cisco Live!

GO BEYOND

#CiscoLiveAPJC

SD-WAN Design Case Studies

Lessons Learned from Cisco's SD-WAN Design Council

Rishika Goel, Technical Marketing Engineer
@ccie33041

BRKENS-2720

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated
by the speaker until November 15, 2024.

https://ciscolive.ciscoevents.com/
ciscolivebot/#BRKENS-2720

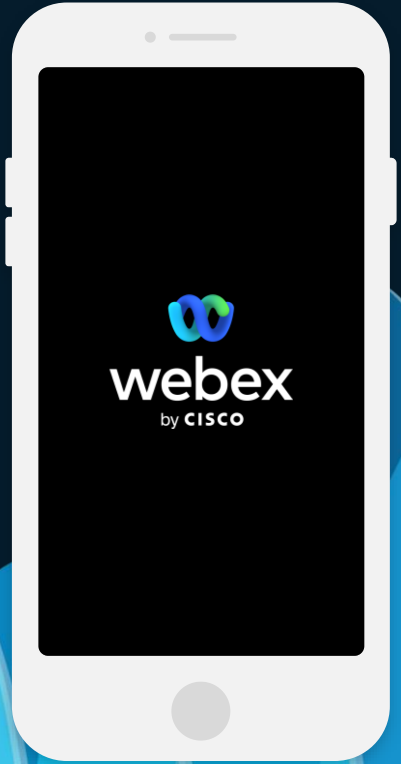

What I do @Cisco

- Started my career in Cisco as a TAC Engineer
- 10 years at Cisco, majority in Jasper and SD WAN BU
- Technical Marketing Engineer, part of Design Council Team
- Participated in various APJC Roadshows/NX Champions

@ccie33041
rigoel@cisco.com
https://www.linkedin.com/in/rishika-goel33041/

Cisco SD-WAN Design Council Introduction

- BU design council includes Cisco members of technical marketing, engineering, product management, and sales.
- Provides guidance for non-standard or undocumented SD-WAN designs

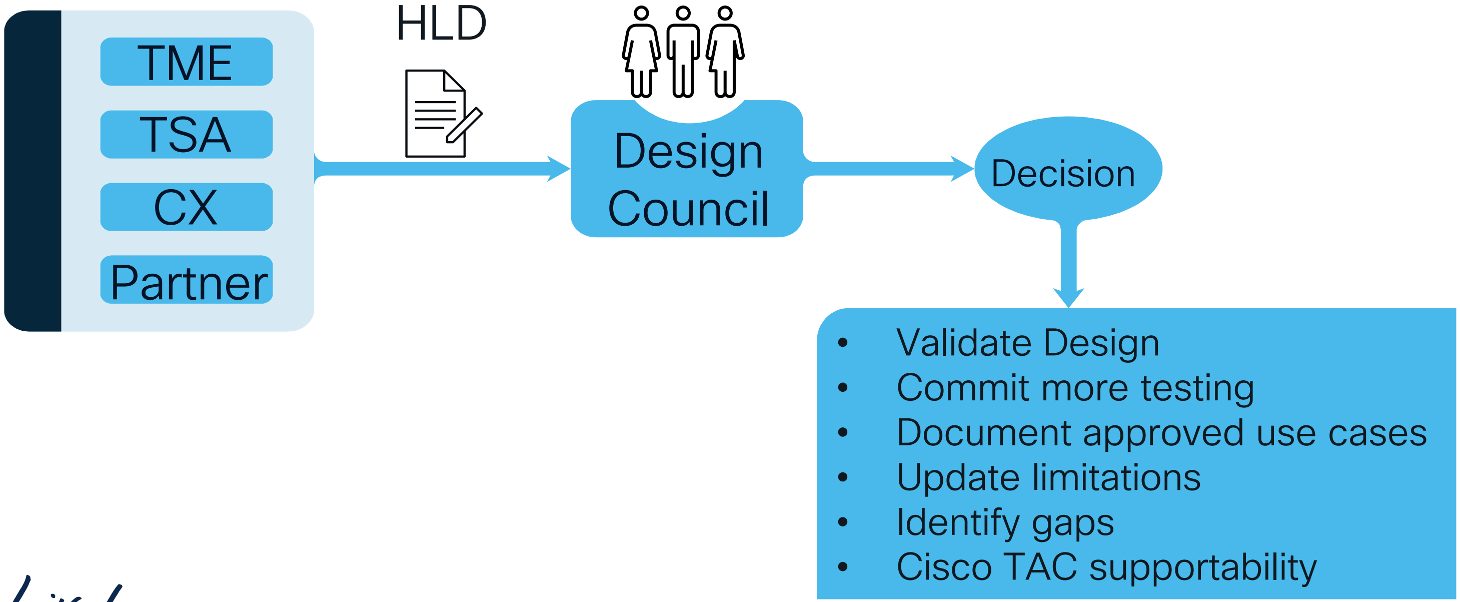

Controller Deployment

Use Case: NMS tools integration with Cloud-Hosted Control Components

Challenge: How to integrate control components with customer internal NMS tools that are RFC1918 privately addressed?

Solution: Extend SD-WAN fabric to the virtual private cloud.

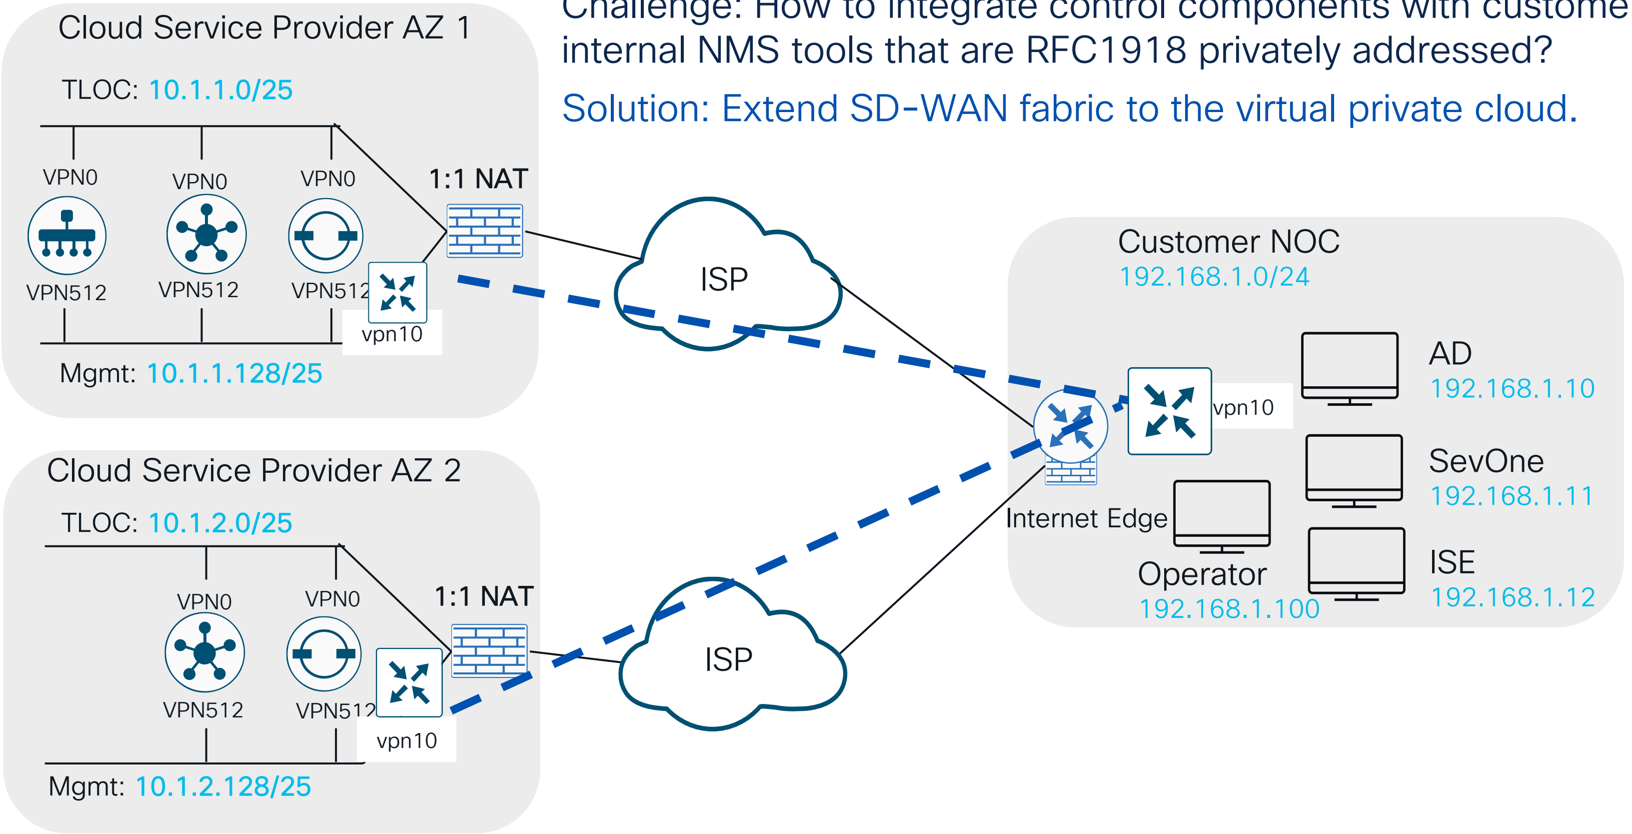

Use Case: Catalyst Controller High Availability
How to to protect against complete outage of a CSP Availability Zone (AZ)

What you want

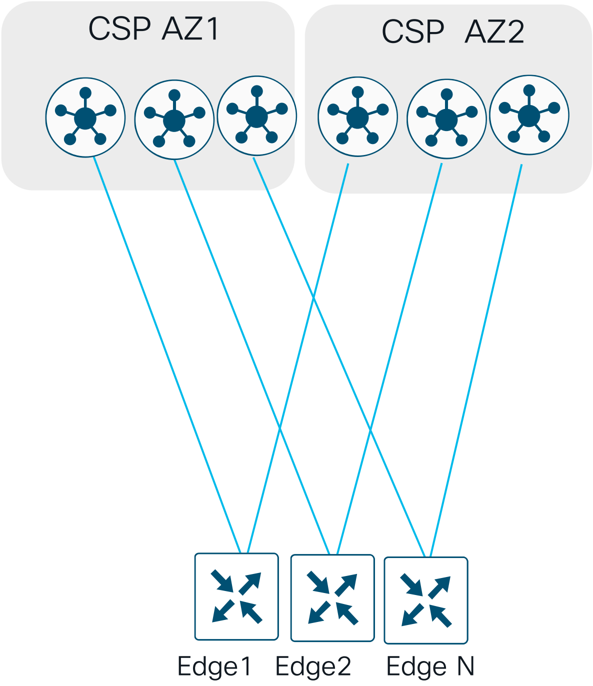


```
system
 max-omp-sessions 2
```

What you might get with default hashing method

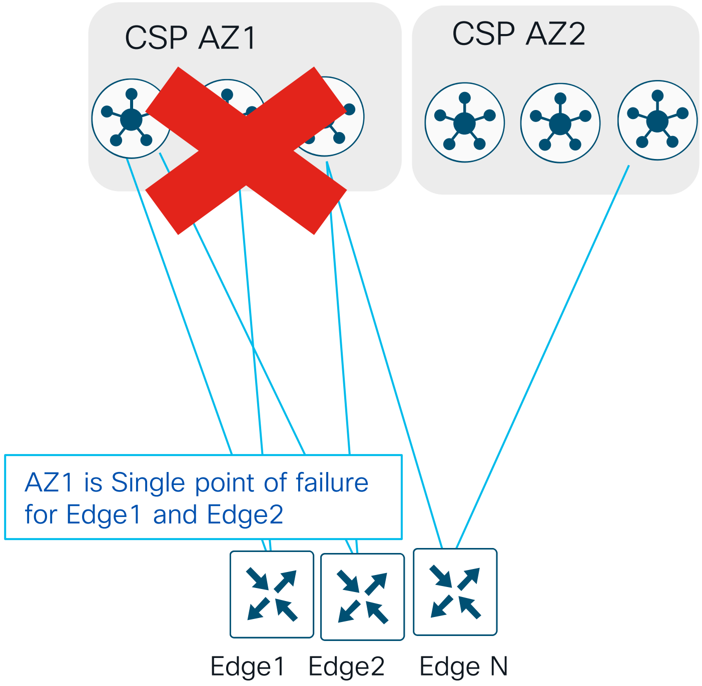


```
system
 max-omp-sessions 2
```


Solution: Deterministic TLS control connections

Controller Groups (CG)

Regionalize controllers with different controller group affinities

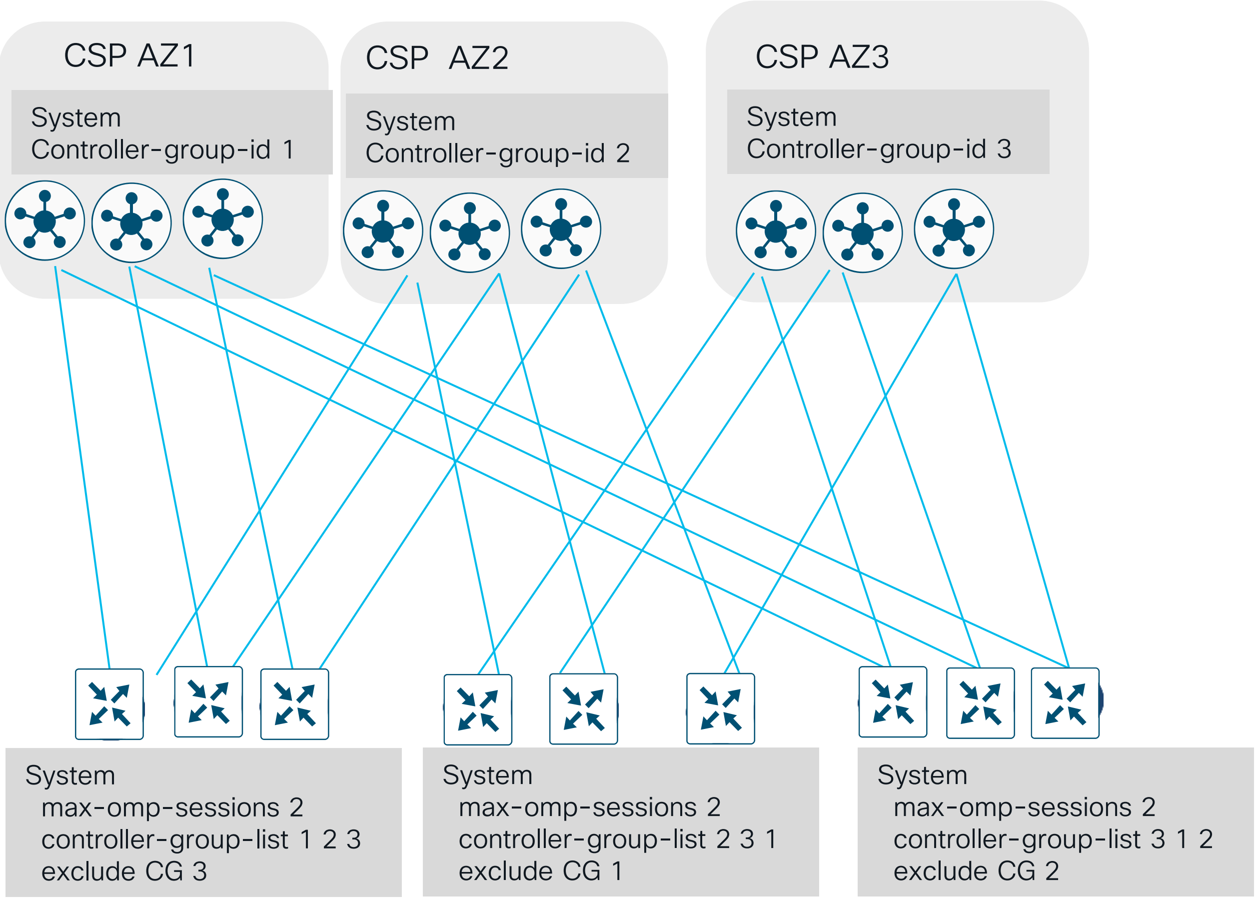

Regionalize WAN Edge with controller-group-lists

Underlay Routing

Use Case: Secure Hub Site design
Requirement: FW protection for controllers, servers and Hub Edge

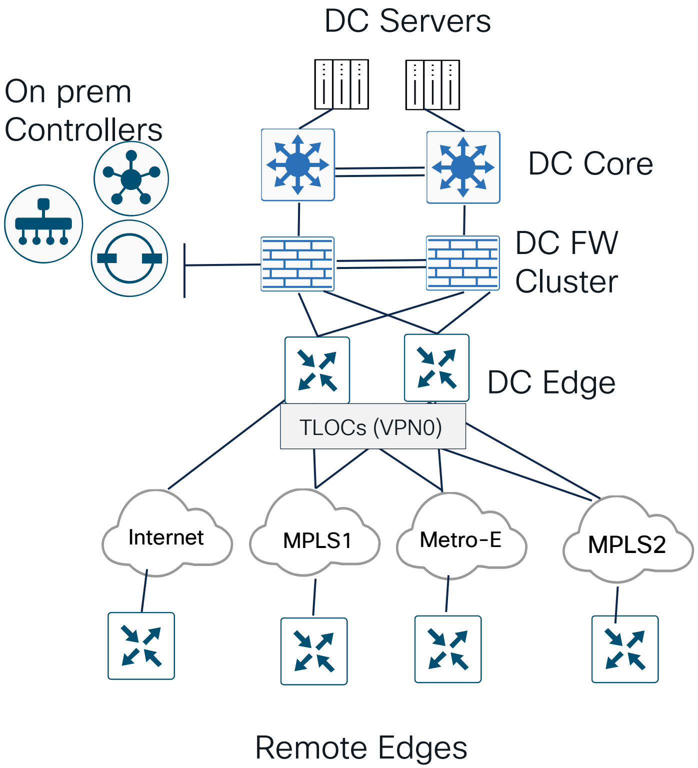

Requirements:

- DC Edge routers must establish TLS connections to controllers over each TLOC (prerequisite for SD-WAN data plane tunnels to form)
- Traffic between WAN Edge and controllers must pass through firewall

Problem1: DC Edge no VPN0 route to controllers via either TLOC

- Remote Edge must establish TLS connections to controllers via Hub

Solution: Collapsed backbone design
L2/L3 aggregation switches for VLAN service chaining

Collapsed Backbone design at Hub Site
All devices and transports connect to L3 switch HA pair

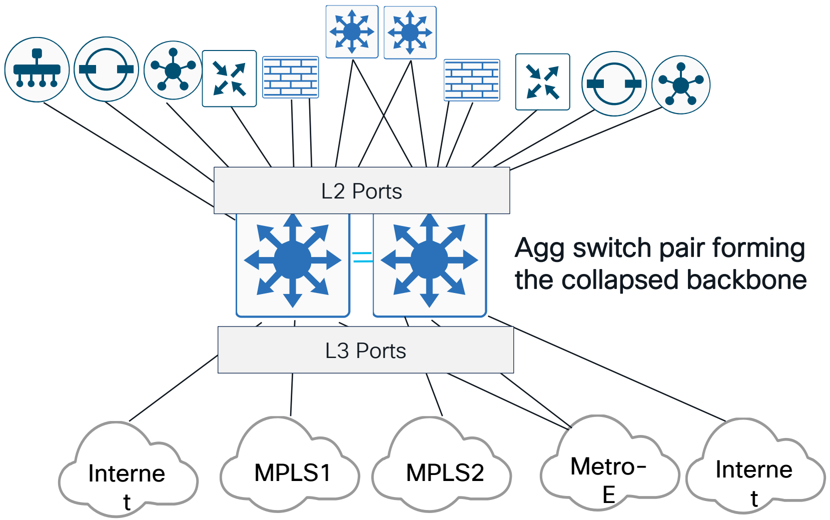

VLAN Service Chain 1
(TLS control connections)
Hub Edge - Firewall - Controllers

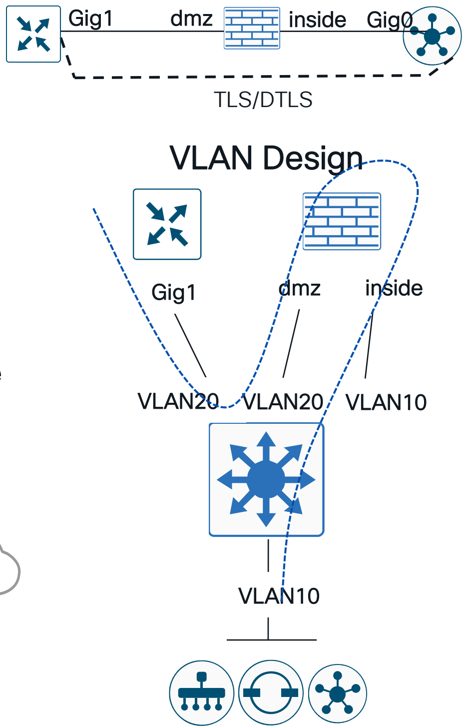

VLAN Service Chain 2
(IPsec over Internet)
Hub Edge - Firewall - Remote Edge

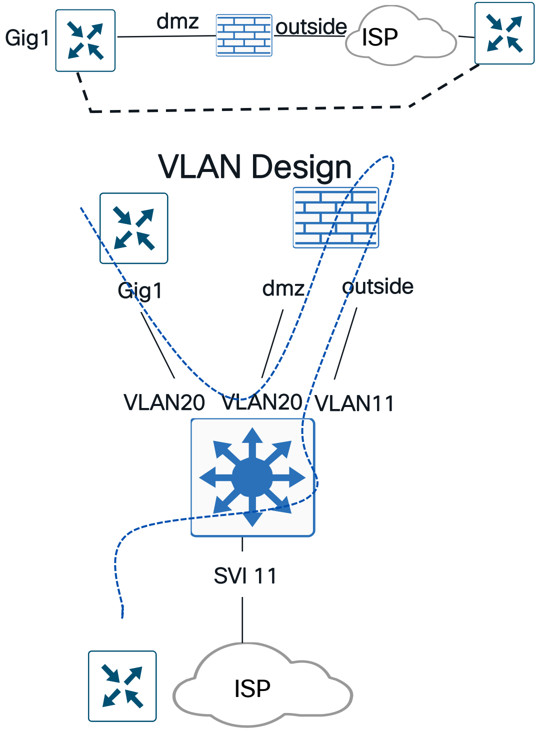

Use Case: Branch router as regional hub for remote branches

Requirement: Hub and spoke via regional Branch

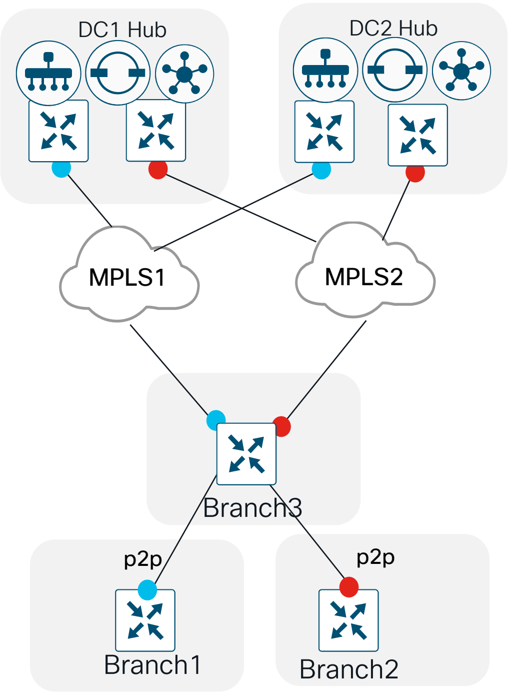

Challenge: Branches1/2 with no MPLS availability
Only point-to-point (p2p) circuits to Branch3

Option1: Add WAN agg switch at regional hub

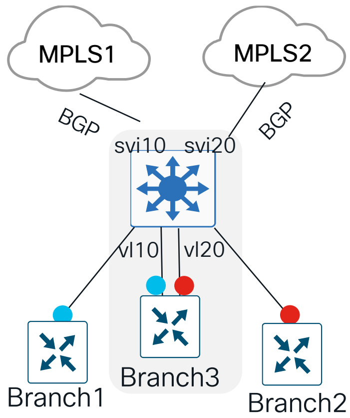

Option2: TLOC Extensions to 'stitch' p2p circuits to MPLS

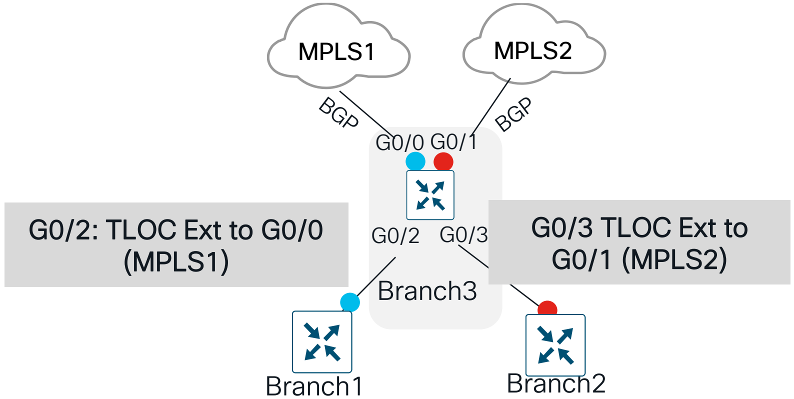

Use Case: On-prem Controllers hosted in a Colo Facility

Design Challenges

Branch Edge cannot form control connections through Hub WAN Edge

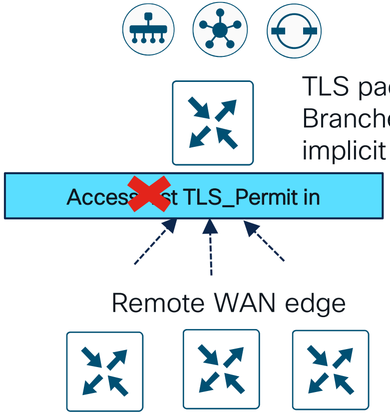


```
Edge2#show platform packet-trace summary
Pkt    Input                          Output                         State   Reason
FWD
5      internal0/0/recycle:0          Gi2                            FWD
6      internal0/0/recycle:0          Gi2                            FWD
7      internal0/0/recycle:0          Gi2                            FWD
8      internal0/0/recycle:0          Gi2                            FWD
9      Gi2                            Gi6                            FWD
10     Gi0                            Gi0                           DROP    479 (SdwanImplicitAclDrop)
```

```
Edge2#show platform packet-trace packet
10 decode
Packet: 10          CBUG ID: 10
Summary
  Input     : GigabitEthernet0
  Output    : GigabitEthernet0
  State     : DROP 479
(SdwanImplicitAclDrop)
  Timestamp
    Start   : 19430074503041 ns
(10/22/2024 13:02:14.185389 UTC)
    Stop    : 19430085645000 ns
(10/22/2024 13:02:14.196531 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet0
    Output      : <unknown>
    Source      : 172.16.0.201
    Destination : 172.16.2.3
    Protocol    : 17 (UDP)
      SrcPort   : 12346
      DstPort   : 12366
```

Create Explicit ACL which permit TLS/DTLS transit through Hub WAN Edge TLOC and apply to tunnel interface

```
policy
access-list Permit_control_traffic
sequence 1
match
source-port 12346 12366 12386 12406 12426
12446 12546 12646 12746 12846 12946 13046
 protocol 17
!
action accept
!
sequence 2
match
source-port 23456 23556 23656 23756 23856
23956 24056 24156
protocol 6
!
action accept
!
default-action drop
sdwan
interface GigabitEthernet0
tunnel-interface
encapsulation ipsec weight 1
color biz-internet
access-list Permit_control_traffic in
!
```

Use Case: Daisy chaining IoT SD-WAN routers
Remote locations with limited transport availability

Physical Topology (underlay)

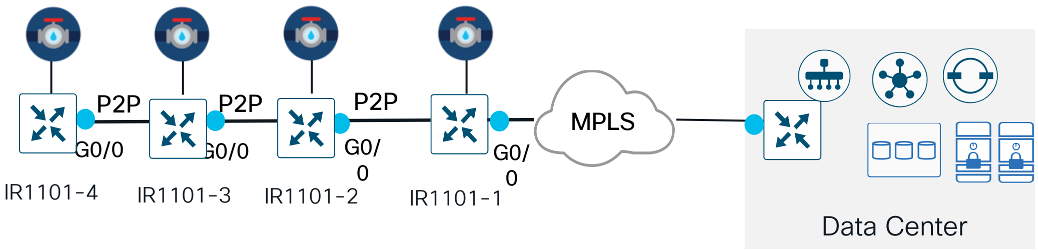

Desired logical topology (overlay)

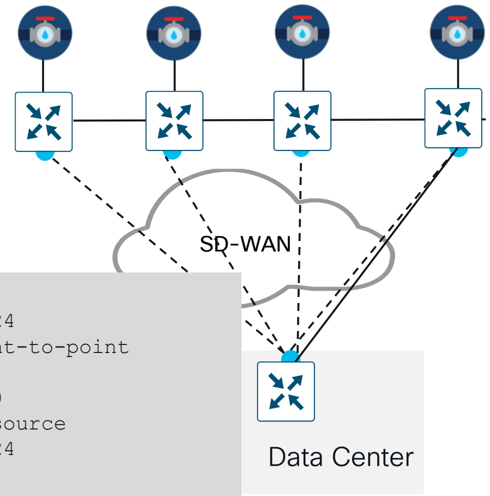

Challenge: Transit routing of control and data plane traffic blocked by implicit ACL (Gig0/0) present on all TLOC interfaces

Solution

- Use Loopback interface as TLOC sources (unbound mode)
- Enable underlay routing protocol (OSPF and BGP)
- Allow OSPF service under tunnel interface

```
interface gig0/0
ip address 1.1.1.1/24
ip ospf network point-to-point
ip ospf 1 area 0
interface loopback10
description tunnel source
ip address 2.2.2.2/24
ip ospf 1 area 0
!
sdwan
interface GigabitEthernet0/0/0
interface loopback10
tunnel interface
color mpls
allow-service ospf
!
router ospf 1
router-id 1.1.1.1
```

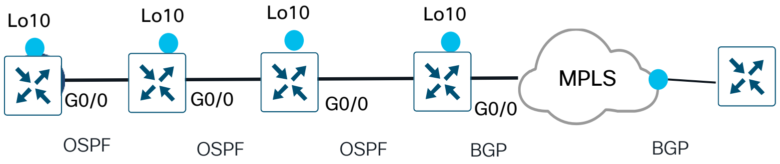

Starlink Satellite

Use Case: Starlink Satellite as SD-WAN Transport

Starlink is a constellation of 6000+ Low Earth Orbit (LEO) satellites offering Internet access via exchange points around the world (plans for 12000)

The Starlink CPE includes a high-performance terminal (dish), PoE injector, ethernet cable, and wifi router that can be replaced by a Cisco or 3rd party SD-WAN router

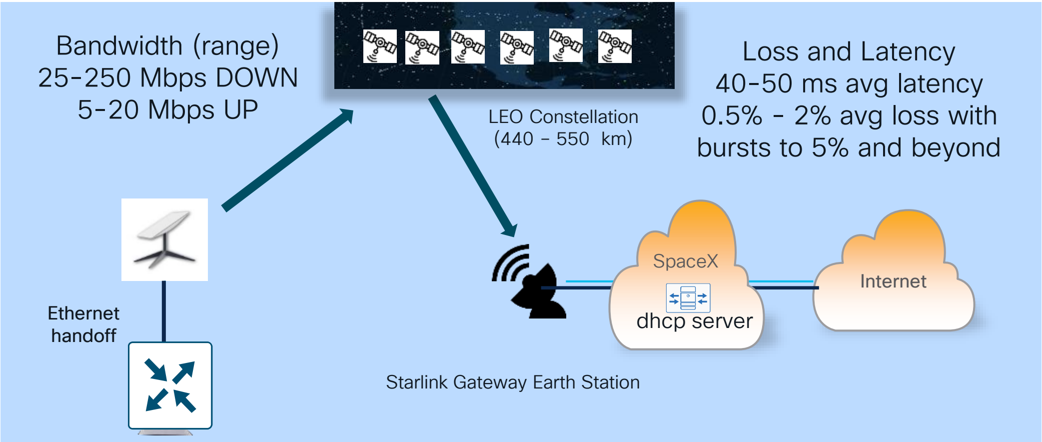

Challenges

- Higher degree of latency than terrestrial transports may impact quality of experience for applications sensitive to delay
- Bandwidth may fluctuate depending on dish placement, radio frequency interference, weather and ground station capacity.
- Packet loss impacts TCP app performance, especially with OS with Reno and Cubic congestion control algorithms.

Recommendations

- Deploy AAR to measure performance and steer latency-sensitive traffic away from Starlink path
- Configure per-tunnel and adaptive QoS (dynamic shaping)
- Enable App-QoE features to mitigate loss and latency

SD-WAN over Starlink configurations

Interface configuration

```
interface GigabitEthernet0/0/0
 description Ethernet to Starlink LEO satellite
 ip dhcp client default-router distance 1
 ip dhcp client route distance 1
 ip address dhcp client-id GigabitEthernet0/0/0
 load-interval 30
 negotiation auto
```

Required for deployments with Starlink and other TLOCs
(ip dhcp client route distance 1 required for Cat8KV)

Tunnel optimizations and adaptive QoS

```
interface GigabitEthernet0/0/0
 tunnel-interface
 encapsulation ipsec
 color custom1
 no last-resort-circuit
 vmanage-connection-preference 1
!
 qos-adaptive
 period 5
 downstream 200000
 downstream range 80000 250000
 upstream 12000
 upstream range 2000 20000
```

Application Aware Routing (AAR)

Prefer MPLS for real time traffic, Starlink for default

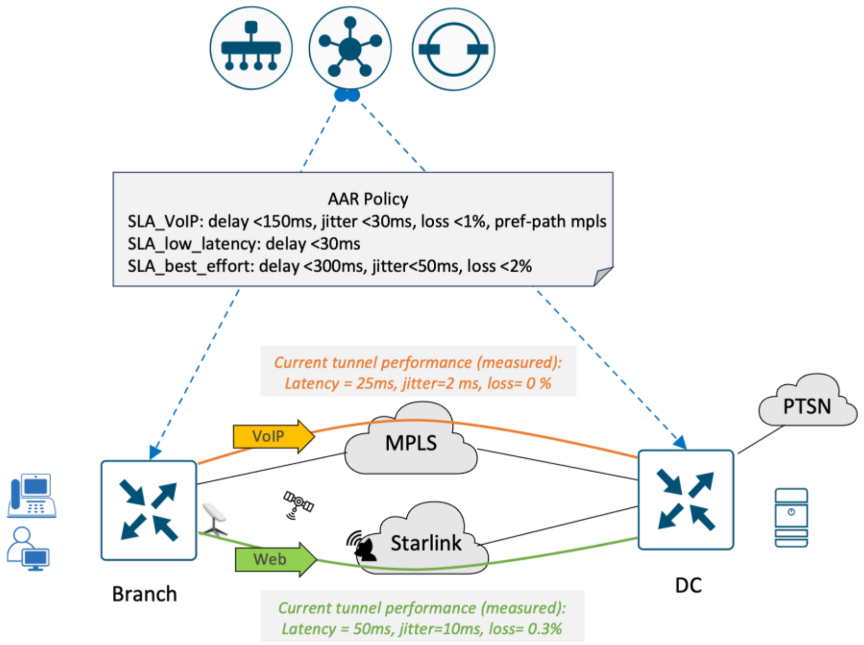


```
sla-class SLA_ClassList-Voice-And-Video
   jitter 30
   latency 45
   loss 2
!
sla-class SLA_ClassList-Default
   jitter 100
   latency 300
   loss 5
  !
app-route-policy _CS1_VPN1_AAR_Default-AAR-Policy
  vpn-list CS1_VPN1
    sequence 1
      match
       app-list APP-List-voip-telephony
       source-ip 0.0.0.0/0
      !
      action
       sla-class SLA_ClassList-Voice-And-Video
        preferred-color mpls
        backup-sla-preferred-color mpls
sequence 161
     match
       dscp 0
       source-ip 0.0.0.0/0
     action
       sla-class Default  preferred-color custom1
```

App-QoE: TCP Optimization

BBR2 increases iPerf/TCP throughput over Starlink by 300%

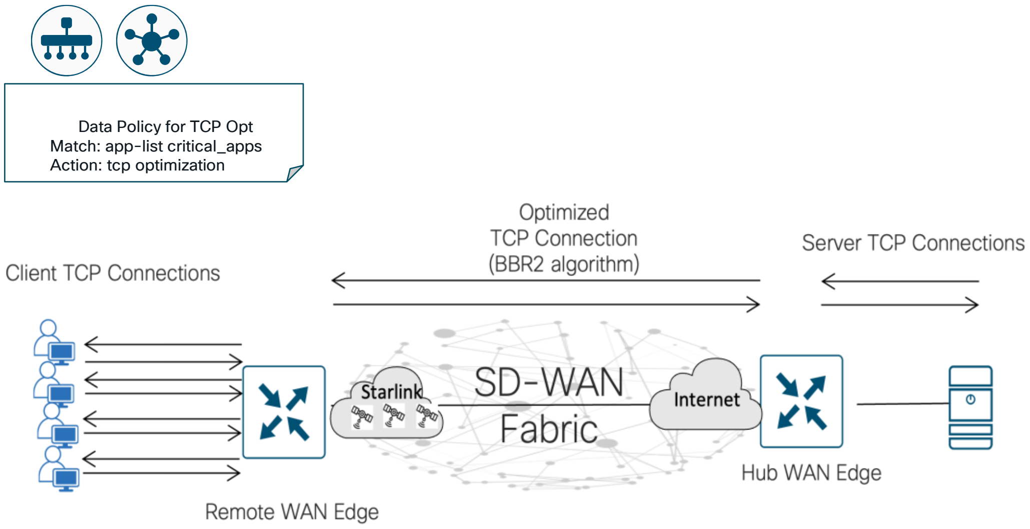


```
Viptela-policy:policy
 data-policy _CS1_VPN1_opt4
  vpn-list CS1_VPN1
    sequence 1
     match
      source-data-prefix-list ubuntu11
      destination-data-prefix-list ubuntu6
     !
     action accept
      tcp-optimization
      service-node-group SNG-APPQOE
      count u11-6_-1323095233

  default-action accept
```

WAN Edge routers proxy TCP sessions between clients and servers and leverage the BBR2 algorithm to improve performance over high latency and lossy links.

App-QoE: Forward Error Correction (FEC)

Reconstructing packet drops over Starlink tunnels

FEC is a mechanism to recover lost packets on a link by sending extra "parity" packet built by calculating XOR value for every group of 4 packets.

If 1 packet is lost in the group of 4, and the parity packet is intact, the receiving WAN edge can reconstruct with another XOR

Data Policy for FEC
Match: app-list critical_apps
Action forward-error-correction

DC router
reconstructs packet 3

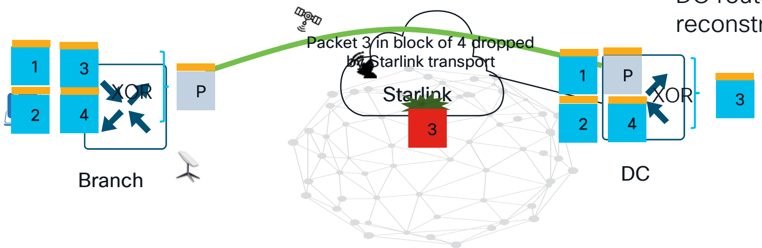


```
viptela-policy:policy
 data-policy _CS1_VPN1_error_correction
  vpn-list CS1_VPN1
    sequence 1
     match
      source-data-prefix-list ubuntu11
      destination-data-prefix-list ubuntu6
     !
     action accept
      count u11-6fec_87880661
      loss-protect fec-adaptive
      loss-protection forward-error-correction adaptive
```

Up to 60% throughput increase over Starlink with FEC
(IPerf3 TCP throughput testing)

App-QoE: Packet Duplication

Proactive loss mitigation by replicating critical flows over redundant tunnels

WAN edge replicates all packets for critical application flows
and forwards copy over redundant SD-WAN tunnel

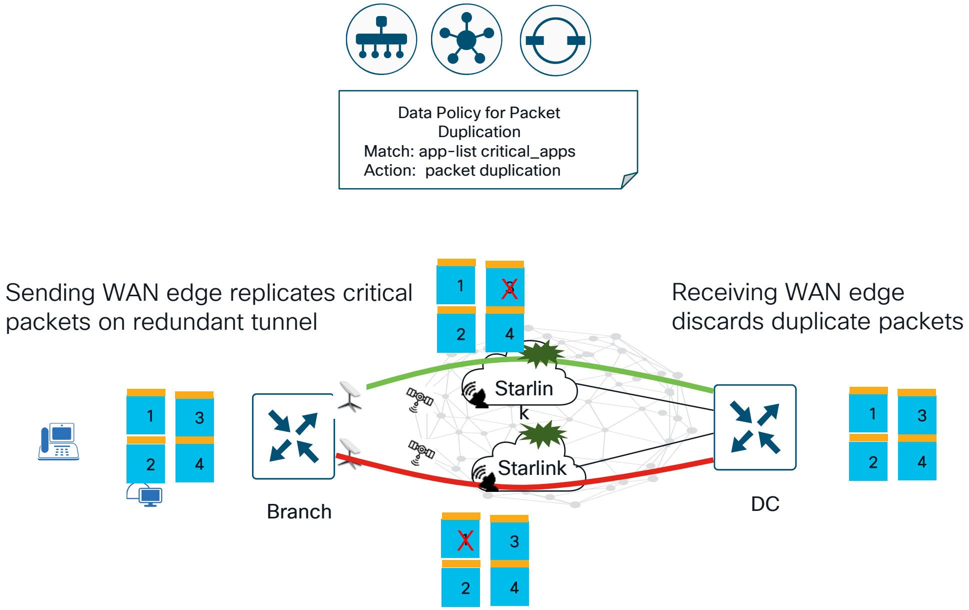


```
viptela-policy:policy
 data-policy _CS1_VPN1_Packetdup
  vpn-list CS1_VPN1
    sequence 1
     match
      source-data-prefix-list ubuntu11
      destination-data-prefix-list ubuntu6
     !
     action accept
      count u11-6fec_1256232839
      loss-protect pkt-dup
      loss-protection packet-duplication
```

Up to 20% throughput increase over Starlink when Packet duplication
(IPerf3 TCP throughput testing)

App-QoE: Data Redundancy Elimination (DRE) with LZ compression

Traffic Reduction over WAN

DRE reduces the amount of WAN traffic by (byte) caching previously seen data patterns

LZ reduces the amount of WAN traffic by doing zip like compression on top of DRE

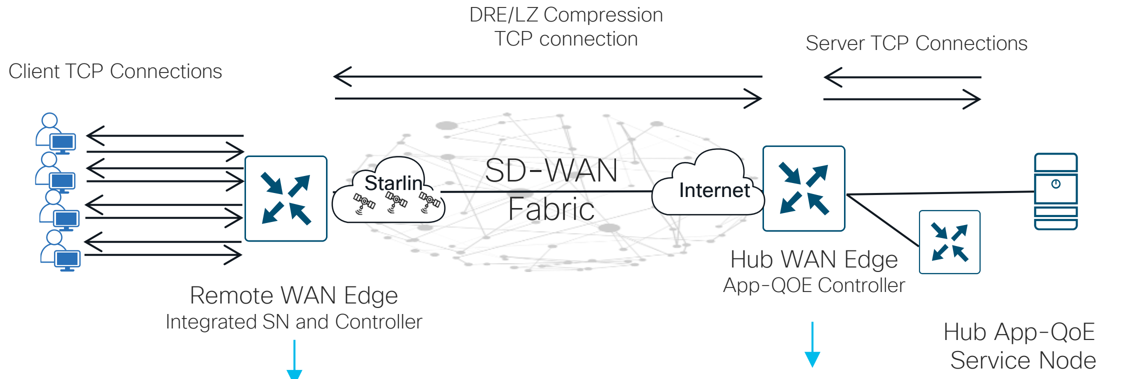

- C8300-1N1S-6T
- C8300-1N1S-4T2X
- C8300-2N2S-6T
- C8300-2N2S-4T2X
- C8200-1N-4T
- C8000V

- C8500L-8S4X
- C8000V

- C8500L-8S4X
- C8000V

```
viptela-policy:policy
 data-policy _CS1_VPN1_opt3
  vpn-list CS1_VPN1
    sequence 1
     match
      source-data-prefix-list ubuntu11
      destination-data-prefix-list ubuntu6
     !
      action accept
      tcp-optimization
      dre-optimization
      service-node-group SNG-APPQOE
      count u11-6_185636872
```

Up to 90% reduction in WAN traffic over Starlink when DRE/LZ is enabled along with TCP opt

(Tested using multiple iterations of FTP downloads of a file hosted on a server at Hub)

Review of App-QoE features with Starlink as SD-WAN transport

Option	App-QoE feature	Throughput increase	WAN Traffic Reduction
A	TCP Optimization	300%	
B	Forward Error Correction	60%	
C	Packet Duplication	20%	
D	Data Redundancy Elimination (DRE) with LZ compression		90%

Recommend TCP Optimization and DRE/LZ for better application experience

NAT

Use Case: Multiple NAT requirements for a site

Requirements

1. Corporate VPN1 **User-1** should **NAT to Pool10** and be load shared across ISP1 and ISP2
2. Corporate VPN1 **User-2** accessing Webex services should **NAT to Interface Gig1** IP address and be pinned to ISP1.
3. Recently integrated Partner VPN2 wants to communicate with other SD-WAN branches.
4. **Web-server VPN3** traffic should be **bypassed from being NATted** (NOT TRANSLATED)
5. All DIA traffic inspected by NGFW

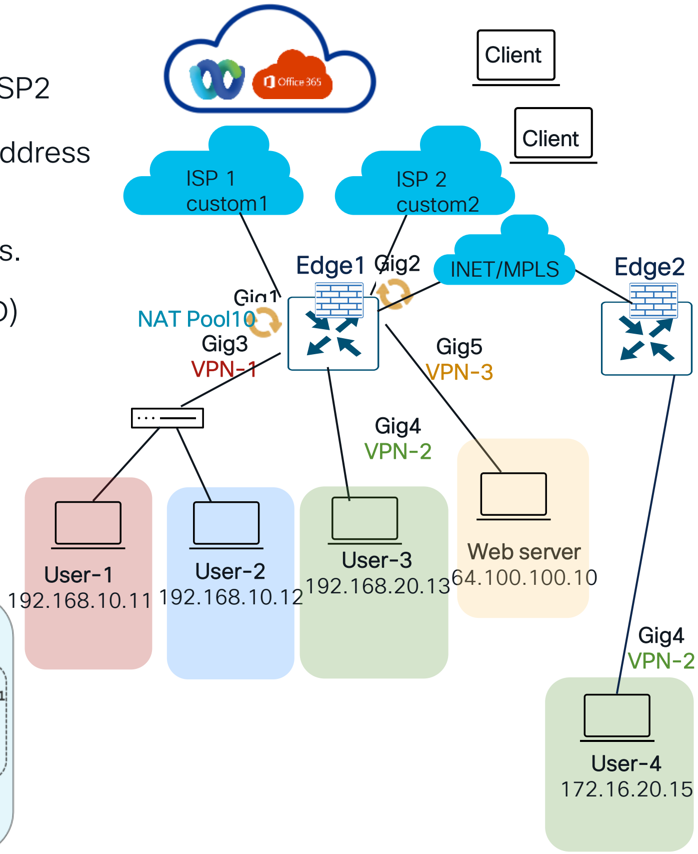

Problem: Single NAT method supported in data policy prior to 17.14

Behavior prior to 17.14

Data-Policy configuration

```
sequence 31
   match
    source-ip 0.0.0.0/0
   action accept
    nat use-vpn 0
    no nat fallback
```

No control to select specific pool or interface

NAT configuration

```
ip nat pool natpool-GigabitEthernet1-0 181.1.1.1 181.1.1.4 prefix-length 27
ip nat inside source list global-list pool natpool-GigabitEthernet1-0 overload egress-interface GigabitEthernet1
                                   (or)
ip nat inside source list nat-dia-vpn-hop-access-list interface GigabitEthernet1 overload
                                   (or)
ip nat inside source list global-list interface Loopback1 overload egress-interface GigabitEthernet1
```

Solution: Multiple DIA NAT methods (17.14)

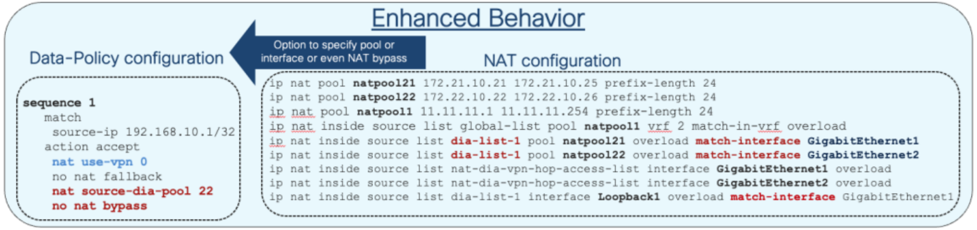

VPN1 (Corporate VPN) Requirements

- User-1 DIA NAT to Pool 10 and load shared across ISP1 and ISP2
- User-2 DIA NAT to Interface Gig1 IP address and forward to ISP1

```
viptela-policy:policy
 data-policy _VPN1_Multiple-DIA-NAT-method
 vpn-list VPN1
 sequence 1
 match
 source-ip 192.168.10.11
!
 action accept
 nat use-vpn 0
 nat source-dia-pool 10
!
sequence 2
 match
 source-ip 192.168.10.12
!
 action accept
 nat use-vpn 0
 set
 local-tloc-list
 color custom1
```

VPN2 (Partner VPN) Requirements

User-3 Service Side NAT to Pool 1 address.

```
from-vsmart data-policy _nat-dia-vpn-1
_SSNAT-Network-Static-8
 direction from-service
 vpn-list nat-dia-vpn-2
  sequence 1
   match
    destination-ip 172.16.20.15
   action accept
nat pool 1
  default-action accept
from-vsmart lists vpn-list nat-dia-vpn-2
 vpn 2
```

VPN3 (Web Server) Requirements

Web Server with Public IP address should bypass NAT outbound to Internet

```
vpn-list VPN3
    sequence 1
     match
      source-ip 64.100.100.10
!
     action accept
      nat use-vpn 0
      nat bypass
```


Problem: Incoming traffic to web server dropped

No route to Web Server in Global VPN0

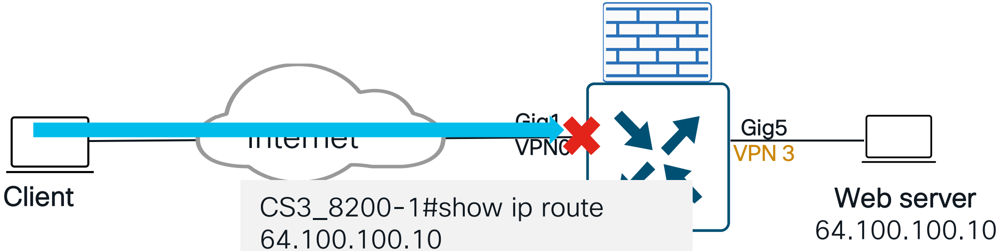

CS3_8200-1#show ip route 64.100.100.10
Routing entry for 0.0.0.0/0
  Known via "static", distance 1, metric 0
  Routing Descriptor Blocks:
  * **<ISP Router IP address>**

```
vpn-list VPN3
 sequence 1
  match
  source-ip 64.100.100.10
!
 action accept
 nat use-vpn 0
 nat bypass
```

Solution: Route leak between Transport and Service VPN

```
vrf definition 3
 description Web Server VRF
 rd 1:10
 address-family ipv4
 route-replicate from vrf global unicast
static
 route-target export 100:10
 route-target import 100:10
 exit-address-family
 !
global-address-family ipv4
 route-replicate from vrf 3 unicast connected
```

CS3_8200-1#show ip route 64.100.100.10
Routing entry for 64.100.100.10/32
  Known via "connected", distance 0, metric 0 (connected)
  Routing Descriptor Blocks:
  * **directly connected, via GigabitEthernet5**
    Route metric is 0, traffic share count is 1

Branch Security

NGFW Redundancy
DIA and SD-WAN use cases

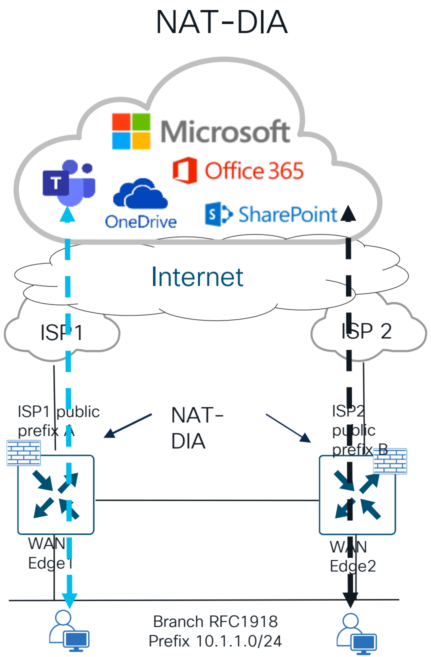

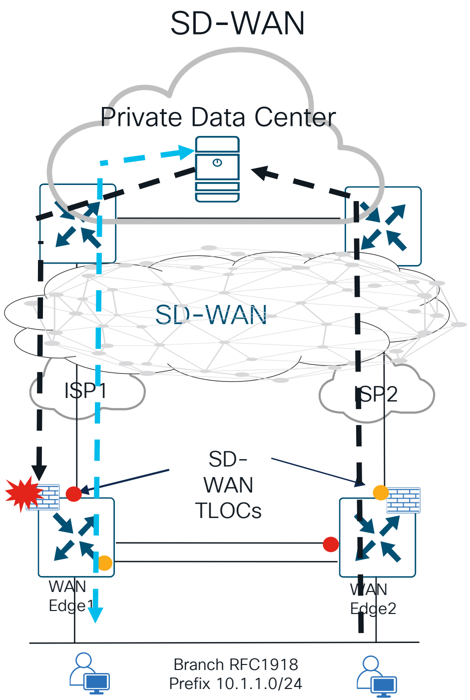

- Dual Edge with NGFW for DIA traffic inspection
- NAT enforces traffic symmetry and guarantees bidirectional traffic through the NGFW

- Dual Edge with NGFW for SD-WAN traffic inspection
- No guarantee for traffic symmetry by default
- Ingress traffic may return on different Edge/NGFW resulting in traffic drops (no NGFW state)

NGFW Redundancy
Solutions for SD-WAN

Primary and backup WAN Edge/NGFW

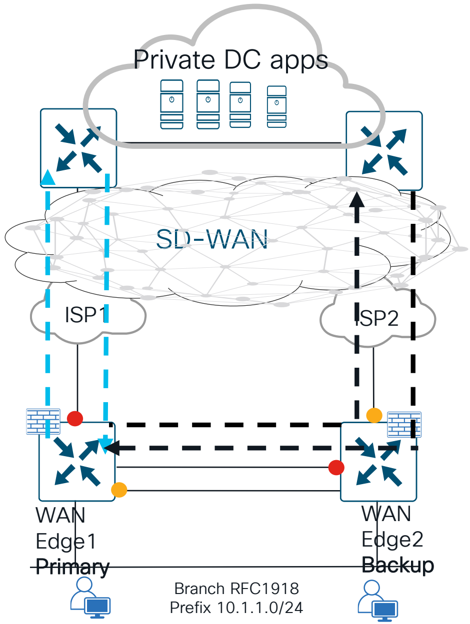


```
WAN Edge Primary
interface gig0
description ISP1 interface
 tunnel-interface
 color custom1 (red)
!
interface gig1
description ISP2 via TLOC EXT
 tunnel-interface
 color custom2 (yellow) encaps
ipsec
!
interface gig2
 vrf forwarding 1
 ip address 10.1.1.2
vrrp 1 address-family ipv4
  priority 105
  address 10.10.1.1 primary
  tloc-change increase-pref 333
```

```
WAN Edge2 Backup
interface gig0
description ISP2 interface
 tunnel-interface
 color custom2(yellow)
 encaps ipsec
!
interface gig1
description ISP1 via TLOC EXT
 tunnel-interface
 color custom1 (red)
!
interface gig2
 vrf forwarding 1
 ip address 10.1.1.3
 vrrp 1 address-family ipv4
 address 10.10.1.1 primary
```

- Enforce symmetry by designating active and backup forwarding nodes
- L2 LAN: VRRP priority set higher on active forwarder
- VRRP tloc preference ensures traffic from DC side is attracted to WAN Edge1
- L3 LAN: Routing metrics tuned to prefer active forwarder during redistribution

Critical Branch High Availability

Branch On Prem Security redundancy (20.15)

17.15/20.15: Redundancy Groups (RG)

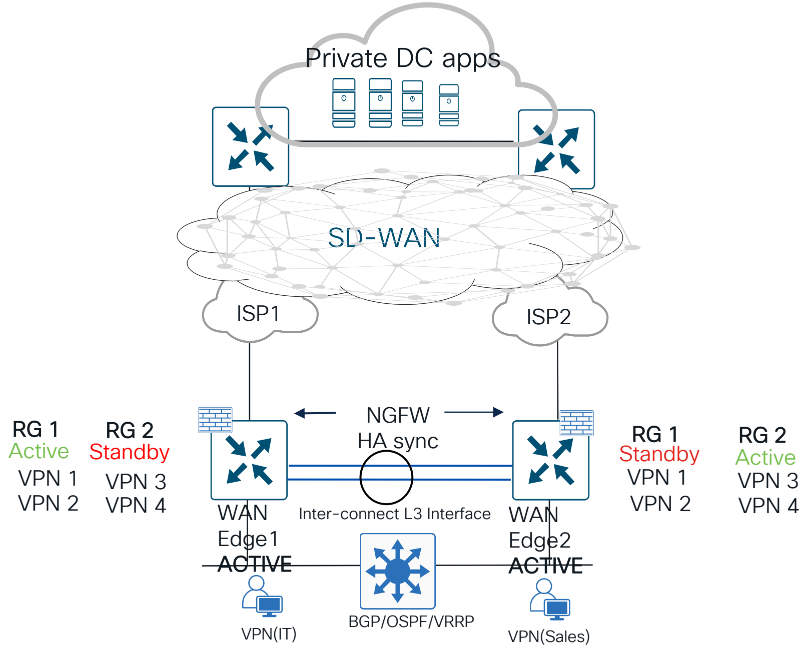

- Traffic symmetry not required
- Active/Active WAN Edge for selected VPN's and NGFW pair
- VPN Homing
- Inter-connect connects two routers in HA mode
  - ✓ Session Sync
  - ✓ Data Traffic (Peer Redirect)
- FW State and NAT Connections are in sync between devices in HA group

Critical Branch High Availability

Branch Cloud security redundancy (20.15)

Option2: Cloud Security (SIG/SSE)
17.15/20.15: Secure Access Context Sharing

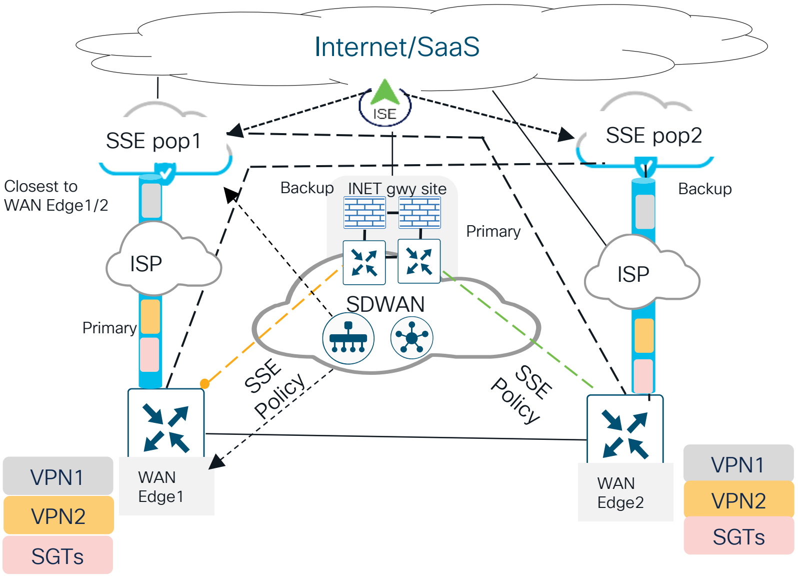

- SSE inspecting the traffic for thin/hybrid sites
- Primary and Backup SSE tunnels for HA
- Context sharing of VPN/SGT identity to enforce identity-based policy at SSE

Scalability and High Availability

Use Case: High Scale Hub and Spoke Deployment
Hub WAN Edge Horizontal Scalability

Requirement 1: 18,000 tunnels at hub

- Catalyst 8500 is highest performing SD-WAN platform supporting up to 10,000 IPSec tunnels
- **How to spread IPSec Tunnels horizontally across 3 different Hub routers in each DC?**

Requirement 2: 300Gbps aggregate bw at Hub

- Aggregate throughput required for all sites exceeds the current throughput capacity of a single Catalyst 8500
- **How to distribute traffic across 3 different hub routers in a horizontal fashion?**

Requirement 3: 600 VRFs at hub site

- Remote sites are different companies that require separation into different VRFs
- 600 VRFs are required to accommodate all partners - Catalyst 8500 supports 300 VRFs
- **How to distribute 600 VRFs across 3 Hub routers?**

Solution: Tunnel Groups for deterministic tunnel placement on hub

Result after Tunnel Group

- VRFs distributed across Hubs
- Tunnels and Traffic distributed across Hubs

```
sdwan
interface GigabitEthernet0
 tunnel-interface
 encapsulation ipsec
 color mpls
 Group 1
interface GigabitEthernet1
 tunnel-interface
 encapsulation ipsec
 color biz-internet
 Group 1
```

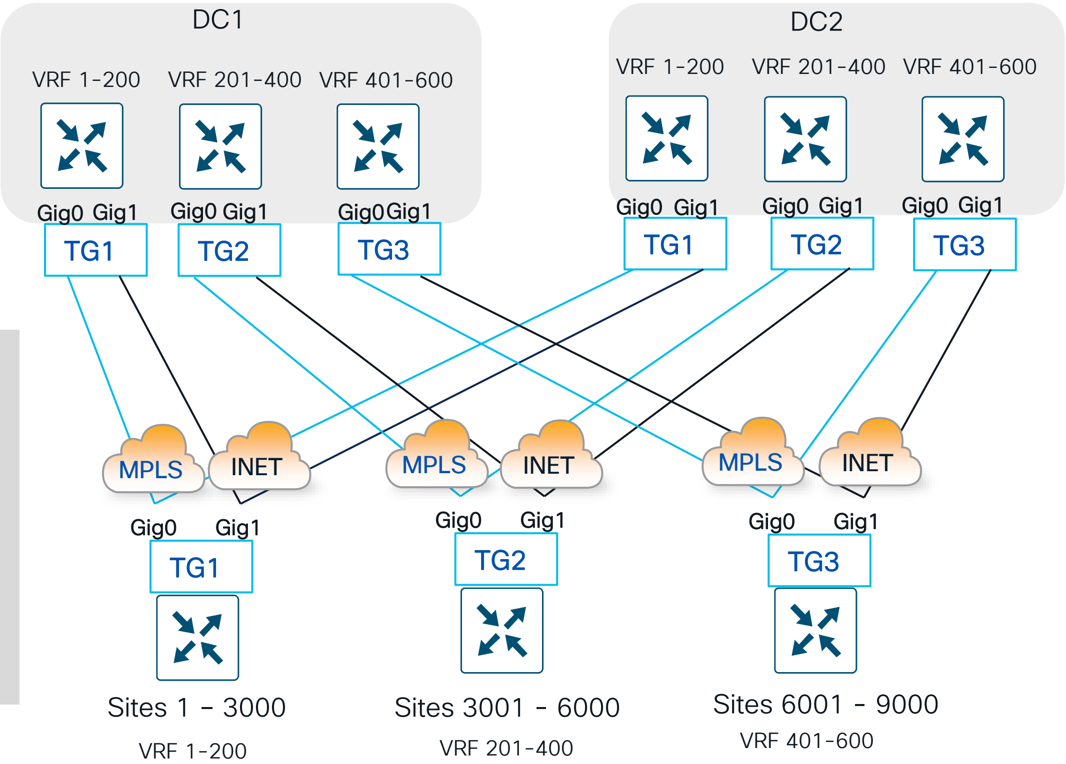

SD-WAN HA Design for Critical Site
Option1: Active/Active Tunnels with TLOC Extensions

Two TLOCs on each router (biz-internet **and** public-internet)

VLANs 100, 200 defined for TLOC extensions on cross-link

Results in 8 Tunnels per site (4 WAN, 4 via TLOC extension)

- 8 OMP paths for DC routes (1 per TLOC)
- ECMP load-balancing across all 8 tunnel paths
- Supports Application Aware Routing (AAR) designs where local access to all TLOCs is a required prerequisite

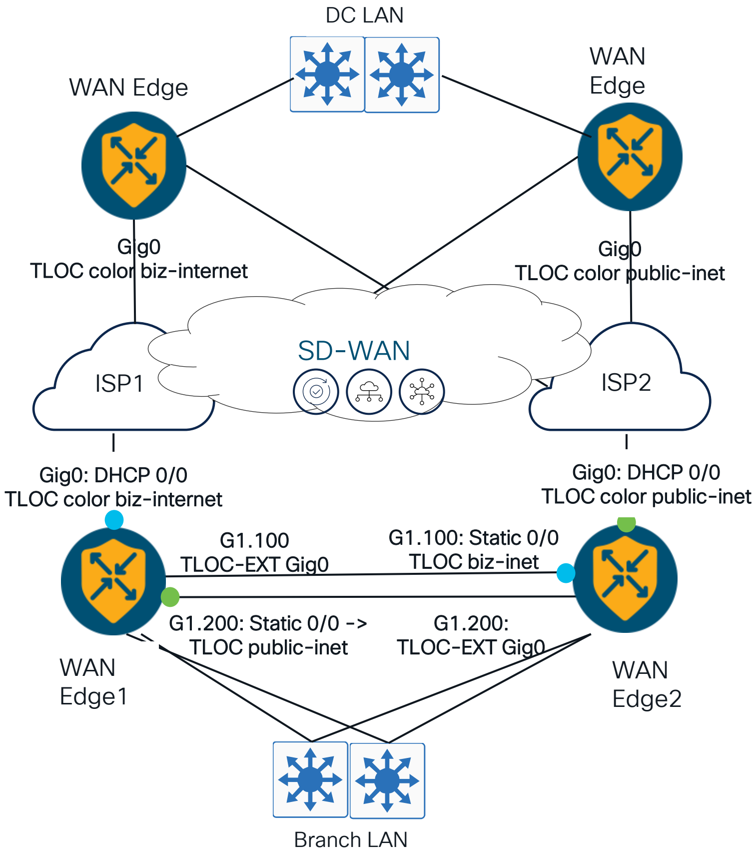

SD-WAN HA Design for Critical Site

Option2: Active/Backup paths with no TLOC Extensions

One TLOC on each router (biz-internet **or** public-internet)

Dedicated VLAN per Service VPN with OSPF enabled
- VLAN100: Service VPN1
- VLAN200: Service VPN2

OMP to OSPF (Service VPN) redistribution for backup path

Results in 4 Tunnels per site (Each router with 2 across WAN)
- LAN switch controls load-balancing of traffic to both WAN edge
- Does not support Application Aware Routing (AAR) designs where each router must have both TLOCs

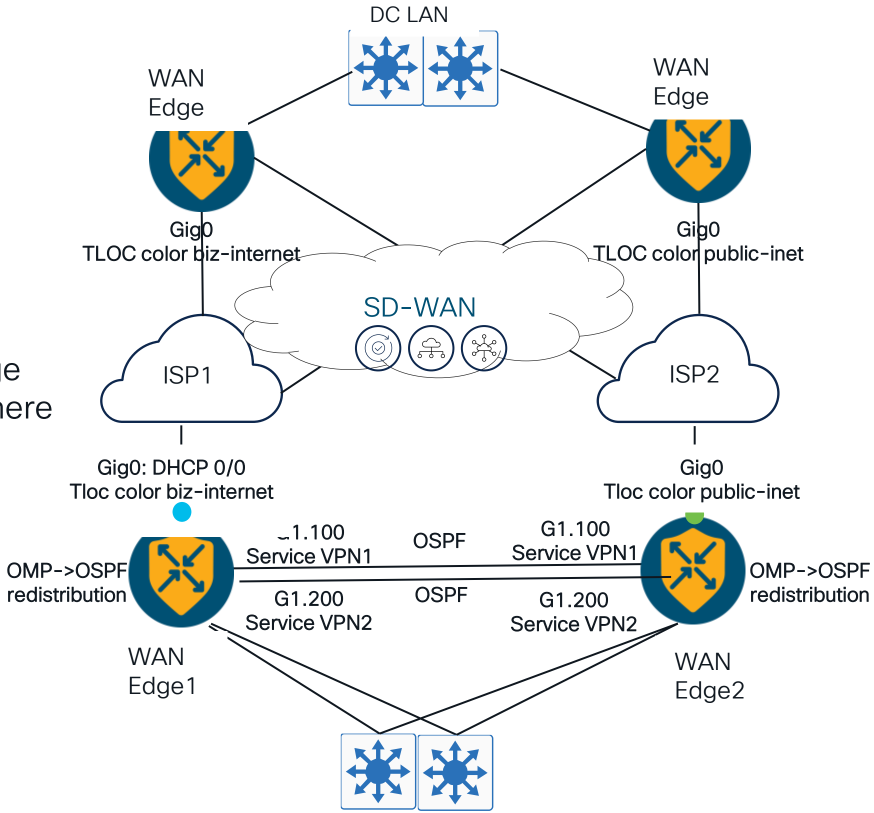

Dual Router Design for Internet High Availability

Option1: Dual DIA using local ISP + TLOC Extension path

Each cEdge with two interfaces for NAT-DIA Internet breakout

- Gig0 local ISP circuit, Gig1.x00 cross-link to peer WAN edge
- Each router with dual default routes
    - DHCP learned 0/0 from local ISP
    - Static 0/0 to peer WAN edge via cross-link

Traffic redirection based on routing or centralized policy

- Guest VPN: Default route in VPN 99 directs all traffic to DIA path

Problem: Some Internet-bound flows looping on cross-link due to dual default routes

- Internet-bound flows hashed to static default route to peer WAN edge may loop on cross-link

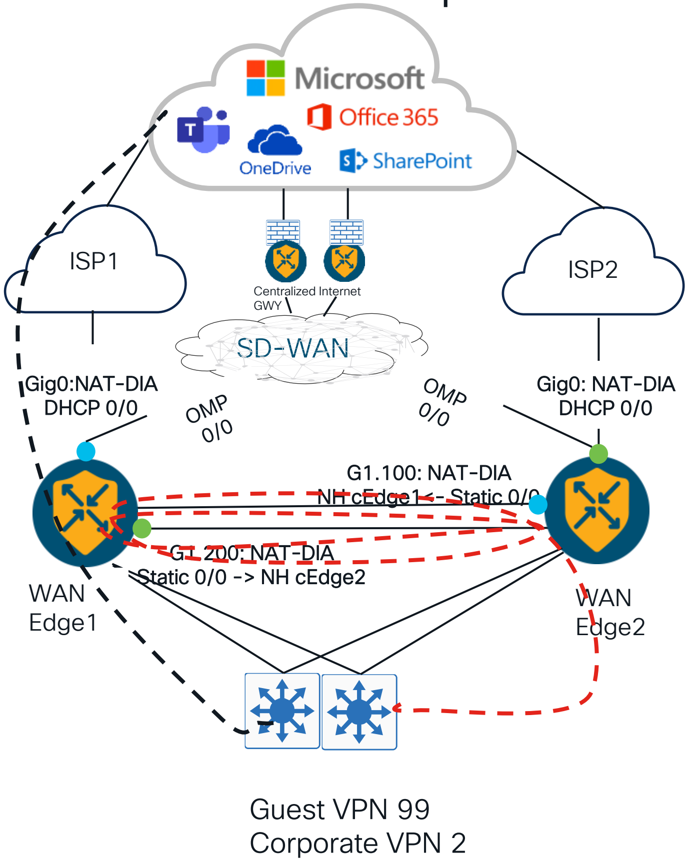

Dual Router Design for Internet High Availability

Option1: Solution for Internet traffic looping on cross-link

Solution: enhanced policy-based route-map (ePBR)

- ePBR is an advanced local data policy that routes traffic based on flexible match criteria such as IP addresses, port numbers, protocols.
- ePBR supports IPSLA object tracking to reduce risk of blackholing
- ePBR configuration requires cli add-on template

Example ePBR on cEdge2 for Guest DIA

```
ip sla 1
  icmp-echo 10.2.1.1
ip sla schedule 1 life forever start-time now
track 1 ip sla 1 state
!
ip access-list extended GUEST
 100 permit ip 10.10.0.0 0.0.0.255 any
class-map match-any DIA-ISP2
 match access-group name GUEST
!
policy-map type epbr DIA-ISP2
 class DIA-ISP2
  set ipv4 next-hop verify-availability 10.2.1.1 10 track 2
!
interface GigabitEthernet1.100
 service-policy type epbr input DIA-ISP2
```

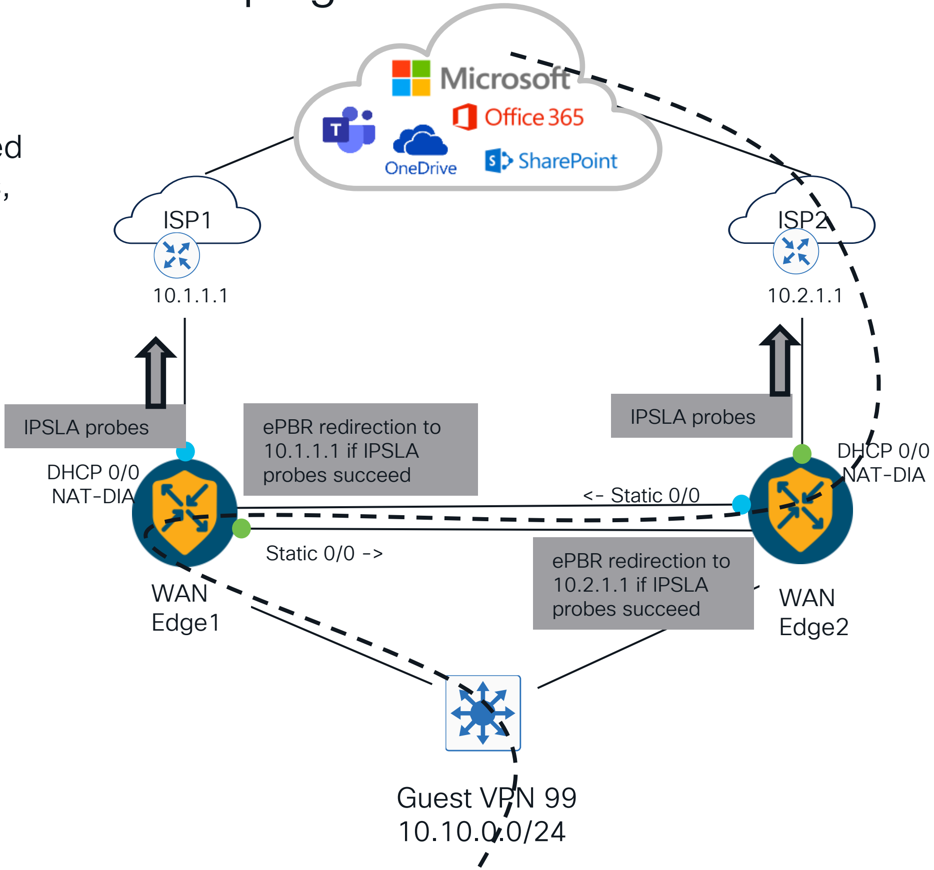

Dual Router Design for Internet High Availability

Option2: Active/Backup paths for DIA

Each WAN edge with single NAT-DIA interface
- Gig0 in Global VPN0 for NAT-DIA to local ISP
- Gig1 in Service VPN 99 (WAN edge 1-2 Cross-link)

DIA path selected as active by the preferred default route in VPN99
- NAT-DIA 0/0 route = admin distance 6

Cross-link path designated as backup due to less preferred distance
- OSPF 0/0 route = admin distance 110
- BGP 0/0 route = admin distance 20 or 200

ECMP load balancing across ISP1/ISP2 controlled by L3 LAN Switch
- OSPF or BGP learned default route from both WAN Edge1 and Edge2

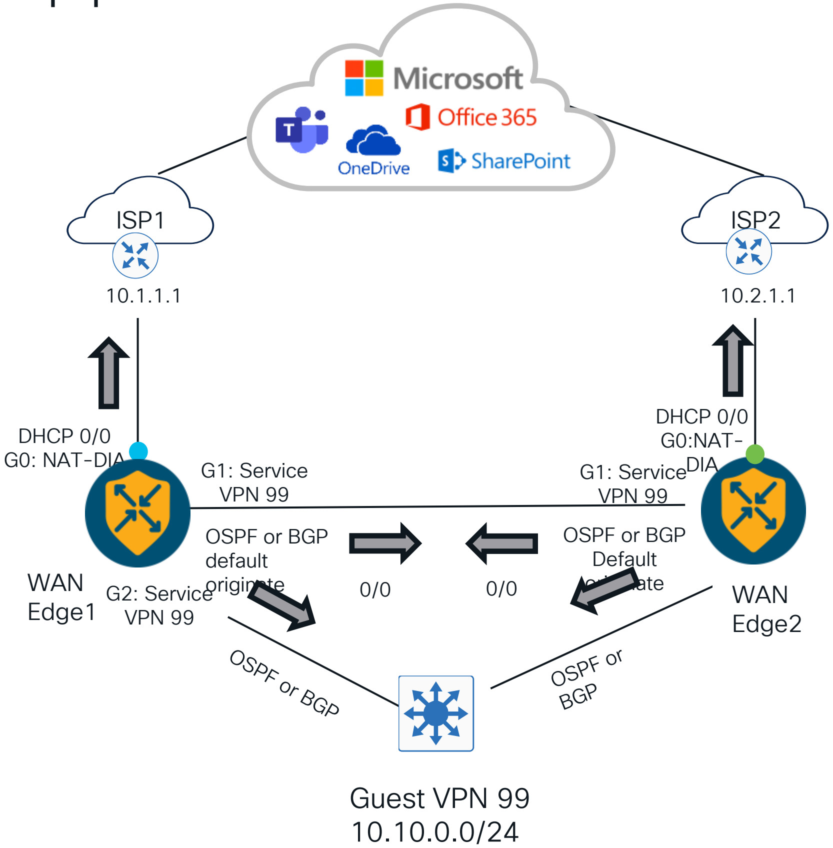

Critical Application High Availability

Requirement: 99.99% uptime SLA for call center applications

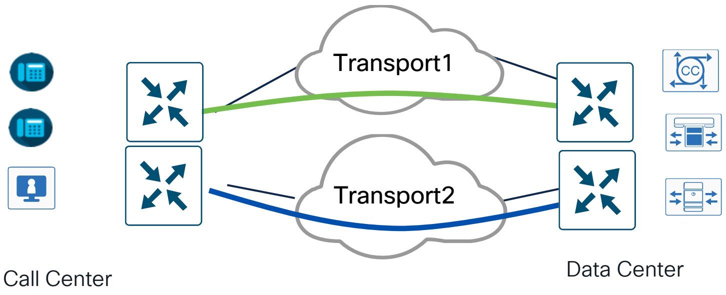

99.99% uptime SLA can only tolerate 52.6 minutes yearly downtime

Hard Circuit Failure can result in up to **7 seconds packet loss**
(Assuming CVD-recommended BFD timers)

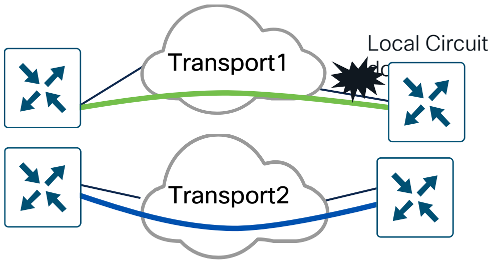

Soft Failure (brownout) may result in up to **12-minute packet loss**
(Assuming CVD-recommended BFD/App-route Timers)

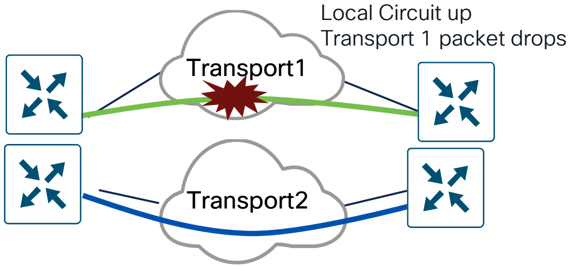

Enhanced AAR (EAAR)
(20.12/17.12)

AAR (Original)

Performance measurements (loss/latency) derived from BFD probe statistics (active measurements)

Soft failures in cloud may take up to 12 minutes to detect and reroute around

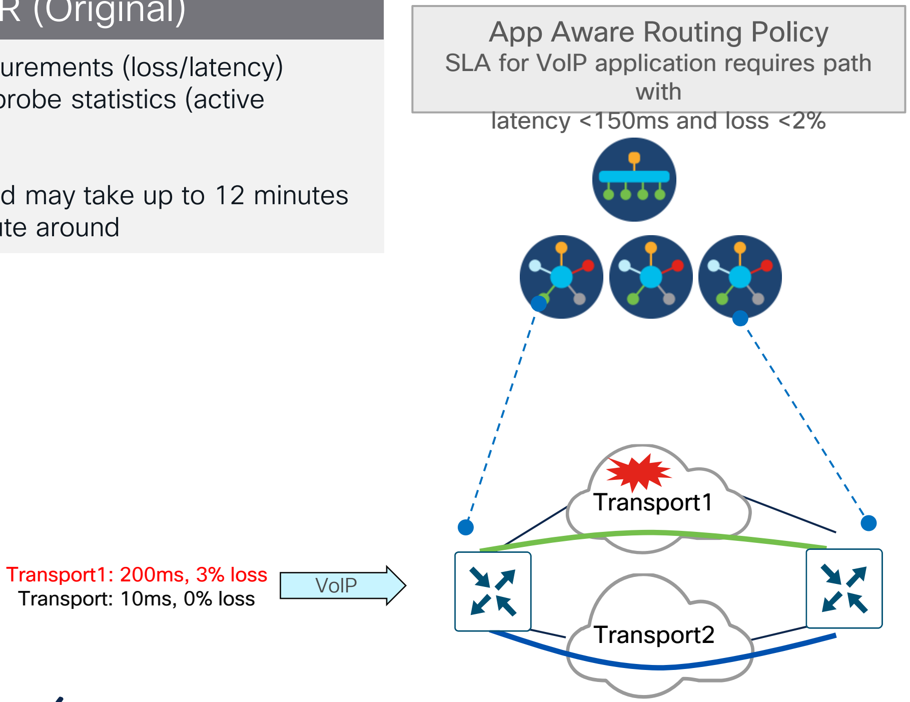

EAAR Improvements

Enhanced (passive) tunnel performance metrics measurements by using **inline data**

Faster tunnel degradation detection and switchover in the order of seconds (minimum 10 sec) to another path when SLA not met

SLA Dampening prevents churn

Critical Application Symmetric Traffic Flow

Requirement: Incoming traffic using AAR to use the same path

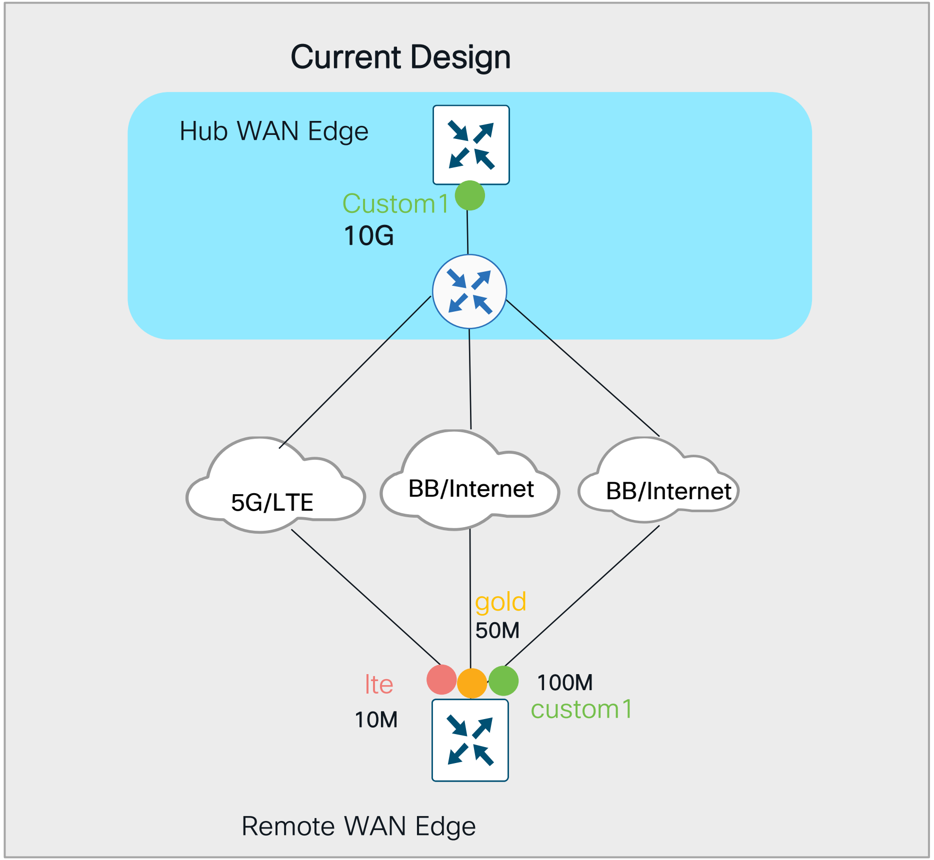

Solution: Remote Color Preference in AAR and Data Policy (20.15/17.15)

AAR (Original)

Able to influence outbound traffic from the branch but could not ensure the return traffic from DC to branch preferring the same path to those critical applications

Branch Site Types with unequal transports can now have symmetrical traffic flow

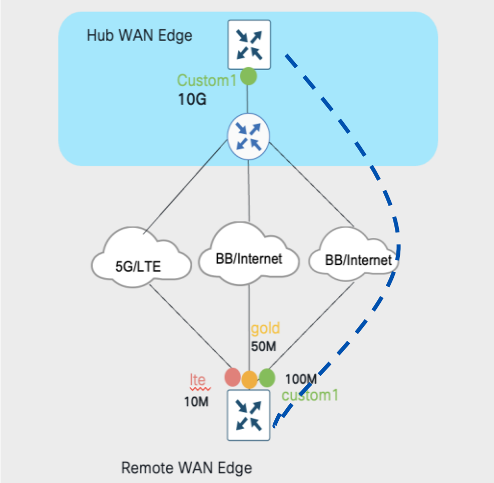

AAR Improvements

- The solution can be simplified by adding preferred remote-color at DC or Hub

  Remote-color preference option available in AAR and Data policy

SD-WAN Overlay High Availability
Goal: Zero downtime

HA Features
- N + 1 Redundant devices and circuits
- Control component redundancy
- Two or more underlay service providers

Design Limitations
- Single instance of OMP
- Does not protect against human error

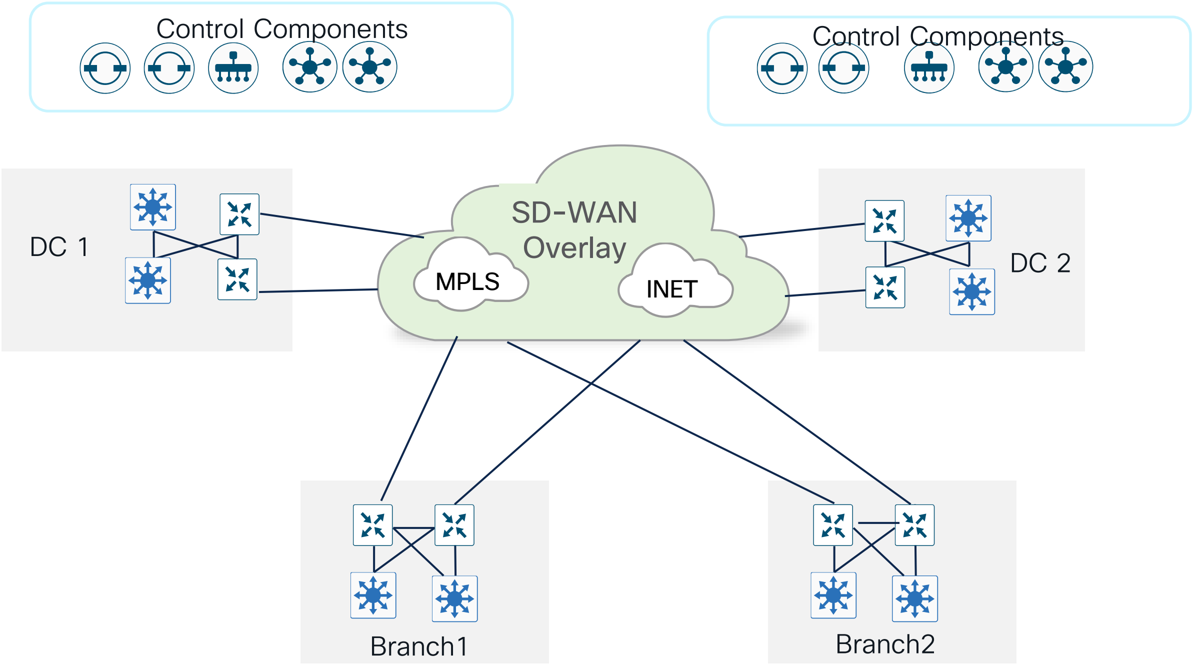

SD-WAN Overlay High Availability
Solution: Dual Overlays (Active/Backup)

Dual overlay Design
- Two unique overlays deployed with redundant WAN edge routers and control components
- Layer 3 LAN switch pair dual homed to each overlay with dynamic routing enabled
- Layer 3 LAN switch pair prefers active overlay based on routing metrics (eg, ospf cost, bgp med)

Caveats
- Two sets of device configurations to manage
- Manual sync of vManage active -> DR databases
- All sites failover to DR Overlay in event of outage on active

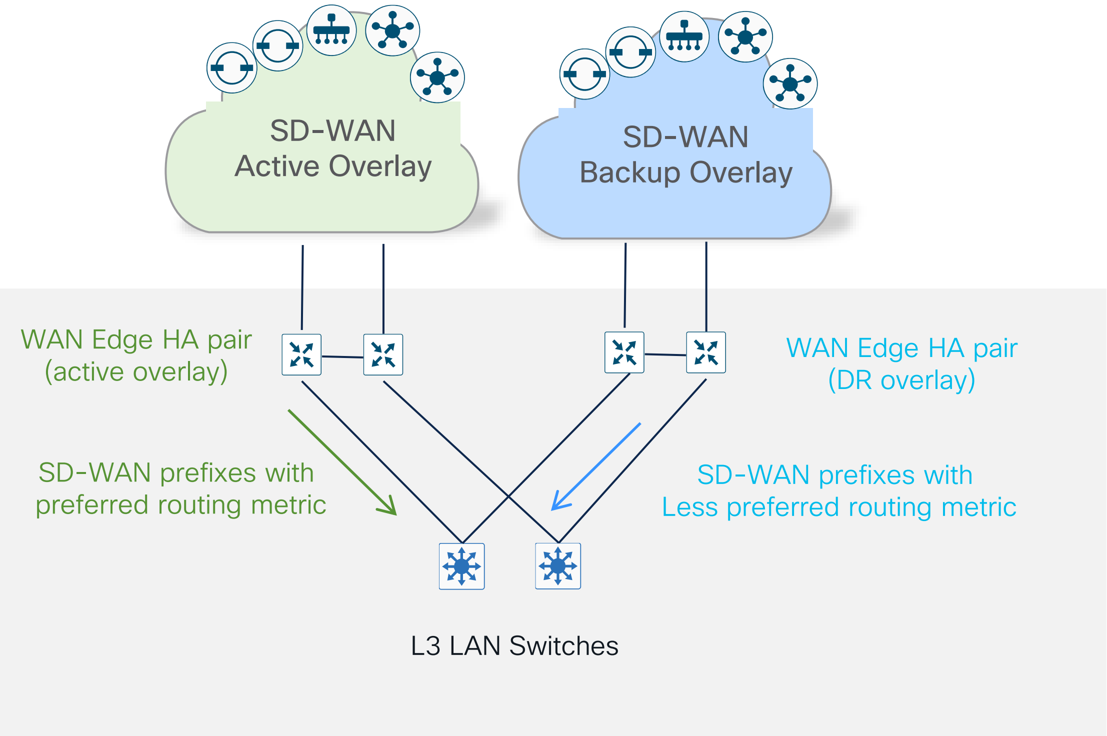

Cloud SaaS Path Optimization

Fast Track SaaS Optimization
Unlocking SaaS App Visibility Through First Packet Match

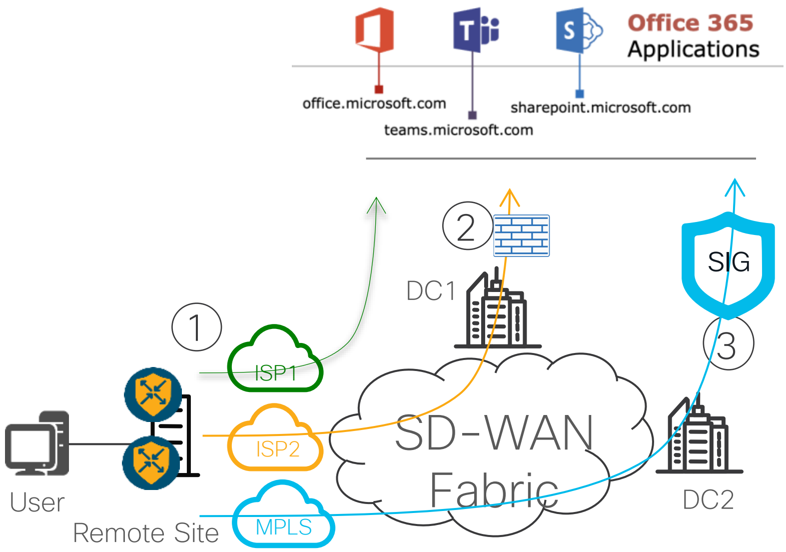

Business intent
- MS Teams forwarded over DIA path via local NGFW
- Other SaaS: Prefer better performing path across SD-WAN overlay to DC1 or DC2

Proposed Solution
- Cloud OnRamp for SaaS with support for SIG (20.3.4)

Problem
- How to achieve first packet match of MS applications for immediate traffic-steering of MS Teams to DIA?

Solution
- CoRSaaS with service-area classification (20.8.1)
- SD-AVC as a service

SDAVC as a Service

Control Plane
SDAVC as a Service

- Separate/decoupled SDAVC backend cloud service (brings agility and auto scale)
- SDAVC Cloud Service pulls M365 URL Categories using M365 web service.
- Dynamically pre-populates Edge router's NBAR cache with M365 IP addresses and URL Categories.
- Easy deployment: enabled by default, with automatic Cloud authentication
- Dynamic Update of built in Protocol Pack

Data Plane
NBAR Agents

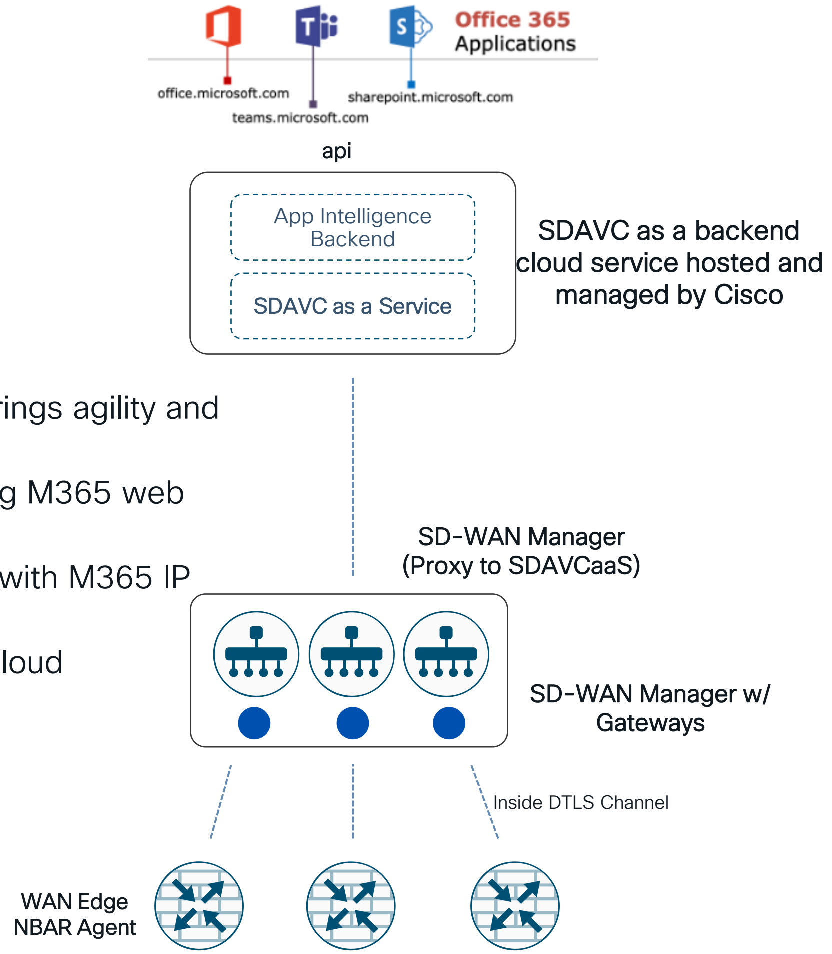

NWPI First Packet Trace

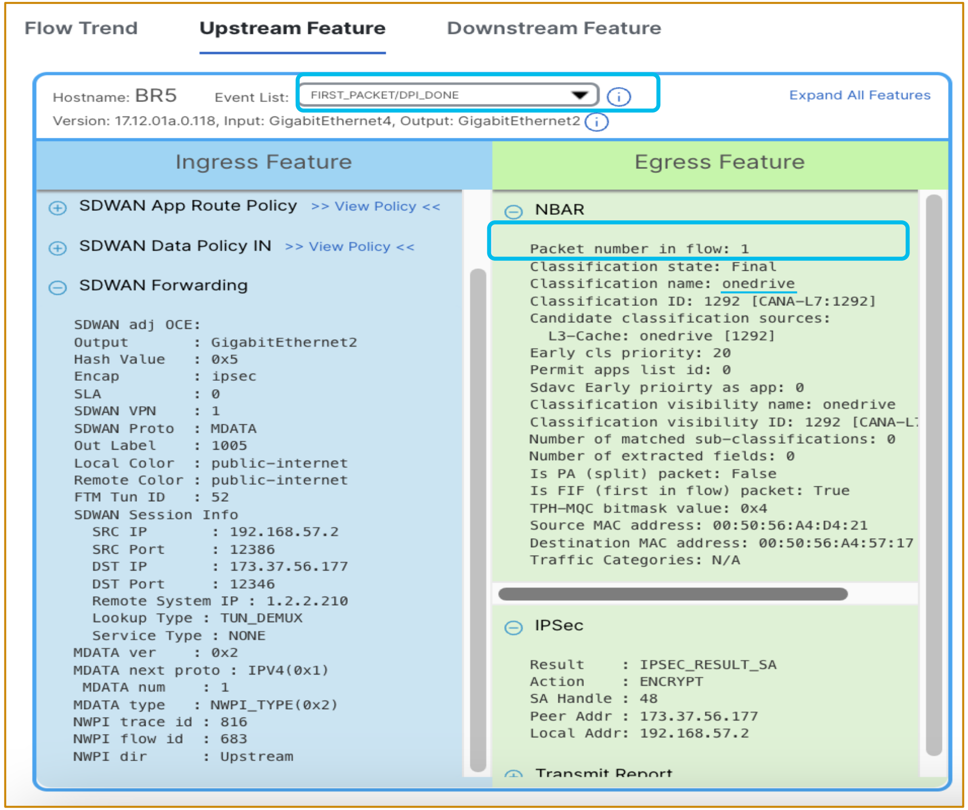

Summary


packet_duplication
tunnel
sd-wan
scalability
controller
overlay
ipsec
underlay
premise
studies
control
adaptive
protection
office
eaar
access
explicit
tools
sd-avc
hosted
cloud
application
design
satellite
application-aware
horizontal
bfd
deployment
case
qos
bandwidth_optimization
council
starlink
groups
sla
cellular
routing
microsoft
saas
list
optimization
nms
integration
daisy-chain
scaling

Catalyst SD-WAN Resources

Design Zone

Design Case Studies

Cisco Catalyst SD-WAN Design Case Studies Introduction
Cisco Catalyst SD-WAN Small Branch Design Case Study
Cisco Catalyst SD-WAN Large Global WAN Design Case Study
Cisco Catalyst SD-WAN Security Sensitive Design Case Study
Cisco Catalyst SD-WAN Remote Access Design Case Study
Cisco Catalyst Cloud First Case Study

Design Guides

Cisco Catalyst SD-WAN Design Guide
Cisco Catalyst SD-WAN Security Design Guide

Deployment Guides

Cisco Catalyst SD-WAN Deployment Guide
Controller Certificates and Authorized Serial Number File Deployment Guide
Secure Direct Cloud Access for IOS-XE SD-WAN Device Deployment Guide
..... And more

YouTube Video Content

Cisco SD-WAN and Cloud Networking

Catalyst SD-WAN Resources

Design Zone

Cisco Learning Network

[Cisco Catalyst SD-WAN optimizations for Starlink]

[NANOG92 Presentation]

slido

Please download and install the Slido app on all computers you use

What are some of the other use cases that aligns within your domain expertise You would like to hear from us in form of case studies?

ⓘ Start presenting to display the poll results on this slide.

Complete Your Session Evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to claim a **Cisco Live T-Shirt**.

Complete your surveys in the **Cisco Live mobile app**.

Continue your education

- Visit the Cisco Stand for related demos
- Book your one-on-one Meet the Expert meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Contact me at: rigoel@cisco.com

Thank you

cisco Live!

GO BEYOND

#CiscoLiveAPJC