454
Outbound SPI : 0x643E23BA, transform : esp-gcm 256
Socket State: Open
Interface: Tunnel200
Session: [0x7F9F8895E290]
Session ID: 12
IKEv2 SA: local 192.0.2.40/4500 remote 192.0.2.10/4500 Active
Capabilities:DNU connid:5 lifetime:23:58:10
Crypto Session Status: UP-ACTIVE
fvrf: IWAN-TRANSPORT-2, Phasel_id: 172.16.1.3
IPSEC FLOW: permit 47 host 192.0.2.40 host 192.0.2.10
Active SAs: 2, origin: crypto map
Inbound: #pkts dec'ed 98 drop 0 life (KB/Sec) 4607985/3490

```

```

Outbound: #pkts enc'ed 130 drop 0 life (KB/Sec) 4607986/3490
Outbound SPI : 0xFC938C6F, transform : esp-gcm 256
Socket State: Open
Interface: Tunnel200
Session: [0x7F9F8895E410]
Session ID: 11
IKEv2 SA: local 192.0.2.40/500 remote 192.0.2.20/500 Active
Capabilities:DU connid:4 lifetime:23:58:02
Crypto Session Status: UP-ACTIVE
fvrf: IWAN-TRANSPORT-2, Phasel_id: 172.16.3.4
IPSEC FLOW: permit 47 host 192.0.2.40 host 192.0.2.20
Active SAs: 2, origin: crypto map
Inbound: #pkts dec'ed 145 drop 0 life (KB/Sec) 4607979/3481
Outbound: #pkts enc'ed 108 drop 0 life (KB/Sec) 4607988/3481
Outbound SPI : 0x8DC164AB, transform : esp-gcm 256
Socket State: Open
Pending DMVPN Sessions:

```

Dynamic peers between spokes are created on demand when traffic flows from spoke to spoke. In this case a ping from LAX to BCN triggers the dynamic tunnel.

```

BR3-LAX-MCBR#ping ip 10.5.100.1 source 10.3.100.1 repeat 100
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 10.5.100.1, timeout is 2 seconds:
Packet sent with a source address of 10.3.100.1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 8/14/93 ms
BR3-LAX-MCBR#
*Jul 10 18:07:54.178: %CRYPTO-4-RECVD_PKT_INV_SPI: decaps: rec'd IPSEC packet has invalid
spi for destaddr=10.30.0.2, prot=50, spi=0x74F38A(7664522), srcaddr=10.50.0.2, input
interface=Tunnel100show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
N - NATed, L - Local, X - No Socket
T1 - Route Installed, T2 - Nexthop-override
C - CTS Capable, I2 - Temporary
# Ent --> Number of NHRP entries with same NBMA peer
NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
UpDn Time --> Up or Down Time for a Tunnel
=====
Interface: Tunnel100, IPv4 NHRP Details
Type:Spoke, NHRP Peers:3,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
2 10.50.0.2 10.6.34.45 UP 00:00:05 DT1
10.6.34.45 UP 00:00:05 DT1
1 10.1.101.21 10.6.34.1 UP 00:10:18 S
1 10.1.102.5 10.6.34.2 UP 00:10:17 S
Interface: Tunnel200, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
1 192.0.2.10 10.6.36.1 UP 00:09:40 S
1 192.0.2.20 10.6.36.2 UP 00:09:50 S
BR3-LAX-MCBR#

```

## 5.2.4 DMVPN Hub Router Configuration walkthrough (DC1-SJC-BR1)

The DC1-SJC-BR1 (MPLS DMVPN Hub) configuration (Crypto, tunnels, QoS and routing) is shown below: This section provides the Hub configuration example (DC1-SJC-BR-1, MPLS hub at DC1). The table below is the DMVPN tunnel and encryption configuration template:

Commands	Description
<pre>vrf definition IWAN-TRANSPORT-2 ! address-family ipv4 exit-address-family</pre>	<p>DMVPN transport routes segmented into a VRF, referred to front door VRF (fVRF)</p> <p>Each DMVPN cloud has dedicated fVRF. In this case transport2 is MPLS</p>
<pre>tunnel source GigabitEthernet3 tunnel mode gre multipoint tunnel key 100 tunnel vrf IWAN-TRANSPORT-2</pre>	<p>DMVPN tunnel source specifies interface, mode, key and tunnel VRF.</p>
<pre>tunnel protection ipsec profile DMVPN-IPSEC-PROFILE</pre>	<p>Tunnel protection encrypts all traffic that traverses the DMVPN tunnel</p>
<pre>no ip redirects ip mtu 1400</pre>	
<pre>nhrp map group RS-GROUP-300MBPS-80 service-policy output RS-GROUP-300MBPS-80-POLICY nhrp map group RS-GROUP-200MBPS-80 service-policy output RS-GROUP-200MBPS-80-POLICY nhrp map group RS-GROUP-100MBPS-80 service-policy output RS-GROUP-100MBPS-80-POLICY nhrp map group RS-GROUP-50MBPS-80 service-policy output RS-GROUP-50MBPS-80-POLICY nhrp map group RS-GROUP-30MBPS-80 service-policy output RS-GROUP-30MBPS-80-POLICY nhrp map group RS-GROUP-20MBPS-80 service-policy output RS-GROUP-20MBPS-80-POLICY nhrp map group RS-GROUP-10MBPS-80 service-policy output RS-GROUP-10MBPS-80-POLICY nhrp map group RS-GROUP-4G-80 service-policy output RS-GROUP-4G-80-POLICY</pre>	<p>NHRP groups defined on hub routers for Per-tunnel QoS traffic shaping towards remote sites.</p>
<pre>ip route vrf IWAN-TRANSPORT-2 0.0.0.0 0.0.0.0 10.1.101.22</pre>	<p>Static default route in the front VRF, next hop is the CE router in the DC.</p>
<pre>ip nhrp authentication cisco123 ip nhrp network-id 100 ip nhrp server-only ip nhrp redirect</pre>	<p>all peers on same overlay must have matching authentication password and network id. Hub routers are NHRP servers for spokes which requires redirect to trigger Phase 3</p>
<pre>ip address 10.6.34.1 255.255.254.0</pre>	<p>All routers on overlay addressed from common subnet</p>
<pre>interface Tunnel100 description MPLS1</pre>	<p>Tunnel 100 connected to MPLS overlay at all hubs and spokes</p>
<pre>interface GigabitEthernet3 description to DC1 MPLS CE router vrf forwarding IWAN-TRANSPORT-2 ip address 10.1.101.21 255.255.255.252</pre>	<p>Physical interface that connects to the WAN and represents tunnel source for DMVPN.</p>



	<p>In case of DC this interface connects to a shared MPLS CE router.</p> <p>The IP subnet assigned to this interface must be announced into the MPLS provider (typically BGP from CE to PE)</p> <p>Physical interface mapped to front VRF</p>
<pre>crypto ipsec security-association replay window-size 1024 ! crypto ipsec transform-set AES256/GCM/TRANSFORM esp-gcm 256 mode transport ! crypto ipsec profile DMVPN-IPSEC-PROFILE set transform-set AES256/GCM/TRANSFORM set ikev2-profile DMVPN-IKEv2-PROFILE</pre>	
<pre>crypto ikev2 proposal AES/GCM/256 encryption aes-gcm-256 prf sha512 group 19</pre>	Crypto ikev2 proposal to peers
<pre>crypto ikev2 profile DMVPN-IKEv2-PROFILE description PSK Profile match fvrf any match identity remote address 0.0.0.0 identity local address 172.16.1.2 authentication remote pre-share authentication local pre-share keyring local DMVPN-KEYRING</pre>	Pre-shared key profile.
<pre>crypto ikev2 policy AES/GCM/256 match fvrf any proposal AES/GCM/256</pre>	crypto ikev2 policy

## 5.2.5 Table: Hub DMVPN Tunnel and Encryption Template Walkthrough

The Hub site is configured with BGP as shown in below table:

Commands	Description
<pre>router bgp 64510 bgp router-id 172.16.1.2 bgp log-neighbor-changes bgp listen range 10.6.34.0/23 peer-group MPLS1- SPOKES neighbor MPLS1-SPOKES peer-group neighbor MPLS1-SPOKES remote-as 64510 neighbor MPLS1-SPOKES description MPLS1 Spoke Route Reflector neighbor MPLS1-SPOKES update-source Tunnel100 neighbor MPLS1-SPOKES timers 20 60 neighbor 10.1.101.10 remote-as 65010 neighbor 10.1.101.10 description DC1-SJC-CORE1 !</pre>	<p>BGP chosen as the routing protocol for DMVPN at hub and spoke sites which are all in the same AS (iBGP peering). Also, as LAN routing protocol for DC/Hub sites</p> <p>Hub routers use dynamic peers (listen range) so that so to avoid pre-defining neighbors. Caveat is the spokes must have tunnel addresses in this range.</p> <p>eBGP to Core router to send/receive site routes</p>
<pre>address-family ipv4 bgp redistribute-internal network 10.0.0.0 network 172.16.1.1 mask 255.255.255.255 redistribute connected neighbor MPLS1-SPOKES activate neighbor MPLS1-SPOKES send-community neighbor MPLS1-SPOKES route-reflector-client neighbor MPLS1-SPOKES next-hop-self all neighbor MPLS1-SPOKES weight 50000 neighbor MPLS1-SPOKES soft-reconfiguration inbound neighbor MPLS1-SPOKES route-map MPLS1-IN in neighbor MPLS1-SPOKES route-map MPLS1-OUT out neighbor 10.1.101.10 activate maximum-secondary-paths ibgp 1 distance bgp 201 19 200 exit-address-family</pre>	<p>Hub routers originate:</p> <p>The enterprise aggregate (10.0.0.0/8 in this case) - to attract first packets from traffic sourced from one branch and sent to another (spoke to spoke). This is required to switch over to phase 3 DMVPN (spoke to spoke) as the hub triggers NHRP signaling to spokes for underlay address.</p> <p>The PfR Hub Master controller address (172.16.1.1/32), which must be reachable to all remote sites.</p> <p>Other eBGP routes are learned by Hub BR's and announced into IBGP towards remote spokes if they pass the conditions of outbound route-map</p> <p>BGP route-maps to and from spokes applied to peer-groups.</p>
<pre>route-map MPLS1-IN deny 10 description All Blocked Prefixes to come IN on BGP match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIXES LOCALDC-PREFIX LOCALMCLOOPBACK TUNNEL-DMVPN ! route-map MPLS1-IN permit 1000 description Allow Everything Else ! route-map MPLS1-OUT permit 10 description All Allowed Prefixes to Go OUT on BGP to Spokes match ip address prefix-list DEFAULT-ROUTE ENTERPRISE-PREFIXES LOCALDC-PREFIX LOCALMCLOOPBACK set local-preference 800 set community 4266393700</pre>	<p>BGP route-map details</p> <p>MPLS1-IN used as a method of loop protection to prevent hubs from learning default, DC routes or enterprise aggregate</p> <p>MPLS1-OUT used to filter routes to only local DC routes, Enterprise aggregate, default and Hub Master controller loopback. Set local-preference is used as a means to prefer particular hubs and transports</p> <p>INET1-IN and INET1-OUT are similarly configured on the DMVPN routers connected to the Internet overlay</p> <p>BGP local-pref used for traffic steering priorities from Branch to Hub</p> <p>DC1 MPLS Hub router with local pref 800 to make it primary</p> <p>DC1 Internet Hub with local pref 700 to make it secondary path</p>

	<p>DC2 MPLS = MPLS Hub router with local pref 600 to make it third</p> <p>DC2 INET = Internet Hub router with local pref 500 to make it forth</p>
<pre>ip prefix-list DEFAULT-ROUTE seq 10 permit 0.0.0.0/0 ! ip prefix-list ENTERPRISE-PREFIXES seq 10 permit 10.0.0.0/8 ip prefix-list ENTERPRISE-PREFIXES seq 20 permit 172.16.0.0/12 ip prefix-list ENTERPRISE-PREFIXES seq 30 permit 192.168.0.0/16 ! ip prefix-list LOCALDC-PREFIX seq 10 permit 10.1.100.0/24 ! ip prefix-list LOCALMCLOOPBACK seq 10 permit 172.16.1.1/32 ! ip prefix-list TUNNEL-DMVPN seq 10 permit 10.6.34.0/23 ip prefix-list TUNNEL-DMVPN seq 20 permit 10.6.36.0/23</pre>	<p>Prefix lists that are referenced by route-maps</p>
<pre>router bgp 64510 bgp router-id 172.16.1.2 bgp log-neighbor-changes bgp listen range 10.6.34.0/23 peer-group MPLS1-SPOKES neighbor MPLS1-SPOKES peer-group neighbor MPLS1-SPOKES remote-as 64510 neighbor MPLS1-SPOKES description MPLS1 Spoke Route Reflector neighbor MPLS1-SPOKES update-source Tunnel100 neighbor MPLS1-SPOKES timers 20 60 neighbor 10.1.101.10 remote-as 65010 neighbor 10.1.101.10 description DC1-SJC-CORE1 !</pre>	<p>BGP chosen as the routing protocol for DMVPN at hub and spoke sites which are all in the same AS (iBGP peering). Also, as LAN routing protocol for DC/Hub sites</p> <p>Hub routers use dynamic peers (listen range) so that so avoid pre-defining neighbors. Caveat is the spokes must have tunnel addresses in this range.</p> <p>eBGP to Core router to send/receive site routes</p>

**Table: Hub BGP routing (DMVPN/WAN and LAN routing)**

The Hub site is configured with following QoS configuration:

Commands	Description
<pre>class-map match-any STREAMING-VIDEO match dscp cs5 match dscp af31 match dscp af32 match dscp af33 class-map match-any INTERACTIVE-VIDEO match dscp cs4 match dscp af41 match dscp af42 match dscp af43 class-map match-any CRITICAL-DATA match dscp cs2 match dscp af11 match dscp af12</pre>	<p>8 class Enterprise QoS model. Expects upstream DSCP marking.</p>

<pre> match dscp af13 match dscp af21 match dscp af22 match dscp af23 class-map match-any VOICE match dscp ef class-map match-any SCAVENGER match dscp cs1 class-map match-any CALL-SIGNALING match dscp cs3 class-map match-any NET-CTRL match dscp cs6 </pre>	
<pre> policy-map WAN class INTERACTIVE-VIDEO bandwidth remaining percent 30 random-detect dscp-based set dscp tunnel af41 class STREAMING-VIDEO bandwidth remaining percent 10 random-detect dscp-based set dscp tunnel af31 class NET-CTRL bandwidth remaining percent 5 set dscp tunnel cs6 class CALL-SIGNALING bandwidth remaining percent 4 set dscp tunnel af21 class CRITICAL-DATA bandwidth remaining percent 25 random-detect dscp-based set dscp tunnel af21 class SCAVENGER bandwidth remaining percent 1 set dscp tunnel af11 class VOICE priority level 1 police cir percent 10 set dscp tunnel ef class class-default bandwidth remaining percent 25 random-detect set dscp tunnel default </pre>	<p>Child policy-map WAN, specifying schedulers and DSCP rewrite to comply with carrier</p>
<pre> policy-map RS-GROUP-30MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 24000000 bandwidth remaining ratio 24 service-policy WAN policy-map TRANSPORT-2-SHAPE-ONLY class class-default shape average 900000000 policy-map RS-GROUP-10MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 8000000 bandwidth remaining ratio 8 service-policy WAN policy-map RS-GROUP-4G-80-POLICY description 80% of RS inbound service rate class class-default </pre>	<p>Parent policy that includes shapers and nested policy (WAN). This is applied to interface.</p>

<pre> shape average 6000000 bandwidth remaining ratio 6 service-policy WAN policy-map RS-GROUP-300MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 240000000 bandwidth remaining ratio 240 service-policy WAN policy-map RS-GROUP-20MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 16000000 bandwidth remaining ratio 16 service-policy WAN policy-map TRANSPORT-1-SHAPE-ONLY class class-default shape average 600000000 policy-map RS-GROUP-50MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 40000000 bandwidth remaining ratio 40 service-policy WAN policy-map RS-GROUP-100MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 80000000 bandwidth remaining ratio 80 service-policy WAN policy-map RS-GROUP-200MBPS-80-POLICY description 80% of RS inbound service rate class class-default shape average 160000000 bandwidth remaining ratio 160 service-policy WAN ! </pre>	
<pre> interface GigabitEthernet3 description to DC1 MPLS CE router vrf forwarding IWAN-TRANSPORT-2 ip address 10.1.101.21 255.255.255.252 negotiation auto no mop enabled no mop sysid service-policy output TRANSPORT-1-SHAPE-ONLY </pre>	WAN facing interface with service-policy applied
<pre> class-map match-any STREAMING-VIDEO match dscp cs5 match dscp af31 match dscp af32 match dscp af33 class-map match-any INTERACTIVE-VIDEO match dscp cs4 match dscp af41 match dscp af42 match dscp af43 class-map match-any CRITICAL-DATA match dscp cs2 match dscp af11 match dscp af12 match dscp af13 match dscp af21 match dscp af22 </pre>	8 class Enterprise QoS model. Expects upstream DSCP marking.

---

<pre>match dscp af23 class-map match-any VOICE match dscp ef class-map match-any SCAVENGER match dscp cs1 class-map match-any CALL-SIGNALING match dscp cs3 class-map match-any NET-CTRL match dscp cs6</pre>	
---	--

***Table: Hub WAN QoS Configurations***

## 5.2.6 The BR3-LAX-MCBR branch DMVPN configuration (Crypto, tunnels, QoS and routing) is shown below:

Commands	Description
<pre>vrf definition IWAN-TRANSPORT-1 ! address-family ipv4 exit-address-family ! vrf definition IWAN-TRANSPORT-2 ! address-family ipv4 exit-address-family !</pre>	Two fVRFs are defined, one for each WAN connection/overlay
<pre>crypto ikev2 proposal AES/GCM/256 encryption aes-gcm-256 prf sha512 group 19</pre>	Crypto ikev2 proposal to peers
<pre>crypto ikev2 policy AES/GCM/256 match fvrf any proposal AES/GCM/256</pre>	crypto ikev2 policy
<pre>crypto ikev2 keyring DMVPN-KEYRING peer ANY address 0.0.0.0 0.0.0.0 pre-shared-key cisco123</pre>	Customer is using pre-shared keys for IPsec authentication rather than certificates.
<pre>crypto ikev2 profile DMVPN-IKEv2-PROFILE description PSK Profile match fvrf any match identity remote address 0.0.0.0 identity local address 172.16.3.1 authentication remote pre-share authentication local pre-share keyring local DMVPN-KEYRING dpd 40 5 on-demand</pre>	Same configuration as Hub, but with dead peer detection (DPD) configured to ensure IPsec liveness. (CPU intensive and not configured on hub)
<pre>class-map match-any STREAMING-VIDEO match dscp af31 af32 af33 cs5 class-map match-any INTERACTIVE-VIDEO match dscp cs4 af41 af42 af43 class-map match-any CRITICAL-DATA match dscp af11 af12 af13 cs2 af21 af22 af23 class-map match-any VOICE match dscp ef class-map match-any SCAVENGER match dscp cs1 class-map match-any CALL-SIGNALING match dscp cs3 class-map match-any NET-CTRL match dscp cs6 ! policy-map WAN class INTERACTIVE-VIDEO bandwidth remaining percent 30 random-detect dscp-based set dscp af41 class STREAMING-VIDEO bandwidth remaining percent 10 random-detect dscp-based set dscp af31</pre>	QoS configs for WAN

<pre> class NET-CTRL bandwidth remaining percent 5 set dscp cs6 class CALL-SIGNALING bandwidth remaining percent 4 set dscp af21 class CRITICAL-DATA bandwidth remaining percent 25 random-detect dscp-based set dscp af21 class SCAVENGER bandwidth remaining percent 1 set dscp af11 class VOICE priority level 1 police cir percent 10 set dscp ef class class-default bandwidth remaining percent 25 random-detect policy-map POLICY-TRANSPORT-1 class class-default shape average 30000000 service-policy WAN policy-map POLICY-TRANSPORT-2 class class-default shape average 50000000 service-policy WAN ! </pre>	
<pre> crypto ipsec security-association replay window- size 1024 ! crypto ipsec transform-set AES256/GCM/TRANSFORM esp-gcm 256 mode transport ! crypto ipsec profile DMVPN-IPSEC-PROFILE set transform-set AES256/GCM/TRANSFORM set ikev2-profile DMVPN-IKEv2-PROFILE </pre>	Crypto (IPsec) parameters
<pre> interface Loopback0 ip address 172.16.3.1 255.255.255.255 </pre>	
<pre> interface Tunnel100 description MPLS1 bandwidth 30000 ip address 10.6.34.41 255.255.254.0 no ip redirects ip mtu 1400 ip pim dr-priority 0 ip pim sparse-mode ip nhrp authentication cisco123 ip nhrp network-id 100 ip nhrp nhs 10.6.34.1 nbma 10.1.101.21 multicast ip nhrp nhs 10.6.34.2 nbma 10.1.102.5 multicast ip tcp adjust-mss 1360 no nhrp route-watch if-state nhrp tunnel source GigabitEthernet2 tunnel mode gre multipoint tunnel key 100 tunnel vrf IWAN-TRANSPORT-1 </pre>	<p>The DMVPN tunnel configuration of a spoke (branch) requires a static underlay to overlay IP definition of the hub router addresses.</p> <p>NHS addresses are the overlay addresses on the tunnels.</p> <p>NBMA addresses are the underlay and must be reachable from each router's VRF routing table (IWAN-TRANSPORT-1 in this case).</p> <pre> ip nhrp nhs 10.6.34.1 nbma 10.1.101.21 multicast ip nhrp nhs 10.6.34.2 nbma 10.1.102.5 multicast </pre>



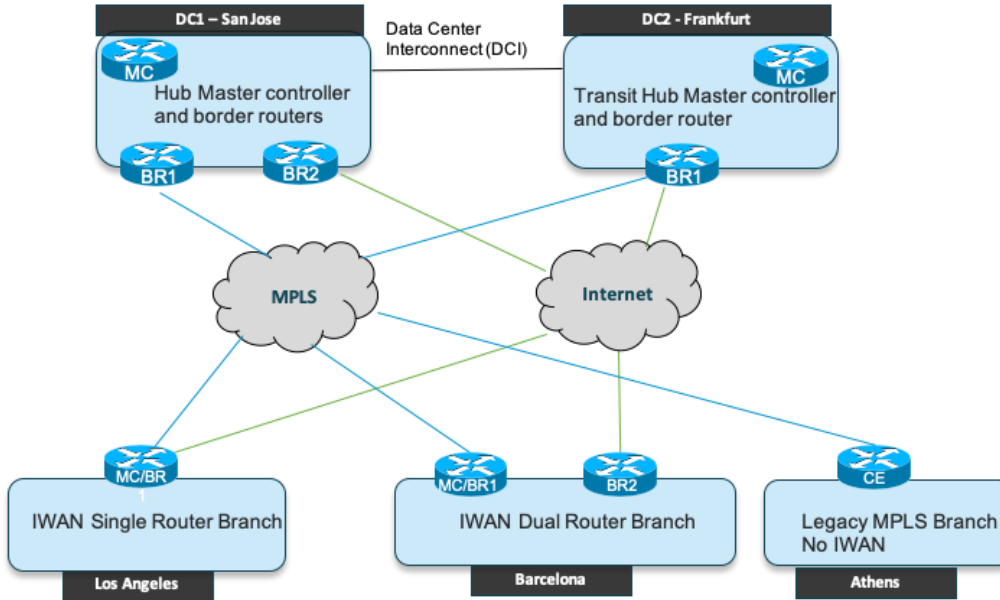
<pre>tunnel protection ipsec profile DMVPN-IPSEC-PROFILE</pre>	
<pre>interface Tunnel200 description INET1 bandwidth 50000 ip address 10.6.36.41 255.255.254.0 no ip redirects ip mtu 1400 ip pim dr-priority 0 ip pim sparse-mode ip nhrp authentication cisco123 ip nhrp network-id 200 ip nhrp nhs 10.6.36.1 nbma 192.0.2.10 multicast ip nhrp nhs 10.6.36.2 nbma 192.0.2.20 multicast ip tcp adjust-mss 1360 no nhrp route-watch if-state nhrp tunnel source GigabitEthernet1 tunnel mode gre multipoint tunnel key 200 tunnel vrf IWAN-TRANSPORT-2 tunnel protection ipsec profile DMVPN-IPSEC-PROFILE</pre>	
<pre>interface GigabitEthernet1 description INET vrf forwarding IWAN-TRANSPORT-2 ip address 192.0.2.40 255.255.255.0 ip access-group ACL-INET-PUBLIC in negotiation auto service-policy output POLICY-TRANSPORT-2 ! interface GigabitEthernet2 description MPLS vrf forwarding IWAN-TRANSPORT-1 ip address 10.30.0.2 255.255.255.252 negotiation auto service-policy output POLICY-TRANSPORT-1 ip route vrf IWAN-TRANSPORT-1 0.0.0.0 0.0.0.0 10.30.0.1 ip route vrf IWAN-TRANSPORT-2 0.0.0.0 0.0.0.0 192.0.2.41  ip access-list extended ACL-INET-PUBLIC permit udp any any eq non500-isakmp permit udp any any eq isakmp permit esp any any permit udp any any eq bootpc permit icmp any any echo permit icmp any any echo-reply permit icmp any any ttl-exceeded permit icmp any any port-unreachable permit udp any any gt 1023 ttl eq 1</pre>	<p>Physical Interface in front door VRFs</p> <p>ACL applied to Internet for edge protection</p> <p>static default routes in each fVRF</p>
<pre>interface GigabitEthernet3 description LAN ip address 10.3.100.1 255.255.255.0 negotiation auto</pre>	
<pre>router bgp 64510 bgp router-id 172.16.3.1 bgp log-neighbor-changes neighbor MPLS1-HUB peer-group</pre>	<p>BGP configuration</p>

<pre> neighbor MPLS1-HUB remote-as 64510 neighbor MPLS1-HUB description To IWAN MPLS1 Hub Router neighbor MPLS1-HUB update-source Tunnel100 neighbor MPLS1-HUB timers 20 60 neighbor INET1-HUB peer-group neighbor INET1-HUB remote-as 64510 neighbor INET1-HUB description To IWAN INET1 Hub Router neighbor INET1-HUB update-source Tunnel200 neighbor INET1-HUB timers 20 60 neighbor 10.6.34.1 peer-group MPLS1-HUB neighbor 10.6.34.2 peer-group MPLS1-HUB neighbor 10.6.36.1 peer-group INET1-HUB neighbor 10.6.36.2 peer-group INET1-HUB neighbor 10.30.0.1 remote-as 65000 neighbor 10.30.0.1 description PE router neighbor 10.30.0.1 shutdown ! address-family ipv4 redistribute connected neighbor 10.6.34.1 activate neighbor 10.6.34.2 activate neighbor 10.6.36.1 activate neighbor 10.6.36.2 activate neighbor 10.30.0.1 activate exit-address-family </pre>	
<pre> vrf definition IWAN-TRANSPORT-1 ! address-family ipv4 exit-address-family ! vrf definition IWAN-TRANSPORT-2 ! address-family ipv4 exit-address-family ! </pre>	<p>Two fVRFs are defined, one for each WAN connection/overlay</p>

## 5.2.7 Performance Routing (PfRv3) Deployment

The Performance Routing architecture elements include the Master Controllers (MC) and Border Routers (BR). Each site has a local MC that is responsible for compiling performance statistics and selecting the outbound path for applications that are bound to an SLA. The border router is the device in the forwarding path that gathers performance of the WAN paths and reports to the MC, forwarding traffic out tunnel interfaces directly to the WAN or indirectly towards an adjacent BR.

In the domain, there is one domain controller that maintains and distributes policies to the other MCs, referred to as the Hub MC. This customer has deployed the Hub MC in DC1.



The PfRv3 design can be summarized as:

- The hub MC is located in the customer DC1, this is where performance policies are configured and distributed to the other MC's in the domain.
- In the case of DC2, the Master controller is considered a transit MC, as traffic from the backbone or other sites may pass through it.
- DC1 and DC2 have dedicated Master Controllers (MCs) that control their Hub Border Routers (BR's) which are also the DMVPN hub router.
- In the LAX branch, the branch MC and BR functions are co-located on the same router.
- In the case of the BCN branch (dual router), one router is both MC/BR and the other router is a standalone BR.

## Performance Routing (PfR) Deployment

Six performance traffic-classes have been defined, with the following characteristics and path-preferences

Class	SLA	Preferred-Path	DSCP Match
<b>VOICE</b>	priority 2 packet-loss-rate threshold 1.0 percent priority 1 one-way-delay threshold 150 msec priority 3 jitter threshold 30000 usec priority 2 byte-loss-rate threshold 1.0 percent	MPLS1	EF
<b>SCAVENGER</b>	priority 2 packet-loss-rate threshold 50.0 percent priority 1 one-way-delay threshold 500 msec priority 2 byte-loss-rate threshold 50.0 percent	INET1	CS1
<b>REAL_TIME_VIDEO</b>	priority 1 packet-loss-rate threshold 1.0 percent priority 1 one-way-delay threshold 500 msec priority 2 one-way-delay threshold 150 msec priority 3 jitter threshold 20000 usec priority 1 byte-loss-rate threshold 1.0 percent	MPLS1	CS4, AF41, AF42, AF43
<b>LOW_LATENCY_DATA</b>	priority 2 packet-loss-rate threshold 5.0 percent priority 1 one-way-delay threshold 100 msec priority 2 byte-loss-rate threshold 5.0 percent	MPLS1	CS2, CS3, AF21, AF22, AF23
<b>DEFAULT</b>	priority 2 packet-loss-rate threshold 10.0 percent priority 1 one-way-delay threshold 500 msec priority 2 byte-loss-rate threshold 10.0 percent	INET1	default
<b>BULK_DATA</b>	priority 2 packet-loss-rate threshold 5.0 percent priority 1 one-way-delay threshold 300 msec priority 2 byte-loss-rate threshold 5.0 percent	INET1	AF11, AF12, AF13

## PfRv3 Verification

The following command output is shown to better understand the Performance routing design.

- `show ip route <prefix>` (check IP routing table prior to optimization)
- `show domain <name> master status` (issue on branch MC)
- `show domain <name> master policy` (issue on branch MC)
- `show domain <name> master traffic-class` (issue on branch MC - once traffic is running)

BR5 BCN output (BR5-BCN-Core is running BGP with BR5-BCN-MCBR1 and BR5-BCN-BR2):

### Check routing from each router at Barcelona branch to prefix in DC1 (10.1.100.0/24)

#### 1. First check routing on MC, which is directly connected to MPLS

```
BR5-BCN-MCBR1#show ip route 10.1.100.0
Routing entry for 10.1.100.0/24
Known via "bgp 64510", distance 200, metric 0
Tag 65010, type internal
Last update from 10.6.34.1 18:45:36 ago
Routing Descriptor Blocks:
* 10.6.34.1, from 10.6.34.1, 18:45:36 ago
Route metric is 0, traffic share count is 1
AS Hops 1
Route tag 65010
MPLS label: none
```

#### 2. Then check routing from BR, which is connected to INET

```
BR5-BCN-BR2#show ip route 10.1.100.0
Routing entry for 10.1.100.0/24
Known via "bgp 64510", distance 200, metric 0
Tag 65010, type internal
Last update from 10.6.36.1 18:46:05 ago
Routing Descriptor Blocks:
* 10.6.36.1, from 10.6.36.1, 18:46:05 ago
Route metric is 0, traffic share count is 1
AS Hops 1
Route tag 65010
MPLS label: none
```

#### 3. Finally, check routing from Legacy Core router, which runs BGP to both IWAN routers. Notice that 10.5.1.1 is preferred, which is the MC connected to MPLS

```
br5-bcn-rtr#show ip bgp 10.1.100.0
BGP routing table entry for 10.1.100.0/24, version 4
Paths: (2 available, best #2, table default)
Advertised to update-groups:
4
Refresh Epoch 1
64510 65010
10.5.2.1 from 10.5.2.1 (172.16.5.2)
Origin incomplete, localpref 100, valid, external
rx pathid: 0, tx pathid: 0
Refresh Epoch 1
64510 65010
10.5.1.1 from 10.5.1.1 (172.16.5.1)
Origin incomplete, localpref 100, valid, external, best
```

```
rx pathid: 0, tx pathid: 0x0
br5-bcn-rtr#
```

#### 4. Traceroute to DC1 prefix (DSCP default) and observe path through MC (10.5.1.1) and MPLS Hub router at DC (10.6.34.1), following routing path

```
br5-bcn-rtr#trace 10.1.100.1
Type escape sequence to abort.
Tracing the route to 10.1.100.1
VRF info: (vrf in name/id, vrf out name/id)
 1 10.5.1.1 2 msec 3 msec 1 msec
 2 10.6.34.1 [AS 64510] 9 msec 15 msec 20 msec
 3 10.1.101.10 [AS 64510] 11 msec 22 msec *
```

#### 5. Now Check Pfrv3 control plane status on Master Controller

```
BR5-BCN-MCBR1#show domain iwan master status
*** Domain MC Status ***
Master VRF: Global
Instance Type: Branch
Instance id: 0
Operational status: Up
Configured status: Up
Loopback IP Address: 172.16.5.1
Load Balancing:
Operational Status: Up
Max Calculated Utilization Variance: 0%
Last load balance attempt: never
Last Reason: Variance less than 20%
Total unbalanced bandwidth:
External links: 0 Kbps Internet links: 0 Kbps
Route Control: Enabled
Transit Site Affinity: Enabled
95% Bandwidth Check: Enabled
Monitor cache usage: 4000 (20%) Auto allocated
Load Sharing: Enabled
Connection Keepalive: 10 seconds
Mitigation mode Aggressive: Disabled
Policy threshold variance: 20
Minimum Mask Length Internet: 24
Minimum Mask Length Enterprise: 24
Syslog TCA suppress timer: 180 seconds
Traffic-Class Ageout Timer: 5 minutes
Minimum Packet Loss Calculation Threshold: 15 packets
Minimum Bytes Loss Calculation Threshold: 1 bytes
Branch to Branch Traffic Control: Enabled
Direct Cloud Access : Disabled
Maximum Traffic Classes Supported: 4000
Minimum Requirement: Met
Borders:
IP address: 172.16.5.2
Version: 2
Connection status: CONNECTED (Last Updated 02:09:04 ago )
Interfaces configured:
Name: Tunnel200 | type: external | Service Provider: INET1 | Status: UP | Zero-SLA: NO |
Path of Last Resort: Disabled
Number of default Channels: 2
Path-id list: 0:2 1:4
Tunnel if: Tunnel0
```

```
IP address: 172.16.5.1
Version: 2
Connection status: CONNECTED (Last Updated 02:09:01 ago )
Interfaces configured:
Name: Tunnel100 | type: external | Service Provider: MPLS1 | Status: UP | Zero-SLA: NO |
Path of Last Resort: Disabled
Number of default Channels: 2
Path-id list: 1:3 0:1
Tunnel if: Tunnel0
```

#### Check BR status on both routers

```
BR5-BCN-MC BR1# show domain iwan border status
Mon Sep 14 17:30:51.658
-----
**** Border Status ****
Instance Status: UP
Present status last updated: 19:59:12 ago
Loopback: Configured Loopback0 UP (172.16.5.1)
Master: 172.16.5.1
Master version: 2
Connection Status with Master: UP
MC connection info: CONNECTION SUCCESSFUL
Connected for: 02:09:08
Route-Control: Enabled
Asymmetric Routing: Disabled
Minimum Mask Length Internet: 24
Minimum Mask Length Enterprise: 24
Connection Keepalive: 10 seconds
Sampling: off
Channel Unreachable Threshold Timer: 4 seconds
Minimum Packet Loss Calculation Threshold: 15 packets
Minimum Byte Loss Calculation Threshold: 1 bytes
Monitor cache usage: 4000 (20%) Auto allocated
Minimum Requirement: Met
Smart Probe Profile:
General Monitor:
Current Provision Level: Master Hub
Master Hub:
Packets per burst: 1
Interval(secs): 1
Quick Monitor:
Current Provision Level: Master Hub
Master Hub:
Packets per burst: 20
Interval(secs): 1
Notification to PD:
add: 1, upd: 0, del: 0
External Wan interfaces:
Name: Tunnel100 Interface Index: 17 SNMP Index: 12 SP: MPLS1 Status: UP Zero-SLA: NO Path of
Last Resort: Disabled Path-id List: 1:3, 0:1
Auto Tunnel information:
Name:Tunnel0 if_index: 18
Virtual Template: Not Configured
Borders reachable via this tunnel: 172.16.5.2
-----

BR5-BCN-BR2#show domain iwan border status
Mon Sep 14 17:30:32.116
-----
**** Border Status ****
```

```

Instance Status: UP
Present status last updated: 19:58:49 ago
Loopback: Configured Loopback0 UP (172.16.5.2)
Master: 172.16.5.1
Master version: 2
Connection Status with Master: UP
MC connection info: CONNECTION SUCCESSFUL
Connected for: 02:08:52
Route-Control: Enabled
Asymmetric Routing: Disabled
Minimum Mask Length Internet: 24
Minimum Mask Length Enterprise: 24
Connection Keepalive: 10 seconds
Sampling: off
Channel Unreachable Threshold Timer: 4 seconds
Minimum Packet Loss Calculation Threshold: 15 packets
Minimum Byte Loss Calculation Threshold: 1 bytes
Monitor cache usage: 4000 (20%) Auto allocated
Minimum Requirement: Met
Smart Probe Profile:
General Monitor:
Sampling: off
Current Provision Level: Master Hub
Master Hub:
Packets per burst: 1
Interval(secs): 1
Quick Monitor:
Current Provision Level: Master Hub
Master Hub:
Packets per burst: 20
Interval(secs): 1
Notification to PD:
add: 1, upd: 0, del: 0
External Wan interfaces:
Name: Tunnel200 Interface Index: 17 SNMP Index: 12 SP: INET1 Status: UP Zero-SLA: NO Path o
f Last Resort: Disabled Path-id List: 0:2, 1:4
Auto Tunnel information:
Name:Tunnel0 if_index: 18
Virtual Template: Not Configured
Borders reachable via this tunnel: 172.16.5.1

```

Now check policy to observe the preferred path and SLA for unmarked traffic (DCSP default). Notice INET1 is the preferred path with MPLS1 backup.

#### Check MC state (control plane state) on MC/BR1

```

BR5-BCN-MC BR1#show domain iwan master traffic-classes
Dst-Site-Prefix: 10.1.100.0/24 DSCP: default [0] Traffic class id:22
Clock Time: 15:29:12 (UTC) 09/14/2020
TC Learned: 00:04:41 ago
Present State: CONTROLLED
Current Performance Status: in-policy
Current Service Provider: INET1 since 00:04:10
Previous Service Provider: Unknown
BW Used: 163 bps
Present WAN interface: Tunnel200 in Border 172.16.5.2
Present Channel (primary): 25 INET1 pfr-label:0:2 | 0:0 [0x20000]
Backup Channel: 28 MPLS1 pfr-label:0:1 | 0:0 [0x10000]
Destination Site ID bitmap: 1
Destination Site ID: 172.16.1.1 (Active)

```



```
Class-Sequence in use: 60
Class Name: DEFAULT using policy best-effort
BW Updated: 00:03:12 ago
Reason for Latest Route Change: Uncontrolled to Controlled Transition
Route Change History:
Date and Time Previous Exit C
urrent Exit Reason
```

#### Check BR state (forwarding plane state) on MC/BR1 and BR2

(notice that Primary interface on MCBR1 is Tunnel0, which is the automatically generated "autotunnel" between MC/BR1 and BR2)

```
BR5-BCN-MCBR1#show domain iwan border traffic-classes
Src-Site-Prefix: ANY Dst-Site-Prefix: 10.1.100.0/24
DSCP: default [0] Traffic class id: 28
TC Learned: 00:01:48 ago
Present State: CONTROLLED
Destination Site ID: 172.16.1.1
If_index: 18
Primary chan id: 25
Primary chan Presence: NEIGHBOR_CHANNEL via border 172.16.5.2
Primary interface: Tunnel0
Backup chan id: 28
Backup chan Presence: LOCAL CHANNEL
Backup interface: Tunnel100
Direct Cloud Access : Disabled
```

#### On BR2, the primary interface is Tunnel200 (DMVPN over INET tunnel)

```
BR5-BCN-BR2#show domain iwan border traffic-classes
Src-Site-Prefix: ANY Dst-Site-Prefix: 10.1.100.0/24
DSCP: default [0] Traffic class id: 28
TC Learned: 00:01:31 ago
Present State: CONTROLLED
Destination Site ID: 172.16.1.1
If_index: 17
Primary chan id: 25
Primary chan Presence: LOCAL CHANNEL
Primary interface: Tunnel200
Primary Nexthop: 10.6.36.1 (BGP)
Backup chan id: 28
Backup chan Presence: NEIGHBOR_CHANNEL via border 172.16.5.1
Backup interface: Tunnel0
Direct Cloud Access : Disabled
```

#### Issue another traceroute from core router and note the traffic is now forwarded from MCBR1 to BR2 and then the DMVPN over INET path

```
br5-bcn-rtr#trace 10.1.100.1
Type escape sequence to abort.
Tracing the route to 10.1.100.1
VRF info: (vrf in name/id, vrf out name/id)
 1 10.5.1.1 2 msec 3 msec 2 msec
 2 172.16.5.2 [AS 64510] 2 msec 6 msec 4 msec
 3 10.6.36.1 [AS 64510] 10 msec 30 msec 11 msec
 4 10.1.101.14 [AS 64510] 15 msec 18 msec 18 msec
 5 10.1.100.1 [AS 65010] 9 msec 10 msec *
br5-bcn-rtr#
```

## PfRv3 Configurations

The PfR Domain Controller / Hub Master Controller (DC1-SJC-MC-1) is configured with the following:

Commands	Description
<pre> domain iwan vrf default master hub source-interface Loopback0 site-prefixes prefix-list DC1-PREFIXES password cisco load-balance enterprise-prefix prefix-list ENTERPRISE- PREFIXES class VOICE sequence 10 match dscp ef policy voice path-preference MPLS1 fallback INET1 class REAL_TIME_VIDEO sequence 20 match dscp cs4 policy real-time-video match dscp af41 policy real-time-video match dscp af42 policy real-time-video match dscp af43 policy real-time-video path-preference MPLS1 fallback INET1 class LOW_LATENCY_DATA sequence 30 match dscp cs2 policy low-latency-data match dscp cs3 policy low-latency-data match dscp af21 policy low-latency-data match dscp af22 policy low-latency-data match dscp af23 policy low-latency-data path-preference MPLS1 fallback INET1 class BULK_DATA sequence 40 match dscp af11 policy bulk-data match dscp af12 policy bulk-data match dscp af13 policy bulk-data path-preference INET1 fallback MPLS1 class SCAVENGER sequence 50 match dscp cs1 policy scavenger path-preference INET1 fallback blackhole class DEFAULT sequence 60 match dscp default policy best-effort path-preference INET1 fallback MPLS1 ! </pre>	<p>Domain policies are configured on the Hub MC only; distributed to all other MC (transit DC and branch) over an EIGRP SAF (overlay).</p> <p>The hub MC and transit Data Center MC must have local prefixes defined in a static prefix-list. Branch prefixes can be learned dynamically.</p> <p>Enterprise prefix lists specify the scope of the PfR deployment, which typically includes all of the RFC 1918 space plus any public space that would be routed internally across DMVPN.</p> <p>Performance Classes defined in sequences.</p> <p>Each class has match criteria for classification, which can be DSCP-based or NBAR-based. In this example all performance traffic classes are matched on DSCP markings.</p> <p>Policies can be defined manually by specifying delay, loss, and/or jitter. Otherwise, pre-defined policy templates can be used as in this example (policy 'voice', policy 'real-time-video', etc....).</p> <p>Path-preference allows deterministic traffic steering on a preferred path. In this case, voice traffic will prefer the MPLS1 path as long as it is within the SLA boundaries specified for 'voice'. (OWD &lt; 150 ms, loss &lt; 1%, jitter &lt; 30ms).</p>
<pre> interface Loopback0 ip address 172.16.1.1 255.255.255.255 ! interface GigabitEthernet1 ip address 10.1.101.1 255.255.255.252  ! interface GigabitEthernet2 ip address 10.1.101.5 255.255.255.252 router bgp 64511 bgp router-id 172.16.1.1 bgp log-neighbor-changes redistribute connected neighbor 10.1.101.2 remote-as 65010 neighbor 10.1.101.2 description DC1-SJC-CORE2 neighbor 10.1.101.6 remote-as 65010 neighbor 10.1.101.6 description DC1-SJC-CORE1 </pre>	<p>Interfaces and routing. The Hub MC only needs layer 3 reachability to the Border Routers and is not in the forwarding path.</p>
<pre> ! ip prefix-list DC1-PREFIXES seq 10 permit </pre>	

<pre> 10.1.100.0/24 ! ip prefix-list ENTERPRISE-PREFIXES seq 10 permit 10.1.100.0/24 ip prefix-list ENTERPRISE-PREFIXES seq 20 permit 10.2.100.0/24 ip prefix-list ENTERPRISE-PREFIXES seq 30 permit 10.3.100.0/24 ip prefix-list ENTERPRISE-PREFIXES seq 40 permit 10.4.100.0/24 ip prefix-list ENTERPRISE-PREFIXES seq 50 permit 10.5.100.0/24 ip prefix-list ENTERPRISE-PREFIXES seq 60 permit 10.6.100.0/24 </pre>	
--	--

**Table: PfR Domain Controller / Hub Master Controller (DC1-SJC-MC-1)**

The Hub Border Router (DC1-Hub1) is configured with below PfR configurations

Commands	Description
<pre> domain iwan vrf default border source-interface Loopback0 master 172.16.1.1 password cisco </pre>	<p>DMVPN Hub router is also PfR Border router.</p> <p>Specify source interface and destination IP address for peering to Hub Master Controller, along with password for simple authentication</p>
<pre> interface Tunnel100 description MPLS1 &lt;snip&gt; domain iwan path MPLS1 path-id 1 </pre>	<p>Hub router includes Path name and id that is sent to all branches over DMVPN. This allows branch router dynamic discovery and is used for PfR path selection policies.</p>

**Table: PfR Hub Border Router configuration (DC1 - Hub1)**

The Branch MC/BR (BR3-LAX-MCBR1) is configured with below PfR configurations:

Commands	Description
<pre> domain iwan vrf default border source-interface Loopback0 master local password cisco master branch source-interface Loopback0 password cisco hub 172.16.1.1 </pre>	<p>The single Branch router at LAX is configured as the local branch MC and BR.</p> <p>The hub MC IP address and password must be specified in order to form the SAF peering to exchange policy information between the branch and Hub MC</p>

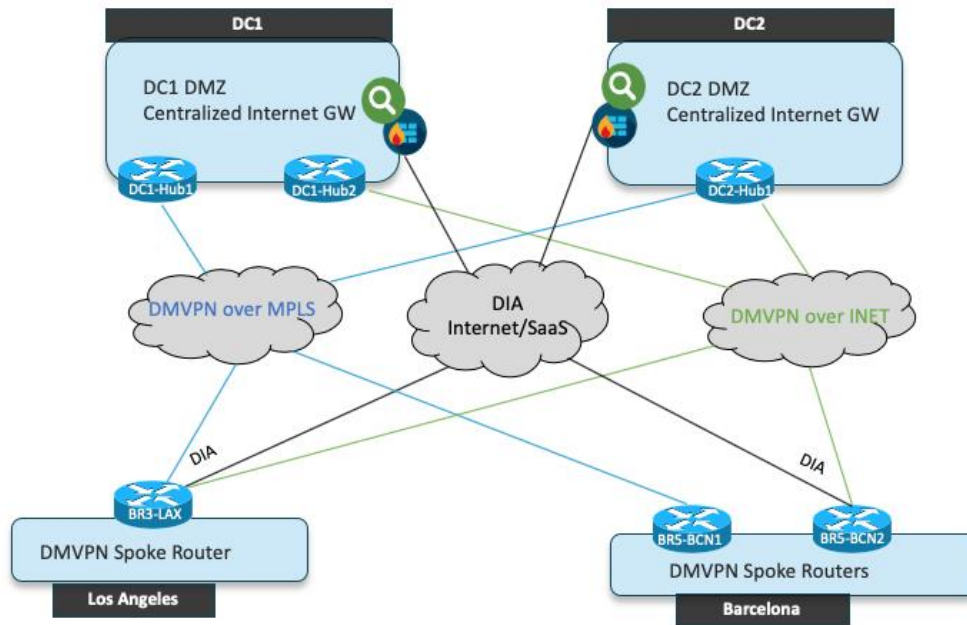
**Table: PfR Branch MC/BR configuration (BR3-LAX-MCBR1)**

## 5.2.8 Direct Internet Access (DIA) Deployment

The customer has deployed Direct Internet Access (DIA) within the IWAN branch sites to allow all Internet-bound traffic and public cloud traffic from the branch to be routed directly to the Internet over the local ISP circuit used as transport for IWAN DMVPN. Enabling DIA can improve performance by removing the latency associated with tunneling the Internet traffic over the WAN to a data center providing Internet access. Zone based firewall is configured for Internet protection on IWAN routers where DIA is configured.

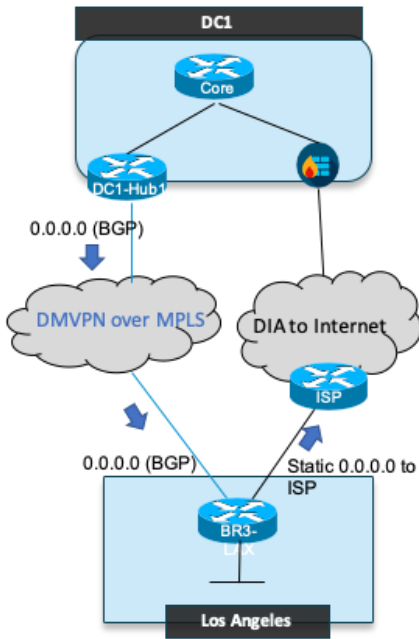
The IWAN requirements for DIA are as follows:

- Primary path for all Internet-bound traffic is the local ISP circuit
- In case of ISP circuit failover or brownout, branch Internet bound traffic rerouted on DMVPN over MPLS overlay to the data center providing centralized Internet access
- Internet traffic never forwarded over DMVPN Internet path for sites having DIA configured
- Sites with no DIA may load-share Internet traffic over both DMVPN overlays



### *DIA Design for Single Router Site (Los Angeles)*

BR3-MCBR1 is the local Internet exit where DIA will occur by means of a static default route towards the ISP. Internet bound traffic reaching this router will be sent directly to the ISP router after being inspected by the Zone Based Firewall. The DIA path will be continuously monitored by router-sourced ICMP probes that will ping Google DNS at 8.8.8.8. In the event of a failure, the static default route will be withdrawn, and traffic will converge onto the DMVPN overlay to the Internet exit at DC1 (following the default route learned by BGP over DMVPN).



### DIA Configurations (Los Angeles)

Commands	Description
<p><b>NAT on Interfaces:</b></p> <pre>interface GigabitEthernet3 description LAN ip address 10.3.100.1 255.255.255.0 ip nat inside ! interface GigabitEthernet1 description INET vrf forwarding IWAN-TRANSPORT-2 ip address 192.0.2.40 255.255.255.0 ip nat outside</pre>	Configure NAT
<p><b>NAT ACL and Route Map</b></p> <pre>ip access-list extended NAT-LOCAL permit ip 10.0.0.0 0.255.255.255 any route-map NAT permit 10 description local Internet NAT for DIA match ip address NAT-LOCAL match interface GigabitEthernet1</pre>	
<p><b>NAT Overload</b></p> <pre>ip nat inside source route-map NAT interface GigabitEthernet1 overload</pre>	
<pre>ip route 0.0.0.0 0.0.0.0 GigabitEthernet1 192.0.2.41 track 1</pre>	Configure Static Default Route for Split tunneling
<pre>ip sla 1 icmp-echo 8.8.8.8 source-interface GigabitEthernet1</pre>	Configure IP SLA object for liveness tracking (ICMP echo probes to 8.8.8.8 used in ip sla 1)

<pre>vrf IWAN-TRANSPORT-2 threshold 2 timeout 1000 frequency 3 ip sla schedule 1 life forever start-time now</pre>	
<pre>zone security default zone security OUTSIDE class-map type inspect match-any INSIDE-TO-OUTSIDE-CLASS match protocol ftp match protocol tcp match protocol udp match protocol icmp policy-map type inspect INSIDE-TO-OUTSIDE-POLICY class type inspect INSIDE-TO-OUTSIDE-CLASS inspect class class-default drop zone-pair security IN_OUT source default destination OUTSIDE service-policy type inspect INSIDE-TO-OUTSIDE-POLICY  interface GigabitEthernet1 description Internet Connection zone-member security OUTSIDE</pre>	<p>Configure Zone Based Firewall to inspect DIA traffic (FTP, TCP, UDP, ICMP)</p>
<pre>ip access-list extended ACL-RTR-IN permit udp any any eq non500-isakmp permit udp any any eq isakmp permit icmp any any echo permit icmp any any echo-reply permit icmp any any ttl-exceeded permit icmp any any port-unreachable permit udp any any gt 1023 ttl eq 1 ! ip access-list extended ACL-RTR-OUT permit udp any any eq non500-isakmp permit udp any any eq isakmp permit icmp any any permit udp any any eq domain ! ip access-list extended DHCP-IN permit udp any eq bootps any eq bootpc ip access-list extended DHCP-OUT permit udp any eq bootpc any eq bootps ! ip access-list extended ESP-IN permit esp any any ip access-list extended ESP-OUT permit esp any any ip access-list extended GRE-IN permit gre any any class-map type inspect match-any INSPECT-ACL-IN-CLASS match access-group name ACL-RTR-IN  class-map type inspect match-any INSPECT-ACL-OUT-CLASS match access-group name ACL-RTR-OUT</pre>	<p>Configure Zone Based Firewall to restrict traffic to the router itself</p> <p>(Note: This method is an advanced alternative to a simple IP ACL for router threat protection from the Internet)</p>

```

class-map type inspect match-any PASS-ACL-IN-
CLASS
match access-group name ESP-IN
match access-group name DHCP-IN
match access-group name GRE-IN

class-map type inspect match-any PASS-ACL-OUT-
CLASS
match access-group name ESP-OUT
match access-group name DHCP-OUT
policy-map type inspect ACL-IN-POLICY
class type inspect INSPECT-ACL-IN-CLASS
inspect
class type inspect PASS-ACL-IN-CLASS
pass
class class-default
drop
policy-map type inspect ACL-OUT-POLICY
class type inspect INSPECT-ACL-OUT-CLASS
inspect
class type inspect PASS-ACL-OUT-CLASS
pass
class class-default
drop

zone-pair security FROM-ROUTER source self
destination OUTSIDE
service-policy type inspect ACL-OUT-POLICY

zone-pair security TO-ROUTER source OUTSIDE
destination self
service-policy type inspect ACL-IN-POLICY

```

## DIA Configurations Verification Single Site Router

Verify DIA at Single Router site (Los Angeles.) Check routing and forwarding path prior to enabling DIA

### Ensure default route learned over DMVPN from MPLS Hub Router in DC1

```

BR3-LAX-MCBR#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
Known via "bgp 64510", distance 200, metric 0, candidate default path
Tag 65010, type internal
Last update from 10.6.34.1 00:22:16 ago
Routing Descriptor Blocks:
* 10.6.34.1, from 10.6.34.1, 00:22:16 ago
Route metric is 0, traffic share count is 1
AS Hops 1
Route tag 65010
MPLS label: none

```

### Traceroute to Internet destination [www.cisco.com](http://www.cisco.com), and observe forwarding across DMVPN MPLS overlay to DC1 prior to DIA

```

BR3-LAX-MCBR#traceroute www.cisco.com source g3
Type escape sequence to abort.
Tracing the route to www.cisco.com (184.25.199.192)
VRF info: (vrf in name/id, vrf out name/id)
 1 10.6.34.1 9 msec 7 msec 3 msec
 2 10.1.101.10 4 msec 8 msec 5 msec
 3 10.1.254.1 4 msec 10 msec 6 msec
 4*

```

```
5*
6. www.cisco.com (184.25.199.192) [AS 65010] 5 msec 5 msec
```

Checks the status of IP SLA ICMP probes to specified destination on the Internet (8.8.8.8) to ensure direct path is alive and not blackholing

```
BR3-LAX-MCBR#show track
Track 1
IP SLA 1 reachability
Reachability is Up
2 changes, last change 00:05:54
Latest operation return code: OK
Latest RTT (milliseconds) 1
Tracked by:
Static IP Routing 0
```

**Check static default route installed with ISP router as next-hop.**

```
BR3-LAX-MCBR#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
Known via "static", distance 1, metric 0, candidate default path
Routing Descriptor Blocks:
* 192.0.2.41, via GigabitEthernet1
Route metric is 0, traffic share count is 1
```

Traceroute to [www.cisco.com](http://www.cisco.com) and note direct hop destination with source IP of LAN interface

```
BR3-LAX-MCBR#traceroute ip www.cisco.com source GigabitEthernet3
Tracing the route to www.cisco.com (184.25.199.192)
VRF info: (vrf in name/id, vrf out name/id)
192.0.2.41 9 msec 7 msec 3 msec
*
www.cisco.com (184.25.199.192) [AS 65010] 3 msec 2 msec 2 msec
```

Check NAT translations

```
BR3-LAX-MCBR#show ip nat translations
Pro Inside global Inside local Outside local Outs
ide global
udp 192.0.2.40:5122 10.3.100.1:49456 184.25.199.192:33436 184.
25.199.192:33436
udp 192.0.2.40:5120 10.3.100.1:49454 184.25.199.192:33434 184.
25.199.192:33434
udp 192.0.2.40:5121 10.3.100.1:49455 184.25.199.192:33435 184.
25.199.192:33435
Total number of translations: 3
```

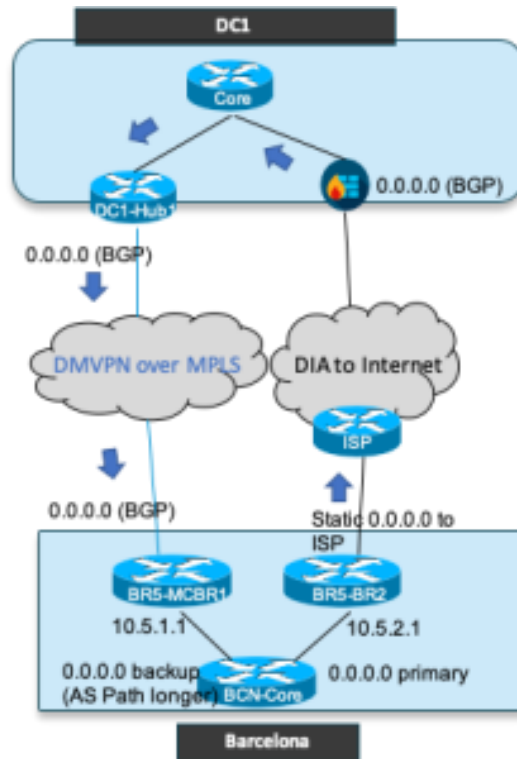
## DIA Design for Dual Router Site (Barcelona)

BR5-BR2 is the local Internet exit where DIA will occur. Internet bound traffic reaching this router will be sent directly over the ISP circuit after being inspected by the Zone Based Firewall.

BR5-MCBR1 is the MPLS connected router with no DIA path. Internet bound traffic reaching this router will be forwarded on the DMVPN overly to DC-HUB1 where it follows the default route to the DC internet exit



Both IWAN routers send default route to BCN-Core L3 router. The path to BR5-BR2 is preferred due to a BGP route-map that raises the local-preference. In the event of a DIA failure (router, circuit or IPSLA tracking), the path will converge to the DMVPN overlay to DC1 via BR5-MCBR1



### DIA Configurations Dual Router Site (Barcelona)

Commands	Description
<p><b>NAT on Interfaces:</b></p> <pre>interface GigabitEthernet1 description link to INET1 vrf forwarding IWAN-TRANSPORT-2 ip address 192.0.2.50 255.255.255.0 ip nat outside ! interface GigabitEthernet2 description link to Core ip address 10.5.2.1 255.255.255.252 ip nat inside</pre>	Configure NAT
<p><b>NAT ACL and Route Map</b></p> <pre>ip access-list extended NAT-LOCAL permit ip 10.0.0.0 0.255.255.255 any route-map NAT permit 10 description local Internet NAT for DIA match ip address NAT-LOCAL match interface GigabitEthernet1</pre>	

<p><b>NAT Overload</b></p> <pre>ip nat inside source route-map NAT interface GigabitEthernet1 overload</pre>	
<pre>ip route 0.0.0.0 0.0.0.0 GigabitEthernet1 192.0.2.51 track 1</pre>	Configure Static Default Route for Split tunneling
<pre>ip sla 1 icmp-echo 8.8.8.8 source-interface GigabitEthernet1 vrf IWAN-TRANSPORT-2 threshold 2 timeout 1000 frequency 3 ip sla schedule 1 life forever start-time now</pre>	Configure IP SLA object for liveness tracking (ICMP to 8.8.8.8)
<pre>route-map BGP2LAN permit 10 match ip address prefix-list DEFAULT set local-preference 1000 ! route-map BGP2LAN permit 20 ! route-map BGPFROMIWAN deny 10 match ip address prefix-list DEFAULT ! route-map BGPFROMIWAN permit 20 ! route-map BGP2IWAN deny 10 match ip address prefix-list DEFAULT ! route-map BGP2IWAN permit 20 ip prefix-list DEFAULT seq 5 permit 0.0.0.0/0</pre>	Configure BGP route maps to control distribution of default route
<pre>router bgp 64510 bgp router-id 172.16.5.2 bgp log-neighbor-changes neighbor INET1-HUB peer-group neighbor INET1-HUB remote-as 64510 neighbor INET1-HUB description To IWAN INET1 Hub Router neighbor INET1-HUB update-source Tunnel200 neighbor INET1-HUB timers 20 60 neighbor 10.5.2.2 remote-as 65054 neighbor 10.5.2.2 description core router neighbor 10.6.36.1 peer-group INET1-HUB neighbor 10.6.36.2 peer-group INET1-HUB ! address-family ipv4 network 172.16.5.1 mask 255.255.255.255 redistribute connected redistribute static neighbor INET1-HUB next-hop-self neighbor INET1-HUB route-map BGPFROMIWAN in neighbor INET1-HUB route-map BGP2IWAN out neighbor 10.5.2.2 activate neighbor 10.5.2.2 route-map BGP2LAN out neighbor 10.6.36.1 activate neighbor 10.6.36.2 activate</pre>	Apply route maps to BGP neighbors (LAN and IWAN peers)

```
default-information originate
exit-address-family
```

## DIA Verification for Dual Router Site (Barcelona)

Default route received from both BCN-NMCBR1 and BCN-BR2 over IBGP. Verify default route on BCN-Core prefers BR5-BR2 (10.5.2.1)

```
br5-bcn-rtr#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
Known via "bgp 65054", distance 20, metric 0, candidate default path
Tag 64510, type external
Last update from 10.5.2.1 01:08:42 ago
Routing Descriptor Blocks:
* 10.5.2.1, from 10.5.2.1, 01:08:42 ago
Route metric is 0, traffic share count is 1
AS Hops 1
Route tag 64510
MPLS label: none
```

Traceroute to [www.cisco.com](http://www.cisco.com) from core router

```
br5-bcn-rtr#traceroute www.cisco.com
Type escape sequence to abort.
Tracing the route to www.cisco.com (184.25.199.192)
VRF info: (vrf in name/id, vrf out name/id)
 1 10.5.2.1 2 msec 2 msec 2 msec
 2 www.cisco.com (184.25.199.192) [AS 64510] 3 msec 3 msec 2 msec
```

Check NAT translations on DIA router (BR5-BCN-BR2)

```
BR5-BCN-BR2#show ip nat translations
Pro Inside global Inside local Outside local Outside global
udp 192.0.2.50:5062 10.5.2.2:49187 184.25.199.192:33437 184.25.199.192:33437
tcp 192.0.2.50:5063 10.5.2.1:21172 52.203.231.173:443 52.203.231.173:443
udp 192.0.2.50:5063 10.5.2.2:49188 184.25.199.192:33438 184.25.199.192:33438
udp 192.0.2.50:5064 10.5.2.2:49189 184.25.199.192:33439 184.25.199.192:33439
Total number of translations: 4
```

Simulate Failover ISP circuit to verify traffic re-routes on MPLS overlay. Verify default route converges to BR5-MCBR1 (10.5.1.1). Then Verify traceroute to destination goes on overlay through DC1 out FW to destination

```
BR5-BCN-BR2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
BR5-BCN-BR2(config)#int g1
BR5-BCN-BR2(config-if)#shut
BR5-BCN-BR2(config-if)#

br5-bcn-rtr#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
Known via "bgp 65054", distance 20, metric 0, candidate default path
```

```

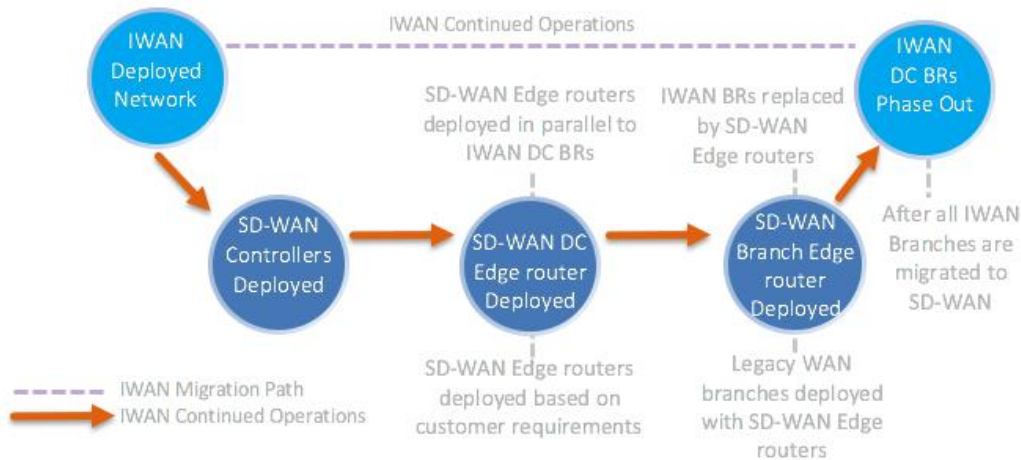
Tag 64510, type external
Last update from 10.5.1.1 00:00:06 ago
Routing Descriptor Blocks:
* 10.5.1.1, from 10.5.1.1, 00:00:06 ago
Route metric is 0, traffic share count is 1
AS Hops 2
Route tag 64510
MPLS label: none

br5-bcn-rtr#traceroute www.cisco.com
Type escape sequence to abort.
Tracing the route to www.cisco.com (184.25.199.192)
VRF info: (vrf in name/id, vrf out name/id)
 1 10.5.1.1 2 msec 1 msec 2 msec
 2 10.6.34.1 [AS 64510] 4 msec 7 msec 5 msec
 3 10.1.101.10 [AS 64510] 5 msec 4 msec 4 msec
 4 10.1.254.1 [AS 64510] 4 msec 7 msec 6 msec
 5 *
www.cisco.com (184.25.199.192) [AS 65010] 9 msec 7 msec
br5-bcn-rtr#

```

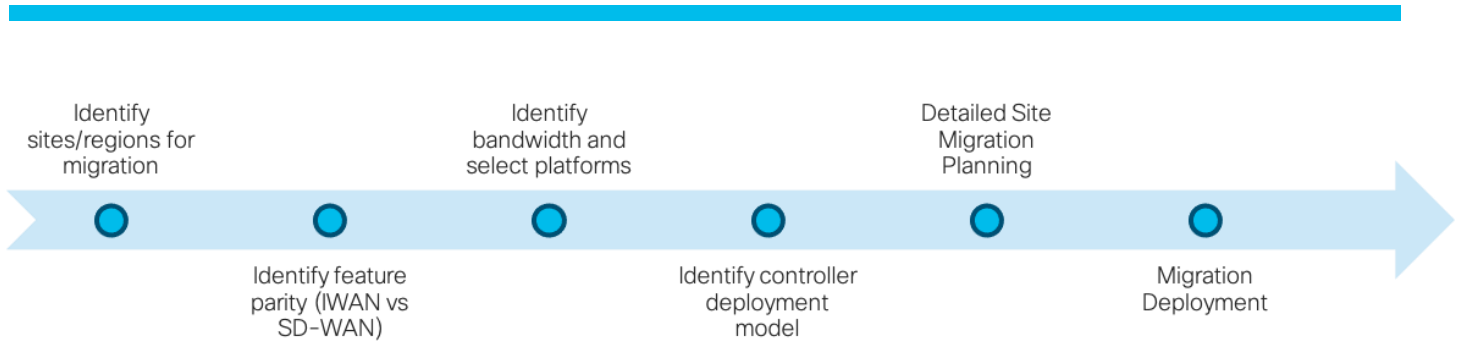
### 5.3 Migration Planning of Case Study

Section 5.1 explains the IWAN Lab setup. In this section let’s go through the guidelines of migration planning. As show in the figure below, when migrating to SD-WAN, Controllers are deployed first. Next SD-WAN Edge devices are installed in the DC(s) parallel to IWAN devices. Then branch SD-WAN Edger routers are deployed. Once all sites are migrated to SD-WAN, we can remove IWAN MC and BR from DC.



**Figure 9: Deployment/Migration Stages**

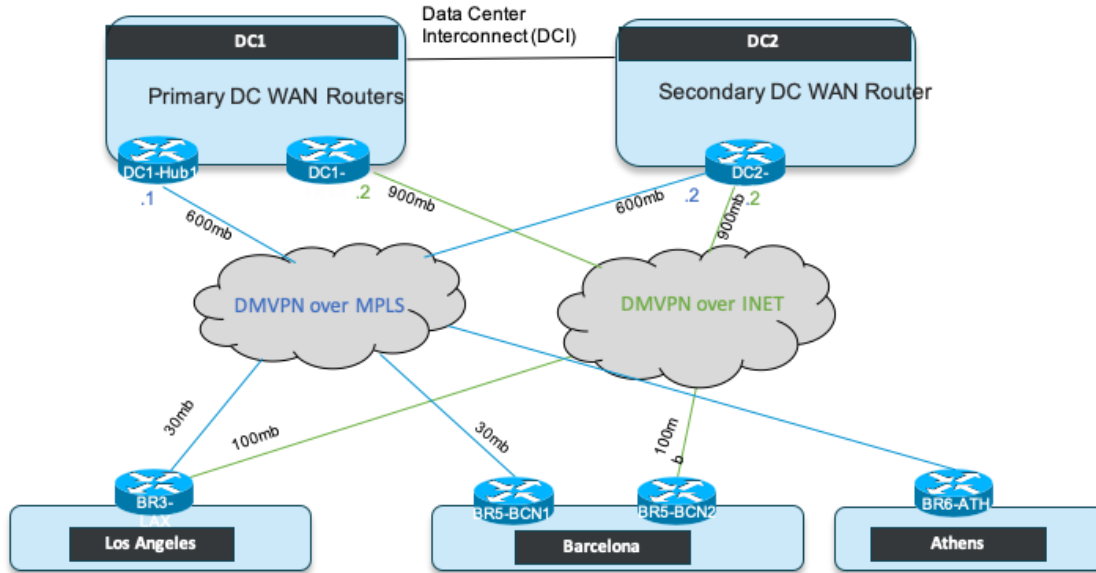
As described in Section 3, the migration planning is critical because moving from IWAN to SD-WAN requires changes to control plane and data plane architecture, design, as well as functional partitioning of the network.



**Figure 10: Migration Planning**

### 5.3.1 Identify Sites/Regions

We plan to migrate DC1, DC2, Los Angeles and Barcelona sites. Athens site is a legacy routing site and does not need migration to Cisco SD-WAN.



**FIGURE #:**

Data center, migration considerations are:

- Migrating existing circuits.
  - DC1: MPLS and Internet terminating on each DC router with TLOC-extension
  - DC2: Single router supporting
- Design is IWAN with DMVPN dynamic tunnels.
- Routing at DCs need to support routing between IWAN, SD-WAN and Legacy sites.
- IWAN use cases are covered in configuration section

Los Angeles and Barcelona site migration considerations:

- Existing circuits.
  - Los Angeles, single router terminates MPLS and Internet
  - Barcelona, dual router site, with TLOC-extension requirement
- Hub and spoke with legacy and IWAN sites.
- Cisco SD-WAN Dynamic tunnels with SD-WAN branches.
- Default route from DC, and specific SD-WAN prefixes
- Policy requirement based on use cases is shown in next section

### 5.3.2 Identify Use Cases and Feature/Configuration Analysis

In this step we are identifying use cases and performing configuration analysis. The configurations of the devices can be created on vManage before the migration of the devices. When migrating to Cisco SD-WAN, the business intent of the IWAN use cases should be deployed with Cisco SD-WAN. Below table shows the use cases that are deployed with IWAN and use cases for Cisco SD-WAN that are planned to be deployed.

Existing IWAN Use Cases	Cisco SD-WAN Use Cases
Secure site-to-site VPN with DMVPN over hybrid transport (MPLS and Internet)	Secure Automated WAN (Dynamic site-to-site IPsec VPN using hybrid MPLS and Internet transport)  Dynamic Tunnels
Application aware routing with PfRv3 with SLA protection and preferred path selection. Details are mentioned in section <b>Performance Routing (PfR) Deployment</b>	Application Performance Optimization (Application aware routing policies that offer SLA protection and preferred path selection)
Secure Direct Internet Access with local Internet exit and Zone Based Firewall at branch	Secure Direct Internet Access at the branch - Local Internet exit with Zone Based Firewall protection and fallback to MPLS overlay path to DC
Per tunnel QoS	Cisco SD-WAN Per Tunnel QoS

#### Feature and Configuration Analysis

In this case study, SD-WAN configurations are only required for migration. Perform analysis of the IWAN features that are deployed in the network and related with the Cisco SD-WAN features. Review [Release Notes](#) on cisco.com to ensure existing features are supported with IOS XE SD-WAN software.

Perform a comparison of IWAN and SD-WAN configurations.

- Perform configuration audit to identify deployed features, for example, routing, QoS, and features outside of IWAN such as Voice.
- Identify the target SD-WAN code version based on required features and platforms to be migrated. Based on the information identified in previous sections, Code image for upgrade will be SD-WAN 17.3/20.3 version. This version supports per-tunnel QoS and dynamic tunnel use case.
- Perform bug scrubbing from release notes and forums  
Lab test the deployment scenario (SVS/partners/customer labs/Cisco dCloud/Cisco Modeling Labs)  
Involve Cisco Teams (Account team, Customer Experience) for additional support if required.

Feature descriptions for each solution	
IWAN	Cisco SD-WAN
Control Plane:	Control Plane:

Master Controller (MC) on every site (software feature on router)	vSmart centralized controller
<b>Data Plane:</b> Border Routers (BR) monitor path quality and enforce path for egress traffic based on policy and SLA	<b>Data Plane:</b> WAN edge monitor path quality and enforce path for egress traffic based on policy and SLA
<b>Path quality monitoring:</b> Proprietary smart probes between PFR Border routers	<b>Path quality monitoring:</b> BFD probes between WAN edge routers
Custom routing with BGP or EIGRP configured on top of DMVPN (ECMP or active/backup with route policy)	Custom routing with overlay management protocol (OMP) policies on vSmart controller. OMP peering is automatically establishes between WAN edge and vSmart controllers and defaults to ECMP load sharing across multiple tunnels to same destination.
PfRv3 for intelligent path selection for applications matched by DSCP or through NBAR deep packet inspection	Application-Aware Routing policies for intelligent path selection matched by DSCP or through NBAR or QOSMOS (vEdge) deep packet inspection
Application optimization through Wide Area Application Services (WAAS). Application visibility with NBAR and NetFlow/IPFIX.	FEC, TCP optimization, packet duplication, QoS, per tunnel QoS, adaptive QoS, intelligent path steering using app-aware, AppNav only (no DRE, caching), Cloud OnRamp for SaaS, SD-AVC
Secure encrypted communications with IPsec over mGRE tunnels. Pre-shared keys or PKI integration for authentication with a variety of strong encryption protocols.	Default SD-WAN IKE-less IPsec and GRE, secure control/data plane using combination of PKI and secure symmetric key exchange, optional pair wise keys, legacy IPsec/GRE, Greatwall UTM (ZBFW, IPS, AMP, URL filtering, Cisco Umbrella), SIG tunnels, segmentation, secure control/data plane, ACL, SSL proxy
cloud-based applications	Cloud OnRamp for IaaS, SaaS and colocation, Umbrella DNS/SIG, Third party SIG
Secure bootstrap, Plug and Play, SUDI	Secure PnP based ZTP, ZTP other methods (one-time password (OTP) with cloud_init, usb boot, on-prem ZTP), ZTP without DHCP (automatic IP), SUDI, TPM support, allow-list for trusted devices serial with Smart account Sync
Cisco IOS (classic) and Cisco IOS XE	Cisco IOS XE and Viptela OS
PKI	Cisco CA, Enterprise CA, VManage as CA
Automation	Full automation support by vManage using templates with REST API support
Manageability APIC-EM with IWAN APP	vManage for complete day0,1,2 configs, monitoring and troubleshooting, vAnalytics for advance analytics
Routing protocols support for (EIGRP, BGP, OSPF)	OMP for Overlay, BGP, OSPF, Static for WAN, BGP, EIGRP, OSPF, Static for LAN.
IPv6	IPv6 (see specific feature details)



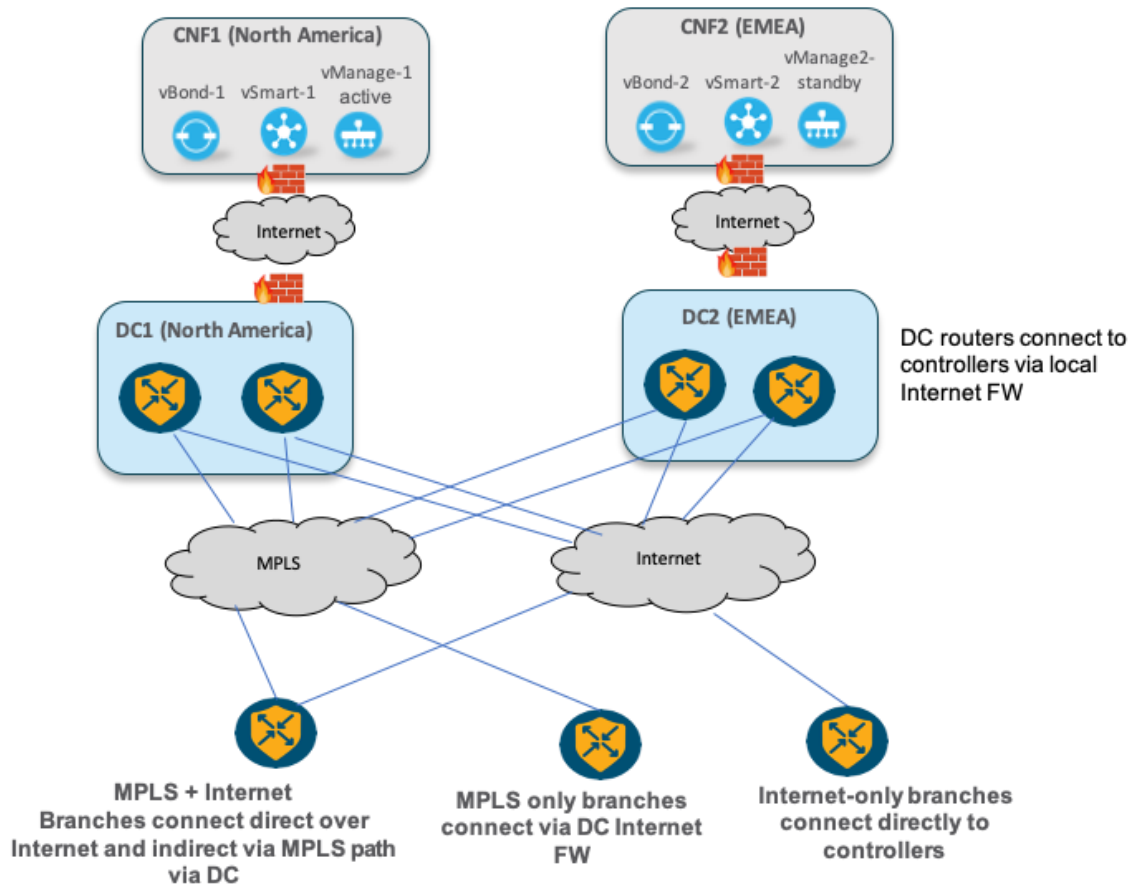
Master Controller	vSmart (centralized controller with redundancy)
Scalability 2000 sites	IKE less IPsec provides scalability of IPsec tunnels, no limit on number of sites with horizontal scaling of controller and routers
Multicast Support	Multicast Support (see Release Notes for details)
NAT	NAT DIA, NAT POOL, Service Side NAT
Segmentation 20 hub, 7 remote VRFs	Multi VPN (VRF) up to 300, inter VRF route leak
Trackers	DIA tracker, SLA monitoring with BFD, Static Route tracker, zScaler L7 health check, VRRP tracker of OMP peering and prefixes
NBAR (N/W Based App Recognition)	NBAR and Customer App Recognition
Direct Cloud Access	Available as Direct Internet Access both from policy and route within Service VPN. Also provide Cloud OnRamp services for optimized SaaS and auto-integration with IaaS
NetFlow v9	Available
VRRP/HSRP	VRRP Supported. vEdge supports up to five VRRP groups per physical/sub-interfaces for Primary and Secondary IP addresses. cEdge support for multiple groups is planned for 17.4.  No HSRP support.
Port Channel interfaces for additional bandwidth capacity and redundancy	SD-WAN routers don't support port link aggregation technology. A different design technique (e.g., ECMP) for high availability should be used.

### 5.3.3 Current Inventory, Bandwidth and Platform Requirement

The current inventory in the case study runs IWAN on IOS XE CSR 1000v at data center. In data center the CSR 1000v will be deployed in parallel with Cisco SD-WAN code. The IOS XE CSR 1000v at Los Angeles and Barcelona sites will be upgraded to Cisco SD-WAN. The bandwidth and scale requirements are according to business requirement. Licensing is checked with the current Cisco SD-WAN bandwidth licensing requirements with the appropriate Level (Essential, Advantage or Premier) and Term of the license.

### 5.3.4 Identify Controller Deployment Model and Requirements

Cisco SD-WAN Controllers can be deployed in two models, On-Prem and Cisco Cloud hosted. This customer chose to self-deploy the SD-WAN controllers in a pair of Carrier Neutral Facility (CNFs) that they were using to host other enterprise shared services including the enterprise CA server used to issue certificates for their mobile access users and IWAN routers. This is considered an on-premise deployment, where the customer is responsible for all of the operational tasks of deployment, monitoring, troubleshooting and backup/restore. The diagram below illustrates the customer's controller deployment and connectivity to the SD-WAN sites.



**Figure#**

The controllers are deployed On-Prem with internet connectivity to all sites. Hardware requirement is identified for the controllers. The customer installed a dedicated UC Server running VMware ESXi in each CNF for controller hosting. The sizing of UCS capacity was based on Cisco guidelines for VM allocations for a medium-sized deployment of 300 SD-WAN sites.

- Hardware: UCS-C240 M5
- Memory: 128GB
- Storage: 1TB SSD
- Hypervisor: VMware ESXi. 6.0

### 5.3.5 Detailed Site Migration Planning

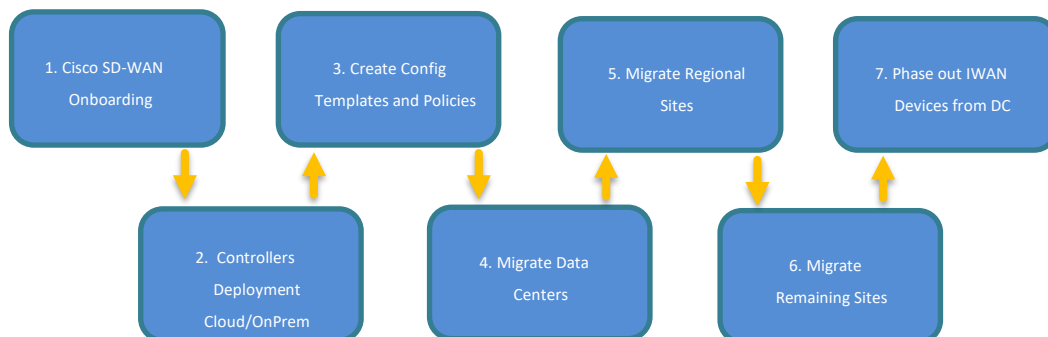
Low-level design details for site migration planning are described in section 3.5. In this case study below information was planned:

- Determined additional port requirement on CEs, firewalls and core devices in data centers to provide connectivity to SD-WAN Edge devices.
- IP addressing is determined for SD-WAN to core connectivity and to CE routers. In addition, VRF segmentation is planned.
- No Controllers redundancy, affinity, clustering or horizontal scaling is required.

- Restore back procedure in case of issues at DCs will be to revert any routing to SD-WAN Edge devices back to IWAN and CEs. At branch location, the restoration will be to downgrade the image and revert back to IWAN.

## 5.4 Migration Deployment Steps for Case Study

The migration follows below steps and sequence



**Figure 11: Migration Sequence**

### 5.4.1 Cisco SD-WAN Onboarding

Using the onboarding process as explained in [section 4.1](#), the controllers and devices are onboarded on SA and VA. Controllers profile is created on VA on PnP and serial file is uploaded on vManage.

### 5.4.2 Deploying Controllers

Controllers are deployed in a colocation facility that was currently offering secure access to Internet and Cloud services. The controllers were deployed on existing UCS servers running VMware, this deployment mode is referred to as “on premise”.

- The process to deploy on premise controllers is described in the Cisco Live presentation [BRKRST-2559](#).
- The process to install enterprise-signed certificates on the controllers is described in the [Cisco SD-WAN Controller Certificates and Authorization File CVD](#).

### Firewall Traffic Requirements

When a WAN Edge router joins the overlay network, it establishes DTLS control plane connections with the controller devices—the Cisco vBond orchestrator, the Cisco vManage, and the Cisco vSmart controller. When initially establishing these DTLS connections, the WAN Edge router uses the base port 12346. If it is unable to establish a connection using this base port, it port-hops through ports 12366, 12386, 12406, and 12426, returning, if necessary, to 12346, until it successfully establishes the DTLS connections with the three controller devices. This same port number is used to establish the IPsec connections and BFD sessions to the other WAN Edge routers in the overlay network.

- Refer to [Firewall Ports for Cisco SD-WAN Deployments](#) for the full details on firewall port requirements.

### High Availability and Scalability of Controllers

A single vManage, vSmart and vBond is deployed in the case study.

---

### 5.4.3 Migrating Data Centers to SD-WAN

NOTE: For the sake of understanding the setup, we will discuss step 4, 5 and 6 of the migration first i.e., the data center and branch inline migration and explain the configurations before step 3 Policies and Advanced configuration use cases. However, you may build the policies and templates first on vManage before migration.

Data centers are migrated first to SD-WAN. This is because the migration of the branch sites is typically gradual and during the migration the data center serves as the transit site for the traffic between IWAN/non-SD-WAN and SD-WAN sites.

Data Center migration prerequisites are verified before actual migration.

1. SD-WAN overlay established with vManage, vBond and vSmart controllers active and onboarded
2. DC WAN edge router details (PID, chassis serial#, device certificate) associated with the customer's SD-WAN virtual account in the Cisco Software Connect - Plug and Play (PnP) Connect Portal
3. WAN edge list including DC WAN edge routers uploaded to customer vManage
4. DC WAN edge Device templates and policies created on vManage
5. vSmart Central policy created on vManage

Below steps are followed to migrate the Data Center

**Step 1.** Baseline current network before any changes

**Step 2.** Pre-stage WAN edge

**Step 3.** Activate central policy

**Step 4.** Attach device templates

**Step 5.** Validate device certificates

**Step 6.** Onboard to vManage

**Step 7.** Validate DC routing

**Step 8.** Verify NetOps

DC SD-WAN cEdge routers are placed, in both San Jose and Frankfurt DCs, in parallel to IWAN routers. For the LAN side connectivity, SD-WAN routers will be connected to the core in the same way IWAN routers are connected i.e., in a full mesh connectivity to core router(s). This is called service side connectivity and unlike in IWAN where the customer connected the LAN side in global VRF routing table, the VRF used in this case is **VRF 1**. This approach provides scalable segmentation and enterprise can add many other isolated segments and use the same SD-WAN overlay for transport instead of creating different **overlays like IWAN per VRF**. The rest of the DC connectivity remains intact including the DCI.

For WAN side, SD-WAN routers will be also be connected to both transport as a best practice for DC, this was not the case in IWAN where auto tunnel was used, and one transport was terminated per IWAN router. Customers have the ability to connect each transport to each SD-WAN edge router, or customers can utilize the TLOC-extension feature (similar to auto tunnel). Using TLOC-extension an SD-WAN router can provide an uplink for the peer router for a transport it is connected to. All SD-WAN routers will form tunnels using both transports provided by northbound CPE (Internet FW and MPLS CE).

Unlike IWAN, Cisco SD-WAN routers keeps both transports in the same VRF, called VPNO, for better routing control and for the ability to use single overlay to carry multiple service VPN segments without creating multiple overlays.

### 5.4.4 Basic Configurations

This section provides an explanation for the basic configuration of one of our data center routers, DC1 San Jose. For more detail and advance configuration explanations, please refer to Cisco SD-WAN documentation and command reference guides.

Configuration Commands	Explanation
<pre>DC1-SJC-cEdge1# show sdwan running-config system system system-ip          1.10.1.1 overlay-id         1 site-id            10 port-offset        0 control-session-pps 300 admin-tech-on-failure sp-organization-name "Viptela-POC-Tool - 19827" organization-name  "Viptela-POC-Tool - 19827" port-hop track-transport track-default-gateway console-baud-rate  115200 no on-demand enable on-demand idle-timeout 10 vbond vbond-test-drive port 12346</pre>	<p>The <b>system</b> configs under SD-WAN are mandatory which helps onboard the devices. These attributes can be filled in a bootstrap file or in a template during onboarding.</p> <p><b>System-ip</b> provides unique identification of the device in overlay and can be analogous to a loopback in legacy routing but it is not routable. The system-ip is unique per device in the SD-WAN overlay</p> <p><b>Site-id</b> provides an identifier to identify a site which can have 1 or more routers. It is used for many purposes including applying policies based on site numbers. Example DC1-San Jose is site id 10.</p> <p><b>Organization-name</b> uniquely identifies overlay name and is key in securely onboarding devices. If the org-name doesn't match to that configured for controllers the device will not be permitted in the overlay.</p> <p><b>vBond</b> FQDN/IP is where the device will reach out first to get permitted to overlay and also get IPs of vManage and vSmart for control communication. vBond also provides NAT detection and determines pre/post NAT IPs of the devices. Port 12346 is the listening port over UDP for vBond by default.</p>
<pre>ip host vbond-test-drive 192.0.2.2</pre>	<p>A host entry exists for vBond FQDN, alternatively DNS can also be used</p>
<pre>DC1-SJC-cEdge1# sh ip vrf Name                Default RD Interfaces 1                   1:1 Gi3 Gi4 65528               &lt;not set&gt; Lo65528 Mgmt-intf           1:512 Gi8  DC1-SJC-cEdge1#  vrf definition 1 description VRF1 for LAN rd             1:1 address-family ipv4</pre>	<p>VRF 1 represent LAN or service side which provides segmentation.</p> <p>MGMT-intf VRF is a default OOB management VRF</p> <p>VPN 0 is the global routing table where transport terminates. This helps create tunnels for the overlay using any transport and provides DIA and other capabilities.</p>

<pre> route-target export 65011:1 route-target import 65011:1 exit-address-family ! address-family ipv6 exit-address-family ! ! vrf definition Mgmt-intf description Mgmt-intf rd 1:512 address-family ipv4 route-target export 1:512 route-target import 1:512 exit-address-family ! address-family ipv6 exit-address-family </pre>	
<pre> interface GigabitEthernet1 description INET WAN interface no shutdown arp timeout 1200 ip address 10.1.254.10 255.255.255.252 no ip redirects ip mtu 1500 ip nat outside load-interval 30 mtu 1500 negotiation auto service-policy output shape_GigabitEthernet1 exit interface GigabitEthernet2 description MPLS WAN interface no shutdown arp timeout 1200 ip address 10.1.1.10 255.255.255.252 no ip redirects ip mtu 1500 load-interval 30 mtu 1500 negotiation auto service-policy output shape_GigabitEthernet2 exit interface GigabitEthernet3 description VPN1 -TO - DC-Core1 no shutdown arp timeout 1200 vrf forwarding 1 ip address 10.1.11.2 255.255.255.252 no ip redirects ip mtu 1500 ip nbar protocol-discovery load-interval 30 mtu 1500 negotiation auto exit interface GigabitEthernet4 description VPN1 -TO - DC-Core2 no shutdown arp timeout 1200 vrf forwarding 1 ip address 10.1.12.2 255.255.255.252 no ip redirects </pre>	<p>Gig1 and Gig2 are part of underlay (global routing/vpn0). Gig1 connects to internet uplink, while Gig2 connects to MPLS CE.</p> <p>Gig3 and Gig4 connect to DC Core and are part of VRF 1 (service VPN or LAN side segment)</p> <p>Gig8 is connected to Out-of-band management network and is not part of overlay.</p>

<pre> ip mtu 1500 ip nbar protocol-discovery load-interval 30 mtu 1500 negotiation auto exit interface GigabitEthernet8 description mgmt interface no shutdown arp timeout 1200 vrf forwarding Mgmt-intf ip address 192.168.150.20 255.255.255.0 no ip redirects ip mtu 1500 load-interval 30 mtu 1500 negotiation auto exit </pre>	
<pre> interface Tunnel1 no shutdown ip unnumbered GigabitEthernet1 no ip redirects ipv6 unnumbered GigabitEthernet1 no ipv6 redirects tunnel source GigabitEthernet1 tunnel mode sdwan exit interface Tunnel2 no shutdown ip unnumbered GigabitEthernet2 no ip redirects ipv6 unnumbered GigabitEthernet2 no ipv6 redirects tunnel source GigabitEthernet2 tunnel mode sdwan </pre>	<p>Tunnel interfaces with <b>tunnel mode sdwan</b> need to be configured and bound to transport physical interfaces.</p> <p>Tunnel 1 is attached to GigabitEthernet1 which is connected to the Internet transport.</p> <p>Tunnel 2 is attached to GigabitEthernet2 which is connected to MPLS transport.</p>
<pre> sdwan service TE vrf global ! interface GigabitEthernet1 tunnel-interface encapsulation ipsec weight 1 no border color biz-internet tunnel-qos hub no last-resort-circuit no low-bandwidth-link no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier default nat-refresh-interval 5 hello-interval 1000 hello-tolerance 12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf </pre>	<p>Transport in Cisco SD-WAN are referred to as TLOCs.</p> <p>Under the <b>sdwan</b> mode, interfaces with TLOCs are configured with colors and other parameters.</p> <p>This configuration of tunnel-interface under sdwan interface triggers the router to reach out to vBond for controller connectivity.</p> <p>The tunnel interface also provided implicit ACL security for the router and doesn't allow any transit traffic. Between tunnel and non-tunnel interfaces.</p> <p>IPsec encapsulation provides security for overlay tunnels. In case of MPLS, GRE can also be used for encapsulation.</p> <p>Weight defines the flow ratio in case ECMP is used between transport and the link has unequal bandwidth.</p> <p>Same tunnels are used to carry all VRF service side traffic by using VPN labels as a differentiator.</p> <p>Cisco SD-WAN routers only communicate and form tunnels to devices which are part of allowed-list and after authentication using PKI infrastructure (see <a href="#">Overlay Bringup documentation</a>)</p> <p><b>Color mpls restrict</b> provides capability for this transport to only form tunnels to same-colored transports. As MPLS does</p>

<pre> no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit rewrite-rule REWRITE-POLICY exit interface GigabitEthernet2 tunnel-interface encapsulation ipsec weight 1 no border color mpls restrict tunnel-qos hub no last-resort-circuit no low-bandwidth-link max-control-connections 0 no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier default nat-refresh-interval 5 hello-interval 1000 hello-tolerance 12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit rewrite-rule REWRITE-POLICY exit appqoe no tcpopt enable !</pre>	<p>not have internet reachability, it will not be able to form full-mesh tunnels. This is the default behavior.</p> <p><b>max control connection 0</b> provides the option to not have any control connection over the MPLS TLOC, as controllers are not reachable using MPLS. By default, all transports attempt to form control connections.</p> <p><b>allow-service</b> provides the capability to open communication for small services which are disabled by default.</p>
<pre> ip route 0.0.0.0 0.0.0.0 10.1.1.9 ip route 0.0.0.0 0.0.0.0 10.1.254.9</pre>	<ul style="list-style-type: none"> <li>• The two default routes provide underlay connectivity using both transports to form tunnels for the overlay using the global routing table.</li> <li>• OSPF/BGP support is also available for transport connectivity.</li> <li>• It is very important to note that p2p interface routers are advertised by uplink CEs to provide tunnel reachability and creation</li> </ul>

The WAN Edge SD-WAN router runs BGP with DC1 core router, which originates the local DC1 prefixes and default route. These routes are installed into VPN (VRF) 1 on the WAN Edge service LAN side, and subsequently redistributed into OMP towards the vSmart controller where they will be advertised to other remote sites. The LAN core is also connected to a data center interconnect (DCI) router that connects to DC2 over a core link used to exchange DC routes. Routing policy is configured on the DCI routers so that DC to DC communications prefer this DCI link over SD-WAN.

When the LA or Barcelona branch advertises their 10.3.x.x or 10.5.xx network respectively, these routes will be propagated to vSmart which will be propagated further to SD-WAN spokes and DCs. Since the route is advertised to both DC1 and DC2 which have a DCI link running BGP, there is a chance that the branch route will be advertised to DC2 from DC1 and vice



versa, thus creating loop. The easiest way to overcome this situation is to use same BGP AS number on both DC cEdge which will block all branch routes learned through DCI and will route traffic through the WAN for branches.

Let us look at the important configurations from both DC1 and DC2 routers:

Configuration output	Description
<pre>DC1-SJC-cEdge1#sh running-config   section bgp router bgp 65011   bgp log-neighbor-changes   distance bgp 20 200 20   !   address-family ipv4 vrf 1     bgp router-id 1.10.1.1     redistribute connected     redistribute omp     propagate-aspath     neighbor 10.1.11.1 remote-as 65010     neighbor 10.1.11.1 description dc core1 router peer   neighbor 10.1.11.1 activate   neighbor 10.1.11.1 send-community both   neighbor 10.1.11.1 maximum-prefix 2147483647 100   neighbor 10.1.12.1 remote-as 65010   neighbor 10.1.12.1 description dc core2 router peer   neighbor 10.1.12.1 activate   neighbor 10.1.12.1 send-community both exit-address-family DC1-SJC-cEdge1#</pre>	<ul style="list-style-type: none"> <li>• EBGp neighborship with DC1 core1 and core 2 routers having remote as 65010 and local-as 65011</li> <li>• Redistribute OMP command redistributes OMP routes learned from other sites into BGP which will be sent to DC core</li> <li>• Propagate-as path can be optionally used to propagate AS information into OMP for loop prevention scenario.</li> <li>• Router-id is kept same as system-ip for serviceability perspective</li> <li>• Distance command provides optional capability to change AD values for internal, local and external routers</li> </ul>
<pre>DC1-SJC-cEdge1#sh sdwan running-config   section omp omp   redistribute omp omp   no shutdown   overlay-as 65555   send-path-limit 16   ecmp-limit 16   graceful-restart   no as-dot-notation   timers     holdtime 60     advertisement-interval 1     graceful-restart-timer 43200     eor-timer 300 exit address-family ipv4 vrf 1   advertise bgp   advertise network 10.0.0.0/8 ! address-family ipv4   advertise bgp   advertise connected   advertise static ! address-family ipv6   advertise connected   advertise static !</pre>	<ul style="list-style-type: none"> <li>• (Optional) <b>overlay-as</b> is configured to identify routes through the overlay for loop prevention</li> <li>• <b>sent-path-limit</b> is the max number of routes sent by edge router to vSmart. Default 4, max 16</li> <li>• <b>ecmp-limit</b> is the max number of paths installed on the edge router for a particular route for ECMP. Default 4, max 16.</li> <li>• <b>advertise network</b> advertises 10.0.0.0/8 as summary route to all sites in order to provide SD-WAN to non-SD-WAN transit routing using hub. The same result could have been achieved from DC core advertising 10.0.0.0/8 network along with default route.</li> <li>• By default, the OMP configuration does not advertise BGP service side routes or OSPF external routes.</li> <li>• <b>advertise bgp</b> advertises BGP routes in OMP so to reach vSmart and other branches.</li> </ul>

<pre>endpoint-tracker static-track1-dci-10.1.100.254 endpoint-ip 10.1.100.254 interval 20 threshold 100 tracker-type static-route track static-track1-dci-10.1.100.254 endpoint-tracker</pre>	<ul style="list-style-type: none"> <li>• A tracker configured tracking DCI route. This is used to ensure cEdge router has connectivity to DCI using DC cores.</li> </ul>
<pre>ip route vrf 1 10.0.0.0 255.0.0.0 10.1.11.1 track name static-track1-dci-10.1.100.254</pre>	<ul style="list-style-type: none"> <li>• A static route is required to generate a summary route through OMP advertisement.</li> <li>• Static route points to next-hop Core routers respectively for different DC routers to track core connectivity to DCI.</li> <li>• The tracker is attached from above config to remove the 10.0.0.0/8 summary route from OMP advertisement if tracker or next-hop becomes unreachable.</li> </ul>
<pre>DC2-FRA-CEDGE1#sh sdwan running-config   section bgp  router bgp 65011 bgp log-neighbor-changes distance bgp 20 200 20 address-family ipv4 unicast vrf 1 bgp router-id 10.2.11.2 neighbor 10.2.11.1 remote-as 65020 neighbor 10.2.11.1 activate neighbor 10.2.11.1 description dc core1 router peer neighbor 10.2.11.1 ebgp-multihop 1 neighbor 10.2.11.1 maximum-prefix 2147483647 100 neighbor 10.2.11.1 route-map set-bgp-as- prepend out neighbor 10.2.11.1 send-community both propagate-aspath redistribute connected redistribute omp exit-address-family !</pre>	<ul style="list-style-type: none"> <li>• DC2 is also using the same AS number to block any branch routes being learned from DC1 through DCI. As OMP AD is 251, DC2 has the potential to install the branch routes from BGP connection to DCI which are learned from DC1 San Jose. Using the same AS will ensure that any routes learned via DCI and originated from DC1 cEdge will be blocked</li> <li>• <b>route-map set-bg-as-prepend-out</b> ensure that all branch routes advertised by DC2 to DCI will be least preferred over DC1 to keep traffic symmetry and active/standby DC scenario.</li> </ul>
<pre>route-map set-bgp-as-prepend permit 1 set as-path prepend 65011 65011 65011 65011 route-map set-bgp-as-prepend permit 65535</pre>	<ul style="list-style-type: none"> <li>• <b>route-map</b> sets policy to prepend AS numbers to routes learned from branches and making DC2 less preferred over DC1 for return traffic to branches.</li> <li>• Note: In order to prefer DC1 over DC2 from branch to DC, we will configure a central policy in vManage (see policy sections).</li> </ul>

After doing basic configuration in DC and Branches, the routes will be propagated end to end. For full configuration of all the routers including DC, DCI and branches refer to appendix section

Below table shows explanation of some important commands to check:

Show outputs	Explanation
<pre>DC1-SJC-cEdge1#sh sdwan control connections  PEER                                PEER CONTROLLER</pre>	<p>This command provides information about control connection to the SD-WAN controllers.</p>

PEER PRIV GROUP TYPE IP PORT	PEER PEER SYSTEM PROT LOCAL	PEER IP COLOR	SITE ID PORT PROXY STATE	DOMAIN ID PUBLIC UPTIME	PEER PUB PRIVATE IP ID
-----					
-----					
-----					
vsmart	dtls	1.1.1.3	1	1	
192.0.2.3				12346	192.0.2.3
12346	biz-internet	No	up	3:01:47:21	0
vbond	dtls	0.0.0.0	0	0	
192.0.2.2				12346	192.0.2.2
12346	biz-internet	-	up	3:01:47:33	0
vmanage	dtls	1.1.1.1	1	0	
192.0.2.4				12446	192.0.2.4
12446	biz-internet	No	up	10:23:37:04	0

Biz-internet is used to reach to the controllers and form DTLS connections to vBond, vSmart and vManage.

The command also provides information about critical network attributes like system-ip, pre/post NAT IPs and ports.

```

DC1-SJC-cEdge1#sh sdwan control local-properties
personality                vedge
sp-organization-name       Viptela-POC-Tool -
19827
organization-name         Viptela-POC-Tool -
19827
root-ca-chain-status       Installed

certificate-status         Installed
certificate-validity        Valid
certificate-not-valid-before Jul 31 02:09:43 2020
GMT
certificate-not-valid-after Jul 29 02:09:43 2030
GMT

enterprise-cert-status     Not-Applicable
enterprise-cert-validity   Not Applicable
enterprise-cert-not-valid-before Not Applicable
enterprise-cert-not-valid-after Not Applicable
dns-name                   vbond-test-drive
site-id                    10
domain-id                  1
protocol                   dtls
tls-port                   0
system-ip                  1.10.1.1
chassis-num/unique-id     CSR-C2F64427-7D0B-
48CD-8D49-A27924ED7D1D
serial-num                 C05C42A0
subject-serial-num         N/A
token                      Invalid
keygen-interval            1:00:00:00
retry-interval             0:00:00:18
no-activity-exp-interval  0:00:00:20
dns-cache-ttl              0:00:02:00
port-hopped                FALSE
time-since-last-port-hop  0:00:00:00
embargo-check              success
number-vbond-peers         0
number-active-wan-interfaces 2
NAT TYPE: E -- indicates End-point independent mapping
A -- indicates Address-port dependent
mapping                    N -- indicates Not learned
Note: Requires minimum two vbonds to learn
the NAT type

PRIVATE                PUBLIC                PUBLIC PRIVATE
MAX RESTRICT/         LAST                SPI TIME        NAT
VM
INTERFACE              IPv4                PORT    IPv4
IPv6                   PORT    VS/VM

```

This command provides information on how the control connection is formed.

It provides details on certificates, serials and local TLOC colors.

<pre> COLOR          STATE CNTRL CONTROL/   LR/LB CONNECTION     REMAINING  TYPE CON  STUN                                     PRF ----- ----- ----- GigabitEthernet1          192.0.2.100    46688 10.1.254.10      :: 12346      1/1 biz-internet    up    2    no/yes/no No/No 0:00:00:14 0:06:26:55 N    5 GigabitEthernet2          10.1.1.10      12346 10.1.1.10      :: 12346      0/0 mpls      up    0    yes/ no/no No/No 12:05:36:35 0:06:26:55 N    5 DC1-SJC-cEdge1# </pre>	
<pre> DC1-SJC-cEdge1#sh sdwan bfd sessions SOURCE TLOC REMOTE TLOC          DST PUBLIC              DST PUBLIC    DETECT TX SYSTEM IP          SITE ID STATE      COLOR COLOR              SOURCE IP          IP PORT               ENCAP  MULTIPLIER INTERVAL(msec) UPTIME TRANSITIONS ----- ----- ----- 172.16.3.1      30      up      mpls mpls            10.1.1.10 10.30.0.2      12346    ipsec 7 1000           10:23:42:00 0 172.16.3.1      30      up      biz-internet biz-internet   10.1.254.10 192.0.2.40     12346    ipsec 7 1000           1:10:24:12 1 172.16.5.1      50      up      mpls mpls            10.1.1.10 10.50.0.2      12346    ipsec 7 1000           10:23:41:54 0 172.16.5.1      50      up      biz-internet biz-internet   10.1.254.10 192.0.2.50     5062    ipsec 7 1000           3:01:55:14 1 172.16.5.2      50      up      mpls mpls            10.1.1.10 10.5.102.2     12346    ipsec 7 1000           10:23:42:00 0 172.16.5.2      50      up      biz-internet biz-internet   10.1.254.10 192.0.2.50     12346    ipsec 7 1000           3:01:55:14 2 DC1-SJC-cEdge1# </pre>	<p>The output shows IPsec tunnels created from DC1-Edge1 perspective.</p> <p>By default, tunnels created in full mesh, but in the output below its clear that MPLS is only forming tunnel to MPLS TLOCs and same for biz-internet.</p> <p>Tunnel endpoints IPs and port are also shown as source destination</p>

### Checking routing on DC routers

Note: Although only DCs has been shown to be migrated as of now, but below output will also show post branch migration output from LA perspective for better understanding the DC routing.

```

DC1-SJC-cEdge1#sh ip route vrf 1
B*    0.0.0.0/0 [20/0] via 10.1.11.1, 01:58:08
      10.0.0.0/8 is variably subnetted, 45 subnets, 5
masks
S     10.0.0.0/8 [1/0] via 10.1.11.1
B     10.1.1.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.1.4/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.1.8/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.1.12/30 [20/0] via 10.1.11.1, 01:58:08
C     10.1.11.0/30 is directly connected,
GigabitEthernet3
L     10.1.11.2/32 is directly connected,
GigabitEthernet3
B     10.1.11.4/30 [20/0] via 10.1.11.1, 01:58:08
C     10.1.12.0/30 is directly connected,
GigabitEthernet4
L     10.1.12.2/32 is directly connected,
GigabitEthernet4
B     10.1.12.4/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.100.0/24 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.4/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.8/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.12/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.20/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.28/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.101.32/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.102.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.102.4/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.102.12/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.254.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.1.254.4/30 [20/0] via 10.1.11.1, 01:58:08
B     10.2.1.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.2.1.8/30 [20/0] via 10.1.11.1, 01:58:08
B     10.2.11.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.2.100.0/24 [20/0] via 10.1.11.1, 01:58:08
m     10.3.100.0/24 [251/0] via 172.16.3.1, 1d10h,
Sdwan-system-intf
m     10.5.1.0/30 [251/0] via 172.16.5.1, 3d02h,
Sdwan-system-intf
m     10.5.2.0/30 [251/0] via 172.16.5.2, 3d02h,
Sdwan-system-intf
m     10.5.100.0/24 [251/0] via 172.16.5.2, 3d02h,
Sdwan-system-intf
      [251/0] via 172.16.5.1, 3d02h,
Sdwan-system-intf
B     10.5.102.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.5.102.4/30 [20/0] via 10.1.11.1, 01:58:08
B     10.6.34.0/23 [20/0] via 10.1.11.1, 01:58:08
B     10.6.36.0/23 [20/0] via 10.1.11.1, 01:58:08
B     10.6.100.0/24 [20/0] via 10.1.11.1, 01:58:08
B     10.10.0.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.20.0.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.30.0.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.40.0.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.40.1.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.50.0.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.60.0.0/30 [20/0] via 10.1.11.1, 01:58:08
B     10.102.1.0/30 [20/0] via 10.1.11.1, 01:58:08
      100.0.0.0/32 is subnetted, 1 subnets
B     100.100.100.100 [20/0] via 10.1.11.1, 01:58:08
      172.16.0.0/32 is subnetted, 5 subnets
B     172.16.1.1 [20/0] via 10.1.11.1, 01:58:08
B     172.16.1.2 [20/0] via 10.1.11.1, 01:58:08
B     172.16.1.3 [20/0] via 10.1.11.1, 01:58:08
B     172.16.2.1 [20/0] via 10.1.11.1, 01:58:08
B     172.16.2.2 [20/0] via 10.1.11.1, 01:58:08
B     192.168.150.0/24 [20/0] via 10.1.11.1, 01:58:08
      200.200.200.0/32 is subnetted, 1 subnets
B     200.200.200.200 [20/0] via 10.1.11.1, 01:58:08
DC1-SJC-cEdge1#

```

Commands shows routes learned in VRF 1 routing table of dc1-ledge1.

Default route is learned via DC-core

Static route is used for advertising summary route to SD-WAN branch routers for SD-WAN to non-SD-WAN communications.

“m” routes represent routes learned from OMP which are the branch routes.

Since we have redistributed OMP routes into BGP and vice versa, the routes will be learned in OMP and BGP process as well

The next-hop in ‘m’ routes points to system-ip of the SD-WAN branch routers

<pre> DC1-SJC-cEdge1#sh sdwan omp routes vpn 1 10.3.100.0 Code: C -&gt; chosen I -&gt; installed Red -&gt; redistributed Rej -&gt; rejected L -&gt; looped R -&gt; resolved S -&gt; stale Ext -&gt; extranet Inv -&gt; invalid Stg -&gt; staged IA -&gt; On-demand inactive U -&gt; TLOC unresolved            PATH          PSEUDO FROM PEER  ID    LABEL  STATUS  KEY    TLOC IP         COLOR  ENCAP  PREFERENCE ----- 1.1.1.3    213   1002   C,I,R   1 172.16.3.1  mpls  ipsec  - 1.1.1.3    214   1002   C,I,R   1 172.16.3.1  biz-internet  ipsec  - DC1-SJC-cEdge1# </pre>	<p>The output explains the routes learned from OMP perspective.</p> <p>The route is learned from LA branch LAN which is shown to be migrated to SD-WAN</p> <p>C, I, R depicts it's a chosen, installed and resolved route.</p> <p>The route is learned using both TLOCs mpls and biz-internet using IPsec encapsulation</p>
<pre> DC1-SJC-cEdge1#sh ip route vrf 1 10.3.100.0 Routing Table: 1 Routing entry for 10.3.100.0/24   Known via "omp", distance 251, metric 0, type omp   Redistributing via bgp 65011   Advertised by bgp 65011   Last update from 172.16.3.1 on Sdwan-system-intf,   1d10h ago   Routing Descriptor Blocks:   * 172.16.3.1 (default), from 172.16.3.1, 1d10h ago,   via Sdwan-system-intf      opaque_ptr 0x7F202071BAC0   Route metric is 0, traffic share count is 1 DC1-SJC-cEdge1# </pre>	<p>DC1 SD-WAN router learns the LA branch router through OMP having AD of 251.</p>
<pre> dc1-sjc-core1#sh ip bgp 10.3.100.0 BGP routing table entry for 10.3.100.0/24, version 65 Paths: (3 available, best #3, table default)   Advertised to update-groups:     1          2          3          5   Refresh Epoch 1   65011 65555     10.1.11.2 from 10.1.11.2 (1.10.1.1)       Origin incomplete, metric 1000, localpref 100,   valid, external       Extended Community: SoO:0:10 RT:65011:1         rx pathid: 0, tx pathid: 0   Refresh Epoch 1   65011 65555, (Received from a RR-client)     10.1.12.6 from 10.1.100.2 (192.168.150.8)       Origin incomplete, metric 1000, localpref 100,   valid, internal         rx pathid: 0, tx pathid: 0   Refresh Epoch 1   65011 65555     10.1.11.6 from 10.1.11.6 (1.10.1.2)       Origin incomplete, metric 1000, localpref 100,   valid, external, best       Extended Community: SoO:0:10 RT:65011:1         rx pathid: 0, tx pathid: 0x0 dc1-sjc-core1# </pre>	<p>The LA route is advertised to core by DC1-edge1 and edge2</p> <p>The route contains AS-PATH information including overlay-as 65555 and local DC1 edge AS of 65011</p> <p>In order to prevent the core from readvertising this route back to DC-1-cedge2, a SoO community is automatically inserted in the routes.</p> <p>A Metric of 1000 is also received for OMP redistributed routes by default</p>

```

vyos@dci-sjc-fra-rtr:~$ sh ip bgp 10.3.100.0
BGP routing table entry for 10.3.100.0/24
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non-peer-group peers:
    10.1.100.2 10.102.1.2
    65010 65011 65555
    10.1.100.2 from 10.1.100.2 (192.168.150.8)
      Origin incomplete, localpref 100, valid, external
      Last update: Thu Oct 29 16:56:56 2020
    65010 65011 65555
    10.1.100.1 from 10.1.100.1 (100.100.100.100)
      Origin incomplete, localpref 100, valid,
external, best
      Last update: Thu Oct 29 16:56:48 2020
vyos@dci-sjc-fra-rtr:~$

```

DCI router receives the LA router preferred from DC1 core routers.

Now that we have seen a sample DC migration from the perspective of DC routers, we take a look at the branch migration steps.

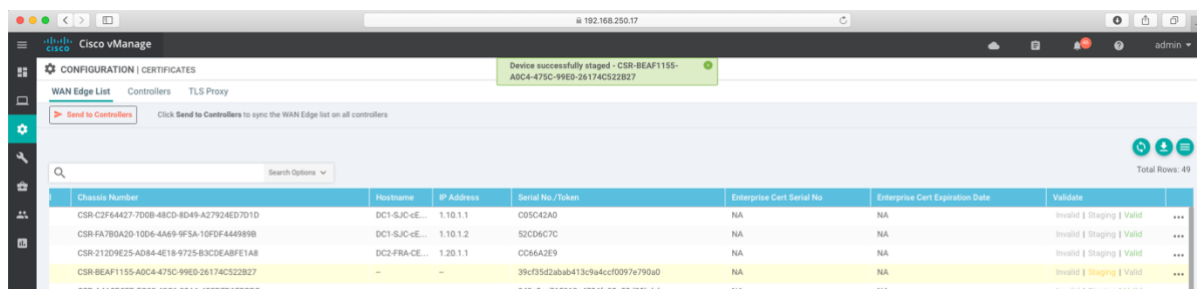
Note: Branch router basic configuration will be similar as explained above. The next section will be showing branch migration using templates, bootstrapping and ZTP.

### 5.4.5 Branch Routers Migration basic config of LA branch missing

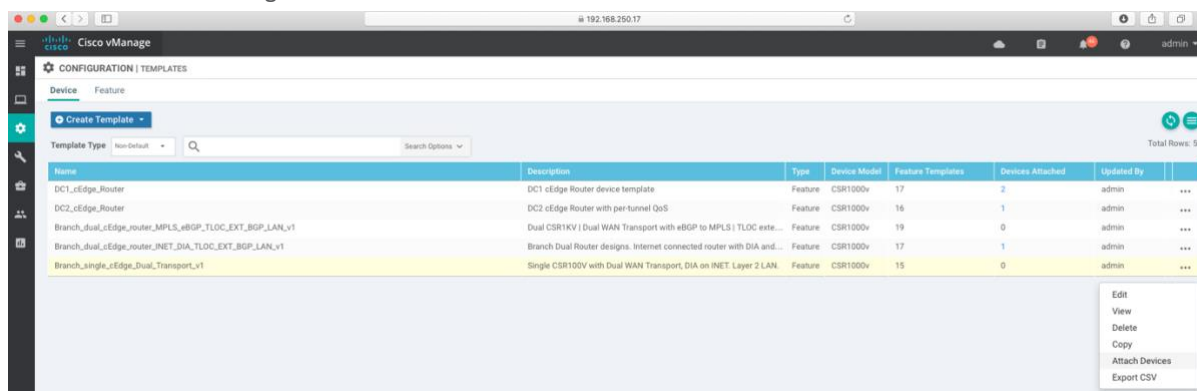
Below are the vManage steps followed for migrating the existing branch router.

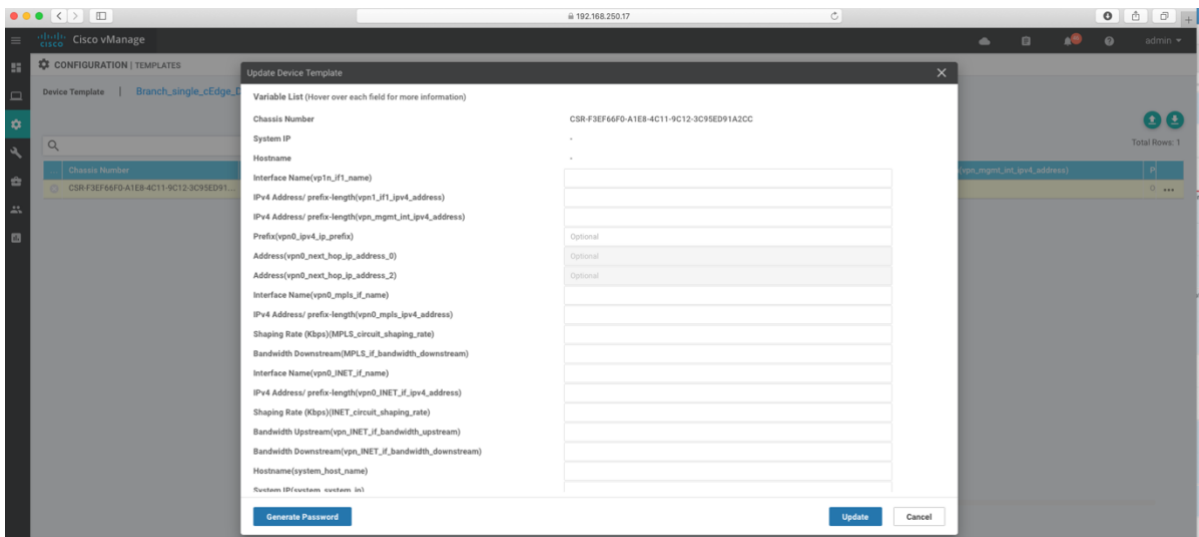
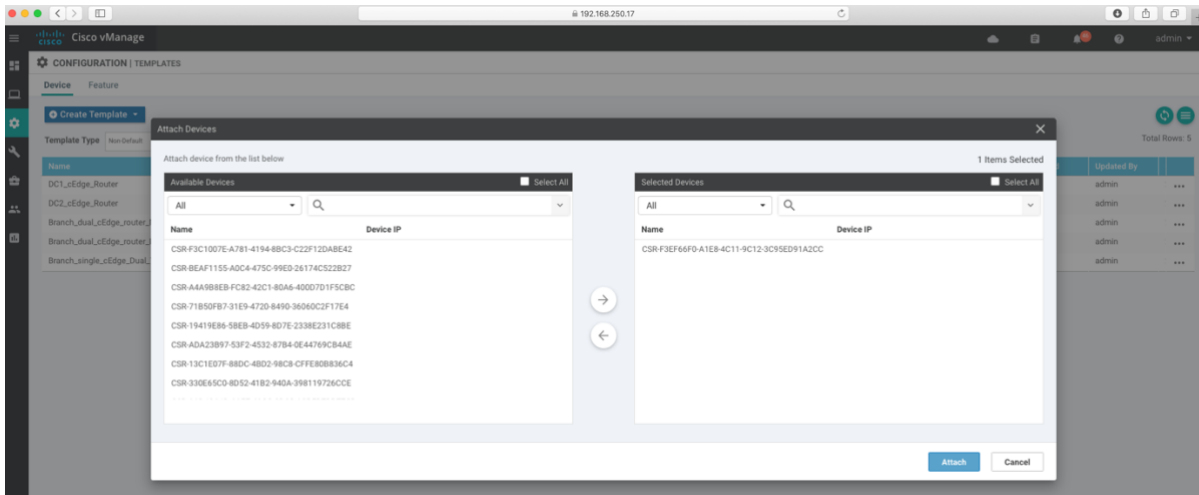
#### vManage NMS steps (Prior to Migration)

1. Place WAN Edge device in staging mode in the devices pane.

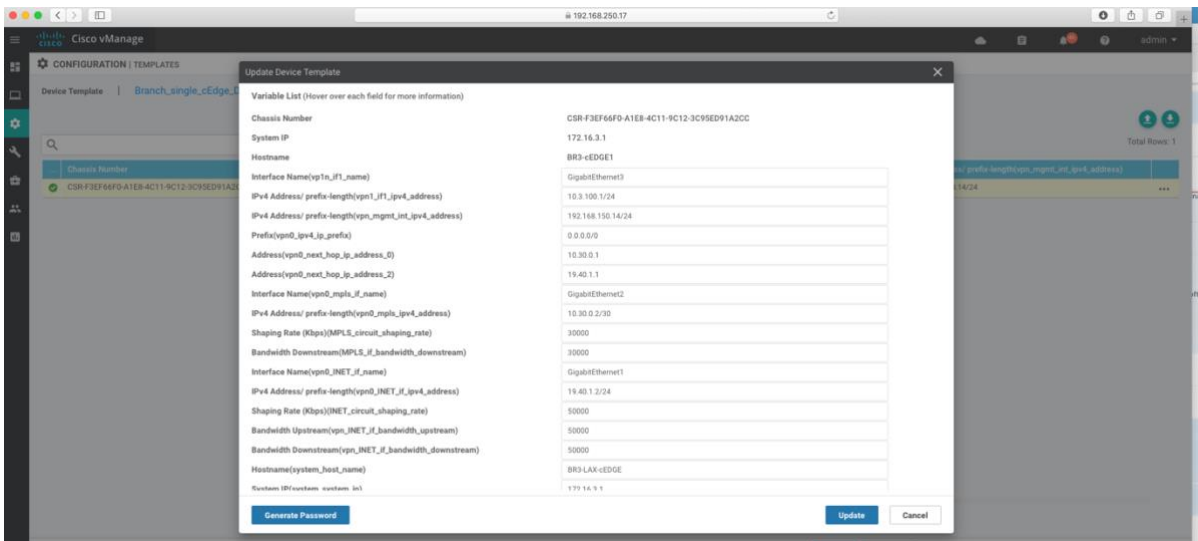
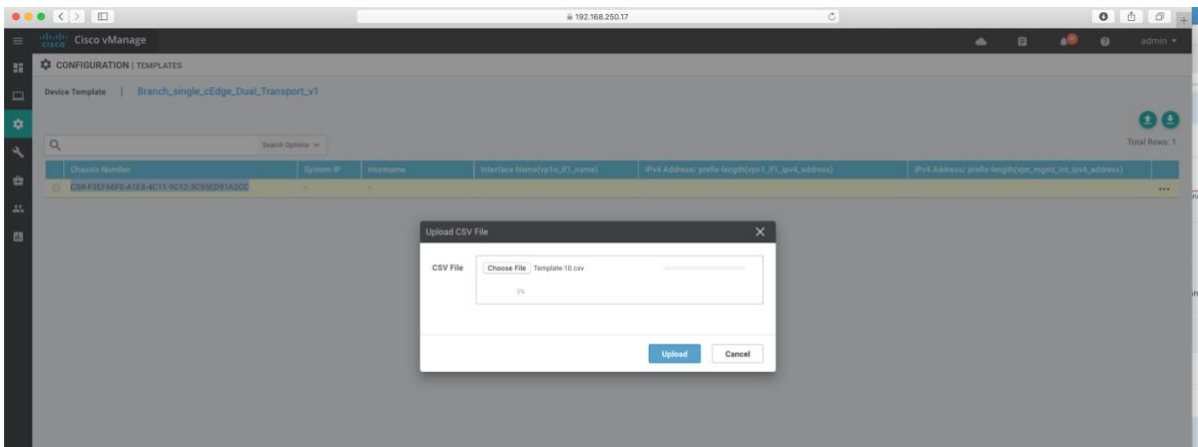
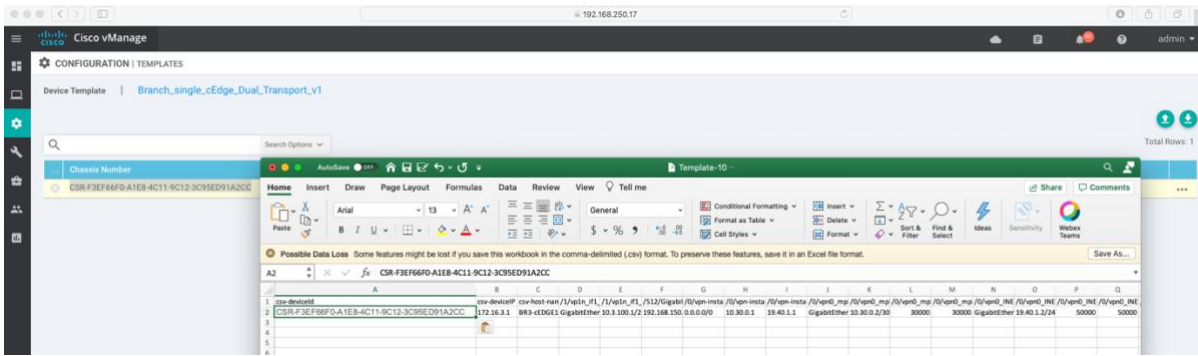


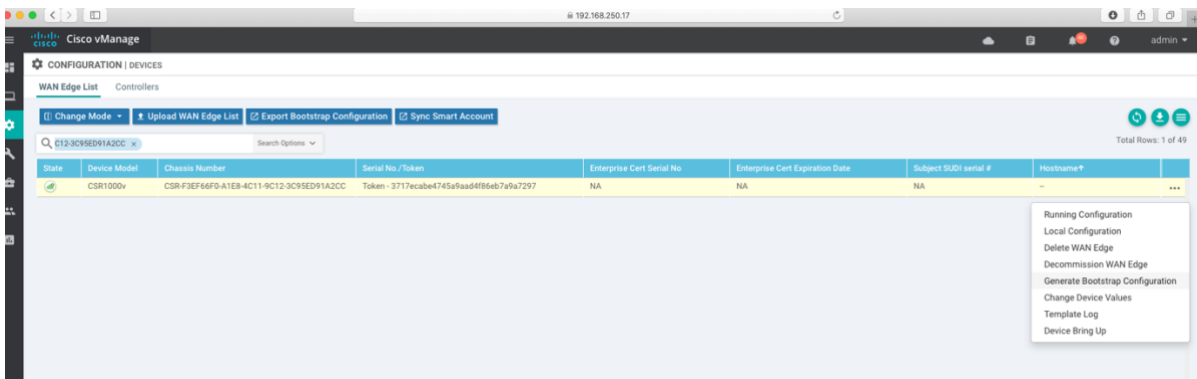
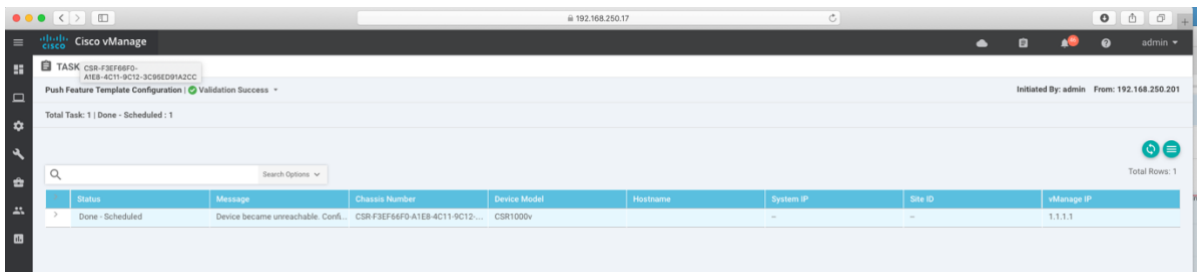
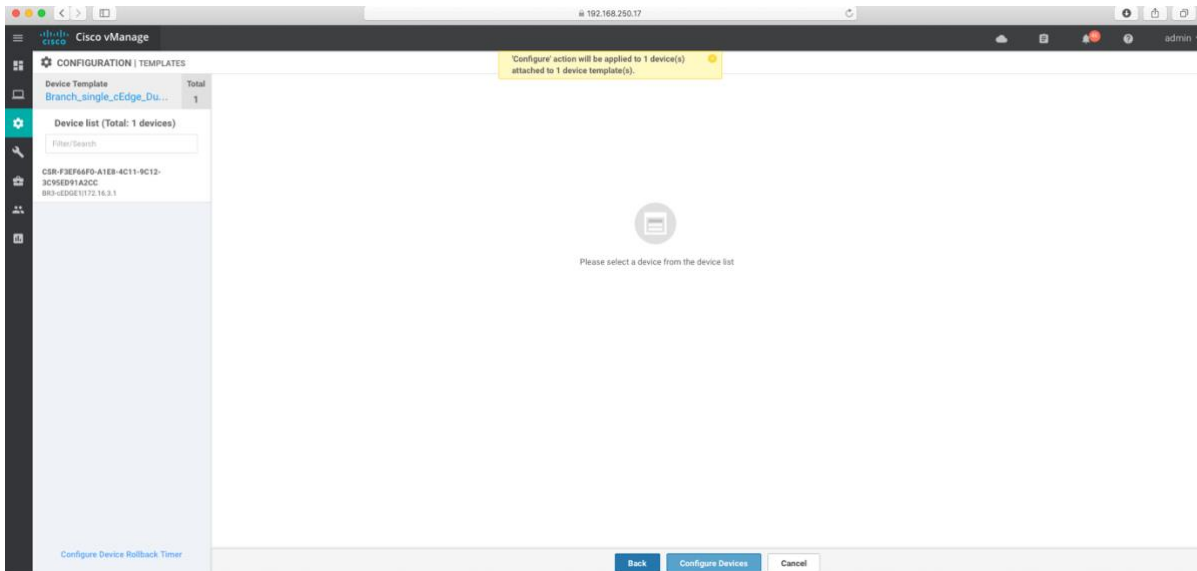
2. Attach the device template to WAN edge device and supply site-specific values for template variables to generate a full device configuration.











### Branch router output (Prior to Migration)

1. Ensure rommon have been upgraded to the minimum level
2. Upload the new IOS XE SD-WAN image onto the router bootflash
3. Branch Router: Copy the bootstrap configuration generated by vManage to the router bootflash
4. Remove existing boot statements and add boot variable that points to the new IOS XE SD-WAN image

5. Save the existing configuration as a named file in the router bootflash

### vManage NMS steps (During Migration)

1. Branch Router: Reboot the router
2. Branch Router: Verify the router comes up on the target image in controller mode
3. Branch Router: If router comes up in autonomous mode, the previous IWAN config will be present - Change to controller mode which will trigger a reboot and reformat
4. Branch Router: Verify control connections formed and WAN edge router in vManage mode
5. vManage NMS: Place WAN Edge device in active mode in the devices pane
6. vManage NMS: Verify WAN Edge control plane and forwarding plane (BFD sessions) formed.

### Backout Procedure

1. Change to autonomous mode (controller-mode disable) which will trigger a reboot and reformat

```
BR3-LAX-cEDGE#controller-mode disable
Disabling controller mode will erase the nvram filesystem, remove all configuration files,
and reload the box!
Ensure the BOOT variable points to a valid image
Continue? [confirm]y
% Warning: Bootstrap config file needed for Day-0 boot is missing
Do you want to abort? (yes/[no]): no
Mode change success
*Sep 2 20:42:54.445: %SYS-7-NV_BLOCK_INIT: Initialized the geometry of nvram
*Sep 2 20:43:00.858: %SYS-5-RELOAD: Reload requested by Exec. Reload Reason: Disabling
controller-mode.
<system reboot messages>
--- System Configuration Dialog ---
Would you like to enter the initial configuration dialog? [yes/no]: no
Would you like to terminate autoinstall? [yes]: yes
No startup-config, starting autoinstall/pnp/ztp...
Autoinstall will terminate if any input is detected on console
Press RETURN to get started!
```

2. Escape from configuration dialog, change the boot statement to original IOS XE image running IWAN and reboot

```
Router#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)#boot system bootflash:csr1000v-universalk9.16.09.05.SPA.bin
Router#wr
Building configuration...
[OK]
Router#reload
Proceed with reload? [confirm]y
*Sep 2 20:48:56.051: %SYS-5-RELOAD: Reload requested by console. Reload Reason: Reload
Command.
```

3. Copy saved IWAN configuration to running configuration

```
Router#copy IWAN-Config-0902 running-config
Destination filename [running-config]?
BR3-LAX-MCBR#conf t
Enter configuration commands, one per line. End with CNTL/Z.
BR3-LAX-MCBR(config)#int g1
BR3-LAX-MCBR(config-if)#no sh
BR3-LAX-MCBR(config-if)#int g2
BR3-LAX-MCBR(config-if)#no sh
```

```
BR3-LAX-MCBBR(config-if)#int g3
BR3-LAX-MCBBR(config-if)#no sh
```

#### 4. Verify that DMVPN has established

```
BR3-LAX-MCBBR#show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
N - NATed, L - Local, X - No Socket
T1 - Route Installed, T2 - Nexthop-override
C - CTS Capable, I2 - Temporary
# Ent --> Number of NHRP entries with same NBMA peer
NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
UpDn Time --> Up or Down Time for a Tunnel
=====
Interface: Tunnel100, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
1 10.1.101.21 10.6.34.1 UP 00:01:12 S
1 10.1.102.5 10.6.34.2 UP 00:01:11 S
Interface: Tunnel200, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
-----
1 192.0.2.10 10.6.36.1 UP 00:01:27 S
1 192.0.2.20 10.6.36.2 UP 00:01:27 S
```

#### 5.4.6 Verification

Below section shows basic configs from Barcelona branch which is a dual router branch. Each router is connected to single transport MPLS or internet and using **TLOC extension feature** to form tunnels using each other's transport. The branch is also running BGP to branch core router.

Configuration output	Description
<pre>BR5-BCN-cEDGE1#sh running-config   section interface interface GigabitEthernet1   description WAN Transport2 Parent Interface   mtu 1504   no ip address   no ip redirects   load-interval 30   negotiation auto   arp timeout 1200   no mop enabled   no mop sysid interface GigabitEthernet1.101   description TLOC Extension Transport1   encapsulation dot1Q 101   ip address 10.5.102.1 255.255.255.252   no ip redirects   ip mtu 1496   arp timeout 1200 interface GigabitEthernet1.102   description INET WAN interface   encapsulation dot1Q 102   ip address 10.5.102.5 255.255.255.252   no ip redirects</pre>	<p>Gig1 provides interconnection between BCN-cedge1 and BCN-cedge2. This interface should be part of VPN 0 when configuring from vManage template</p> <p>Gig1.101 provides tagged sub-interface p2p connectivity to BCN-cedge2 for biz-internet TLOC. The subinterface MTU is adjusted from default 1500 to 1496 to adjust 4-byte VLAN tag.</p> <p>Gig1.102 provides tagged sub-interface p2p connectivity to BCN-cedge1 to form tunnels using MPLS TLOC of BCN-cedge2 MPLS transport. The subinterface MTU is adjusted from default 1500 to 1496 to adjust 4-byte VLAN tag.</p> <p>Gig 2 connects to MPLS WAN on BCN edge-1</p>

<pre> ip mtu 1496 ip nat outside service-policy output shape_GigabitEthernet1.102 interface GigabitEthernet2 description MPLS WAN interface ip address 10.50.0.2 255.255.255.252 no ip redirects load-interval 30 negotiation auto arp timeout 1200 no mop enabled no mop sysid service-policy output shape_GigabitEthernet2 interface GigabitEthernet3 description VPN1 -TO - DC-Core1 vrf forwarding 1 ip address 10.5.1.1 255.255.255.252 no ip redirects ip nbar protocol-discovery load-interval 30 negotiation auto arp timeout 1200 no mop enabled no mop sysid interface GigabitEthernet8 description mgmt interface vrf forwarding Mgmt-intf ip address 192.168.150.18 255.255.255.0 no ip redirects load-interval 30 negotiation auto arp timeout 1200 no mop enabled no mop sysid </pre>	<p>Gig 3 connects to Barcelona LAN core router in VRF 1</p> <p>Gig 8 provides OOB mgmt</p>
<pre> interface Tunnel2 ip unnumbered GigabitEthernet2 no ip redirects ipv6 unnumbered GigabitEthernet2 no ipv6 redirects tunnel source GigabitEthernet2 tunnel mode sdwan interface Tunnel102001 ip unnumbered GigabitEthernet1.102 no ip redirects ipv6 unnumbered GigabitEthernet1.102 no ipv6 redirects tunnel source GigabitEthernet1.102 </pre>	<p>Tunnel2 terminates biz-internet</p> <p>Tunnel102001 terminates mpls TLOC</p>
<pre> BR5-BCN-cEDGE1#sh sdwan running-config sdwan <b>interface GigabitEthernet1.101</b> <b>tloc-extension GigabitEthernet2</b> exit <b>interface GigabitEthernet1.102</b> <b>tunnel-interface</b> <b>encapsulation ipsec weight 1</b> <b>no border</b> <b>color biz-internet</b> tunnel-qos spoke no last-resort-circuit no low-bandwidth-link no vbond-as-stun-server </pre>	<p>Gig1.101 provides TLOC-extension to BCN-cedge2 for terminating biz-internet tunnels</p> <p>Gig1.102 forms tunnel using mpls TLOC using TLOC-extension provided by BCN-cedge2</p> <p>Note: similar configs will be done on BCN-CEDGE2 but in reverse</p>

<pre> vmanage-connection-preference 5 port-hop carrier                                default nat-refresh-interval                    5 hello-interval                          1000 hello-tolerance                         12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit bandwidth-downstream 50000 rewrite-rule REWRITE-POLICY exit interface GigabitEthernet2 tunnel-interface encapsulation ipsec weight 1 no border color mpls restrict tunnel-qos spoke no last-resort-circuit no low-bandwidth-link max-control-connections                 0 no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier                                default nat-refresh-interval                    5 hello-interval                          1000 hello-tolerance                         12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit bandwidth-downstream 30000 rewrite-rule REWRITE-POLICY exit appqoe no tcpopt enable !</pre>	
<pre> omp no shutdown overlay-as                             100</pre>	<p>OMP configs remain similar to DC.</p>

<pre> send-path-limit 16 ecmp-limit 16 graceful-restart no as-dot-notation timers   holdtime 60   advertisement-interval 1   graceful-restart-timer 43200   eor-timer 300 exit address-family ipv4 vrf 1   advertise bgp ! address-family ipv4   advertise bgp   advertise connected   advertise static ! address-family ipv6   advertise connected   advertise static ! ! ! </pre>	
---	--

### 5.4.7 Mapping IWAN to SD-WAN Configuration Policies and Advanced Use Cases

At this stage of the migration, the IWAN deployed uses cases are captured and configuration analysis has been performed. Note that the routing information at each site is critical for configuration and policies creation and must be captured while planning during the low-level design phase. See the Appendix for full configurations and details.

#### DC Preference

##### Intent:

Customer have 2 DCs and wants to make DC1- San Jose active vs DC2- Frankfurt as Standby Data center for any branch to DC communication.

##### Solution:

As shown earlier DC1 and DC2 prefixes are advertised from DC1 and DC2 respectively including a common default route and a 10.0/8 summary route to the overlay branches. In order for branches to prefer default and summary route from DC1 instead of DC2, a centralized control policy is configured and pushed to branch routers to prefer OMP routes from DC1 over DC2. This will take care from branch to DC preference. For the return traffic, BGP as-path prepend is used from DC2 on LAN side to advertise sub-optimal branch prefixes so that DC can prefer DC1 as exit point for reaching branch LAN.

Configuration Output	Explanation
<pre> control-policy Central-topology-policy-v2 sequence 1   match route   site-list DC1   prefix-list _AnyIpv4PrefixList ! action accept   set   preference 400 </pre>	<p>Centralized control policy created</p> <p>Sequence 1 matches on any route originating with site-id list assigned to DC1.</p> <p>Setting OMP preference of routes originating from DC1.</p> <p>Sequence 11 matches on any route originating with site-id list assigned to DC2.</p>





R	installed	1.1.1.3	162	1002
ipsec	200	1.20.1.1	biz-internet	
C,I,R	installed	1.1.1.3	263	1002
ipsec	400	1.10.1.1	mpls	
C,I,R	installed	1.1.1.3	264	1002
ipsec	400	1.10.1.1	biz-internet	

## Blocking DC to DC tunnels

### Intent:

Customer always wants data center sites to use DCI for any Inter-DC communication

### Solution:

There are multiple ways to enforce inter-DC communication through DCI. For simplicity perspective, blocking any tunnels between DCs will ensure all DC to DC traffic will have DCI as the only available path.

Configuration Output	Explanation
<pre>control-policy DC-block-tunnels sequence 1   match tloc     site-list DC_Group   !   action reject   !   ! default-action accept !  site-list DC_Group site-id 10 site-id 20 ! apply-policy  site-list DC_Group control-policy DC-block-tunnels out !  ! !</pre>	<p><b>Control-policy</b> created with name DC-block-tunnels</p> <p><b>Sequence 1</b> matches TLOCs (tunnels) originating from data centers.</p> <p><b>Action reject</b> anything matching within the current sequence. In this case, this will block TLOCs tunnel formation with sites in DC_Group.</p> <p><b>Default-action accept</b> allows TLOCs (tunnels) from anywhere else other than DCs</p> <p>Site-list matching DC sites</p> <p>Apply policy to both DC1 and DC2 sites WAN edges. Both DC1 and DC2 will block tunnels from each other and only allow tunnels from branches</p>
<pre>DC1-SJC-cEdge1#sh sdwan bfd sessions                                  SOURCE TLOC      REMOTE TLOC DST PUBLIC                                DST PUBLIC DETECT    TX SYSTEM IP      SITE ID STATE          COLOR COLOR              SOURCE IP IP                                PORT ENCAP MULTIPLIER  INTERVAL(msec) UPTIME TRANSITIONS ----- -----</pre>	<p>From output on DC1, DC1 rejects all tunnels from DC2 system IPs and allows branch tunnels using multiple TLOCs</p>

-----			
-----			
-----			
172.16.3.1	30	up	mpls
mpls	10.1.1.10		
10.30.0.2			12346
IPsec 7	1000		11:03:35:44
0			
172.16.3.1	30	up	biz-
internet	biz-internet		10.1.254.10
192.0.2.40			12346
ipsec 7	1000		1:14:17:55
1			
172.16.5.1	50	up	mpls
mpls	10.1.1.10		
10.50.0.2			12346
ipsec 7	1000		11:03:35:37
0			
172.16.5.1	50	up	biz-
internet	biz-internet		10.1.254.10
192.0.2.50			5062
ipsec 7	1000		3:05:48:57
1			
172.16.5.2	50	up	mpls
mpls	10.1.1.10		
10.5.102.2			12346
ipsec 7	1000		11:03:35:44
0			
172.16.5.2	50	up	biz-
internet	biz-internet		10.1.254.10
192.0.2.50			12346
ipsec 7	1000		3:05:48:57
2			

## Dynamic Tunnels

### Intent:

Dynamic Tunnel for branch tunnel scalability for low end platforms

### Solution:

Unlike IWAN which uses DMVPN Phase 3 for building dynamic spoke-to-spoke tunnels, Cisco SD-WAN configurations allow permanent tunnels to be created with default configurations (shown in previous sections). This provides better visibility into circuit using bfd probe data and provides efficient path for data plane. However, in large deployments there can be thousands of sites forming permanent tunnels and with a low-end branch router, this can become a challenge (based on the number of tunnels supported per platform). Therefore, the enterprise wants to keep the same behavior of having dynamic tunnels for its spoke to spoke traffic. In order to configure dynamic tunnels first a backup path must be established using hub or any sites (in this cases DC Hub1) to provide packet forwarding path for initial flow until dynamic tunnels are spawned between spokes. When the interesting traffic is detected between spokes bidirectionally, dynamic tunnel will be created.

Note: Cisco SD-WAN provides capability to inter-operate between dynamic tunnel configured branches and static tunnel configured branches.

Configuration output	Explanation
<pre>control-policy Central-topology-policy-v2 sequence 21 match route site-list Dynamic-tunnel-branches prefix-list _AnyIpv4PrefixList ! action accept set tloc-action backup tloc-list DC-TLOCs ! ! ! default-action accept !</pre>	<p>Prepare centralized control policy</p> <p>Sequence 21 matches any route originating from dynamic tunnel branches matched using site-ids</p> <p>Set next-hop TLOC to DC TLOC as backup. This means that if the direct path is not yet established between branches, hub will be utilized as next hop to reach branch. Once dynamic tunnels are created, this hub based backup path is removed</p>
<pre>! site-list Dynamic-tunnel-branches site-id 30 site-id 50 ! tloc-list DC-TLOCs tloc 1.10.1.1 color mpls encap ipsec tloc 1.10.1.1 color biz-internet encap ipsec !</pre>	<p><b>Site-list</b> matches all the dynamic tunnel branches LA and BCN</p> <p><b>TLOC-list</b> defines next hop router and transport TLOCs. DC1-SANJOSE-CEDGE1. (in this case). The list can contain any other routers or sites as well but, in this example, only one hub router is taken as backup path for dynamic tunnel in the list.</p>
<pre>apply-policy site-list Branch_Group control-policy Central-topology-policy-v2 out !</pre>	<p>Policy applied to all branches participating in dynamic tunnels</p>
<pre>service TE vrf global</pre>	<p>This is required at HUBS which are providing backup paths to branch during tunnel establishment.</p> <p>Configure under VPN 0 template in vManage</p>
<pre>BR3-LAX-cEDGE#sh sdwan running-config system system system-ip          172.16.3.1 overlay-id         1 site-id            30 port-offset        0 control-session-pps 300 admin-tech-on-failure sp-organization-name "Viptela-POC-Tool - 19827" organization-name  "Viptela-POC-Tool - 19827" port-hop track-transport track-default-gateway console-baud-rate  115200 on-demand enable on-demand idle-timeout 10 vbond vbond-test-drive port 12346</pre>	<p>The spokes LA and Barcelona will be configured with <b>on-demand enable knob</b> enabled under system configs</p>
<pre>BR5-BCN-cEDGE1#sh sdwan system on-demand remote- system SITE-ID    SYSTEM-IP    ON-DEMAND    STATUS IDLE-TIMEOUT-EXPIRY(sec)</pre>	<p>On demand tunnel enabled branches show up as <b>'yes'</b>, from BCN-1 its shows LA and BCN-2 routers</p> <p><b>Inactive</b> shows tunnels not yet established</p>

<pre> ----- 10      1.10.1.1      no      - - 10      1.10.1.2      no      - - 20      1.20.1.1      no      - - 30      172.16.3.1    yes     inactive - 50      172.16.5.2    yes     inactive - BR5-BCN-cEDGE1# </pre>	
<pre> BR5-BCN-cEDGE2#sh sdwan system on-demand remote- system SITE-ID   SYSTEM-IP   ON-DEMAND   STATUS IDLE-TIMEOUT-EXPIRY(sec) ----- 10      1.10.1.1      no      - - 10      1.10.1.2      no      - - 20      1.20.1.1      no      - - 30      172.16.3.1    yes     inactive 50      172.16.5.1    yes     inactive </pre>	<p>On demand tunnel enabled branches show up as 'yes', from BCN-2 its shows LA and BCN-1 routers</p> <p><b>Inactive</b> shows tunnels not yet established</p>
<pre> BR5-BCN-cEDGE1#sh sdwan bfd session SYSTEM IP   SITE ID   STATE   COLO ----- 1.10.1.1    10       up      biz-internet 1.10.1.1    10       up      mpls 1.10.1.2    10       up      biz-internet 1.10.1.2    10       up      mpls 1.20.1.1    20       up      biz-internet 1.20.1.1    20       up      mpls </pre>	<p><b>show bfd output</b> only shows DC sites with which permanent tunnels are created</p>
<pre> viptela@ubuntu:~\$ tracepath -n 10.3.100.10 1?: [LOCALHOST] pmtu 1500 1:  10.5.100.1 1.700ms 1:  10.5.100.1 1.316ms 2:  10.5.1.1 1.350ms 3:  10.5.1.1 10.056ms pmtu 1434 <b>3:  10.1.254.10</b> <b>42.089ms</b> 4:  10.3.100.10 14.853ms reached Resume: pmtu 1434 hops 4 back 4 viptela@ubuntu:~\$ </pre>	<p>Initial tracepath shows ping from Barcelona to LA branch server taking the backup DC router, DC1-SJC-1 router (in bold).</p> <p>At this point dynamic tunnel establishment is triggered after bidirectional traffic detection</p>
<pre> viptela@ubuntu:~\$ tracepath -n 10.3.100.10 1?: [LOCALHOST] pmtu 1500 1:  10.5.100.1 1.397ms 1:  10.5.100.1 1.613ms 2:  10.5.1.1 1.213ms 3:  10.5.1.1 1.141ms pmtu 1434 <b>3:  192.0.2.40</b> <b>4.025ms</b> </pre>	<p>Subsequent trace shows path now is created and takes direct BCN to LA tunnels</p>

<pre>4: 10.3.100.10 6.953ms reached Resume: pmtu 1434 hops 4 back 4 viptela@ubuntu:~\$</pre>	
<pre>BR5-BCN-cEDGE1#sh sdwan system on-demand remote-system SITE-ID      SYSTEM-IP    ON-DEMAND    STATUS IDLE-TIMEOUT-EXPIRY(sec) ----- 10          1.10.1.1      no           - - 10          1.10.1.2      no           - - 20          1.20.1.1      no           - - 30          172.16.3.1    yes          active 72 50          172.16.5.2    yes          inactive - BR5-BCN-cEDGE1#</pre>	<p>Now tunnel status shows active with LA</p>
<pre>BR5-BCN-cEDGE1#sh sdwan bfd sessions SOURCE TLOC REMOTE TLOC SYSTEM IP    SITE ID STATE    COLOR COLOR        SOURCE IP ----- 1.10.1.1     10      up       biz-internet biz-internet 10.5.102.5 1.10.1.1     10      up       mpls mpls         10.50.0.2 1.10.1.2     10      up       biz-internet biz-internet 10.5.102.5 1.10.1.2     10      up       mpls mpls         10.50.0.2 1.20.1.1     20      up       biz-internet biz-internet 10.5.102.5 1.20.1.1     20      up       mpls mpls         10.50.0.2 <b>172.16.3.1</b>   <b>30</b>    <b>up</b>     <b>biz-</b> <b>internet</b>   <b>biz-internet</b> <b>10.5.102.5</b> <b>172.16.3.1</b>   <b>30</b>    <b>up</b>     <b>mpls</b> <b>mpls</b>       <b>10.50.0.2</b></pre>	<p>After Dynamic Tunnel triggered, bfd session from Barcelona show tunnels up with LA using both tunnels</p>

### Application Aware Routing (AAR)

**Intent:**

Match PFRv3 policy and intelligently route customer traffic using different transport based on tunnel health.

**Solution:**

Cisco SD-WAN solution, by default, utilizes custom bfd packet-based probing to calculate loss, latency and jitter values between tunnel endpoints. The same is used for peer liveliness detection. The probes are sent using DSCP 48 markings to calculate tunnel health over user configured time interval. Future release of Cisco SD-WAN will provide capability to send probes using custom DSCP and custom QoS queues in order to get better visibility of actual data behavior.

AAR policy is configured using centralize data policy and sent to all the sites using a site list. The policy below provides the closest alternative to the previous pfrv3 policy configured under IWAN sections.

Configuration Output	Explanation
<pre>viptela-policy:policy sla-class BULK_DATA1   latency 300   loss 5 ! sla-class DEFAULT1   latency 500   loss 10 ! sla-class LOW_LATENCY_DATA1   latency 100   loss 5 ! sla-class REAL_TIME_VIDEO1   latency 150   loss 1   jitter 20 ! sla-class SCAVENGER1   latency 500   loss 50 ! sla-class VOICE-ITU1   latency 150   loss 1   jitter 30 !  app-route-policy_VPN1-corporate_AAR-Policy-v2 vpn-list VPN1-corporate sequence 1 match   app-list mybusinessapp-list   source-ip 0.0.0.0/0 ! action   sla-class BULK_DATA1 preferred-color mpls ! ! sequence 11 match   dscp 46   source-ip 0.0.0.0/0 ! action   sla-class VOICE-ITU1 preferred-color mpls ! ! sequence 21 match   dscp 32 34 36 38   source-ip 0.0.0.0/0 ! action   sla-class REAL_TIME_VIDEO1 preferred- color mpls ! ! sequence 31 match   dscp 16 18 20 22 24</pre>	<p>The AAR policy is configured with the same intent as for PFR policy.</p> <p>SLA classes are defined for different traffic types with loss, latency, jitter values matching the business intent</p> <p>(Default, Bulk Data, Voice, Video, Low Latency Data and Scavenger).</p> <p>The traffic is classified using different DSCP markings in VPN 1 and associated with SLA traffic classes.</p> <p><b>preferred color mpls</b> ensures that MPLS is always preferred for traffic classes voice, video and low latency data. unless MPLS does not meet SLA, in that case, biz-internet will be preferred.</p> <p>Similarly, for traffic classes bulkdata, scavenger and default class the biz-internet color is preferred.</p> <p>The strict keyword ensures if none of the tunnel meeting SLA, scavenger traffic can be dropped.</p>

<pre> source-ip 0.0.0.0/0 ! action sla-class LOW_LATENCY_DATA1 preferred- color mpls ! ! sequence 41 match dscp 10 12 14 source-ip 0.0.0.0/0 ! action sla-class BULK_DATA1 preferred-color biz- internet ! ! sequence 51 match dscp 8 source-ip 0.0.0.0/0 ! action sla-class SCAVENGER1 strict preferred- color biz-internet ! ! sequence 61 match dscp 0 source-ip 0.0.0.0/0 ! action sla-class DEFAULT1 preferred-color biz- internet ! ! ! </pre>	
<pre> ! data-policy _VPN1-corporate_Traffic-data-policy vpn-list VPN1-corporate sequence 1 match dscp 8 source-ip 0.0.0.0/0 ! action accept set local-tloc-list color biz-internet encap ipsec restrict ! ! ! default-action accept ! </pre>	<p>When SLA class for scavenger doesn't meet SLA on internet, the AAR policy will move traffic to MPLS, which is not the intent for this type of traffic. In order for scavenger traffic to not select MPLS link, a traffic data policy is configured with a restrict keyword.</p> <p>The policy simply complements the AAR policy and implements blackholing for scavenger traffic if internet doesn't meet SLA by restricting the TLOC to only biz-internet</p>
<pre> lists site-list ALL-SITES site-id 10 site-id 20 site-id 30 </pre>	<p>List of All sites used for the policies implementation</p>

<pre> site-id 50 !</pre>	
<pre> apply-policy site-list ALL-SITES data-policy _VPN1-corporate_Traffic-data- policy from-service app-route-policy _VPN1-corporate_AAR-Policy-v2</pre>	Applying policy from service side to all sites for VPN 1 traffic.

## Custom-App and SD-AVC

### Intent:

Customer have a business web-app which works on TCP port 8500. Due to criticality of the business app, customer wants first packet match capability for detecting the app and use app-route policy to send it to MPLS link for better App performance

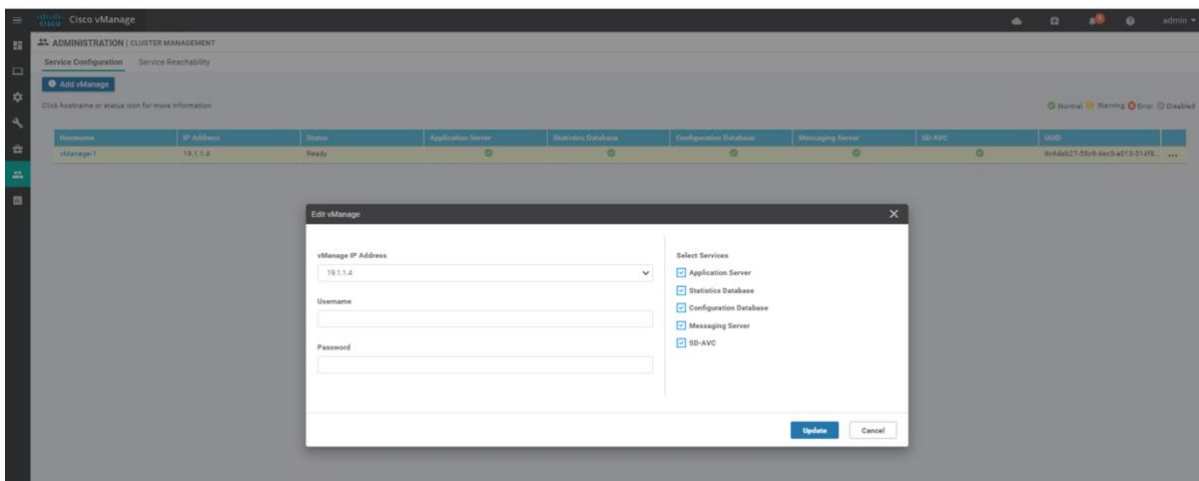
Note: This is a hypothetical customer scenario which is configured in this customer use case to show Cisco SD-WAN SD-AVC and custom-app matching capabilities. Apart from TCP/UDP ports current Cisco SD-WAN solution provides capability to match on domain-names using regex. Future release of Cisco SD-WAN will provide capability to match URLs as well (see roadmap). In this customer scenario web-app is hosted on port 8500 on SANJOSE DC server

### Solution:


[Cisco SD-AVC](#) uses Cisco NBAR2 and other components that operate on devices in the network to provides recognition of network application traffic for visibility, analytics, application-aware routing, and application-based policies, such as QoS and application-based firewall policy. It also provides analytics at the network level.

In addition to the standard protocols provided in a Protocol Pack, you can define protocols, called custom applications, to identify internet traffic, often for uncommon network applications that are of specific interest to their organization.

[Custom applications](#) augment the protocols provided in a Protocol Pack. SD-AVC is enabled under vManage cluster settings. Once configured SDAVC service runs over the network.







Configuration	Explanation
<pre>BR5-BCN-cEDGE2#sh sdwan running-config policy policy app-visibility flow-visibility</pre>	<p><b>app-visibility</b> local policy enables IP NBAR protocol discovery capabilities under service side interfaces</p>
<pre>BR5-BCN-cEDGE2#sh run int GigabitEthernet 2 Building configuration... Current configuration : 251 bytes ! interface GigabitEthernet2 description VPN1 -TO - DC-Core1 vrf forwarding 1 ip address 10.5.2.1 255.255.255.252 no ip redirects <b>ip nbar protocol-discovery</b> load-interval 30 negotiation auto arp timeout 1200 no mop enabled no mop sysid end BR5-BCN-cEDGE2#</pre>	<p><b>ip nbar</b> is enabled under service interfaces.</p>
<pre>from-vsmart lists app-list mybusinessapp-list app mybusinessapp</pre>	<p>A custom business app is created under app-list. vManage provides interface to create custom app definitions</p>
<pre>BR5-BCN-cEDGE2#sh ip nbar protocol-id mybusinessapp Protocol Name          id type ----- mybusinessapp         3147 PPDK LOCAL BR5-BCN-cEDGE2#</pre>	<p>App id is assigned to custom app</p>
<pre>BR5-BCN-cEDGE2#sh ip nbar protocol-attribute mybusinessapp Protocol Name : mybusinessapp   encrypted : encrypted=yes   tunnel : tunnel=yes   category : browsing   sub-category : enterprise-transactional-apps application-group : other p2p-technology : p2p-tech-no traffic-class : transactional-data business-relevance : default application-set : general-browsing application-family : web</pre>	<p>General custom-app attribute can be seen</p>
<pre>from-vsmart app-route-policy _VPN1-corporate_AAR-Policy-v2 vpn-list VPN1-corporate sequence 1 match source-ip 0.0.0.0/0 app-list mybusinessapp-list action</pre>	<p>App route policy matches on mybusinessapp (custom-app) and prefers MPLS</p>

<pre>sla-class BULK_DATA1 no sla-class strict sla-class preferred-color mpls</pre>	
<pre>viptela@ubuntu:~\$ wget http://10.1.100.10:8500 --2020-11-09 20:01:02-- http://10.1.100.10:8500/ Connecting to 10.1.100.10:8500... connected. HTTP request sent, awaiting response... 200 OK Length: 1558 (1.5K) [text/html] Saving to: 'index.html.1' index.html.1 100%[=====&gt;] 1.52K -- .-KB/s in 0s 2020-11-09 20:01:03 (106 MB/s) - 'index.html.1' saved [1558/1558] viptela@ubuntu:~\$</pre>	<p>Client does a HTTP get request on business app hosted on SANJOSE DC over port 8500</p>
<pre>BR5-BCN-cEDGE2#show ip nbar protocol-discovery protocol GigabitEthernet2 Last clearing of "show ip nbar protocol-discovery" counters 01:36:24  Output Input ----- Protocol Packet Count Packet Count Byte Count Byte Count 30sec Bit Rate (bps) 30sec Bit Rate (bps) 30sec Max Bit Rate (bps) 30sec Max Bit Rate (bps) ----- ----- dns 11546 0 882668 0 1000 0 1000 0 icmp 0 6275 0 653316 0 1000 0 1000 0 <b>mybusinessapp 6</b> 6 0 2118 0 0 0 1000 0 bgp 0 2392 0 152923 0</pre>	<p>Ip nbar protocol discovery detects the app and shows counters</p>
<pre>BR5-BCN-cEDGE2#show sdwan app-fwd cflowd flows vpn 1 format table 1 10.1.100.10 10.5.100.10 8500 60154 0 6 27 0 6 2034 Tue Nov 10 01:14:20 2020 GigabitEthernet2 GigabitEthernet3.101 mybusinessapp web No Drop 0 0 0 0 2</pre>	<p>Cflowd output shows mybusinessapp under VPN 1 utilizing MPLS for egress as per app-route policy</p>

<pre>BR3-LAX-cEDGE# sh ip nbar classification cache sync import last   include mybus  4        0.0.0.0/0        8500          mybusinessapp  5        :::/0        8500          mybusinessapp</pre>	App definitions learned on LA branch through SD-AVC service.
--	--

## NAT DIA and ZBFW

### Intent:

- Provide internet access to branches using local internet exit. If local internet goes down, use data centers to exit to internet.
- Only allow web and ICMP traffic to exit to internet.

### Solution:

Spoke branches needs to access internet with minimal latency and without choking DC bandwidth. If internet goes down on the local branch, branches can use default route from DC for backup internet routing. Therefore, NAT DIA is configured on all the spoke branches with tracking functionality. Apart from configuring DIA, enterprise wanted only web and ICMP traffic to go out for internet and block rest of the traffic which is achieved using zone based firewall security policy under Greatwall security suite of Cisco SD-WAN. (see detail documentation for NAT DIA for other methods for DIA).

Looking at the configuration on branch sites:

Configuration Commands	Explanation
<pre>ip nat route VRF 1 0.0.0.0 0.0.0.0 global</pre>	A static default route configured under VRF 1 exiting through VPN 0 routing table.
<pre>BR3-LAX-cEDGE#sh running-config interface GigabitEthernet 1 Building configuration... Current configuration : 299 bytes ! interface GigabitEthernet1 description INET WAN interface ip address 192.0.2.40 255.255.255.0 no ip redirects <b>ip nat outside</b> load-interval 30 negotiation auto endpoint-tracker track-internet arp timeout 1200 no mop enabled no mop sysid service-policy output shape_GigabitEthernet1 end</pre>	IP NAT outside required for DIA to work
<pre><b>endpoint-tracker track-internet</b> <b>endpoint-ip 8.8.8.8</b> <b>tracker-type interface</b> threshold 100 interval 20 multiplier 2 ! interface GigabitEthernet1 description INET WAN interface</pre>	<p>An endpoint tracker is configured, which removes the NAT default route from VPN 1 if destination is not reachable. These are HTTP probes which are sent to destination.</p> <p>The tracker is applied for internet facing NAT interface</p>

<pre> no shutdown arp timeout 1200 ip address 192.0.3.50 255.255.255.0 no ip redirects ip mtu 1500 ip nat outside load-interval 30 mtu 1500 <b>endpoint-tracker track-internet</b> negotiation auto service-policy output shape_GigabitEthernet1 </pre>	
<pre> policy zone-based-policy ent-sec-pol sequence 1 seq-name internet-allow match source-data-prefix-list LAN-prefix <b>destination-port 443 80 2164 53 90</b> protocol 6 17 <b>protocol-name https http ddns-v3 dns dnsix</b> ! <b>action inspect</b> ! ! sequence 11 seq-name icmp-allow match source-data-prefix-list LAN-prefix protocol 1 17 ! action inspect ! ! <b>default-action drop</b> ! zone vpn0-zone vpn 0 ! zone vpn1-zone vpn 1 ! zone-pair ZP_vpn1-zone_vpn0-zo_-2092444543 source-zone vpn1-zone destination-zone vpn0-zone zone-policy ent-sec-pol ! lists data-prefix-list LAN-prefix ip-prefix 10.0.0.0/8 ! ! zone-to-nozone-internet deny ! </pre>	<p>ZBFW configuration with packet matching DNS, HTTP, HTTPS, ICMP traffic inspected.</p> <p>Zone pair created between VPN 1 to VPN 0 for traffic security</p> <p>Everything else is denied between the zones</p>
<pre> BR3-LAX-cEDGE#sh ip route vrf 1 Routing Table: 1 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS- IS level-2 ia - IS-IS inter area, * - candidate default, U - per-user static route H - NHRP, G - NHRP registered, g - NHRP registration summary o - ODR, P - periodic downloaded static route, l - LISP a - application route + - replicated route, % - next hop override, p - overrides from Pfr </pre>	<p>A NAT route with AD 6 installed in the VRF 1 routing table.</p>

<pre>&amp; - replicated local route overrides by connected  Gateway of last resort is 0.0.0.0 to network 0.0.0.0 n*Nd 0.0.0.0/0 [6/0], 07:42:21, Null0</pre>	
<pre>BR3-LAX-cEDGE#sh sdwan OMP routes Code: C -&gt; chosen I -&gt; installed Red -&gt; redistributed Rej -&gt; rejected L -&gt; looped R -&gt; resolved S -&gt; stale Ext -&gt; extranet Inv -&gt; invalid Stg -&gt; staged IA -&gt; On-demand inactive U -&gt; TLOC unresolved            PATH ATTRIBUTE VPN  PREFIX          FROM PEER      ID LABEL STATUS  TYPE          TLOC IP        COLOR ENCAP PREFERENCE ----- ----- 1      0.0.0.0/0          1.1.1.3        41      1002 C,I,R  installed  1.10.1.2      mpls ipsec  400 C,I,R  installed  1.1.1.3        42      1002 ipsec  400 C,I,R  installed  1.10.1.2      biz-internet R      installed  1.20.1.1      mpls ipsec  200 R      installed  1.1.1.3        161     1002 ipsec  200 R      installed  1.20.1.1      biz-internet ipsec  200 C,I,R  installed  1.1.1.3        263     1002 ipsec  400 C,I,R  installed  1.10.1.1      mpls ipsec  400 C,I,R  installed  1.1.1.3        264     1002 ipsec  400 C,I,R  installed  1.10.1.1      biz-internet ipsec  400</pre>	<p>Default route from both DCs also received though OMP.</p> <p>Only DC1 routes are installed due to higher OMP preference.</p>
<pre>BR3-LAX-cEDGE#sh ip route vrf 1 Routing Table: 1 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS- IS level-2 ia - IS-IS inter area, * - candidate default, U - per-user static route H - NHRP, G - NHRP registered, g - NHRP registration summary o - ODR, P - periodic downloaded static route, l - LISP a - application route + - replicated route, % - next hop override, p - overrides from PFR &amp; - replicated local route overrides by connected Gateway of last resort is 0.0.0.0 to network 0.0.0.0 n*Nd 0.0.0.0/0 [6/0], 00:01:34, Null0</pre>	<p>Although default received using OMP but due to OMP AD of 251, NAT DIA route will be preferred.</p> <p>When NAT DIA route goes down due to TLOC or tracker down, OMP default route will be installed pointing to DC1.</p>
<pre>BR5-BCN-cEDGE1#sh ip route vrf 1 Routing Table: 1 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2</pre>	<p>For redundant branch like Barcelona running TLOC extension, the WAN-EDGE1 doesn't terminate internet link, therefore NAT default is not configured on this router, and only OMP default route is installed.</p>

<pre> E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS- IS level-2 ia - IS-IS inter area, * - candidate default, U - per-user static route H - NHRP, G - NHRP registered, g - NHRP registration summary o - ODR, P - periodic downloaded static route, l - LISP a - application route + - replicated route, % - next hop override, p - overrides from PFR &amp; - replicated local route overrides by connected Gateway of last resort is 1.20.1.1 to network 0.0.0.0 m* 0.0.0.0/0 [251/0] via 1.20.1.1, 6d20h, Sdwan-system-intf [251/0] via 1.10.1.1, 6d20h, Sdwan-system-intf </pre>	<p>This same route is sent to Barcelona core after OMP to BGP redistribution which adds a metric of 1000 by default.</p> <p>A network 0/0 is configured under BGP to propagate default to neighbors.</p>
<pre> BR5-BCN-cEDGE2#sh ip route vrf 1 Routing Table: 1 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default, U - per-user static route H - NHRP, G - NHRP registered, g - NHRP registration summary o - ODR, P - periodic downloaded static route, l - LISP a - application route + - replicated route, % - next hop override, p - overrides from PFR &amp; - replicated local route overrides by connected Gateway of last resort is 0.0.0.0 to network 0.0.0.0 n*Nd 0.0.0.0/0 [6/0], 5d10h, Null0 </pre>	<p>Barcelona internet terminating wan edge-2 uses a NAT default route.</p> <p>OMP routes not installed in this router due to OMP AD 251 vs NAT AD 6.</p> <p>Same default route is advertised using BGP network statement to BCN core.</p>
<pre> BR5-BCN-cEDGE2#sh sdwan running-config   section bgp router bgp 64510   bgp log-neighbor-changes   distance bgp 20 200 20   address-family ipv4 unicast vrf 1     bgp router-id 172.16.5.2     neighbor 10.5.2.2 remote-as 65054     neighbor 10.5.2.2 activate     neighbor 10.5.2.2 description BCN Core Router     neighbor 10.5.2.2 ebgp-multihop 1     neighbor 10.5.2.2 send-community both     network 0.0.0.0 mask 0.0.0.0   redistribute connected   redistribute omp   redistribute nat-route dia   exit-address-family ! </pre>	<p>Both BCN-1 and BCN-2 also advertises this default route to BCN-core using network 0.0.0.0 statement.</p>
<pre> br5-bcn-rtr#sh ip route Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default, U - per-user static route o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP a - application route + - replicated route, % - next hop override, p - overrides from PFR Gateway of last resort is 10.5.2.1 to network 0.0.0.0 B* 0.0.0.0/0 [20/0] via 10.5.2.1, 5d10h  br5-bcn-rtr#sh ip bgp 0.0.0.0 BGP routing table entry for 0.0.0.0/0, version 90 Paths: (2 available, best #1, table default)   Advertised to update-groups:     3   Refresh Epoch 1   64510 </pre>	<p>BCN core only installs route from BCN-2 having DIA which is the desired behavior.</p> <p>This behavior is achieved because BCN-1 is redistributing OMP default which adds high metric value of 1000 while BCN-2 have a local static default DIA route which BGP simply advertise using network statement without adding metric value.</p> <p>If tracker is down, both BCN WAN edges will advertise default using OMP route. BCN core can chose any route or do ECMP.</p>

<pre> 10.5.2.1 from 10.5.2.1 (172.16.5.2)   Origin IGP, metric 0, localpref 100, valid, external, best   Extended Community: RT:64510:1   rx pathid: 0, tx pathid: 0x0 Refresh Epoch 1 64510 10.5.1.1 from 10.5.1.1 (172.16.5.1)   Origin IGP, metric 1000, localpref 100, valid, external   Extended Community: SoO:0:50 RT:64510:1   rx pathid: 0, tx pathid: 0 br5-bcn-rtr# </pre>	
<pre> viptela@ubuntu:~\$ wget cisco.com --2020-11-09 20:51:59-- http://cisco.com/ Resolving cisco.com (cisco.com)... 184.25.199.192 Connecting to cisco.com (cisco.com) 184.25.199.192 :80... connected. HTTP request sent, awaiting response... 200 OK Length: 530 [text/html] Saving to: 'index.html.10' index.html.10      100%[=====&gt;]          530  -- .-KB/s    in 0s 2020-11-09 20:51:59 (106 MB/s) - 'index.html.10' saved [530/530] viptela@ubuntu:~\$ </pre>	<p>Testing HTTP traffic to cisco.com provides DIA over the branch</p>
<pre> viptela@ubuntu:~\$ tracepath cisco.com 1?: [LOCALHOST] pmtu 1500  1:  10.3.100.1 0.735ms  1:  10.3.100.1 0.501ms  2:  cisco.com 2.315ms reached Resume: pmtu 1500 hops 2 back 1 viptela@ubuntu:~\$ </pre>	<p>Tracing path to cisco.com (184.25.199.192) takes the traffic to internet using DIA configurations instead of routing it through DC</p>

## QoS Configurations

### Intent:

- Customer need different DSCPs tagged traffic to be mapped to different queues
- In order to prevent the HUB from choking the spoke branches due to difference of shaper, per-tunnel QoS configuration is required to intelligently shape traffic based on spoke download capacity.

### Solution:

Customer’s IWAN network runs the QoS configuration with advance support for per-tunnel QoS. Following section provides the mapping of SD-WAN QoS configuration covering per tunnel QoS with basic QoS configuration (for more details, read the detail documentation on SD-WAN [QoS config](#)). SD-WAN makes it easy to do QoS configuration using vManage GUI and all relevant cli configurations are automatically generated once configured using GUI.

### QoS config at DC-HUB

Configuration	Explanation
---------------	-------------

```

access-list QOS-ACL
sequence 1
match
  dscp 26 28 30 40
!
action accept
  class STREAMING-VIDEO
!
!
sequence 11
match
  dscp 32 34 36 38
!
action accept
  class INTERACTIVE_VIDEO
!
!
sequence 21
match
  dscp 10 12 14 16 18 20 22
!
action accept
  class CRITICAL-DATA
!
!
sequence 31
match
  dscp 46
!
action accept
  class VOICE
!
!
sequence 41
match
  dscp 8
!
action accept
  class SCAVENGER
!
!
sequence 51
match
  dscp 24
!
action accept
  class CALL-SIGNALING
!
!
sequence 61
match
  dscp 48
!
action accept
  class NET-CTRL
!
!
sequence 71
match
  dscp 0
!
action accept
  class DEFAULT

```

Using the same criteria as in IWAN, SD-WAN supports total of 8 queues (7+default). By default, queue 2 is the default queue but in this config, in order to match IWAN QoS config, the default queue is mapped to queue 3.

A local policy ACL is created to map the already marked DSCP traffic received by the customer traffic and mapped to the relevant QoS classes. Example DSCP 46 is mapped to class VOICE.

Note: The same could have been achieved using central data policy which can map application or DSCP and map it to QoS classes.



<pre> ! ! default-action drop ! ! </pre>	
<pre> class-map class Queue0 queue 0 class VOICE queue 0 class CRITICAL-DATA queue 1 class Queue1 queue 1 class Queue2 queue 2 class SCAVENGER queue 2 class DEFAULT queue 3 class Queue3 queue 3 class INTERACTIVE_VIDEO queue 4 class Queue4 queue 4 class NET-CTRL queue 5 class Queue5 queue 5 class Queue6 queue 6 class STREAMING-VIDEO queue 6 class CALL-SIGNALING queue 7 class Queue7 queue 7 ! </pre>	<p>Classes are mapped to queues.</p> <p>Note: These configs are auto-generated using GUI local policy configuration for QoS.</p>
<pre> class-map match-any CALL-SIGNALING match QoS-group 7 ! class-map match-any CRITICAL-DATA match qos-group 1 ! class-map match-any DEFAULT match qos-group 3 ! class-map match-any INTERACTIVE_VIDEO match qos-group 4 ! class-map match-any NET-CTRL match qos-group 5 ! class-map match-any Queue0 match qos-group 0 ! class-map match-any Queue1 match qos-group 1 ! class-map match-any Queue2 match qos-group 2 ! class-map match-any Queue3 match qos-group 3 ! class-map match-any Queue4 match qos-group 4 ! class-map match-any Queue5 match qos-group 5 ! class-map match-any Queue6 match qos-group 6 ! class-map match-any Queue7 match qos-group 7 ! </pre>	<p>Classes mapped to QoS-group as scheduling is applied using QoS-group</p>

<pre> class-map match-any SCAVENGER   match qos-group 2 ! class-map match-any SDWAN_underlay   match any ! class-map match-any STREAMING-VIDEO   match qos-group 6 ! class-map match-any VOICE   match qos-group 0 ! </pre>	
<pre> policy-map WAN-QOS   class Queue0     priority level 1     police rate percent 10   !   !   class Queue1     bandwidth remaining ratio 25     random-detect precedence-based   !   class class-default     bandwidth remaining ratio 1     random-detect precedence-based   !   class Queue3     bandwidth remaining ratio 15     random-detect precedence-based   !   class Queue4     bandwidth remaining ratio 30     random-detect precedence-based   !   class Queue5     bandwidth remaining ratio 5     random-detect precedence-based   !   class Queue6     bandwidth remaining ratio 10     random-detect precedence-based   !   class Queue7     bandwidth remaining ratio 4     random-detect precedence-based   ! ! </pre>	<p>The QoS scheduler defines the bandwidth and scheduling for each class of traffic. It also defines congestion avoidance mechanism like red-drop or tail drop for specific class.</p> <p>By default, all control traffic is sent queue0 which is a Priority queue.</p> <p>Voice traffic is also mapped to the same priority queue.</p> <p>Bandwidth remaining ratio provides granularity and control on traffic distribution based on ratios.</p>
<pre> policy-map per_tunnel_qos_policy_GigabitEthernet1   class SDWAN_underlay     <b>bandwidth remaining percent 10</b>     service-policy WAN-QOS   ! ! policy-map per_tunnel_qos_policy_GigabitEthernet2   class SDWAN_underlay     <b>bandwidth remaining percent 10</b>     service-policy WAN-QOS   ! ! </pre>	<p>Nested QoS defines bandwidth remaining percent between underlay and overlay.</p> <p>When configured from vManage, the tunnel bandwidth percent is defined to be used for overlay, so the remaining percent is automatically applied for underlay. 90% is reserved for overlay so 10% is remaining percentage for underlay.</p> <p>The underlay bandwidth can be anything from TLOC-extension, DIA or traffic not going in overlay. Scheduler WAN-QOS is also applied which is same for overlay.</p>

<pre> policy-map shape_GigabitEthernet1 class class-default service-policy per_tunnel_qos_policy_GigabitEthernet1 shape average 900000000 ! ! policy-map shape_GigabitEthernet2 class class-default service-policy per_tunnel_qos_policy_GigabitEthernet2 shape average 600000000 ! ! </pre>	<p>Parent shaper applied accordingly</p>
<pre> interface GigabitEthernet1 description INET WAN interface ip address 10.1.254.10 255.255.255.252 no ip redirects ip nat outside load-interval 30 negotiation auto arp timeout 1200 no mop enabled no mop sysid <b>service-policy output shape_GigabitEthernet1</b> ! interface GigabitEthernet2 description MPLS WAN interface ip address 10.1.1.10 255.255.255.252 no ip redirects load-interval 30 negotiation auto arp timeout 1200 no mop enabled no mop sysid <b>service-policy output shape_GigabitEthernet2</b> </pre>	<p>Policy-map applied under interfaces</p>
<pre> sdwan service TE vrf global ! interface GigabitEthernet1 tunnel-interface encapsulation ipsec weight 1 no border color biz-internet <b>tunnel-qos hub</b> no last-resort-circuit no low-bandwidth-link no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier default nat-refresh-interval 5 hello-interval 1000 hello-tolerance 12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp </pre>	<p>Hub sites are configured with tunnel-QoS hub command. Also tunnel interfaces are configured with DSCP re-write rule</p>

<pre> no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit rewrite-rule REWRITE-POLICY exit interface GigabitEthernet2 tunnel-interface encapsulation ipsec weight 1 no border color mpls restrict <b>tunnel-qos hub</b> no last-resort-circuit no low-bandwidth-link max-control-connections      0 no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier                       default nat-refresh-interval          5 hello-interval                1000 hello-tolerance               12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit rewrite-rule REWRITE-POLICY </pre>	
<pre> rewrite-rule REWRITE-POLICY class CALL-SIGNALING low dscp 10 layer-2-cos 1 class CRITICAL-DATA low dscp 18 layer-2-cos 2 class DEFAULT low dscp 10 layer-2-cos 1 class INTERACTIVE_VIDEO low dscp 10 layer-2-cos 1 class NET-CTRL low dscp 10 layer-2-cos 1 class SCAVENGER low dscp 10 layer-2-cos 1 class STREAMING-VIDEO low dscp 10 layer-2-cos 1 class VOICE low dscp 46 layer-2-cos 2 ! </pre>	<p>Optional re-write policy to match SP markings for DSCP and CoS</p>
<pre> policy-map SDWANPolicy4210704 class class-default shape average 3000000 bandwidth remaining ratio 3 service-policy WAN-QOS policy-map SDWANPolicy4210707 class class-default shape average 5000000 bandwidth remaining ratio 5 service-policy WAN-QOS policy-map SDWANPolicy4210719 class class-default </pre>	<p>Auto generated shapers for different tunnels based on branch downstream bandwidth configuration</p>

<pre> shape average 5000000 bandwidth remaining ratio 5 service-policy WAN-QOS policy-map SDWANPolicy4210717 class class-default shape average 30000000 bandwidth remaining ratio 3 service-policy WAN-QOS policy-map SDWANPolicy4210714 class class-default shape average 50000000 bandwidth remaining ratio 5 service-policy WAN-QOS policy-map SDWANPolicy4210713 class class-default shape average 30000000 bandwidth remaining ratio 3 service-policy WAN-QOS !</pre>	
--	--

### Branch QoS sample config at LA

Configuration	Explanation
<pre> access-list QOS-ACL sequence 1 match dscp 26 28 30 40 ! action accept class STREAMING-VIDEO ! ! sequence 11 match dscp 32 34 36 38 ! action accept class INTERACTIVE_VIDEO ! ! sequence 21 match dscp 10 12 14 16 18 20 22 ! action accept class CRITICAL-DATA ! ! sequence 31 match dscp 46 ! action accept class VOICE ! ! sequence 41 match</pre>	<p>Using the same criteria as in IWAN, SD-WAN supports total of 8 queues (7+default). By default, queue 2 is the default queue but in this config, in order to match IWAN QoS config, the default queue is mapped to queue 3.</p> <p>A local policy ACL is created to map the already marked DSCP traffic received by the customer traffic and mapped to the relevant QoS classes. Example DSCP 46 is mapped to class VOICE.</p> <p>Note: The same could have been achieved using central data policy which can map application or DSCP and map it to QoS classes.</p>

<pre> dscp 8 ! action accept class SCAVENGER ! sequence 51 match dscp 24 ! action accept class CALL-SIGNALING ! ! sequence 61 match dscp 48 ! action accept class NET-CTRL ! ! sequence 71 match dscp 0 ! action accept class DEFAULT ! ! default-action drop ! !</pre>	
<pre> class-map class Queue0 queue 0 class VOICE queue 0 class CRITICAL-DATA queue 1 class Queue1 queue 1 class Queue2 queue 2 class SCAVENGER queue 2 class DEFAULT queue 3 class Queue3 queue 3 class INTERACTIVE_VIDEO queue 4 class Queue4 queue 4 class NET-CTRL queue 5 class Queue5 queue 5 class Queue6 queue 6 class STREAMING-VIDEO queue 6 class CALL-SIGNALING queue 7 class Queue7 queue 7 !</pre>	<p>Classes are mapped to queues.</p> <p>Note: These configs are auto-generated using GUI local policy configuration for QoS.</p>
<pre> class-map match-any CALL-SIGNALING match qos-group 7 ! class-map match-any CRITICAL-DATA match qos-group 1 ! class-map match-any DEFAULT match qos-group 3 ! class-map match-any INTERACTIVE_VIDEO match qos-group 4 !</pre>	<p>Classes mapped to QoS-group</p>

<pre> class-map match-any NET-CTRL   match qos-group 5 ! class-map match-any Queue0   match qos-group 0 ! class-map match-any Queue1   match qos-group 1 ! class-map match-any Queue2   match qos-group 2 ! class-map match-any Queue3   match qos-group 3 ! class-map match-any Queue4   match qos-group 4 ! class-map match-any Queue5   match qos-group 5 ! class-map match-any Queue6   match qos-group 6 ! class-map match-any Queue7   match qos-group 7 ! class-map match-any SCAVENGER   match qos-group 2 ! class-map match-any STREAMING-VIDEO   match qos-group 6 ! class-map match-any VOICE   match qos-group 0 ! </pre>	
<pre> policy-map WAN-QOS   class Queue0     priority level 1     police rate percent 10   !   class Queue1     bandwidth remaining ratio 25     random-detect precedence-based   !   class class-default     bandwidth remaining ratio 1     random-detect precedence-based   !   class Queue3     bandwidth remaining ratio 15     random-detect precedence-based   !   class Queue4     bandwidth remaining ratio 30     random-detect precedence-based   !   class Queue5     bandwidth remaining ratio 5     random-detect precedence-based   !   class Queue6 </pre>	<p>The QoS scheduler defines the bandwidth and scheduling for each class of traffic. It also defines congestion avoidance mechanism like red-drop or tail drop for specific class.</p> <p>By default, all control traffic is sent queue0 which is a Priority queue.</p> <p>Voice traffic is also mapped to the same priority queue.</p> <p>Bandwidth remaining ratio provides granularity and control on traffic distribution based on ratios.</p>

<pre>bandwidth remaining ratio 10 random-detect precedence-based ! class Queue7 bandwidth remaining ratio 4 random-detect precedence-based ! !</pre>	
<pre>rewrite-rule REWRITE-POLICY class CALL-SIGNALING low dscp 10 layer-2-cos 1 class CRITICAL-DATA low dscp 18 layer-2-cos 2 class DEFAULT low dscp 10 layer-2-cos 1 class INTERACTIVE_VIDEO low dscp 10 layer-2-cos 1 class NET-CTRL low dscp 10 layer-2-cos 1 class SCAVENGER low dscp 10 layer-2-cos 1 class STREAMING-VIDEO low dscp 10 layer-2-cos 1 class VOICE low dscp 46 layer-2-cos 2 !</pre>	<p>Optional re-write policy provides re-marking to match SP markings</p>
<pre>policy-map shape_GigabitEthernet2 class class-default shape average 3000000 service-policy WAN-QOS policy-map shape_GigabitEthernet1 class class-default shape average 5000000 service-policy WAN-QOS</pre>	<p>Policy map is applied to interfaces for both internet and MPLS</p>
<pre>sdwan interface GigabitEthernet1 tunnel-interface encapsulation ipsec weight 1 no border color biz-internet tunnel-qos spoke no last-resort-circuit no low-bandwidth-link no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier default nat-refresh-interval 5 hello-interval 1000 hello-tolerance 12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit <b>bandwidth-downstream 50000</b> rewrite-rule REWRITE-POLICY exit interface GigabitEthernet2 tunnel-interface</pre>	<p>Bandwidth downstream dictates the per tunnel shaper value for every TLOC when sending traffic from hub to branch. The bandwidth downstream is communicated to the HUB using OMP protocol so that HUB can appropriately apply shaper for that tunnel.</p>



---

<pre>encapsulation ipsec weight 1 no border color mpls restrict tunnel-qos spoke no last-resort-circuit no low-bandwidth-link max-control-connections      0 no vbond-as-stun-server vmanage-connection-preference 5 port-hop carrier                       default nat-refresh-interval         5 hello-interval               1000 hello-tolerance              12 allow-service all no allow-service bgp allow-service dhcp allow-service dns allow-service icmp no allow-service sshd no allow-service netconf no allow-service ntp no allow-service ospf no allow-service stun allow-service https no allow-service snmp no allow-service bfd exit <b>bandwidth-downstream 30000</b> rewrite-rule REWRITE-POLICY exit appqoe no tcpopt enable !</pre>	
--	--

## 5.4.8 Day2 Monitoring and Serviceability

There are various ways you can manage and monitor a router. Management interfaces provide access to devices in the Cisco SD-WAN overlay network, allowing you to collect information from the devices in an out-of-band fashion and to perform operations on the devices, such as configuring and rebooting them. For details on serviceability topics refer to [Cisco Live session on serviceability](#)

### SNMP

The Simple Network Management Protocol (SNMP) allows you to manage all Cisco SD-WAN devices in the overlay network. The Cisco SD-WAN software supports SNMP v2c. You can configure basic SNMP properties—device name, location, contact, and community—that allow the device to be monitored by an SNMP Network Management System (NMS). You can configure trap groups and SNMP servers to receive traps. The object identifier (OID) for the internet port of the SNMP MIB is 1.3.6.1. Refer to [SNMP guide](#) for more detail.

### SYSLOG

System logging operations use a mechanism similar to the UNIX syslog command to record system-wide, high-level operations that occur on the Cisco SD-WAN devices in the overlay network. The log levels (priorities) of the messages are the same as those in standard UNIX commands, and you can configure the priority of the syslog messages that should be logged. Messages can be logged to files on the Cisco SD-WAN device or to a remote host. Logging can be configured using [Config guide](#).

### FNF

Cflowd traffic flow monitoring is equivalent to Flexible NetFlow (FNF). Cflowd monitors traffic flowing through Cisco IOS XE SD-WAN devices in the overlay network and exports flow information to a collector, where it can be processed by an IPFIX analyzer. For a traffic flow, Cflowd periodically sends template reports to flow collector. These reports contain information about the flows and the data is extracted from the payload of these reports. (Traffic Flow Monitoring with [Cflowd](#)). Additionally, FNF can be monitored locally on the device by enabling it through local-policy- flow-visibility command.

Configuration	Explanation
<pre>BR5-BCN-cEDGE2#sh sdwan running-config policy policy   app-visibility   flow-visibility</pre>	<p>The app visibility commands, once configured using vManage local policy, expands to FNF IOS XE configurations.</p> <p>The flow output is used as one of the major tools for tracing flows and observing policies impacts</p> <p>The flow can also be sent to external NetFlow analyzer like LiveAction</p>
<pre>BR5-BCN-cEDGE2#sh running-config   section flow flow record sdwan_flow_record   description flow_and_application_visibility records   match ipv4 destination address   match ipv4 protocol   match ipv4 source address   match routing vrf service   match transport destination-port   match transport source-port collect application name collect connection id long collect counter bytes long</pre>	<p>In the current customer environment, only local policy is configured to view NetFlow data instead of external collector.</p> <p>The policy visibility expands to this IOS XE configurations automatically.</p>

<pre> collect counter bytes sdwan dropped long collect counter packets long collect counter packets sdwan dropped long collect flow end-reason collect interface input collect interface output collect ipv4 dscp collect overlay session id input collect overlay session id output collect timestamp absolute first collect timestamp absolute last collect transport tcp flags collect drop cause id collect sdwan sla-not-met collect sdwan preferred-color-not-met collect sdwan qos-queue-id flow exporter sdwan_flow_exporter_0 description export flow and application visibility records to vManage destination local sdwan transport udp 5458 export-protocol ipfix option application-table option drop-cause-table option application-attributes flow monitor sdwan_flow_monitor description monitor flows for vManage and external collectors exporter sdwan_flow_exporter_0 cache timeout inactive 10 cache timeout active 60 cache entries 250000 record sdwan_flow_record ip visibility global flow monitor sdwan_flow_monitor input BR5-BCN-cEDGE2# </pre>	
<pre> BR5-BCN-cEDGE2#show sdwan app-fwd cflowd flows Generating output, this might take time, please wait ... app-fwd cflowd flows <b>vpn 1 src-ip 10.5.100.10</b> <b>dest-ip 8.8.8.8 src-port 33413 dest-port 53 dscp</b> <b>0 ip-proto 17</b> tcp-cntrl-bits      27 icmp-opcode        0 total-pkts         2 total-bytes        152 start-time         "Tue Nov 10 00:14:17 2020" egress-intf-name   GigabitEthernet1 ingress-intf-name  GigabitEthernet2 application        dns family            network-service drop-cause         "No Drop" drop-octets        0 drop-packets       0 sla-not-met        0 color-not-met      0 queue-id           2 fec-d-pkts         0 fec-r-pkts         0 pkt-dup-d-pkts-orig 0 pkt-dup-d-pkts-dup 0 pkt-dup-r-pkts     0 pkt-cxp-d-pkts     0 </pre>	<p>The Cflowd output provides per VRF source destination IP/ports/protocol/DSCP information including ingress/egress interfaces and advanced statistics.</p>

---

traffic-category	0
ssl-read-bytes	0
ssl-written-bytes	0
ssl-en-read-bytes	0
ssl-en-written-bytes	0
ssl-de-read-bytes	0
ssl-de-written-bytes	0
ssl-service-type	0
ssl-traffic-type	0
ssl-policy-action	0

### 5.4.9 API

The Cisco SD-WAN software provides a REST API, which is a programmatic interface for controlling, configuring, and monitoring the Viptela devices in an overlay network. You access the REST API through the vManage web server. Explore [SD-WAN APIs](#) on Cisco's DevNet site.

---

## 6 Cisco SD-WAN Advanced Use-cases

There are many other advanced capabilities and use cases which Cisco SD-WAN can provide in order to meet the needs of today and tomorrow's next generation wide area networks. The section below provides brief summaries and links to the documentation for each feature. Please contact your Cisco account team or representative for additional information.

### 6.1 AppQoE Features

#### FEC

Forward Error Correction (FEC) is a mechanism to recover lost packets on a link by sending extra "parity" packets for every group (N) of packets. As long as the receiver receives a subset of packets in the group (at-least N-1) and the parity packet, up to a single lost packet in the group can be recovered. The ideal use case for FEC is incase when customer have small transactions which needs to be protected from packet loss. Refer to FEC documentation for further details.

#### Packet Duplication

Packet duplication sends copies of packets on alternate available paths to reach Cisco IOS XE SD-WAN devices. If one of the packets is lost, a copy of the packet is forwarded to the server. Receiving Cisco IOS XE SD-WAN devices discard copies of the packet and forward one packet to the server. Packet duplication is suitable for edges with multiple access links. Refer to [Packet Duplication documentation](#) for further details

#### TCP Optimization

TCP optimization fine tunes the processing of TCP data traffic to decrease round-trip latency and improve throughput using TCP BBR algorithm. For end-host clients with old TCP/IP stack TCP optimization can provide efficient transport optimization. Refer to [TCP optimization documentation](#) for further details

#### WAAS Integration

Many customers have integrated WAAS with IWAN solution for optimization of the application traffic. Cisco SD-WAN allows the WAAS deployment using AppNav capability to support the WAAS migration from IWAN to Cisco SD-WAN network designs. Refer to the [Cisco SD-WAN WAAS Deployment and Migration Guide](#) for complete details on how WAAS can be migrated to Cisco SD-WAN.

#### Adaptive QoS

Enterprise networks are increasingly using the Internet as a form of WAN transport. Therefore, QoS models need to adapt accordingly. QoS works effectively when deployed in a service-level agreement (SLA) environment, like Multiprotocol Label Switching (MPLS) networks. The available bandwidth on the Internet at a given time can vary. It can often be much lesser than the actual bandwidth that is offered by the service provider. In a non-SLA environment, QoS has limitations because it can't predict the changing bandwidth on the link.

---

With [adaptive QoS](#), the shapers at the edge of the enterprise (WAN interface shaper and per-tunnel shaper) can adapt to the available WAN bandwidth, both Internet and Long-term Evolution (LTE). Thus, adaptive QoS can control differentiated drops at the enterprise edge and reduce the packet drops in the Internet core. When the adaptive QoS capability is not available, shapers that are applied as part of the egress QoS policy are static in value. They are configured based on the service provider bandwidth offering and don't change with time, thus they don't reflect the actual available Internet bandwidth.

### SD-AVC and Custom App Support

[Cisco SD-AVC](#) uses Cisco NBAR2 and other components that operate on devices in the network to provides recognition of network application traffic for visibility, analytics, application-aware routing, and application-based policies, such as QoS and application-based firewall policy. It also provides analytics at the network level.

In addition to the standard protocols provided in a Protocol Pack, you can define protocols, called custom applications, to identify internet traffic, often for uncommon network applications that are of specific interest to their organization.

[Custom applications](#) augment the protocols provided in a Protocol Pack.

### Cloud OnRamp for IaaS and SaaS

Cloud OnRamp for SaaS (formerly called CloudExpress service) addresses these challenges. It enables you to select specific SaaS applications and interfaces, and to let Cisco SD-WAN determine the best performing path for each SaaS application, using the specified interfaces.

Cloud OnRamp for IaaS extends the fabric of the Cisco SD-WAN overlay network into public clouds, allowing branches with Cisco CSR1000V Cloud Services routers to connect directly to public-cloud application providers. By eliminating the need for a physical data center, Cloud OnRamp for IaaS improves the performance of IaaS applications.

With Cisco SD-WAN Cloud OnRamp for CoLocation solution built specifically for colocation facilities, the traffic is routed to the best-permissible path from branches and remote workers to where those applications are hosted. The solution also allows distributed enterprises to have an alternative to enabling direct internet access at the branch and enhance their connectivity to infrastructure-as-a-service (IaaS) and software-as-a-service (SaaS) providers.

See details of Cloud OnRamp solutions of Cisco SD-WAN in the [guide](#).

## 6.2 Security Use Cases

Security is a critical element of today's networking infrastructure. Network administrators and security officers are hard pressed to defend their network against attacks and breaches. As a result of hybrid clouds and remote employee connectivity, the security perimeter around networks is disappearing. The Cisco SD-WAN solution takes a fundamentally different approach to security, see the details in [Security Configuration guide](#)

### AAA

Use the Manage Users screen to add, edit, or delete users and user groups from the vManage NMS. Only a user logged in as the **admin** user or a user who has Manage Users write permission can add, edit, or delete users and user groups from the vManage NMS. Check [Configure User Access and Authentication guide](#). Refer to guide '[How to Check and Verify Single Sign On](#)'

---

## SAML 2.0

vManage supports SSO. SSO allows a user to login to vManage by authenticating against an external Identity Provider (IP). This feature supports SAML 2.0 specification for SSO.

## IPsec / GRE

To securely transfer traffic from the overlay network to a service network, you can configure IPsec tunnels that run the Internet Key Exchange (IKE) protocol. IKE-enabled IPsec tunnels provide authentication and encryption to ensure secure packet transport. You create an IKE-enabled IPsec tunnel by configuring an IPsec interface. IPsec interfaces are logical interfaces, and you configure them just like any other physical interface. You [configure IKE](#) protocol parameters on the IPsec interface, and you can configure other interface properties.

## TLS-proxy

The [SSL/TLS Proxy feature](#) allows you to configure an edge device as a transparent SSL/TLS proxy. Such proxy devices can then decrypt incoming and outgoing TLS traffic to enable their inspection by Unified Threat Defense (UTD) and identify risks that are hidden by end-to-end encryption.

## Firewall

The [Enterprise Firewall with Application Awareness](#) uses a flexible and easily understood zone-based model for traffic inspection, compared to the older interface-based model.

A firewall policy is a type of localized security policy that allows stateful inspection of TCP, UDP, and ICMP data traffic flows. Traffic flows that originate in a given zone are allowed to proceed to another zone based on the policy between the two zones. A zone is a grouping of one or more VPNs. Grouping VPNs into zones allows you to establish security boundaries in your overlay network so that you can control all data traffic that passes between zones.

## IPS/IDS

This feature enables [Intrusion Prevention System \(IPS\) or Intrusion Detection System \(IDS\)](#) for branch offices on Cisco SD-WAN. It is delivered using a virtual image on Cisco IOS XE SD-WAN devices. This feature uses the Snort engine to provide IPS and IDS functionalities.

Snort is an open source network IPS that performs real-time traffic analysis and generates alerts when threats are detected on IP networks. It can also perform protocol analysis, content searching or matching, and detect a variety of attacks and probes (such as buffer overflows).

## URL Filtering

The [URL Filtering feature](#) enables the user to provide controlled access to Internet websites or Intranet sites by configuring the URL-based policies and filters on the device. The user can configure the URL Filtering profiles to manage the web access. The URL Filtering feature is implemented using the security virtual image similar to the IPS feature.

## Umbrella DNS

The [SD-WAN Umbrella](#) Integration feature enables cloud-based security service by inspecting the Domain Name System (DNS) query that is sent to the DNS server through the device. The security administrator configures policies on the Umbrella portal to either allow or deny traffic towards the fully qualified domain name (FQDN). The router acts as a DNS forwarder on the network edge, transparently intercepts DNS traffic, and forwards the DNS queries to the Umbrella cloud.

---

## Cloud Security

Cisco SD-WAN branch routers can support SD-WAN, routing, security, and other LAN access features that can be centrally managed. On high-end devices, all these features can be provided at a higher scale and performance as is often required by large enterprises. However, on lower-end devices not all security features can be enabled simultaneously without degrading performance. These routers can integrate with [Secure Internet Gateways \(SIG\)](#) which do the majority of the processing to secure enterprise traffic. When the SIG is set up, all client traffic, based on routing or policy, is forwarded to the SIG. In addition, the SIG can also protect roaming users, mobile users, and BYOD use-cases.

### Advance DIA guide

One of the many ways to overcome these challenges within an organization is to use Direct Internet Access (DIA) with Cisco Software Defined WAN (SD-WAN). DIA is a component of the Cisco SD-WAN architecture in which certain Internet-bound traffic or public cloud traffic from the branch can be routed directly to the Internet, thereby bypassing the latency of tunneling Internet-bound traffic to a central site. Refer to [DIA CVD](#) for details.

## 6.3 Unified Communications

This feature lets you use feature templates and voice policies to enable [Cisco Unified Communications \(UC\)](#) voice services for supported routers. When Cisco UC voice services are enabled, routers can process calls for various endpoints, including voice ports, POTS dial peers, SIP dial peers, and phone profiles in Cisco Unified SRST mode.

You can configure items for UC voice services from the Feature tab and the Voice Policy page for a supported device.

Configuring UC voice services for Cisco Unified Communications requires that Cisco vManage runs Cisco SD-WAN Release 20.1.1.

This feature is supported on Cisco 4000 Series Integrated Services Routers.



---

# Appendix

## Configurations and Policies

This section provides reference to [GITHUB REPO](#) which will cover the following content, used in the above customer case study.

- Configurations
- vManage templates
- Central policies
- Local policies
- Security policies
- Bootstrap file example