



You make **possible**



Next-Gen SD-WAN (Viptela)

Design, Deployment and Best Practices

Stefan Olofsson, Technical Solutions Architect
Jean-Marc Barozet, Principal TME

TECRST-2191

CISCO *Live!*

Barcelona | January 27-31, 2020



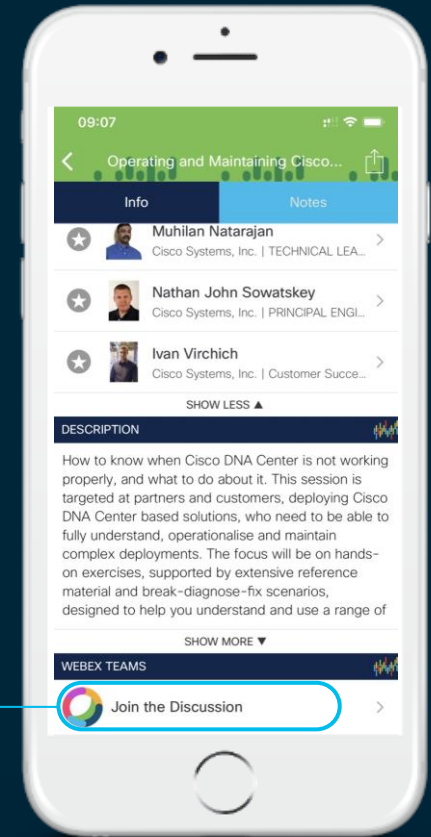
Cisco Webex Teams

Questions?

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

How

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



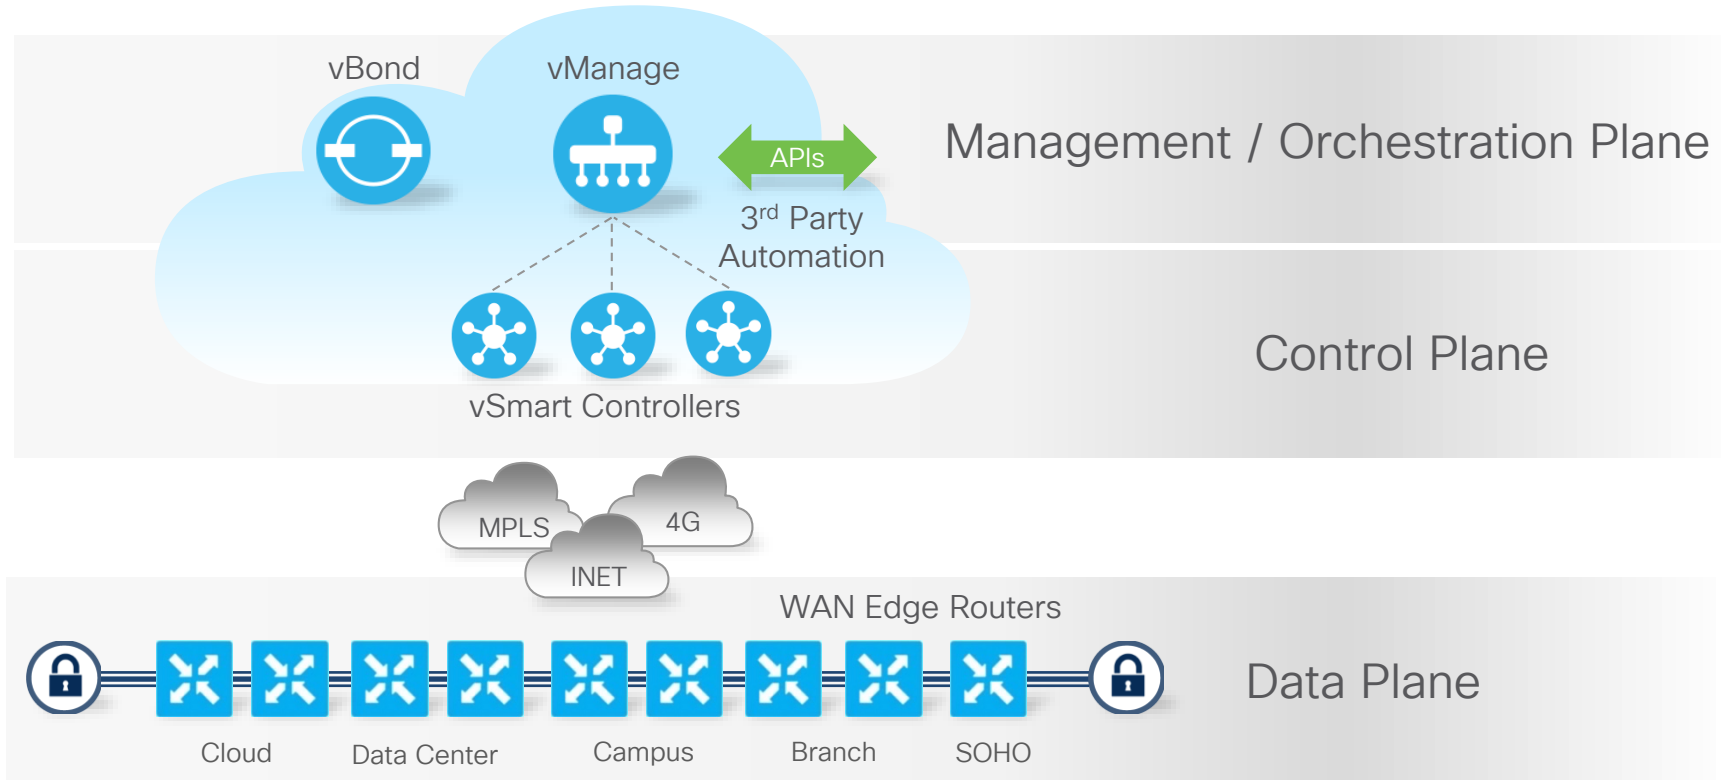
Agenda

- Introduction and Cisco SD-WAN Architecture Review
- SD-WAN Controller Deployment
- SD-WAN Control Plane and Design
- SD-WAN Data Plane and Design
- Policy Framework Introduction
- Overlay Network Design and Services
- Site Design
- Recommended Settings and Operational Best Practices

Cisco SD-WAN – Architecture Review

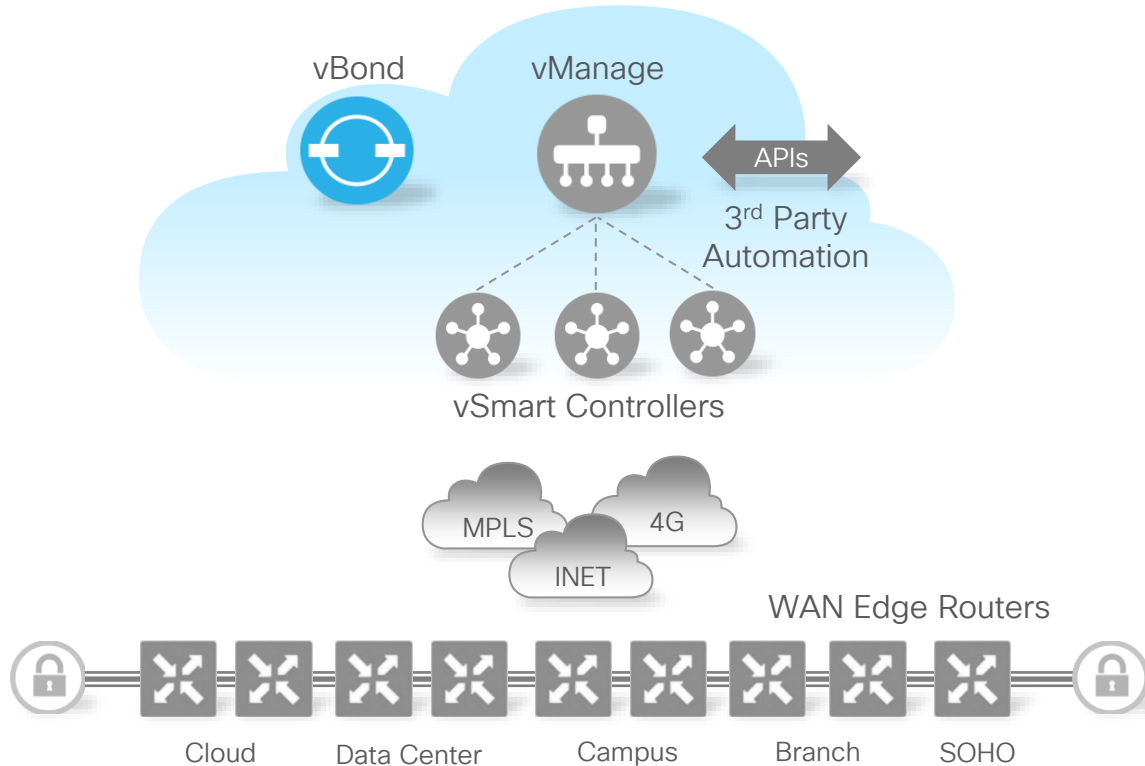
Cisco SD-WAN Architecture Overview

Applying SDN Principles Onto The Wide Area Network



Cisco SD-WAN Architecture Overview

Orchestration Plane - vBond

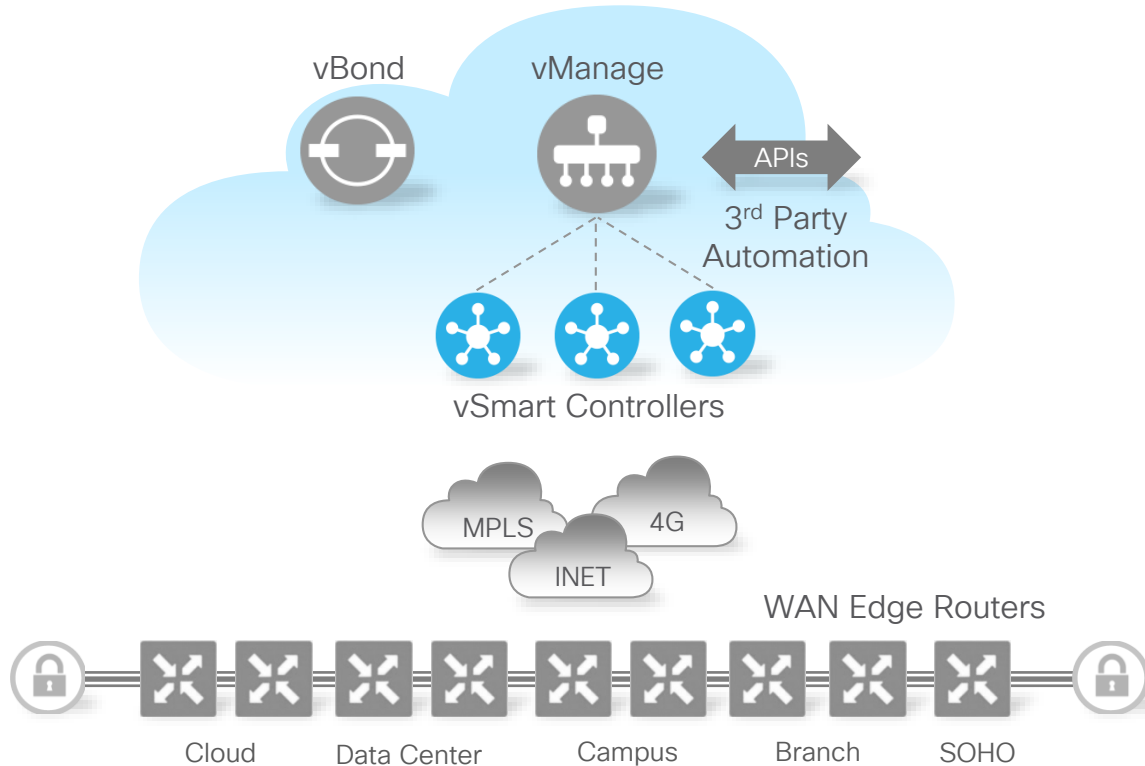


Characteristics

- Orchestrates control and management plane
- First point of authentication
- Distributes list of vSmarts/ vManage to WAN Edge routers
- Facilitates NAT traversal
- Requires universally reachable IP-Address/Port [Can reside behind Port-Forwarding]
- Independent and resilient layer
- Multitenant or single tenant

Cisco SD-WAN Architecture Overview

Control Plane - vSmart

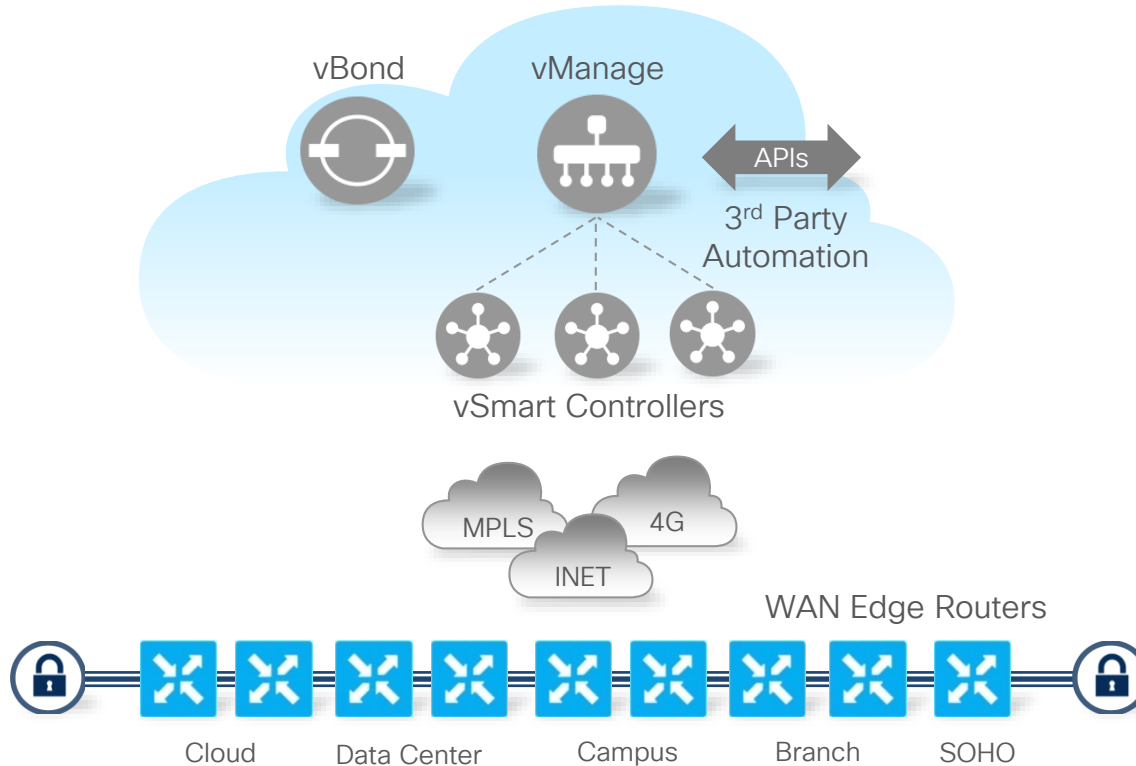


Characteristics

- Runs the Overlay Management Protocol (OMP)
- Facilitates fabric discovery
- Disseminates control plane information between Edges
- Distributes Data and App-Aware Routing policies to WAN Edge routers
- Implements control plane policies
- Enables simple and scalable hub-and-spoke control plane
- Independent and resilient layer

Cisco SD-WAN Architecture Overview

Data Plane - WAN Edge

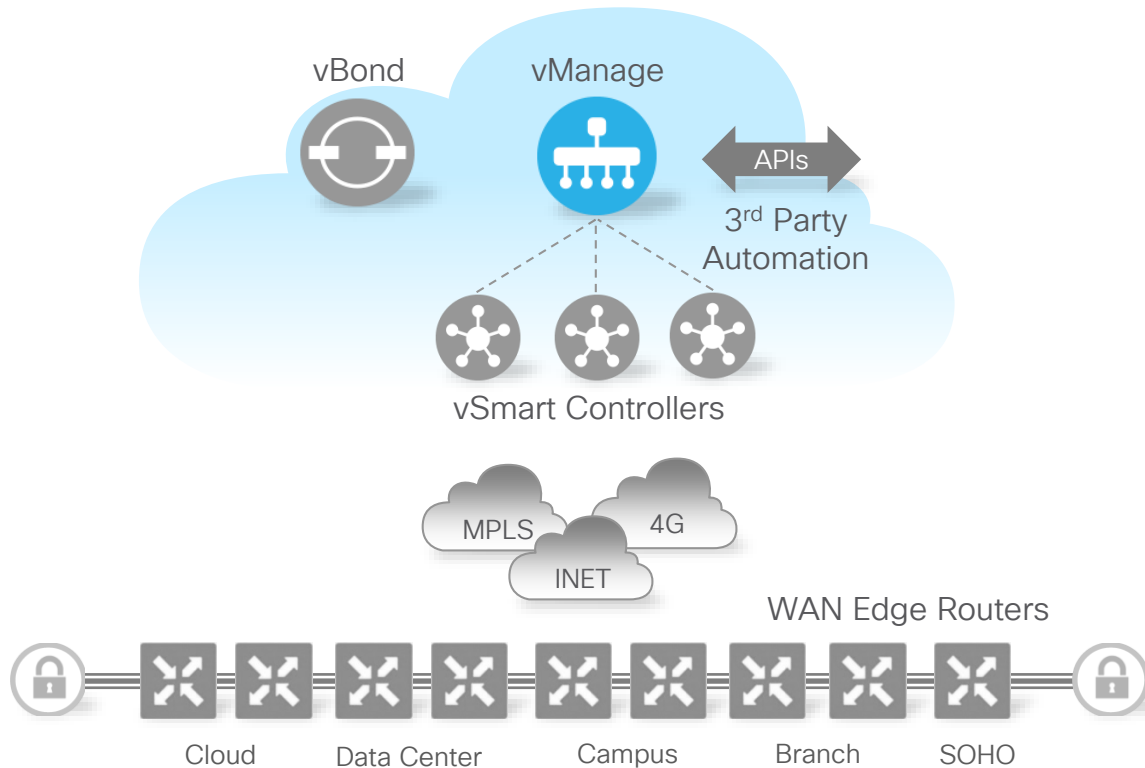


Characteristics

- Full stack WAN Edge router
- Secure data plane with remote WAN Edge routers
- Establishes secure control plane with vSmart controllers (OMP)
- Implements Data and App-Aware routing policies
- Collects and exports SLA and performance statistics
- Supports traditional routing protocols (OSPF, BGP) and First-hop Redundancy (VRRP)
- Supports Zero Touch Deployment
- Physical and Virtual form factor

Cisco SD-WAN Architecture Overview

Management Plane - vManage



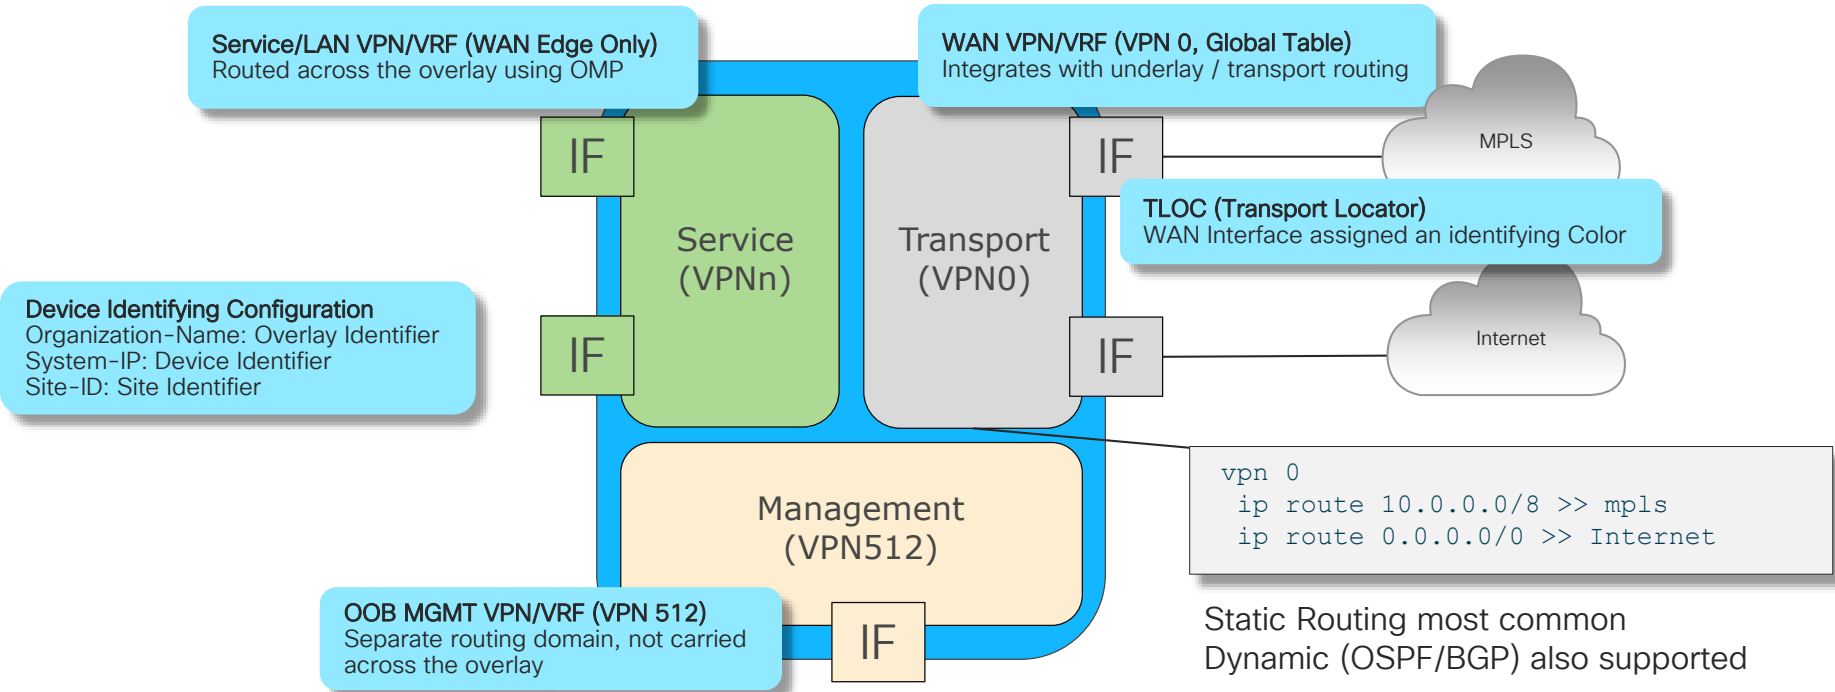
Characteristics

- Single pane of glass for Day0, Day1 and Day2 operations
- Centralized provisioning
- Multitenant or single tenant
- Policies and Templates
- Troubleshooting and Monitoring
- Software upgrades
- GUI with RBAC
- Programmatic interfaces (REST, NETCONF)
- Independent and resilient layer

Cisco SD-WAN – Terminology and key functions

SD-WAN Edge Device Architecture

Connecting the WAN/Transport VRF with the Underlay



Cisco SD-WAN Terminology

- Transport Side – Controller or WAN Edge Interface connected to the underlay/WAN network
 - Always VPN 0 (i.e. Global Table)
 - Traffic typically tunneled/encrypted, unless split-tunneling is used
- Service Side – WAN Edge interface attaching to the LAN
 - VPN 1-511 (512 Reserved for OOB Mgmt)
 - Traffic forwarded as is from original source
- TLOC – Collection of entities making up a transport side connection
 - System-IP: IPv4 Address (non-routed identifier)
 - Color: WAN Interface identifier on local WAN Edge
 - Encryption Key: The encryption key used for traffic destined to the originating TLOC
 - Private TLOC: IP Address on interface sitting on inside of NAT
 - Public TLOC: IP Address on interface sitting on outside of NAT
 - Private/Public can be the same if connection is not subject to NAT
- vRoute – Routes Carried in OMP for destinations reachable across the overlay
 - vRoute tagged with attributes as it is picked up by OMP

Cisco SD-WAN Terminology

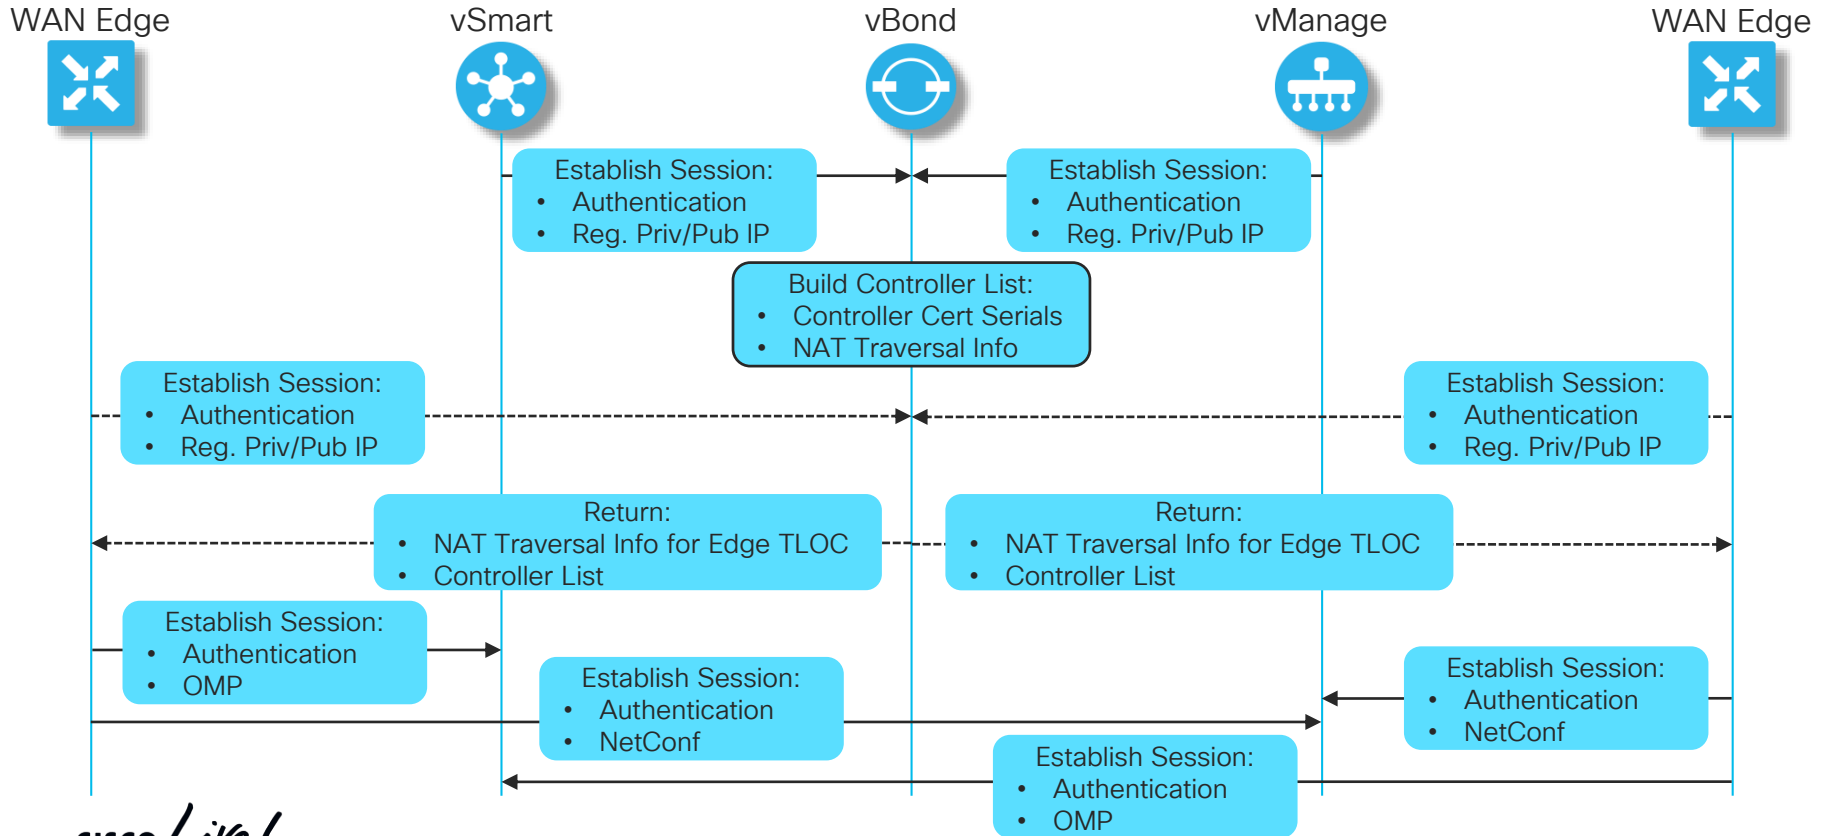
- OMP – Overlay Management Protocol
 - Dynamic Routing Protocol managing the Overlay domain
 - Integrated mechanism for distribution Routing, Encryption and Policies
- Site-ID – Identifies the Source Location of an advertised prefix
 - Configured on every WAN Edge, vSmart and vManage
 - Does not have to be unique, but then assumes same location
 - Required configuration for OMP and TLOC to be brought up
- System-IP – Unique identifier of an OMP Endpoint
 - 32 Bit dot decimal notation (an IPv4 Address)
 - Logically a VPN 0 Loopback Interface, referred to as “system”
 - The system interface is the termination point for OMP
- Organization-Name – Defines the OU to match in the Certificate Auth Process
 - OU carried in both directions for authentication b/t control and WAN Edge nodes
 - Can be set to anything as long as it’s consistent across the Cisco SD-WAN domain

Cisco SD-WAN – Network Bring-up

Cisco SD-WAN Network Bring-up

Control Plane Establishment

Permanent Session —————>
Temporary Session - - - - ->



WAN Edge Controller Discovery

Finding vBond

- Controllers and WAN Edges must know about vBond upfront
 - Via Configuration or Zero Touch Provisioning
- Controllers will attach to every known vBond
 - Ensures that every vBond knows of every controller
- WAN Edge will try vBonds one-by-one on every TLOC
 - Controller connections are established on every TLOC by default
- Both Controllers and WAN Edges find vBonds in the same way:
 - Locally configured IP-address (for a single vBond) or FQDN (for multiple vBonds)
 - FQDN can be resolved via DNS or locally (host statements)
 - In case of ZTP and need for local resolution, an IP-address can be pushed initially and host statements put in place when template configuration is applied

WAN Edge Controller Discovery

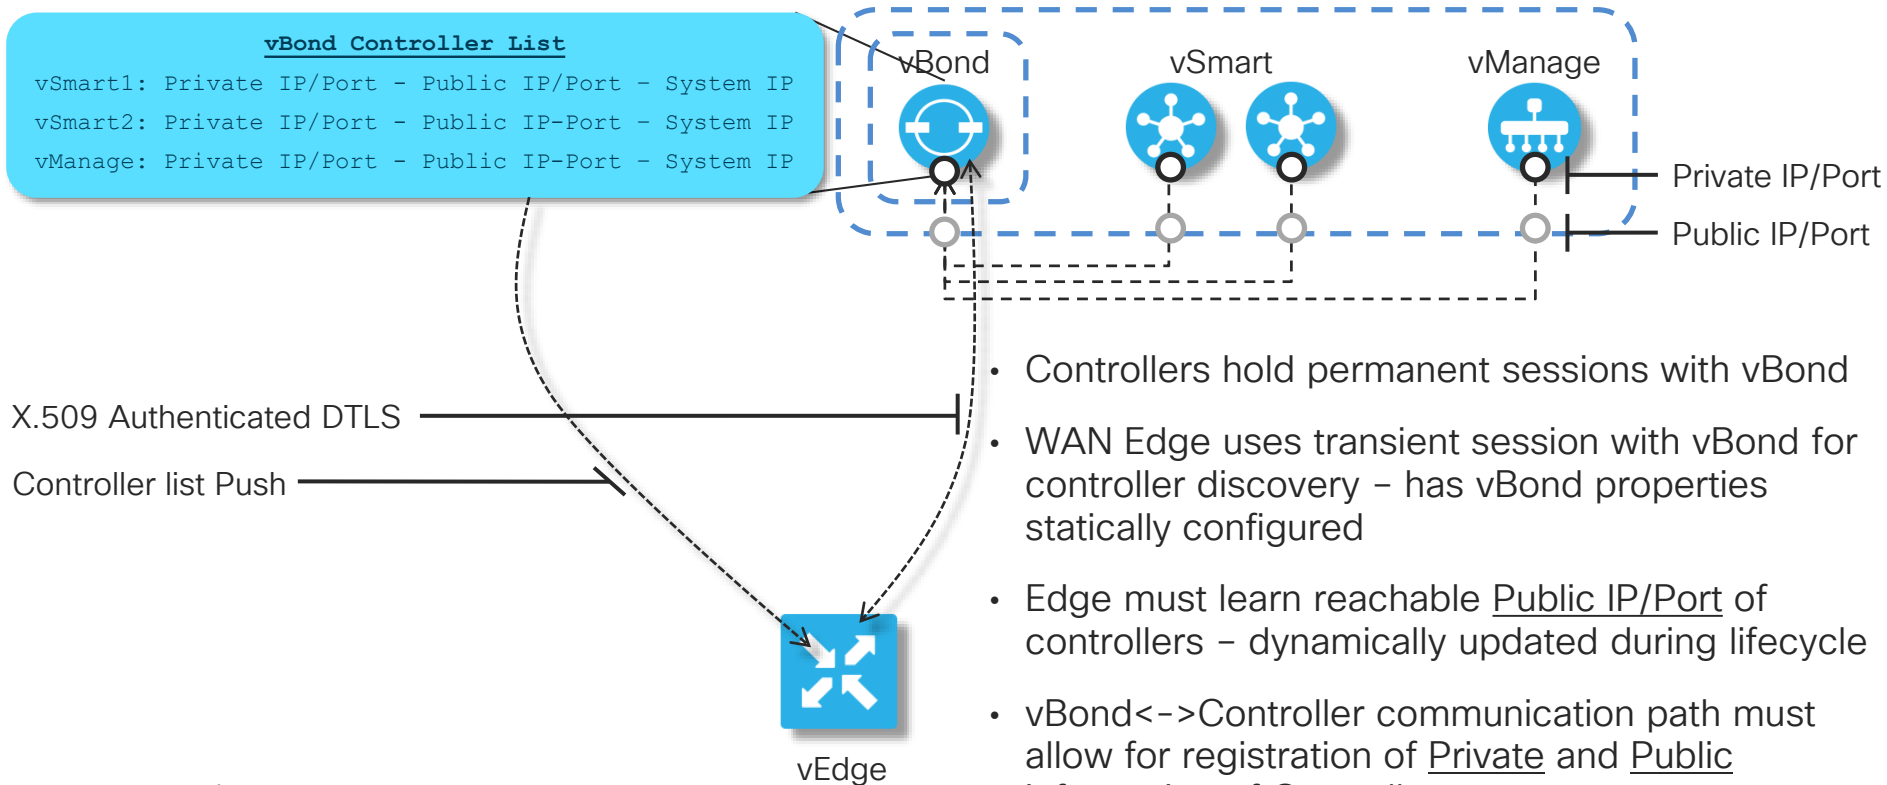
Using vBond

vBond Controller List

vSmart1: Private IP/Port - Public IP/Port - System IP

vSmart2: Private IP/Port - Public IP-Port - System IP

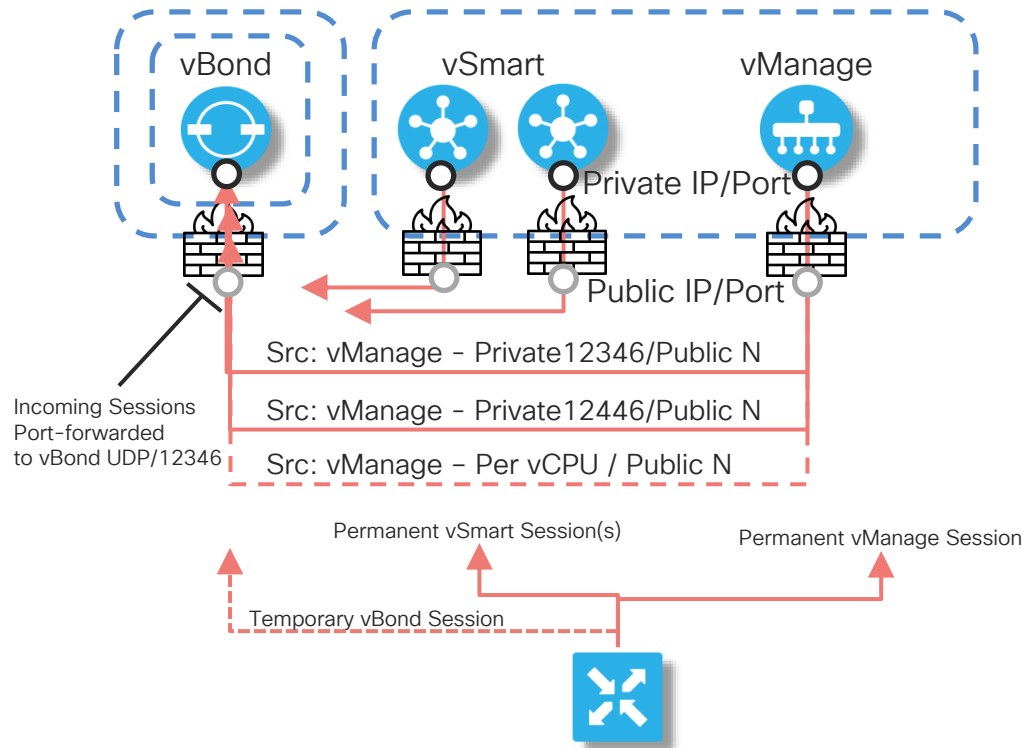
vManage: Private IP/Port - Public IP-Port - System IP



- Controllers hold permanent sessions with vBond
- WAN Edge uses transient session with vBond for controller discovery – has vBond properties statically configured
- Edge must learn reachable Public IP/Port of controllers – dynamically updated during lifecycle
- vBond<->Controller communication path must allow for registration of Private and Public information of Controllers

Controller Firewall Traversal

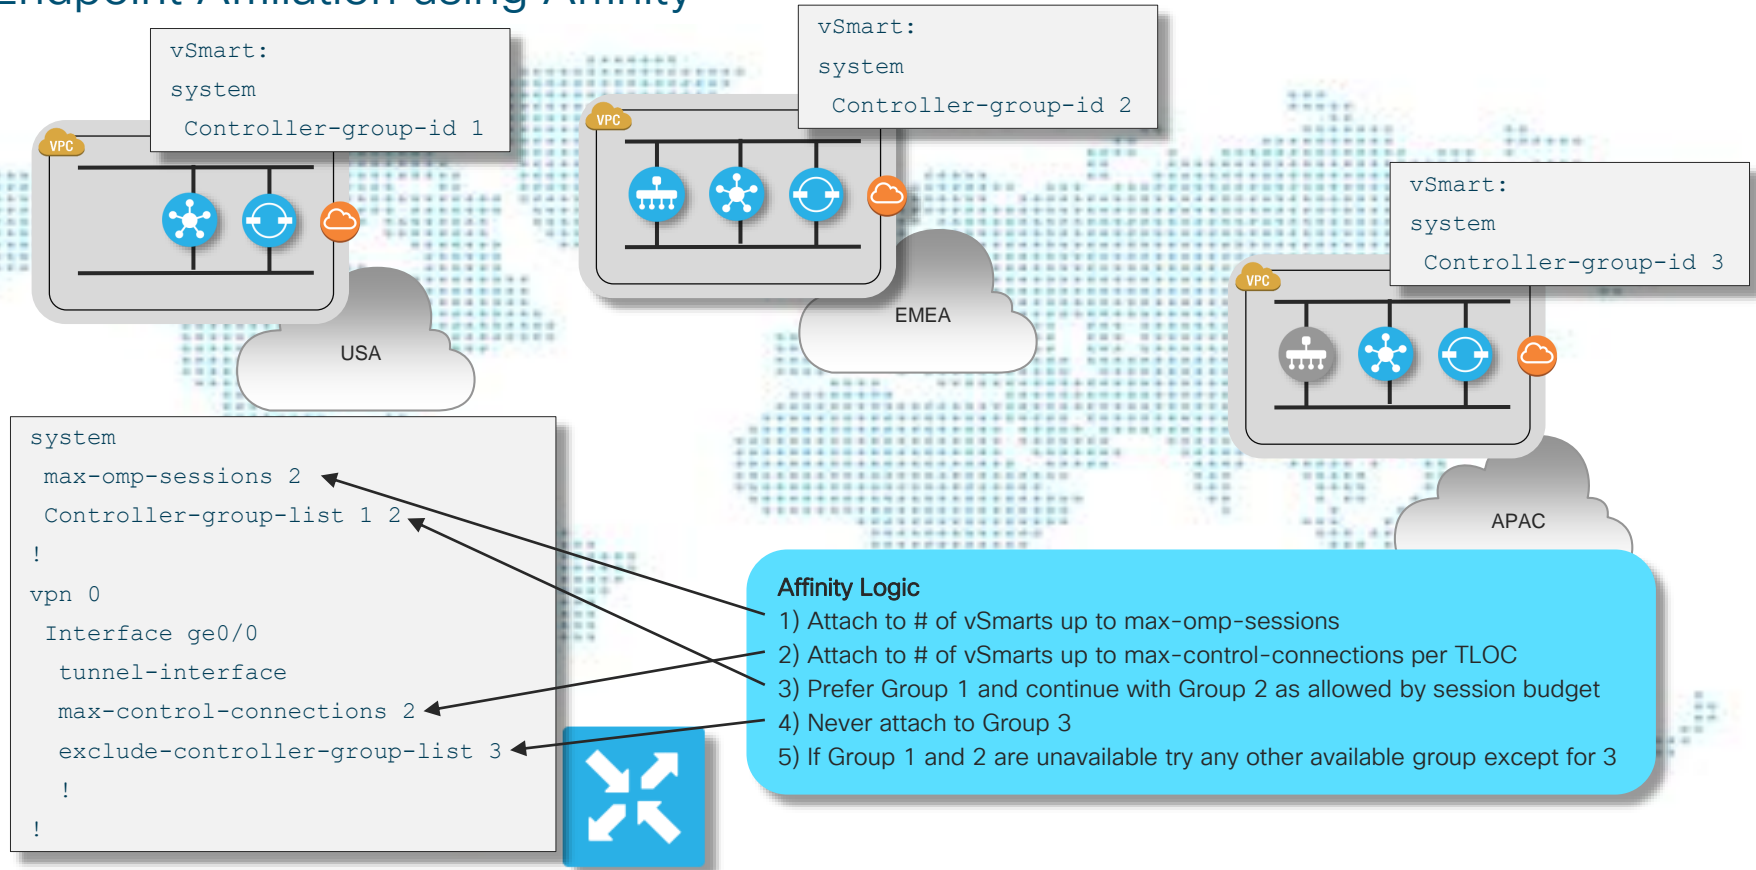
Port Numbers and Session Establishment Direction



- vBond Service: UDP/12346
Not tied to a local IP-address
- vSmart/vManage behavior is identical
Will establish permanent sessions with every vBond
One session established per vCPU for load-sharing
- vSmart: UDP/12346 + 100 per vCPU
12346, 12446, 12546, etc
- vManage: UDP/12346 + 100 per vCPU
12346, 12446, 12546, etc
- WAN Edge: UDP/12346 + 20 up to 12426
12346, 12366, 12386, 12406, 12426*
Source-Port will periodically change if port-hop is enabled* - 12346 is default
Base port can move +1-19 using port offset**

SD-WAN Controller Large Scale Deployment

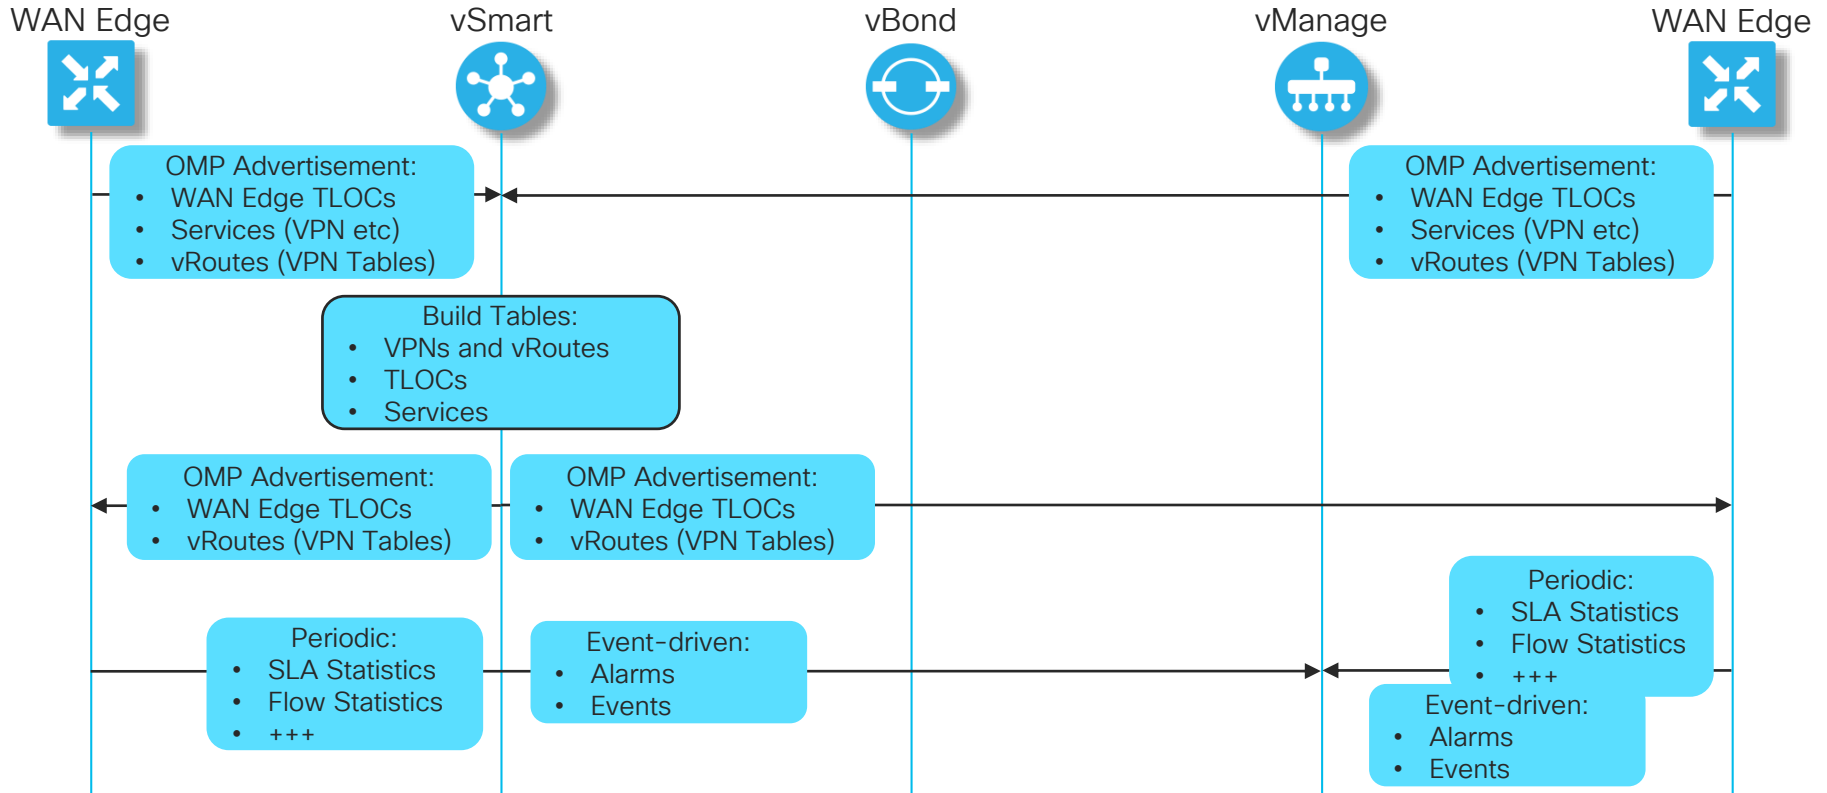
Endpoint Affiliation using Affinity



Cisco SD-WAN Network Bring-up

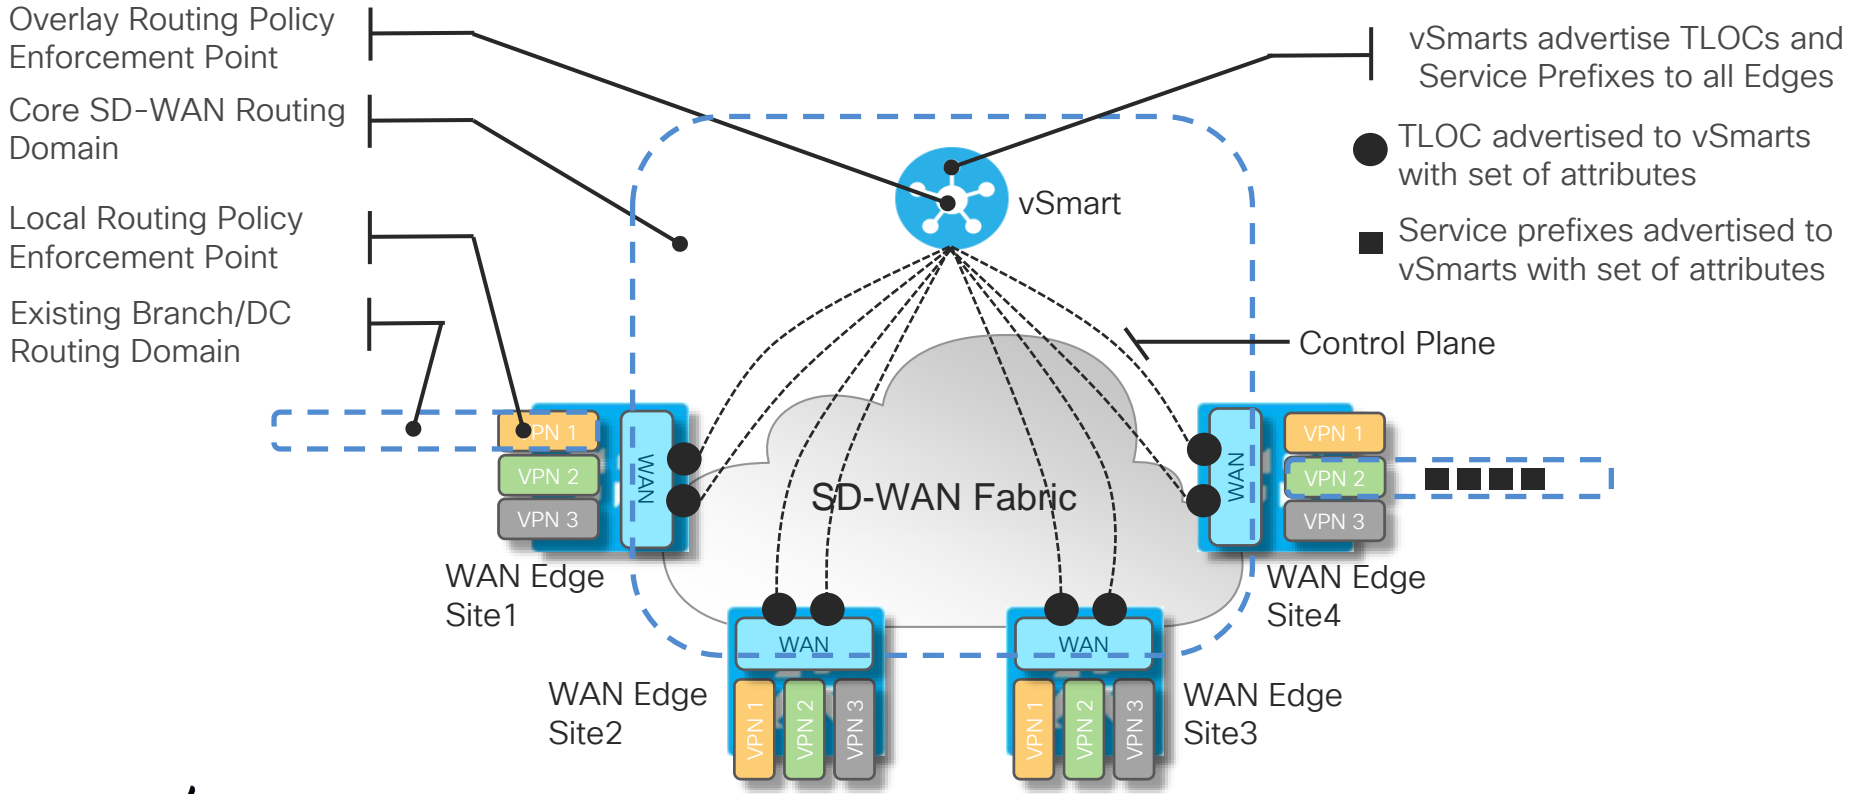
Control Plane Exchanges

Permanent Session —————>
Temporary Session - - - - ->



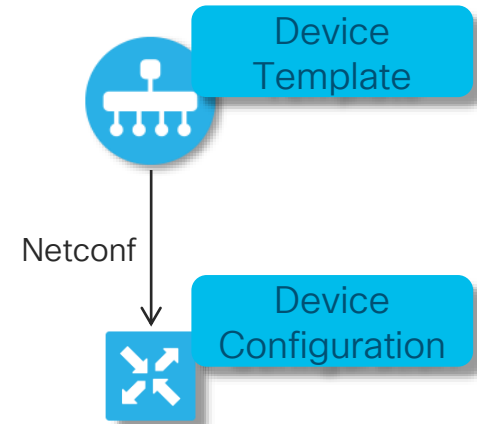
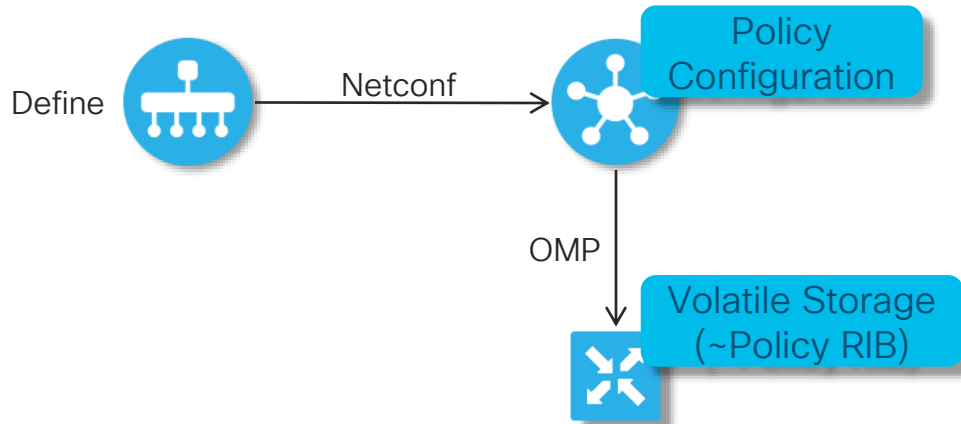
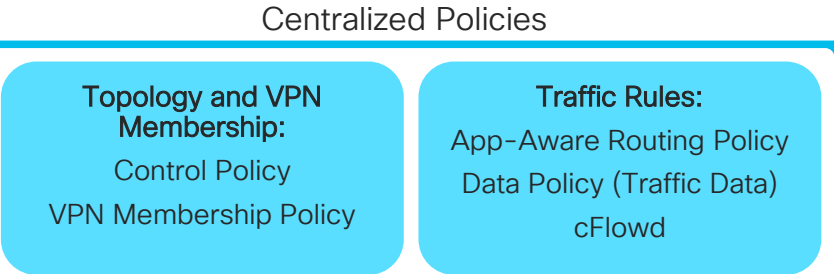
Cisco SD-WAN Overlay Routing

Multi-domain Routing Fabric



Cisco SD-WAN Policy Architecture

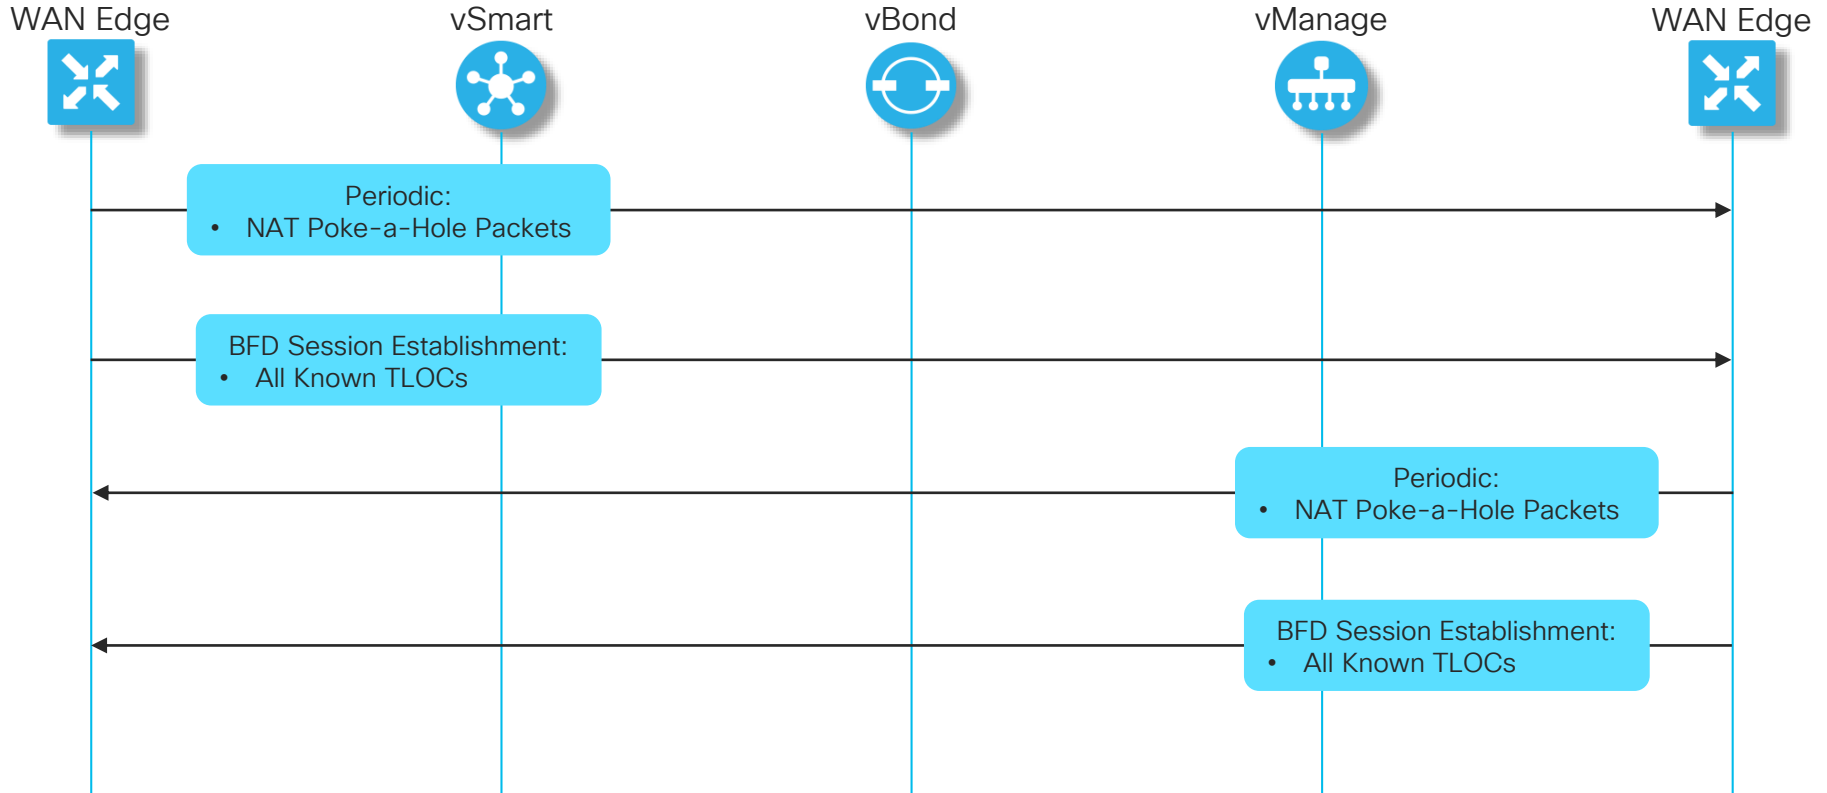
Policy Categories



Cisco SD-WAN Network Bring-up

Permanent Session —————>
Temporary Session - - - - ->

Data Plane Establishment



TLOCs, Colors, Site-IDs and Carriers

Definitions

- TLOC Color used as static identifier for:
 - TLOC Interface on WAN Edge device
 - Underlay network attachment
- The specific color used is categorized as Private or Public
 - Private Colors [mpls, private1-6, metro-ethernet]
 - All other colors are public [red, blue,..., public-ethernet,...]
- Private vs Public color is highly significant
- Color setting applies to:
 - WAN Edge to WAN Edge Communication
 - WAN Edge to Controller Communication

TLOCs, Colors, Site-IDs and Carriers

Private and Public Color Significance

1 Private Color to Private Color



2 Private Color to Public Color




3 Public Color to Public Color



TLOCs, Colors, Site-IDs and Carriers


Color Contention Resolution

- If Site-IDs are identical and colors public:
 - Use Private information
- Carrier setting is final influencer to decide on Private/Public IP/Port
 - Use if two endpoints are using private colors and you need session between them to be established between their Public IP/Port



Control and Data Plane Establishment

Significance of TLOC Color



- MPLS uses Private Color, Internet uses Public Color
- Connectivity optimized within and across domains

Cisco SD-WAN – Controller Deployment

Cisco SD-WAN Controllers


Deployment Options and Designs

- Design is highly influenced by:
 - On or Off-prem hosting
 - Security Requirements
 - Need to manage Data Plane attachment to the overlay in combination with NAT
- Controller Design must handle any combination of data plane and choice of hosting model
- Key concepts, such as TLOC, Colors and Carrier settings are crucial for a fully functional overlay

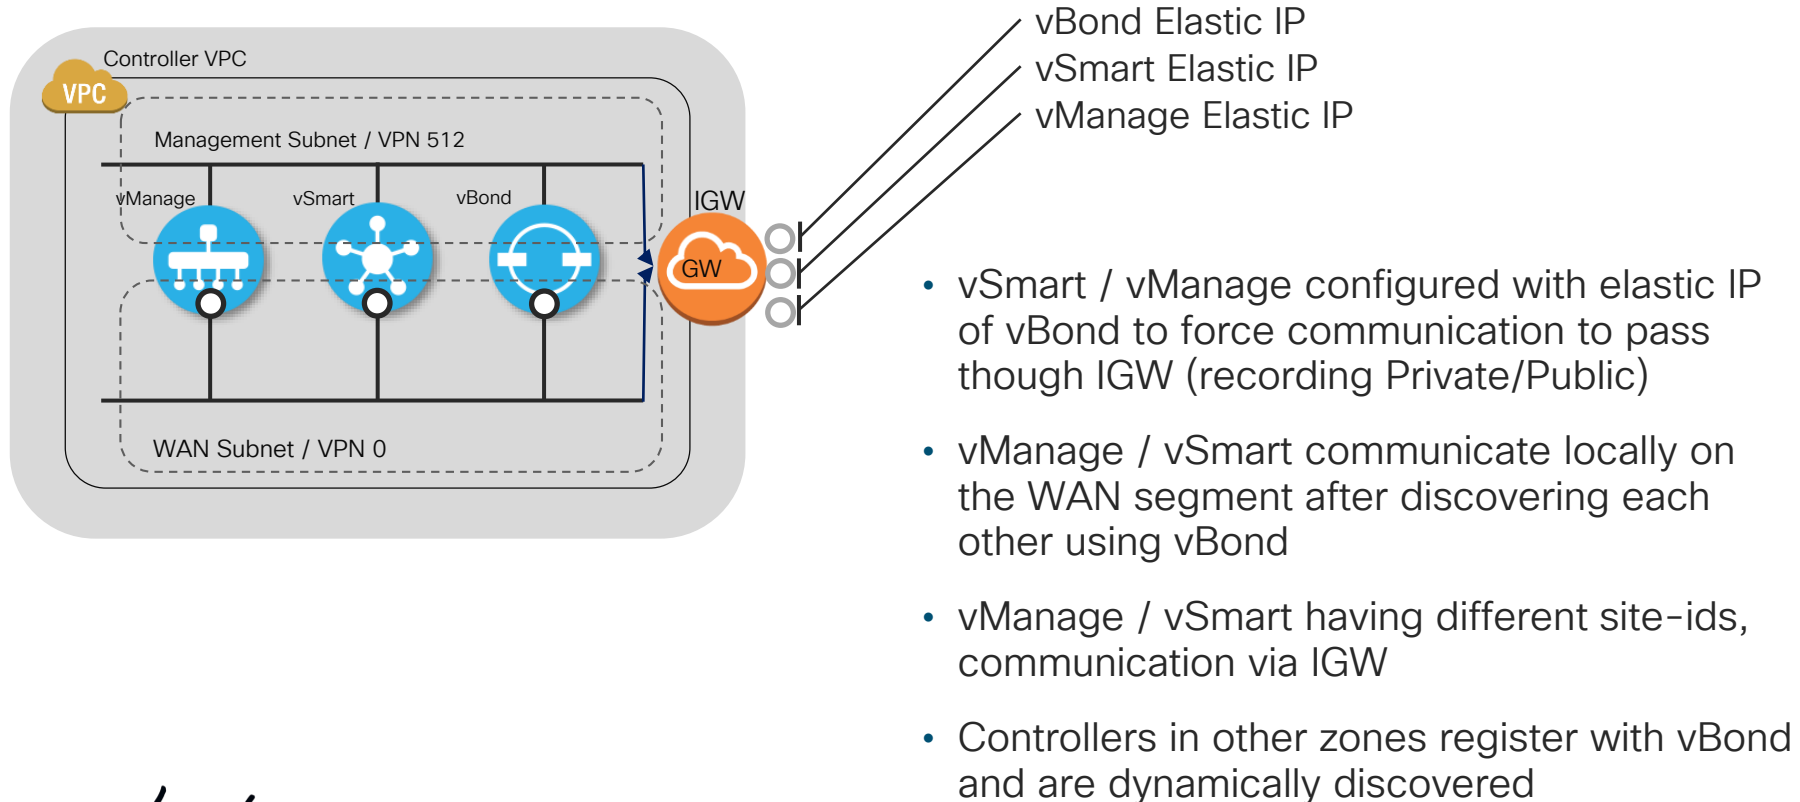
Cisco SD-WAN – Controller Deployment Models

Controller Deployment Models

- Cloud hosted
 - AWS or Azure
 - Single or Multiple Availability Zones
 - Recommended Model
- Cloud hosted + On prem
 - Public Cloud, Private Cloud and/or DC
 - IP connectivity between domains required
 - Not Supported
- On prem only
 - Private Cloud or DC
 - Public and Private transport still supported
 - Specific design considerations required




Cloud hosted Deployment - Recommended




Cloud hosted Deployment – Recommended

Control and Data Plane Establishment

Interconnected data plane –
Contiguous connectivity



Separate underlays –
Disjointed connectivity



Cloud-hosted Deployment

Summary

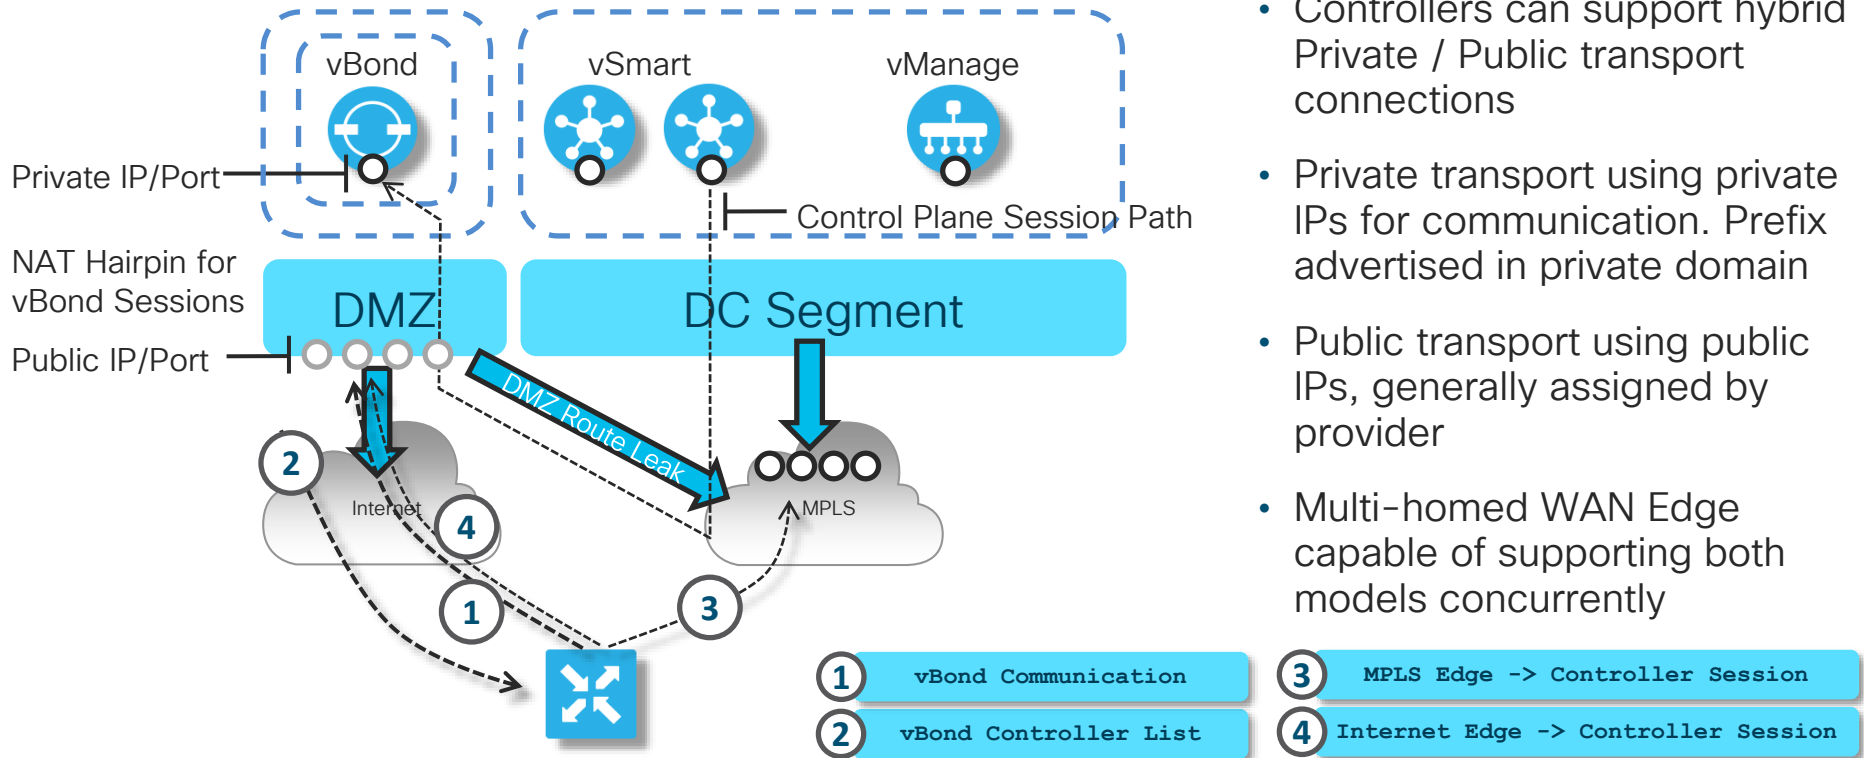
- Recommended mode of deployment
 - Ease of deployment – Cisco orchestrated
 - No On-Prem design considerations
 - Easy to scale and to deliver redundancy / HA
- Requirements
 - Internet connectivity from every site (unless using DirectConnect)
 - If using MPLS Transport, Internet breakout required for Control Plane
- Challenge
 - With a single Internet connection, no DirectConnect or Internet Breakout from MPLS – No Controller Redundancy

On-prem Deployment Considerations

- Supporting NAT Traversal
 - vBond supporting Private + Public Discovery
- Supporting Hybrid Environments
 - Interconnected MPLS and Internet Domains
 - Separate MPLS and Internet Domains
- Redundancy
- Firewall Traversal

On-Prem Deployment – vBond / NAT Traversal

Controllers accessible via Private and Public Transport

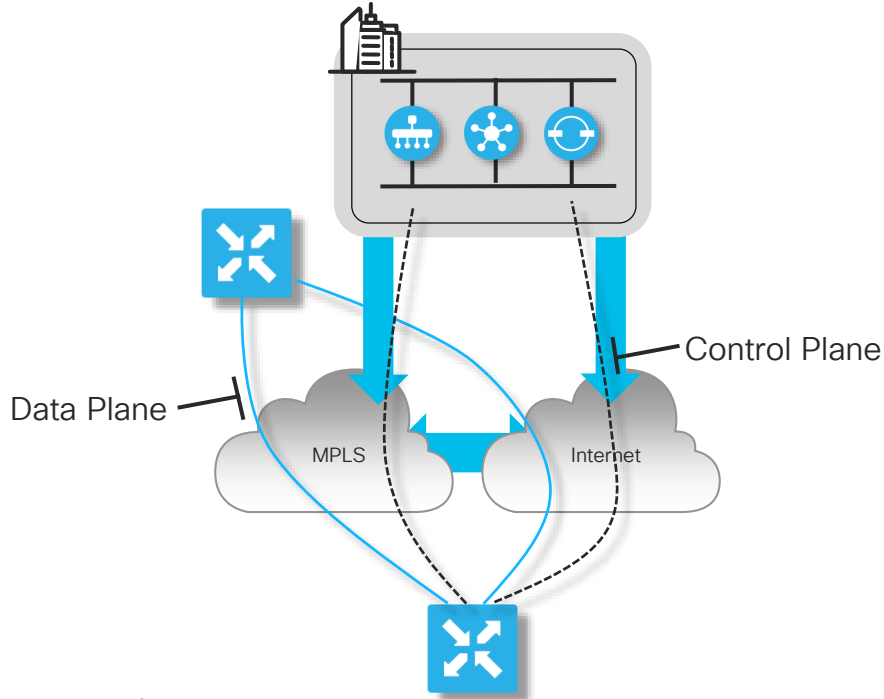


- Controllers can support hybrid Private / Public transport connections
- Private transport using private IPs for communication. Prefix advertised in private domain
- Public transport using public IPs, generally assigned by provider
- Multi-homed WAN Edge capable of supporting both models concurrently

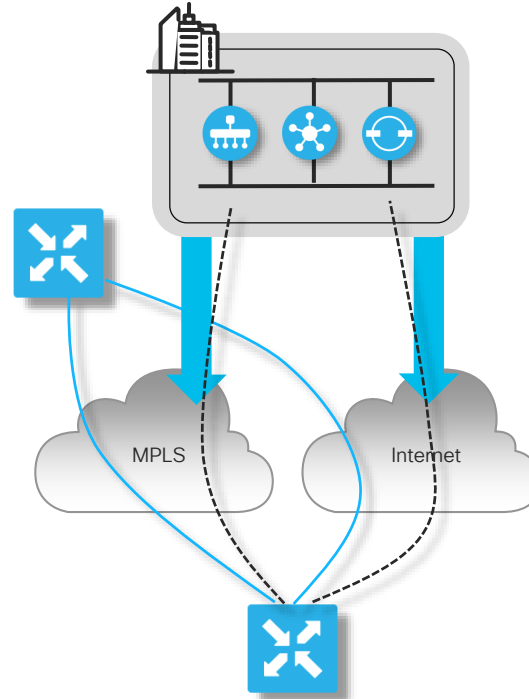
On-Prem Deployment

Control and Data Plane Establishment

Interconnected data plane –
Contiguous connectivity

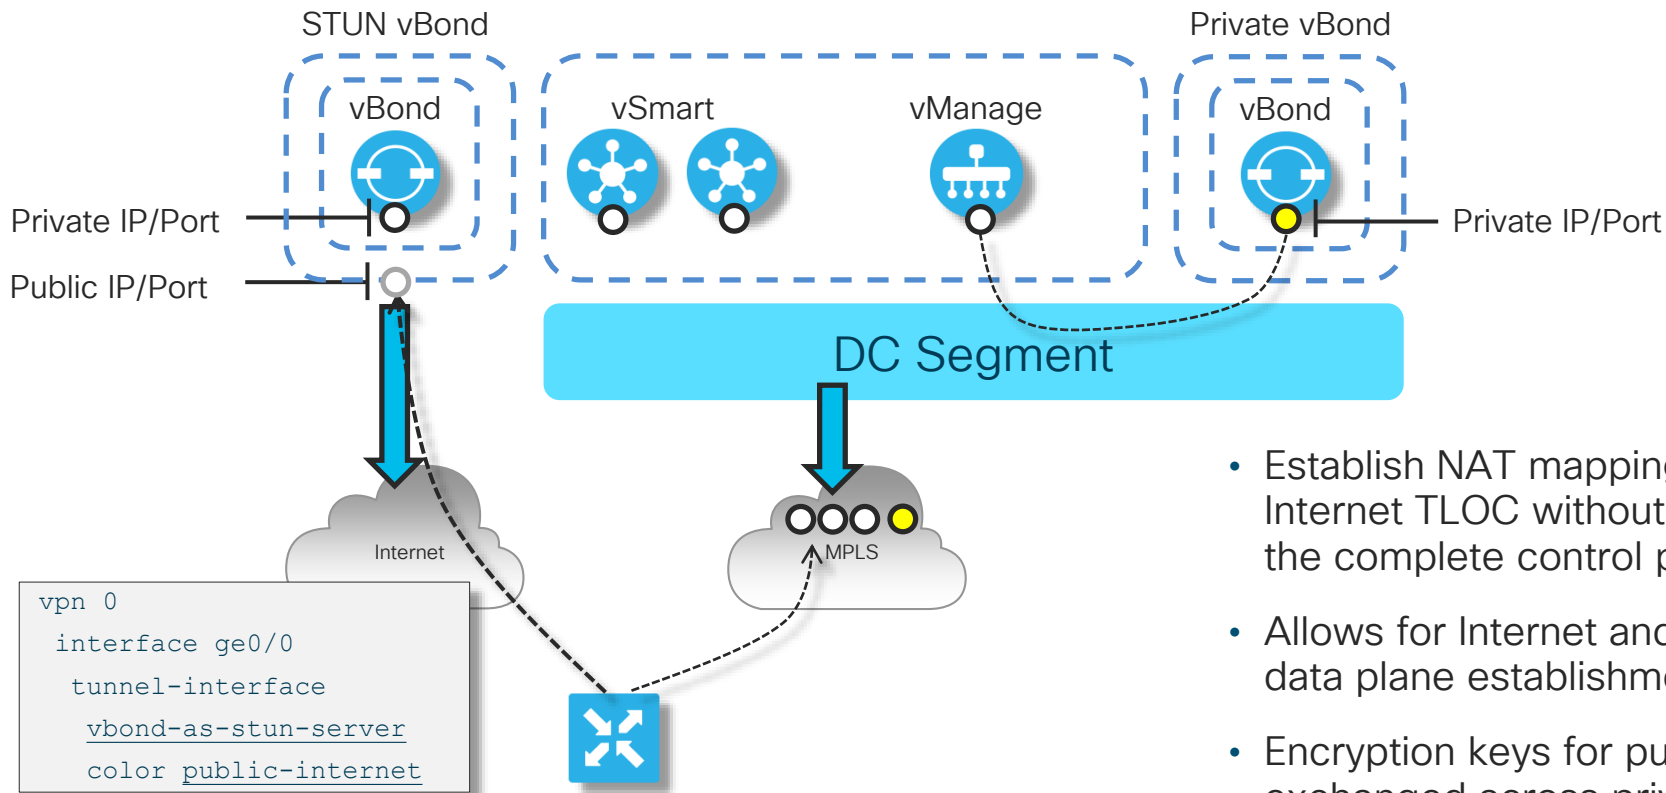


Separate underlays –
Disjointed connectivity



vBond-as-Stun-Server (vEdge Only)

Controllers accessible via Private Transport Only

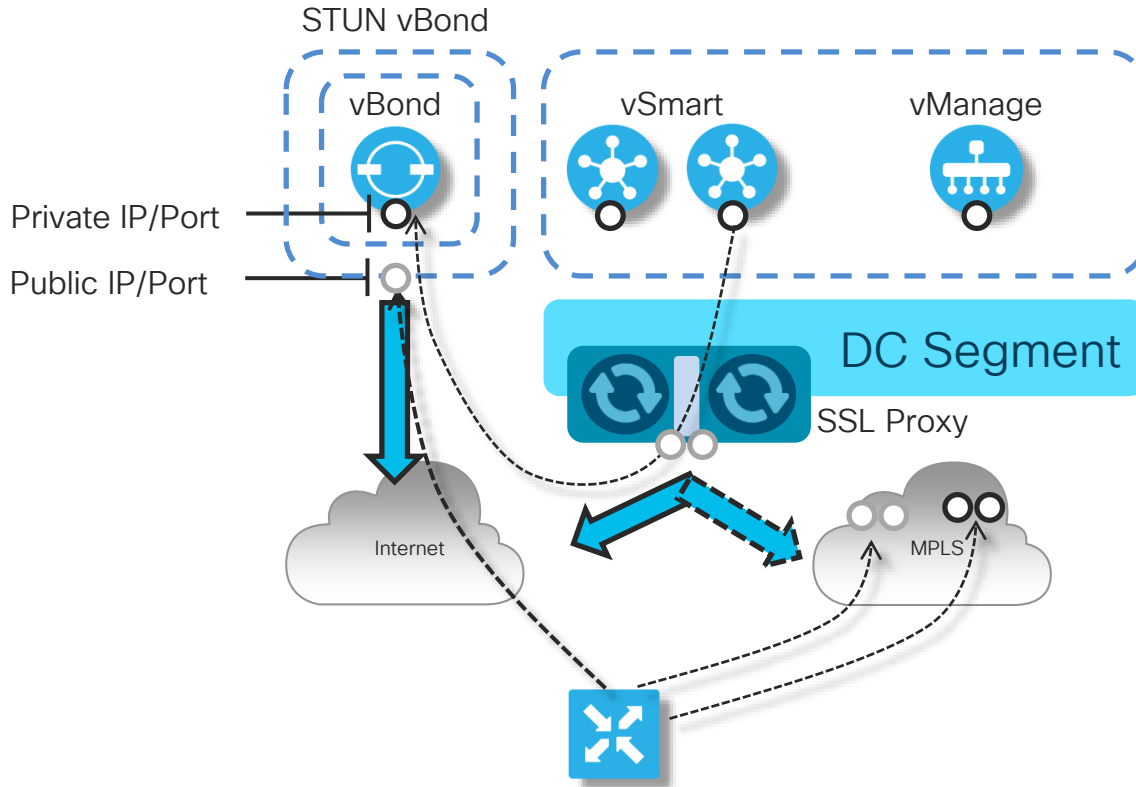


```
vpn 0
interface ge0/0
  tunnel-interface
  vbond-as-stun-server
  color public-internet
```

- Establish NAT mappings for Internet TLOC without operating the complete control plane
- Allows for Internet and MPLS data plane establishment
- Encryption keys for public exchanged across private only

Controller Proxy Access

Use SSL Proxy to maintain secure isolation of controllers



- Private IP/Port information never exposed outside of DC
- Ports assigned on Proxy are configured on vManage/vSmart
- Proxy needs signed cert to allow for controller/edge authentication
- Ports maps are established with vBond using standard mechanism

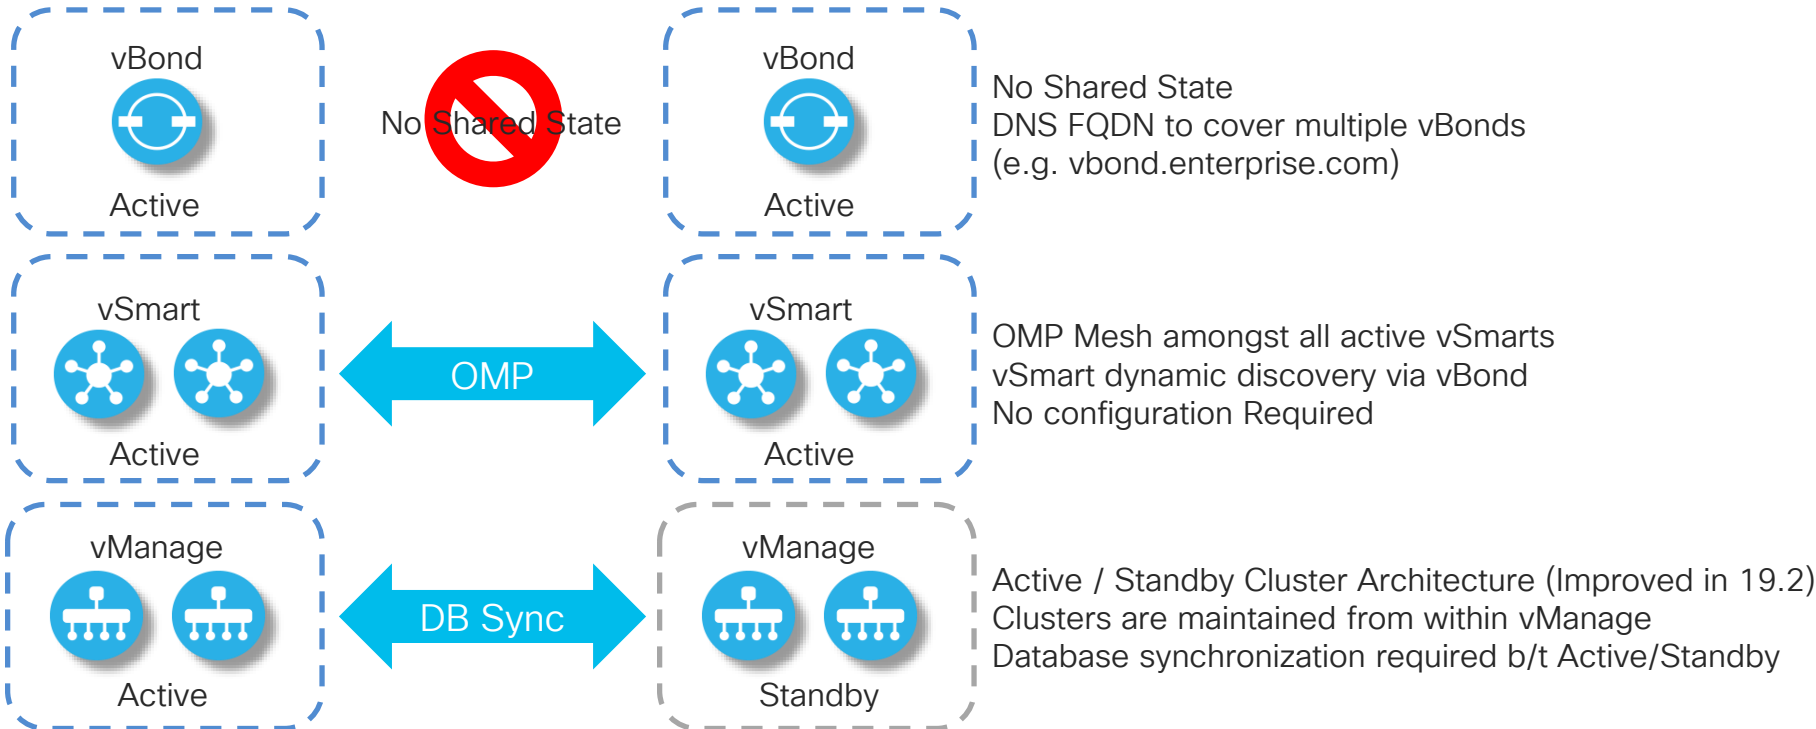
Cisco SD-WAN – Controller Redundancy

SD-WAN Controller Redundancy

High Level Description





Same Principles Apply for Cloud and On-Prem



vManage Controller Redundancy

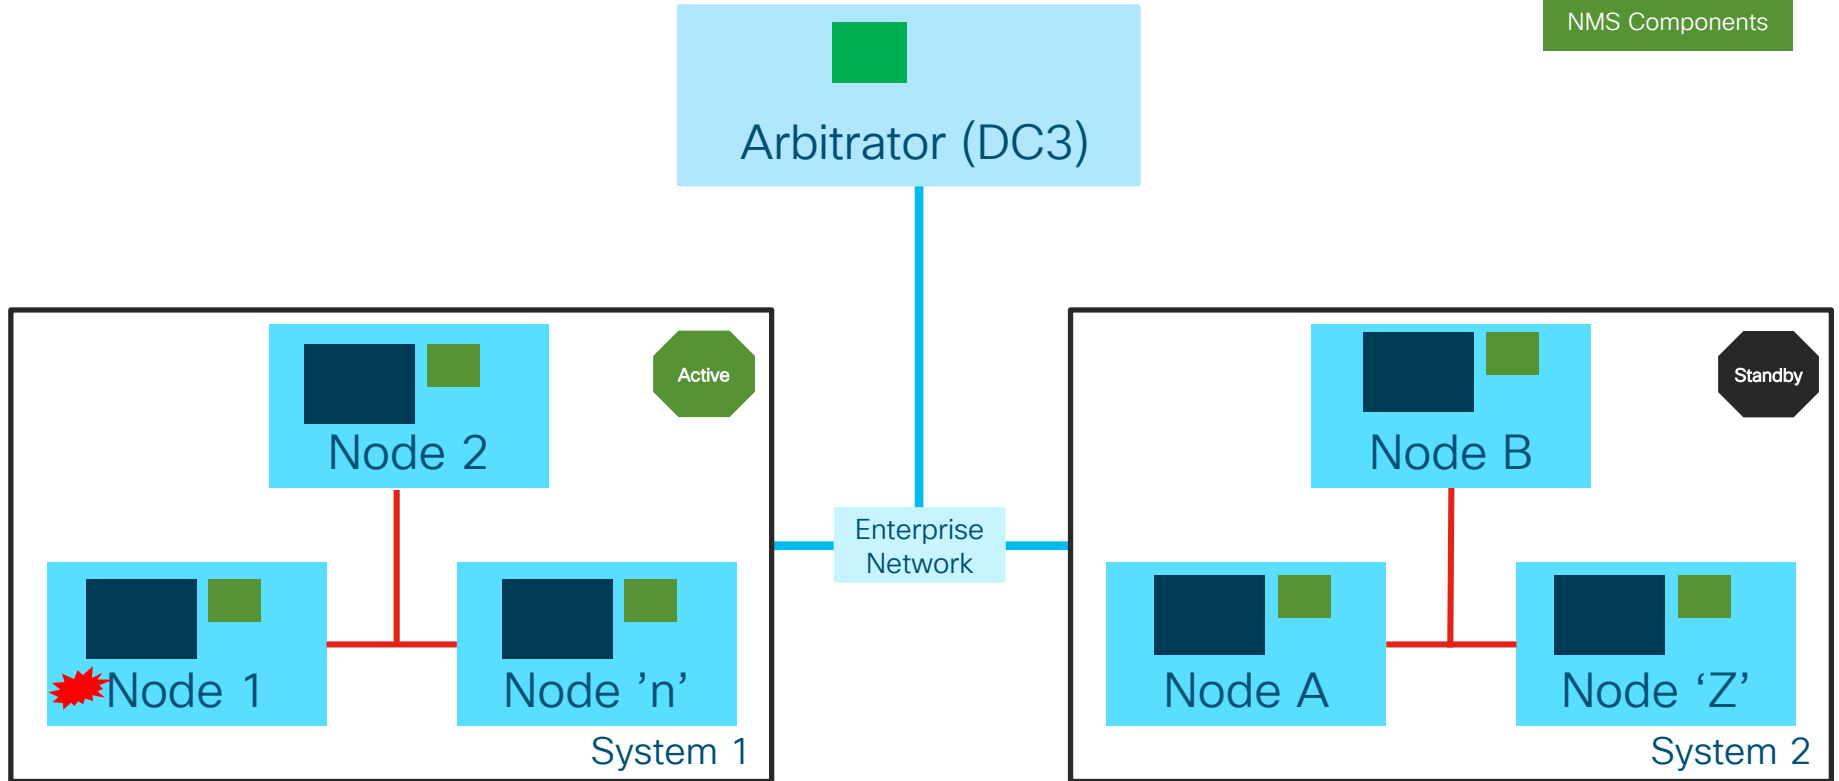

Passive Redundancy (=>R19.2)



- vManage scales horizontally using Clustering
- Add more vManage nodes to cluster in DC for Scale and local HA
- Inter-Cluster Redundancy with DC local Clusters
 - Additional vManage interface used for cluster connectivity

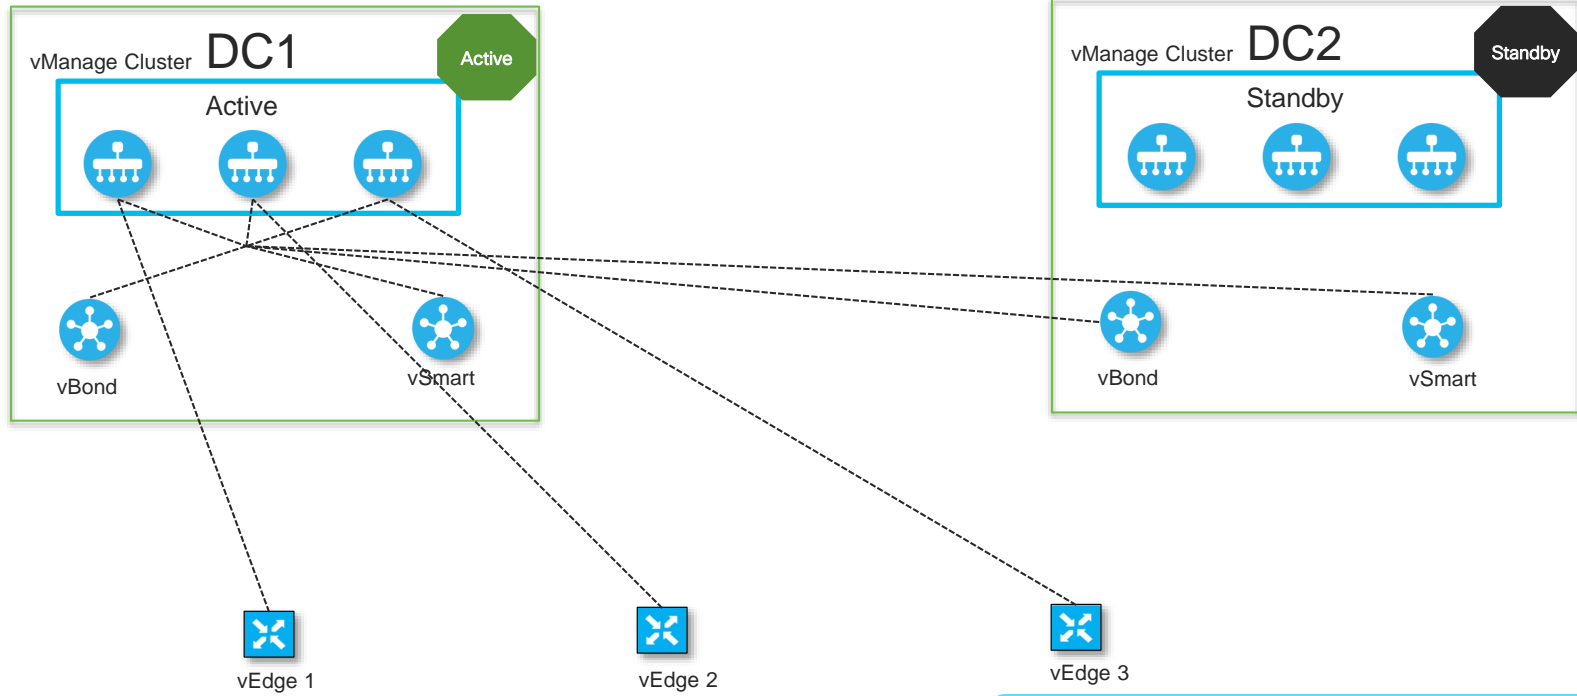
vManage Controller Redundancy

Active Redundancy (=>19.2)



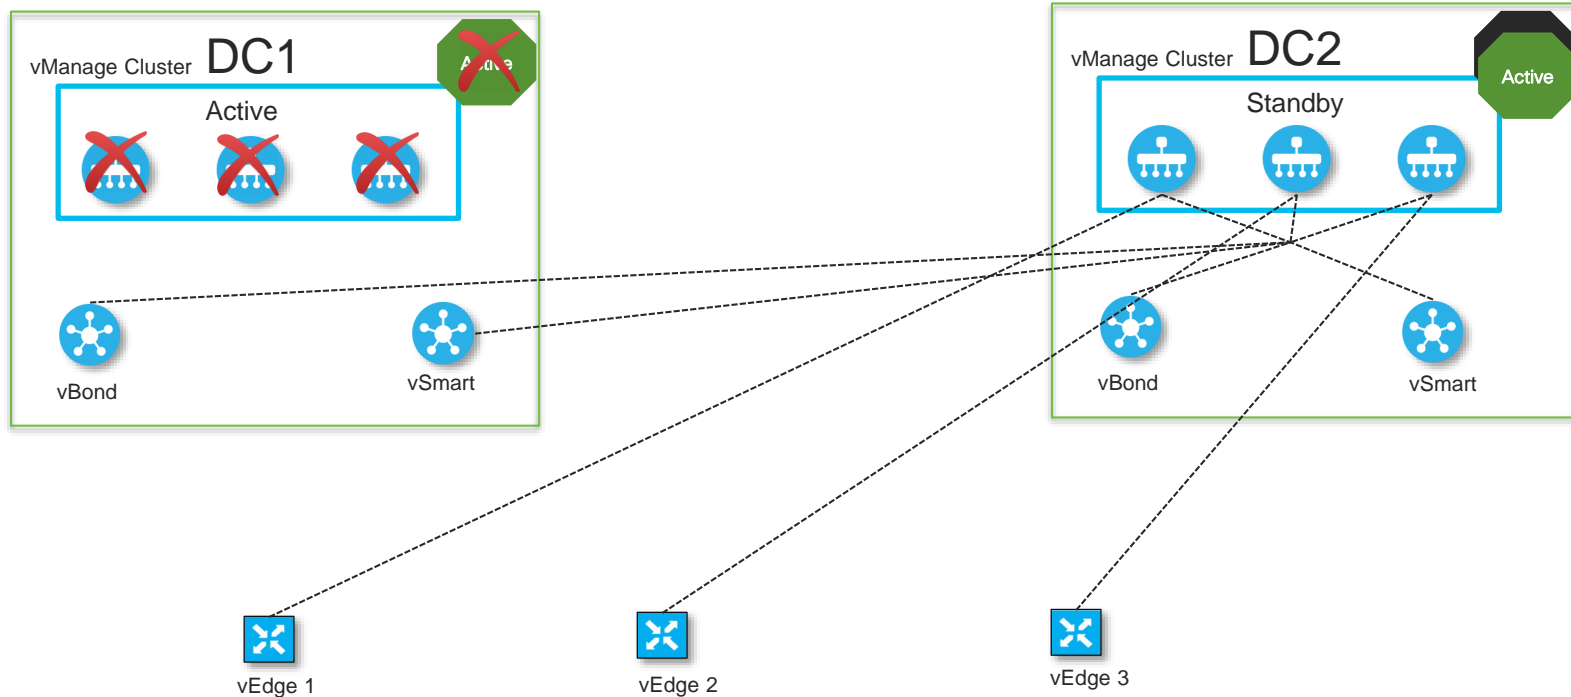
vManage Cluster – Failover

Steady State – Primary Cluster Operational




vManage Cluster – Failover

Post Failover State– Standby Cluster Operational



Passive vManage Cluster Redundancy

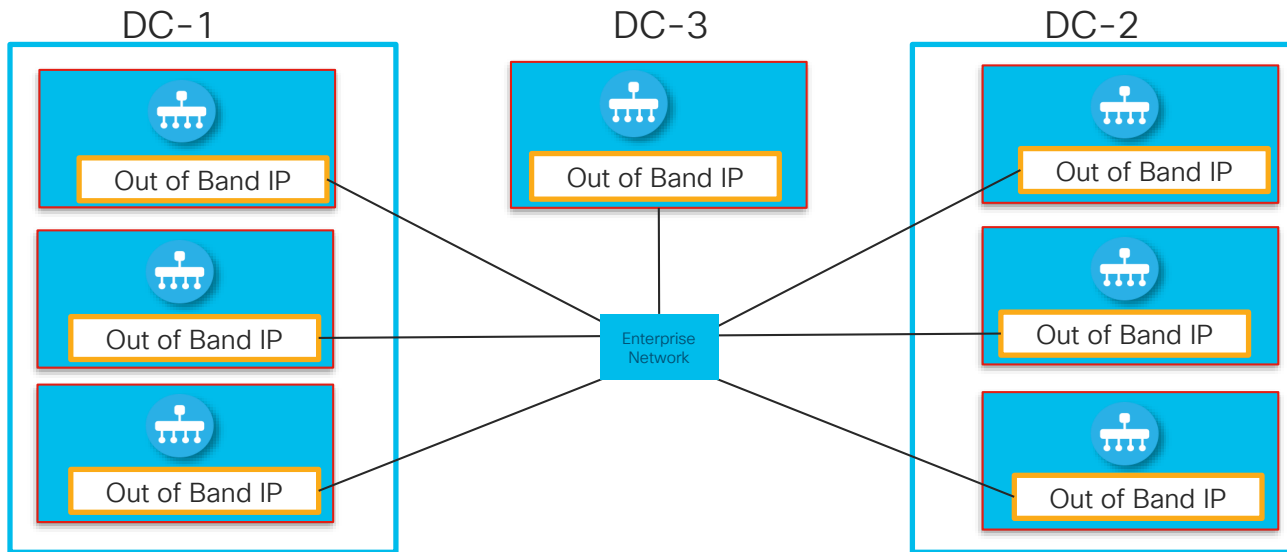
Database Synchronization



vManage Service	Traffic Direction	Protocol	Port
Application Server	Bidirectional	TCP	443
Netconf	Bidirectional	TCP	830

Active vManage Cluster Redundancy

Database Synchronization and Cluster Monitoring



vManage Service	Traffic Direction	Protocol	Port
Application Server	Bidirectional	TCP	443
Netconf	Bidirectional	TCP	830
Consul	Bidirectional	TCP	18600, 18500, 18501, 18301, 18302, 18300

vManage Redundancy Configuration

Configure Standby Mode and DB Synchronization

The screenshot displays the vManage Administration | Disaster Recovery interface. The main content area is titled "Cluster Status" and shows two clusters. Cluster 1 is active and has four nodes, with Node 2 and Node 3 in a failed state. Cluster 2 is in standby mode and also has four nodes, with Node 2 and Node 3 in a failed state. A "Make Active" button is visible next to Cluster 2. To the right, a "Details" panel provides information about the last replication, time to replicate, size of data, last switch, reason for switch, start time, replication interval, and delay threshold.

Cluster Status

Cluster 1 (Active: IP Address)

Node	IP Address	State
Node 1	172.22.112.22	✓
Node 2	172.22.112.22	✗
Node 3	172.22.112.22	✗
Node 3	172.22.112.22	✓
Node 4	172.22.112.22	✓

Cluster 2 [Make Active](#)

Node	IP Address	State
Node 1	172.22.112.22	✓
Node 2	172.22.112.22	✗
Node 3	172.22.112.22	✗
Node 3	172.22.112.22	✓
Node 4	172.22.112.22	✓

Arbitrator

Details

Last Replicated: 04 Oct 2018 10:00:06 AM PDT

Time to Replicated: 04 Oct 2018 10:00:06 AM PDT

Size of Data: 2000 gb

History

Last Switch: 04 Oct 2018 10:00:06 AM PDT

Reason for Switch: Auto (Node Down / Control Down)

Schedule

Start Time: 12 AM

Replication Interval: 15 mins

Delay Threshold: 0 mins

Cisco SD-WAN The Control Plane

Overlay Management Protocol

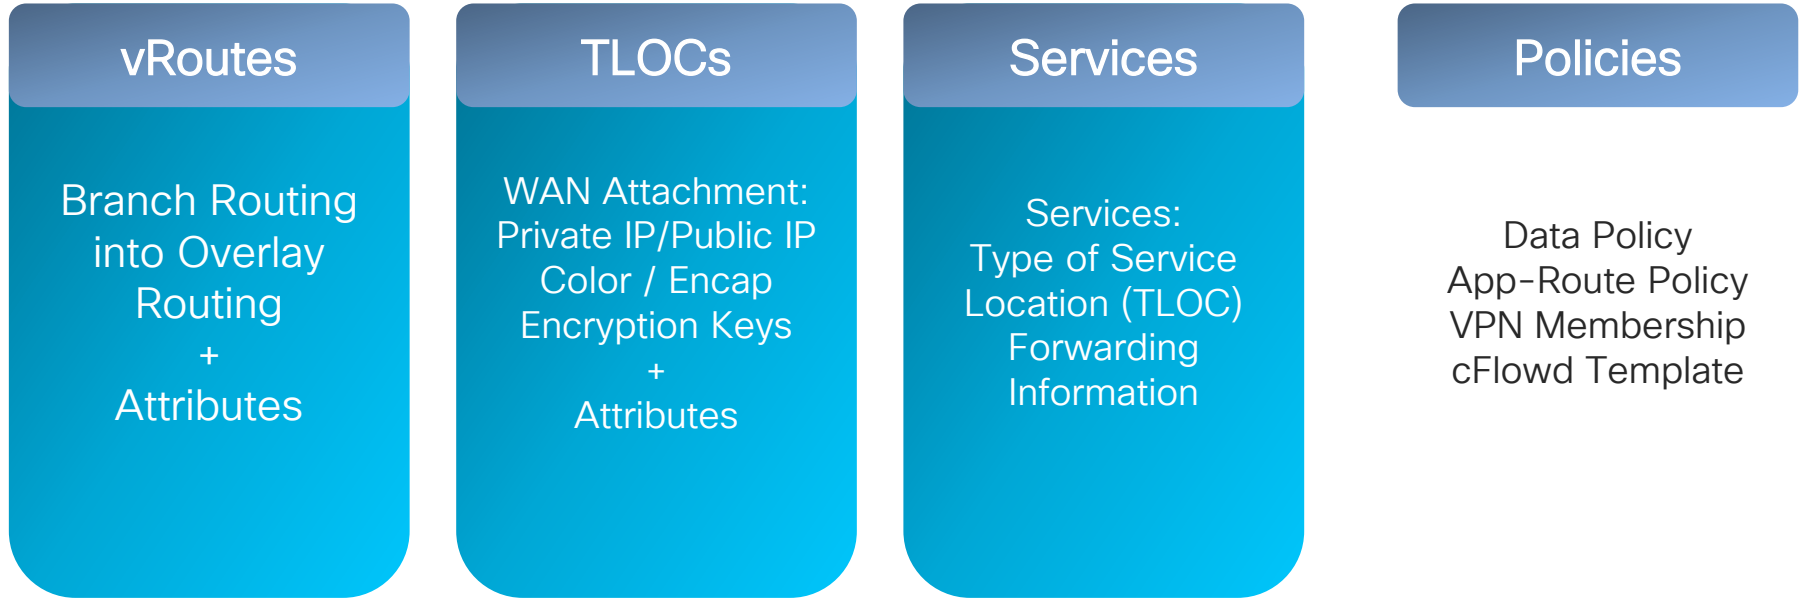
High Level Description

- Path Vector Routing Protocol specifically designed for overlay networks
- Natively Multiprotocol, Multipath and VPN/Segment Aware
- Peer Auto-discovery w/ Zero line config for basic operation
- Inherent Route Target Constraint Capability
- Automatic Distribution of targeted local routing
- Overlay and Legacy Domain Loop Avoidance capabilities
- Reliable and Secure Transport (SSL)
- Broad Attribute Support
 - Preference
 - Identification
 - Legacy Source Protocol Information
- Consistent Routing and Encryption Synchronization
- Multi-domain capable



Overlay Management Protocol

Distribution of Routing Information for Topology-driven Routing

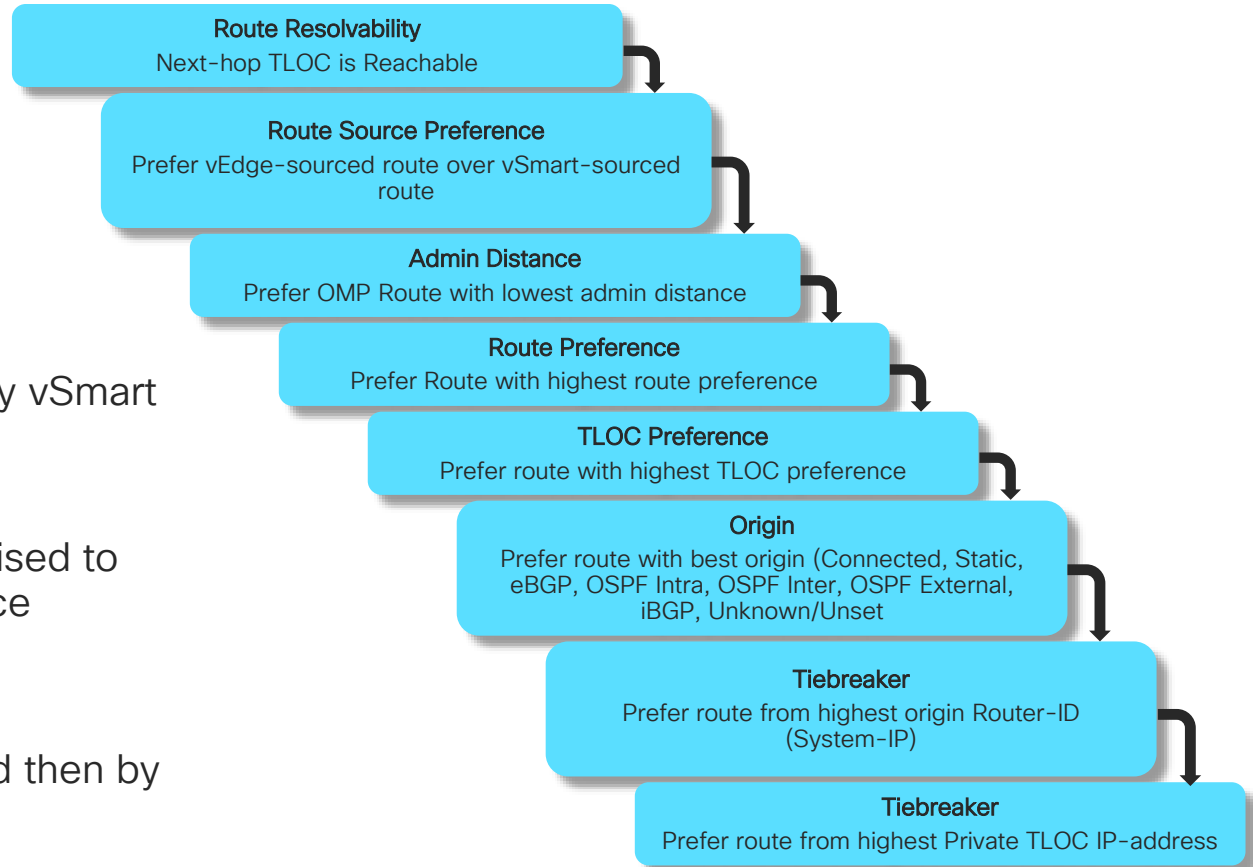


Distribution of Routing Information and Policies subject to endpoint push

Updates sent only on changes – Routing engine operates as with existing protocols (BGP)

Overlay Management Protocol

Path Selection



- Default: 4 paths advertised by vSmart omp

```
Send-path-limit [1-16]
```

- Backup routes can be advertised to vEdges for faster convergence


omp

```
Send-backup-paths
```

- Origin by Admin Distance and then by Protocol Cost / Metric

Dynamic Routing for VPN Segments

OMP Overlay Routing in relation to local Routing



Global: Generic setting for node
VPN Specific: Overrides Global

Node Global OMP Configuration

```
omp
no shutdown
graceful-restart
advertise bgp
advertise connected (default)
advertise ospf external
advertise static (default)
```

VPN Specific OMP Configuration

```
vpn 1
omp
advertise aggregate <prefix>
advertise bgp
advertise connected
advertise network <prefix>
advertise ospf external
advertise static
```

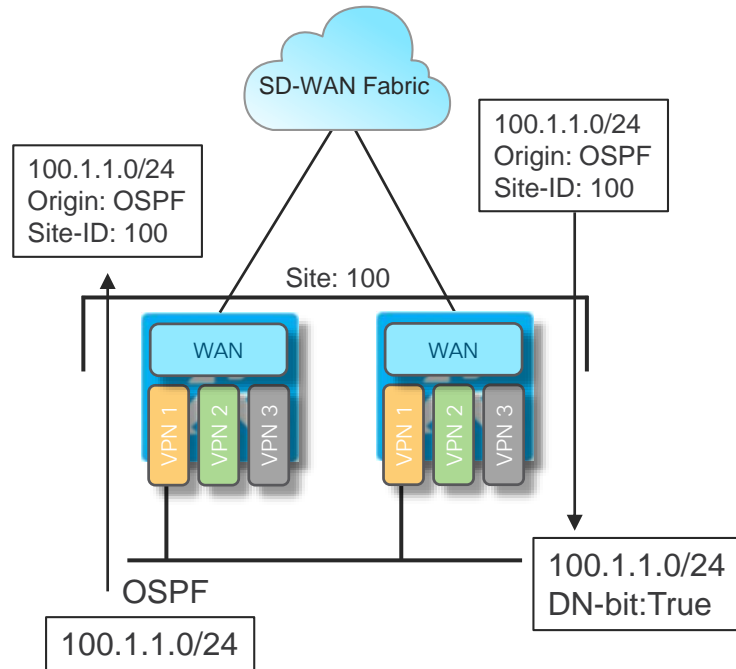
```
vpn 1
router
ospf
redistribute omp
area 0
interface ge0/3
exit
exit
!
```

```
vpn 3
router
bgp 123
address-family ipv4-unicast
redistribute omp
!
neighbor 1.1.1.1
no shutdown
remote-as 321
!
```

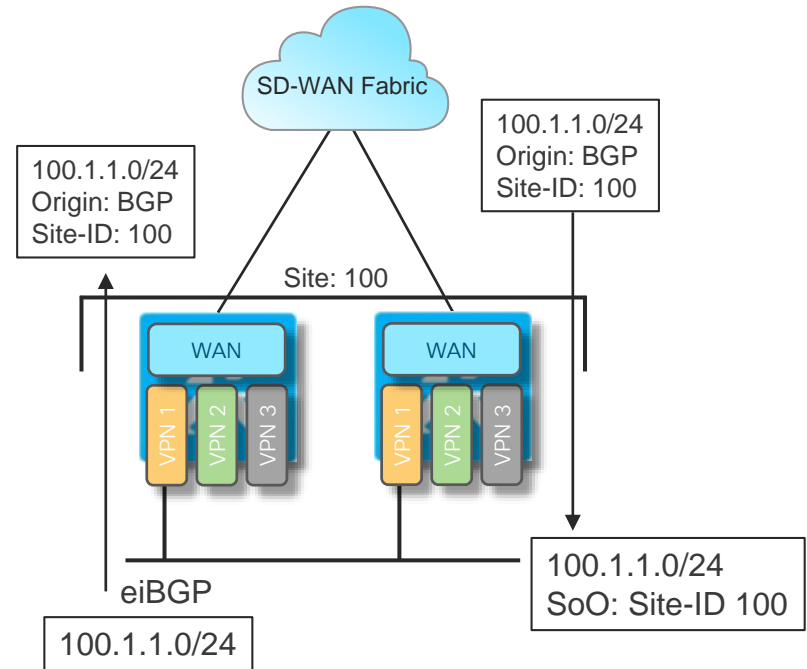
No redistribution of OMP into OSPF/BGP is enabled by default

SD-WAN Overlay Routing

LAN Side Routing Loop Avoidance



OSPF:
Set DN-bit for any redistributed prefix from OMP to OSPF

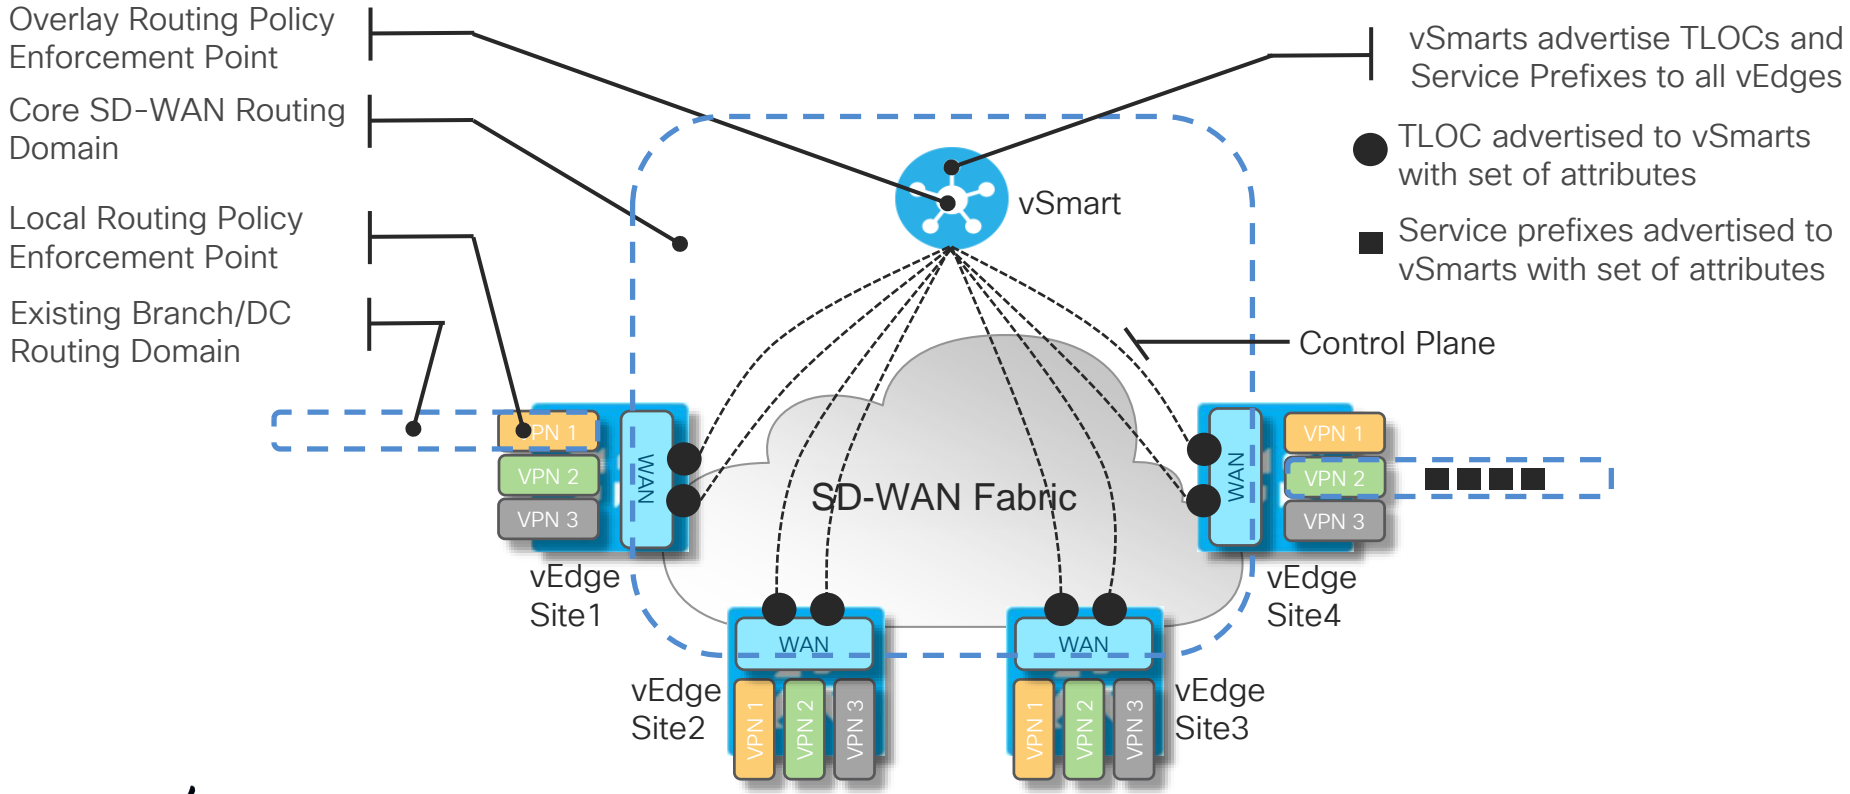


BGP:
Set SoO to indicate Site-of-Origin for OMP to BGP

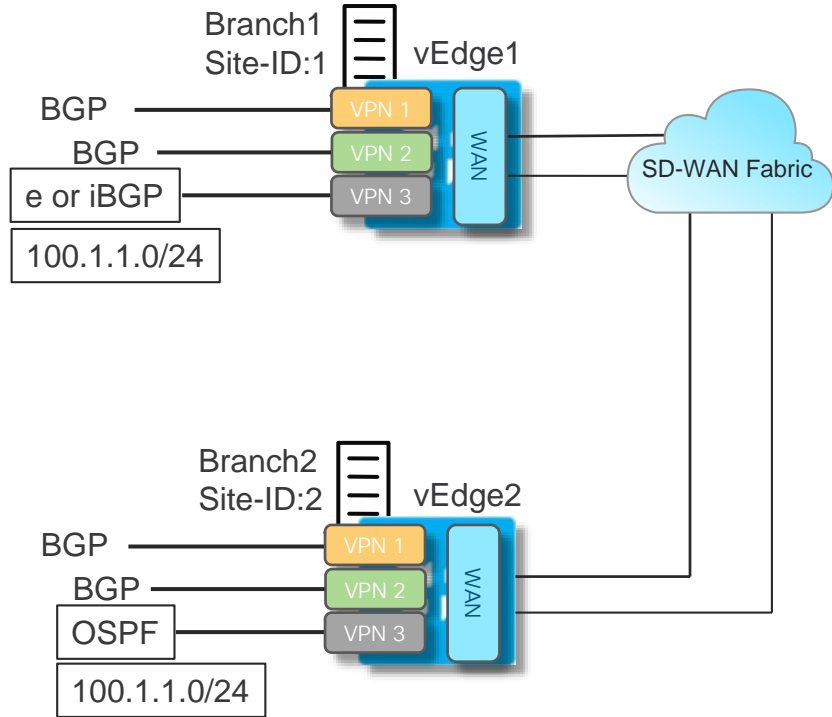
Cisco SD-WAN Overlay Routing - Design Considerations

Cisco SD-WAN Overlay Routing

Multi-domain Routing Fabric



Overlay Routing – Path Selection



vSmart



Vpn	Prefix	Peer	Status
3	100.1.1.0/24	vEdge?	C,R

iBGP vs OSPF

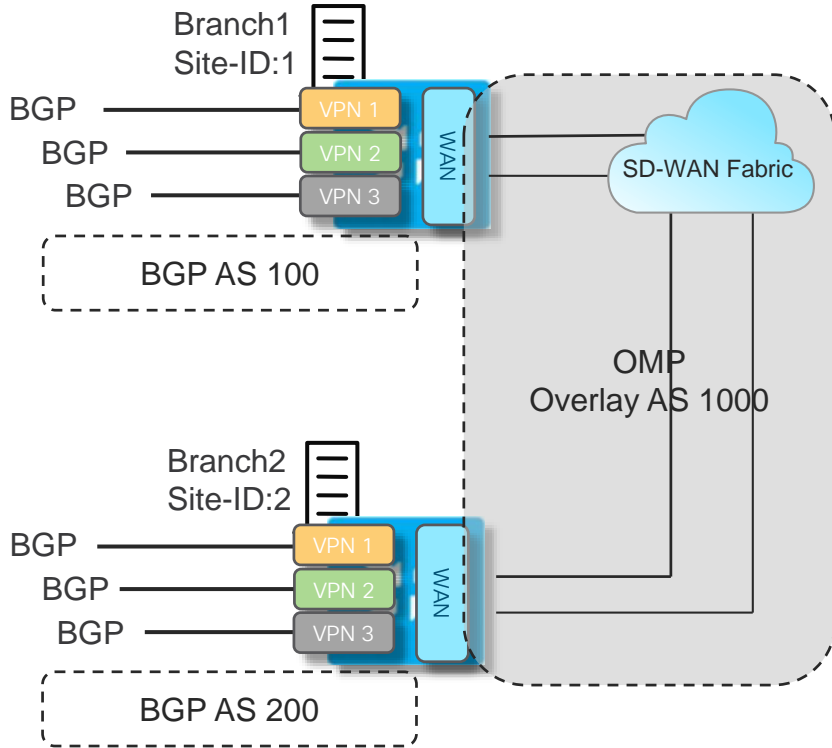
Vpn	Prefix	Peer	Status
3	100.1.1.0/24	vEdge2	C,R

eBGP vs OSPF

Vpn	Prefix	Peer	Status
3	100.1.1.0/24	vEdge1	C,R

SD-WAN Overlay Routing

Overlay AS / Propagate AS



```
omp
overlay-as 1000
```

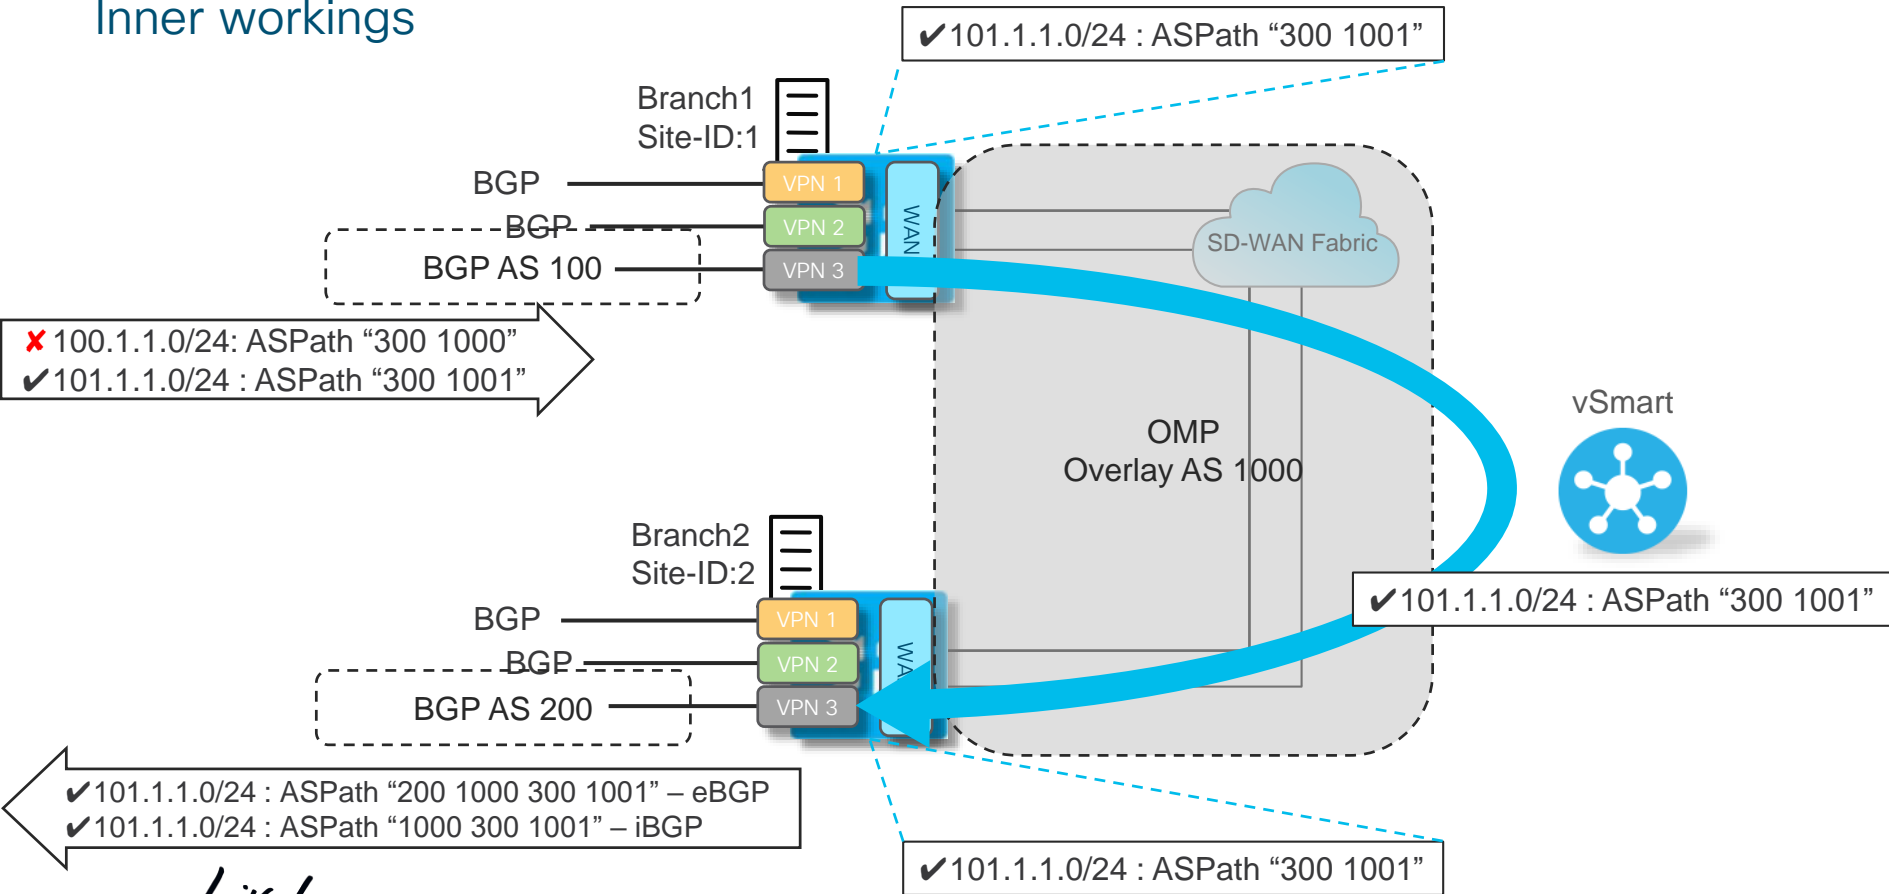
- Overlay AS used to assign an AS to the OMP Domain
- OMP AS is inserted into the AS-path when advertised downstream (iBGP + eBGP)
- Effective tool to ensure loop prevention by using existing BGP behavior

```
vpn 1
router bgp <as>
propagate-aspath
```

- Propagate-aspath enables propagation of AS-path for BGP prefixes across OMP domain

Overlay Routing – Overlay AS / Propagate AS

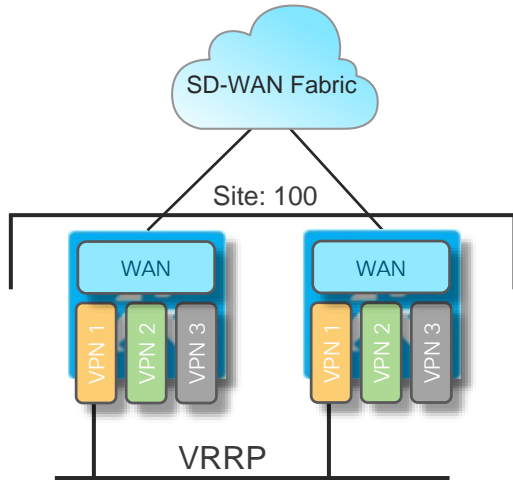
Inner workings



Site Design

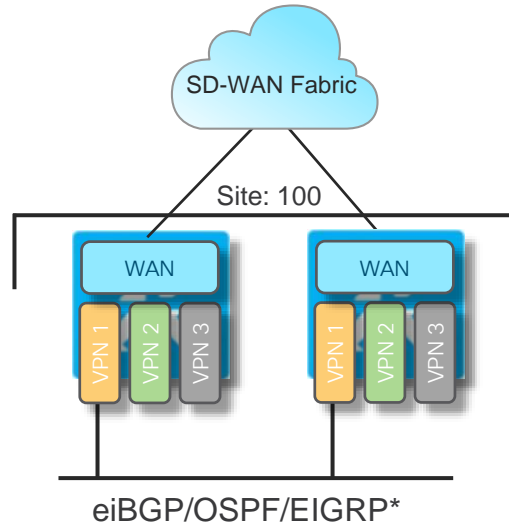
Site Redundancy

Bridged vs Routed LAN



- Layer 2 Site uses VRRP per Segment
- RFC5798 Compliant
- OMP and Prefix upstream tracking


- Layer 3 Site uses L3 dynamic routing per Segment
- *EIGRP is IOS-XE Only
- Mutual OMP/LAN Distribution must be enabled if required



- Loop Avoidance via OSPF DN or BGP SoO
- Multi-pathing as provided by Routing
- WAN Multi-pathing enabled by default
- Site-ID should be identical across same site routers
- Control Policy to manage WAN path priority
- Local Policy for LAN Routing requirements

Site Redundancy

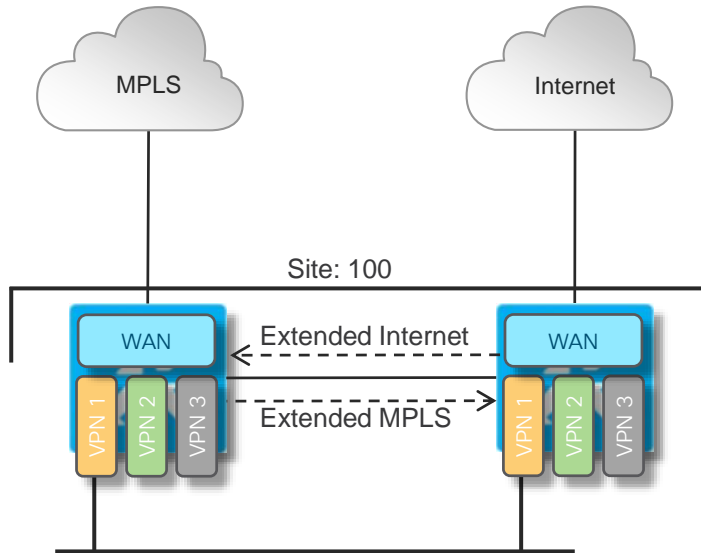
Meshed WAN Transport Attach



- All Transports directly attached to every node in branch
- Each node manages WAN control and data plane independently
- Only simple static default routing required for each transport
- WAN Dynamic Routing also supported is required (usually for certain MPLS providers)

Site Redundancy

Extended WAN Transport Attach aka TLOC Extension




- Directly attached Transport extended to neighboring node
- Each node manages WAN control and data plane independently
- Routing through neighbor for extended TLOC
- L2 or L3 (IOS-XE Only) extension
 - vEdge: L2
 - cEdge: L2 or L3 (GRE)
- Discrete links or 802.1Q can be used between neighbors

Site Redundancy using TLOC Extension

Configuration Example

```
vpn 0
interface ge0/0
  tunnel-interface
  color mpls
  !
  !
interface ge0/1.10
  tunnel-interface
  color biz-internet
  !
  !
interface ge0/1.20
  tloc-extension ge0/0
  !
ip route 0.0.0.0/0 <nhop on ge0/0>
ip route 0.0.0.0/0 <nhop on ge0/1.10>
```



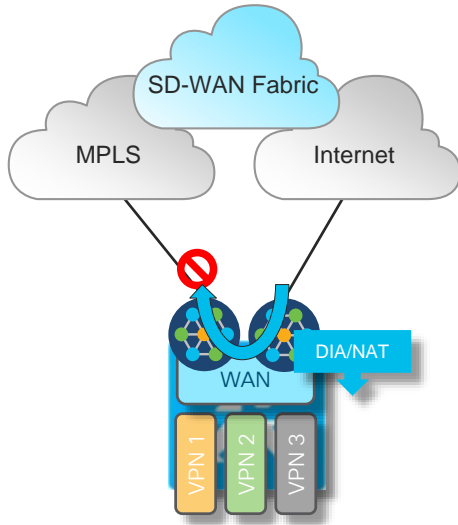
```
vpn 0
interface ge0/0
  nat
  tunnel-interface
  color biz-internet
  !
  !
interface ge0/1.10
  tloc-extension ge0/0
  !
  !
interface ge0/1.20
  tunnel-interface
  color mpls
  !
  !
ip route 0.0.0.0/0 <nhop on ge0/0>
ip route 0.0.0.0/0 <nhop on ge0/1.20>
```

Return traffic is handled by advertising the TLOC subnet (ge0/1.10) to MPLS

Return traffic is handled by NAT on ge0/0

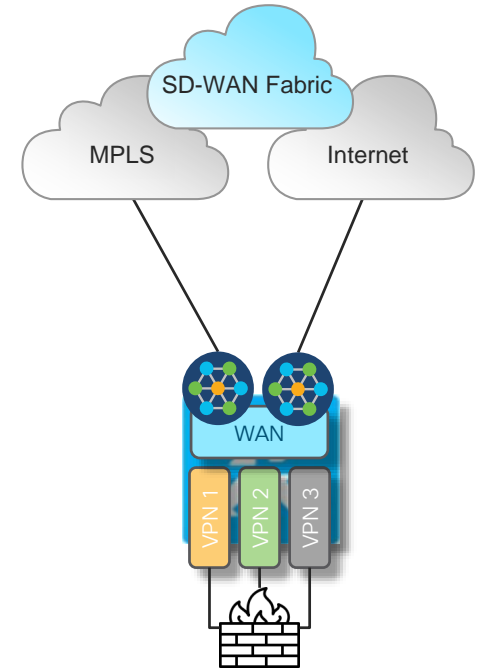
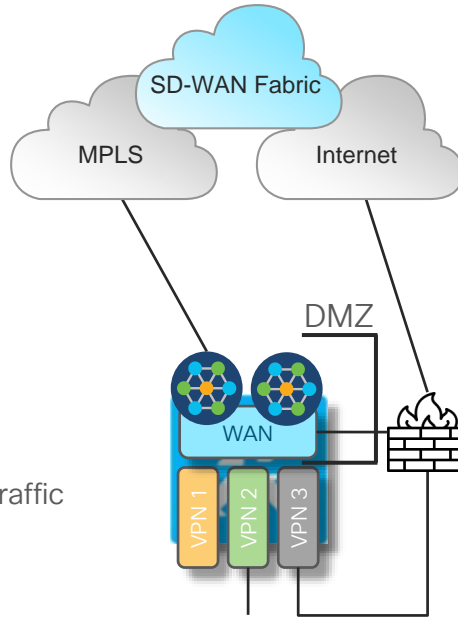
SD-WAN Edge Security Site Design

Various Integration Options



External Firewall for Internet Circuits

- FW provides NAT / vEdge does NAT traversal
- FW could be In / Offline for internal connectivity



Fabric Security

- Secures TLOC Attachment
- Prevents vEdge Transit / Inter-TLOC Traffic



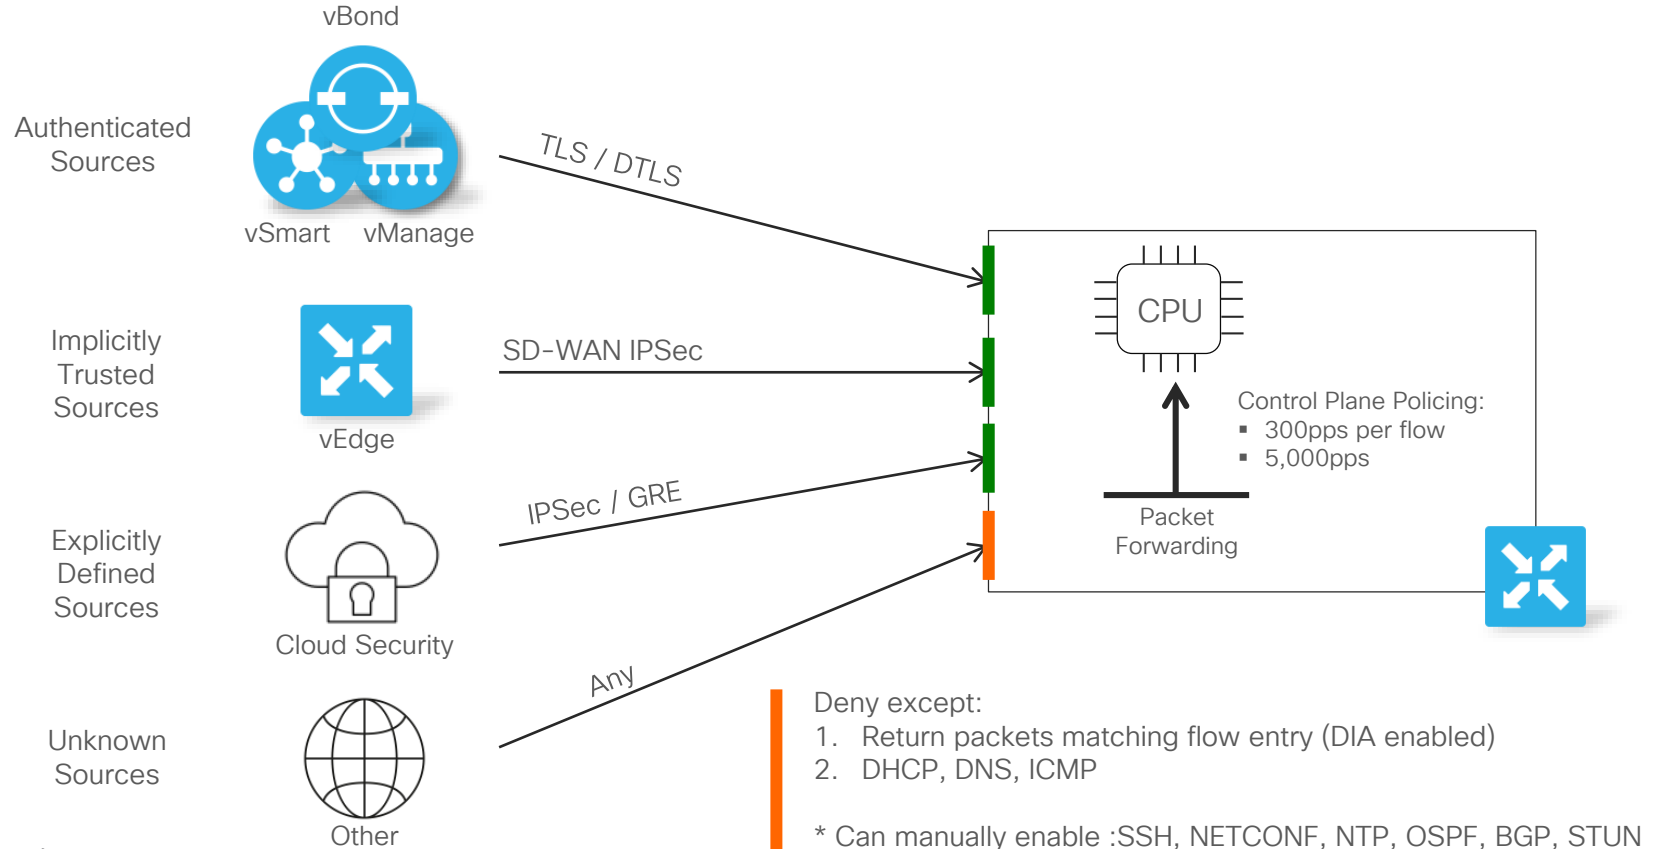
Local vEdge NAT for DIA

- Port Restricted NAT
- vEdge Zone Based Firewall Protection

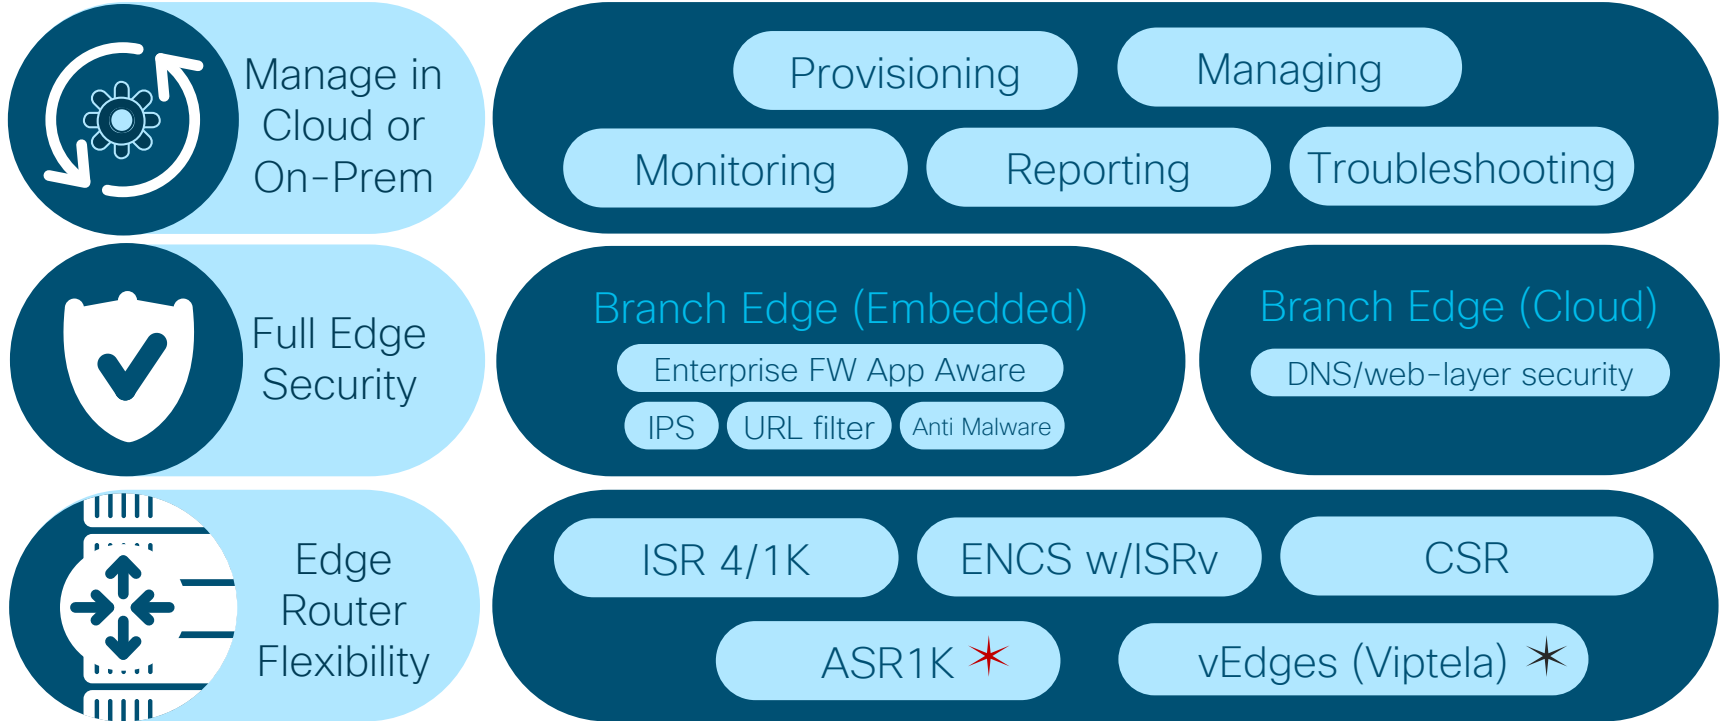
Service-side Firewall

- Built-in ZBFW / App-Layer FW equivalent
- External FW Securing the internal network

WAN Edge Fabric Security Capabilities



SD-WAN Security



LTE Design and Deployment Options

LTE Design and Deployment Options

- Active LTE as a standard TLOC
 - Reduce control and management traffic
- Active LTE as a low bandwidth circuit
 - Reduce control/management traffic and synchronize for quite period to meet M2M requirements
- LTE as a gateway of Last Resort
 - LTE circuit is down and brought up when all other transports are down
- LTE as the ONLY circuit
 - Reduce control and management traffic plus disable many stats collection from vManage

Generic Guidelines for LTE

Reduction in control traffic

1. Control Plane Hello Interval

`vpn 0`

`interface <name>`

`tunnel-interface hello-interval <ms>`

Range: 100-600000 milliseconds (10 minutes)

Default: 1000 milliseconds (1 second)

2. BFD hello timer can be extended to 5 minutes

`bfd`

`color <color> hello-interval <ms>`

Range: 100-300000 milliseconds (5 minutes)

Default: 1000 milliseconds (1 second)

3. Ensure LTE is not preferred for vManage communication

`vpn 0`

`interface <name>`

`tunnel-interface`

`vmanage-connection-preference <n>`

Range: 0 through 8

Default: 5

For an LTE-only device

Ensure vManage collection of (high-volume) statistics is disabled

LTE Operation

Last Resort Circuit

- LTE interface is kept down when other transport interfaces are up

```
vpn 0
```

```
interface <name>
```

```
tunnel-interface
```

```
last-resort-circuit
```

- Convergence time - Bring-up triggered when all other WAN links are down
 - Time for bfd failure detection on other transports [7 sec]
 - Time to bring LTE up and operational (IP assigned)
 - Bring up control connections on LTE [20-30 sec]
 - Bring up BFD/IPSec connections [5-10 sec]

LTE Operation

Low-Bandwidth Circuit

- LTE interface synchronizes the control and BFD traffic
- Provides long quiet periods to meet M2M requirements
- LTE circuit releases the channel during that time

[vpn](#) 0

[interface](#) <name>

[tunnel-interface](#)


[low-bandwidth-link](#)

- This configuration command is relevant only for a spoke vEdge router in a hub-and-spoke deployment scenario
- On the spoke router only
- on such links, application-aware routing data is collected only when user data is transmitted from the LAN to the WAN, to reduce BFD traffic on the link

Cisco SD-WAN Data Plane - Design Considerations

Cisco SD-WAN Data Plane


Default Mode of Operation




- Assuming Interconnected underlays, full mesh between all TLOCs
- Nodes will attempt to form full mesh by default, regardless of TLOC color

Customer design - Site Types and Connectivity


Data Plane Design Example




Single




Dual




DC/Regional



Single
Redundant



Dual
TLOC-Extension



DC/Regional
TLOC-Extension

BFD / Tunnel Consumption by Site Type

Customer Deployment Example

Site Type	Devices	Internet Links	MPLS Links	TLOCs	TLOC Extension	Colors
Single	1	1	0	1	No	1
Single Redundant	1	1	1	2	No	2
Dual	2	2	2	4	No	4
Dual TLOC Extension	2	2	2	8	Yes	4
DC	2	4	2	6	No	6
DC TLOC Extension	2	4	2	12	Yes	6

Source: Cisco Internal i.e. Stefan

BFD / Tunnel Consumption by Site Type

Consumption with Default Configuration

Site Type	Single	Single Redundant	Dual	Dual TLOC Extension	DC	DC TLOC Extension
Single	1/1	2/2	4/4	8/8	6/6	12/12
Single Redundant	2/2	4/4	8/8	16/16	12/12	24/24
Dual	4/4	8/8	16/16	32/32	24/24	48/48
Dual TLOC Extension	8/8	16/16	32/32	64/64	48/48	96/96
DC	6/6	12/12	24/24	48/48	36/36	72/72
DC TLOC Extension	12/12	24/24	48/48	96/96	72/72	144/144

Source: Cisco Internal i.e. Stefan

SD-WAN Data Plane Design

Tools and Techniques for defining data plane connectivity

Color Restrict

- Limit Data Plane Establishment to TLOCs of the Same Color
- Simple configuration knob that effectively limits data plane connectivity

TLOC Groups (new in 19.1 (vEdge) / 16.11.2 (cEdge))

- Configure groups that determine actual data plane topology
- Can be combined with restrict
- Define groups based on site type and connectivity requirements for scalable rollout

Control Policy

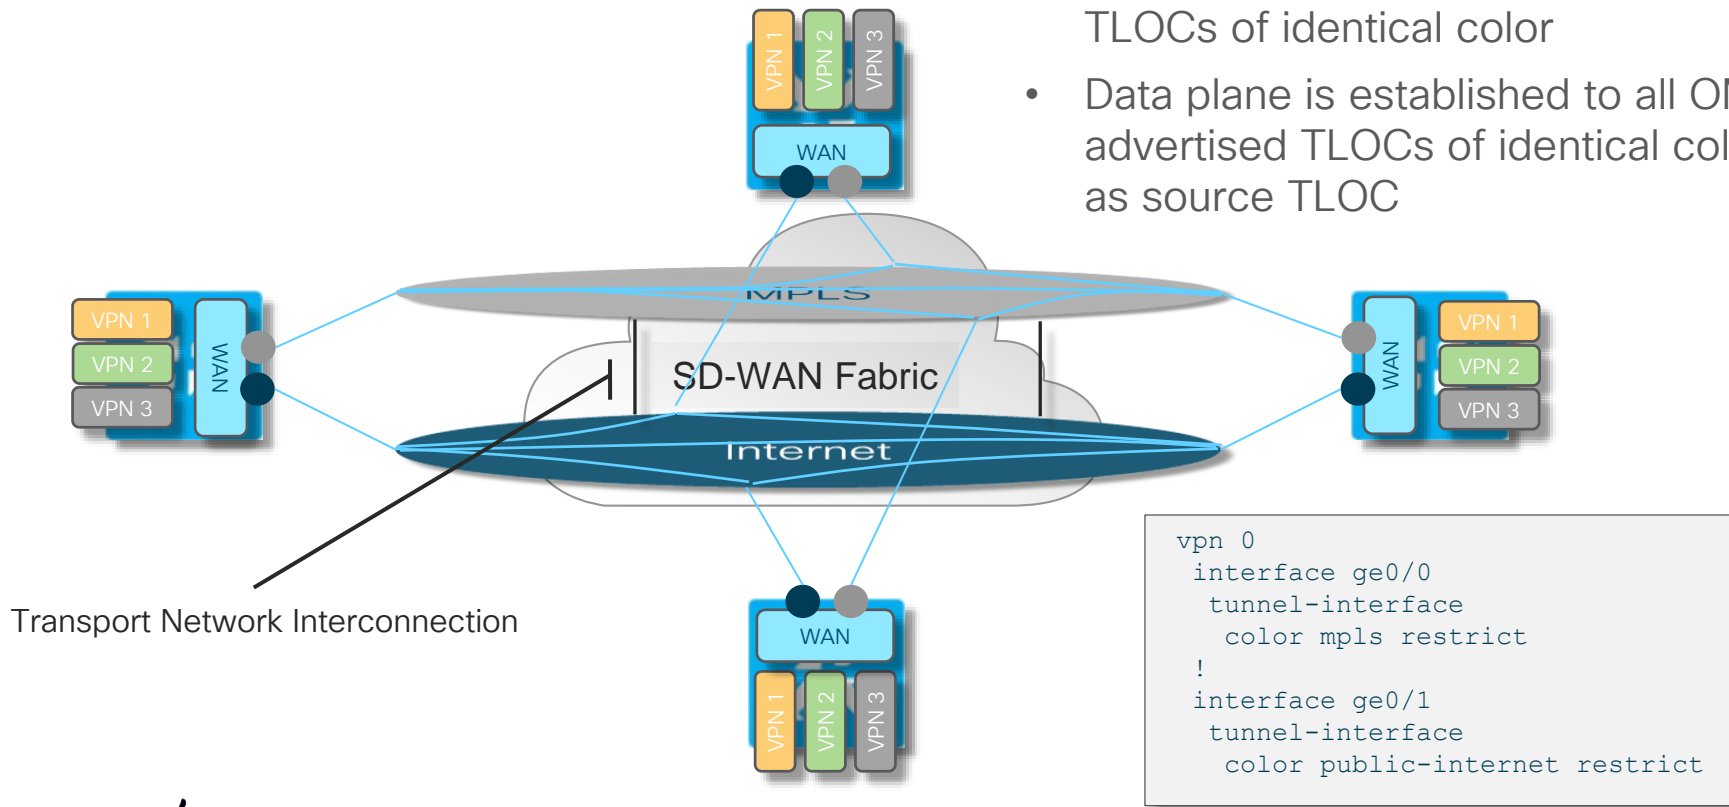
- A control policy allows for TLOC filtering and reassignments
- Can be combined with TLOC Groups and Restrict if needed
- Ultimate control over TLOC distribution, visibility, preference and connectivity model

Cisco SD-WAN Data Plane

Color Restrict



- Limits data plane connectivity to TLOCs of identical color
- Data plane is established to all OMP advertised TLOCs of identical color as source TLOC



BFD / Tunnel Consumption by Site Type

Consumption with Restrict Configuration

Site Type	Single	Single Redundant	Dual	Dual TLOC Extension	DC	DC TLOC Extension
Single	1/1	1/1	1/1	2/2	1/1	2/2
Single Redundant	1/1	2/2	2/2	4/4	2/2	4/4
Dual	1/1	2/2	4/4	8/8	4/4	8/8
Dual TLOC Extension	2/2	4/4	8/8	8/8	8/8	8/8
DC	1/1	2/2	4/4	8/8	6/6	12/12
DC TLOC Extension	2/2	4/4	8/8	8/8	12/12	12/12

Source: Cisco Internal i.e. Stefan

Cisco SD-WAN Data Plane

TLOC Groups

- Limits data plane connectivity to identified group on each TLOC
- Provides the flexibility required to build any model data plane

● MPLS TLOC — Physical Link
● Internet TLOC — Tunnel

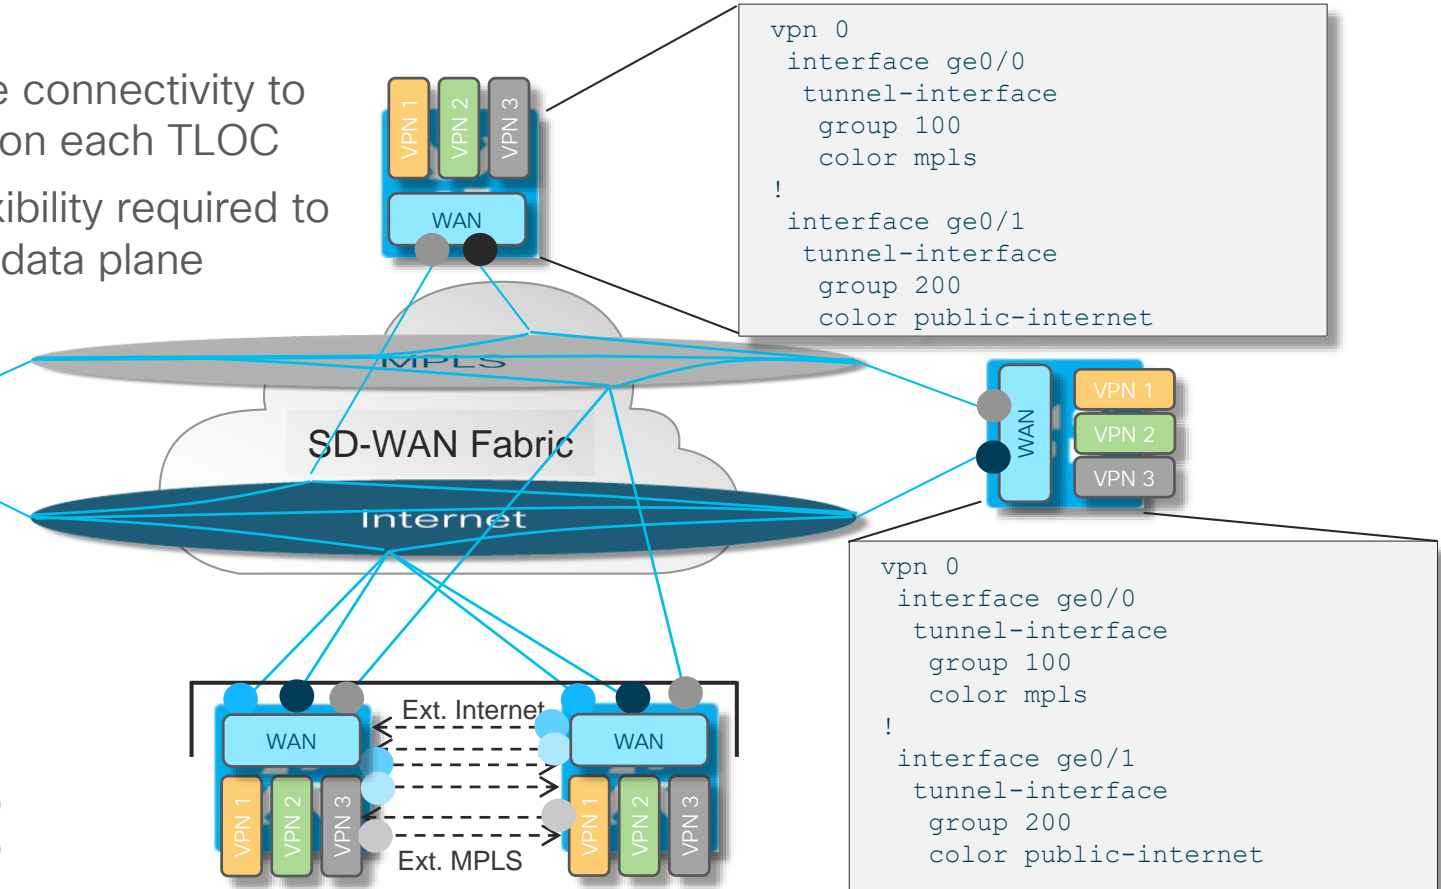
```

vpn 0
 interface ge0/0
  tunnel-interface
  group 100
  color mpls
!
 interface ge0/1
  tunnel-interface
  group 200
  color public-internet
    
```

```

vpn 0
 interface ge0/0
  tunnel-interface
  group 100
  color mpls
!
 interface ge0/1
  tunnel-interface
  group 200
  color public-internet
    
```

- MPLS - Group 100
- Internet - Group 200
- MPLS 2 - Group 100
- Internet 2 - Group 200
- Internet 3 - Group 200




cisco *Live!*


Cisco SD-WAN Data Plane

TLOC Groups and Restrict


Default Behavior – Full Mesh among all colors




MPLS and MPLS 2 with Group 100
Internet and Internet 2 with Group 200




Restrict Configured for all Colors



Restrict for all Colors
MPLS and MPLS 2 with Group 100
Internet and Internet 2 with Group 200
Hub with No Group - Promiscuous



Hub



Cisco SD-WAN Data Plane

TLOC Groups - When to Use

- In a given WAN Edge, a color can be used only once
- This makes for a challenge in networks with multiple MPLS or Internet connections where:
 - Centralized resources are required across branches
 - Those resources aren't connected to all underlay transports and consequently isn't present across all colors
- Restrict becomes too restrictive - limits data plane to single color
- Control Policy Filtering doesn't help overcome a requirement for per-color meshing combined with inter-color connectivity

Cisco SD-WAN Data Plane

Control Policy

- Control Policy can be used to:
 - Filter TLOCs to limit data plane peers per color
 - Re-assign TLOCs for vRoutes to adapt VPN connectivity to data plane
 - Assign priorities to influence path selection
- A Control Policy allows for controlling the distribution of TLOCs across the network along with attribute settings
 - TLOC present means data plane is established
 - TLOC attributes control load-balancing and path selection
- Control Policy Examples coming up

Cisco SD-WAN Data Plane

Using TLOC Attributes

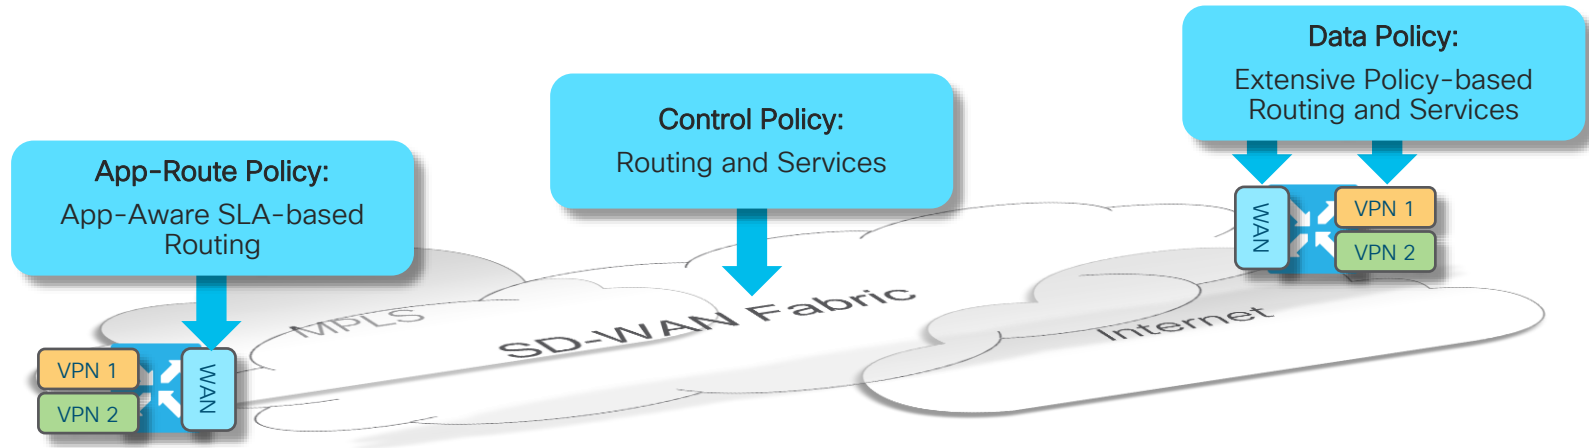
- Every TLOC is advertised with some quite powerful attributes
 - Encapsulation (IPsec and/or GRE) - can be referred to in Policy
 - Weight (Relative bandwidth to other TLOCs on the originating node)
Edge applies local and remote TLOC weight when load-balancing
 - Preference (Used in path selection)

```
vpn 0
 interface ge0/0
  tunnel-interface
  encapsulation gre preference 100 weight 10
  encapsulation ipsec preference 100 weight 10
  !
  !
  !
```

Cisco SD-WAN Policy Framework - Overview

Cisco SD-WAN Policy Architecture

Suite of Policies to address different functional domains



- Control Policies are applied at vSmart: Tailors routing information advertised to vEdge endpoints
- App-Route Policies are applied at vEdge: SLA-driven path selection for applications
- Data Policies are applied at vEdge: Extensive Policy driven routing

Control Policies

Overlay Management Protocol Routing Policies

- Control policies are applied and executed on vSmart to influence routing in the Overlay domain
- Control policies filter or manipulate OMP Routing information to:
 - Enable services
 - Influence path selection
- Control Policies controls the following services:
 - Service Chaining
 - Traffic Engineering
 - Extranet VPNs
 - Service and Path affinity
 - Arbitrary VPN Topologies
 - and more ...
- The Control Policy is one of the centralized and powerful tools in the Cisco SD-WAN toolbox

Data Policies

Policy-driven Routing and Service Enablement

- Data policies:
 - Applied on vSmart
 - Advertised to and executed on WAN Edge
- A Data policy acts on an entire VPN and is not interface-specific
- Data Policies are used to enable the following functions and services:
 - Application Pinning
 - NAT/DIA
 - Classification, Policing and Marking
 - and more ...
- Use a Data Policy for any type of data plane centered traffic management

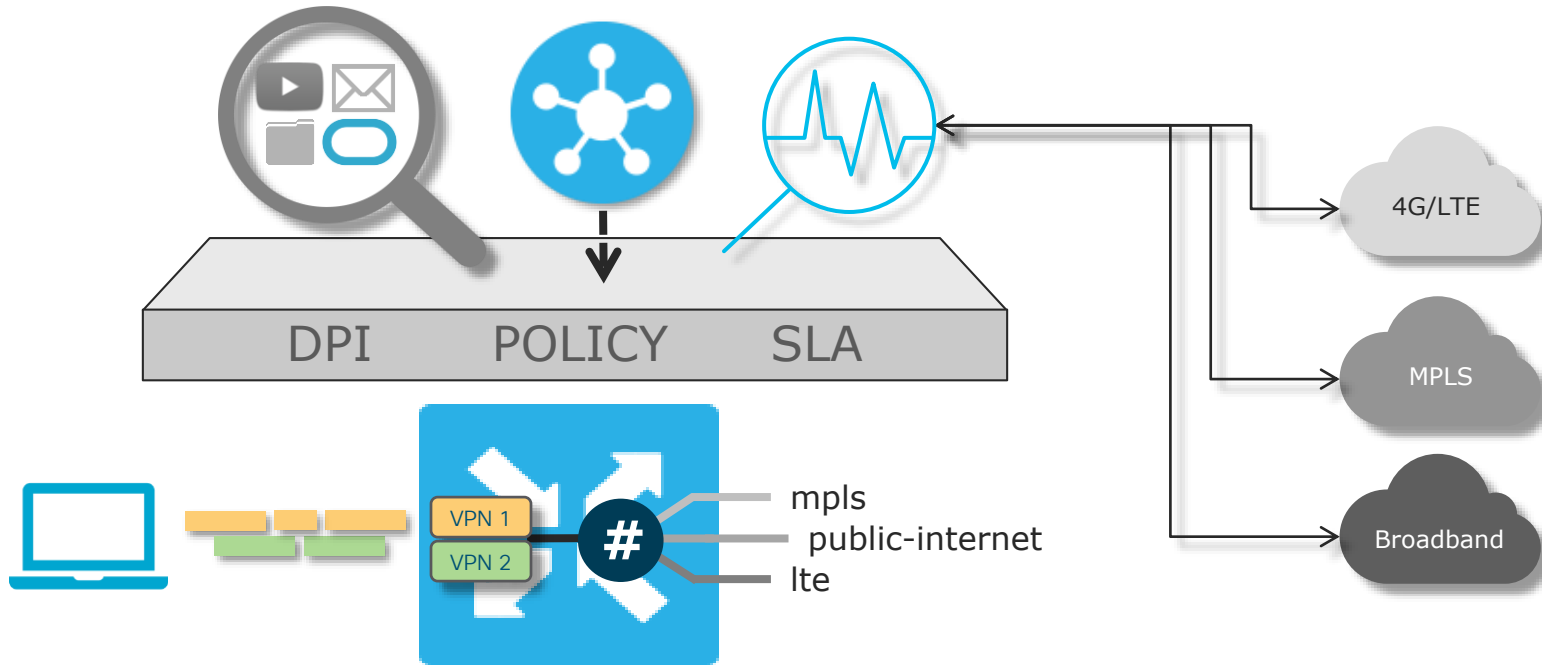
App-Route Policies

Centralized Policy for enabling SLA-driven routing on WAN Edge endpoints

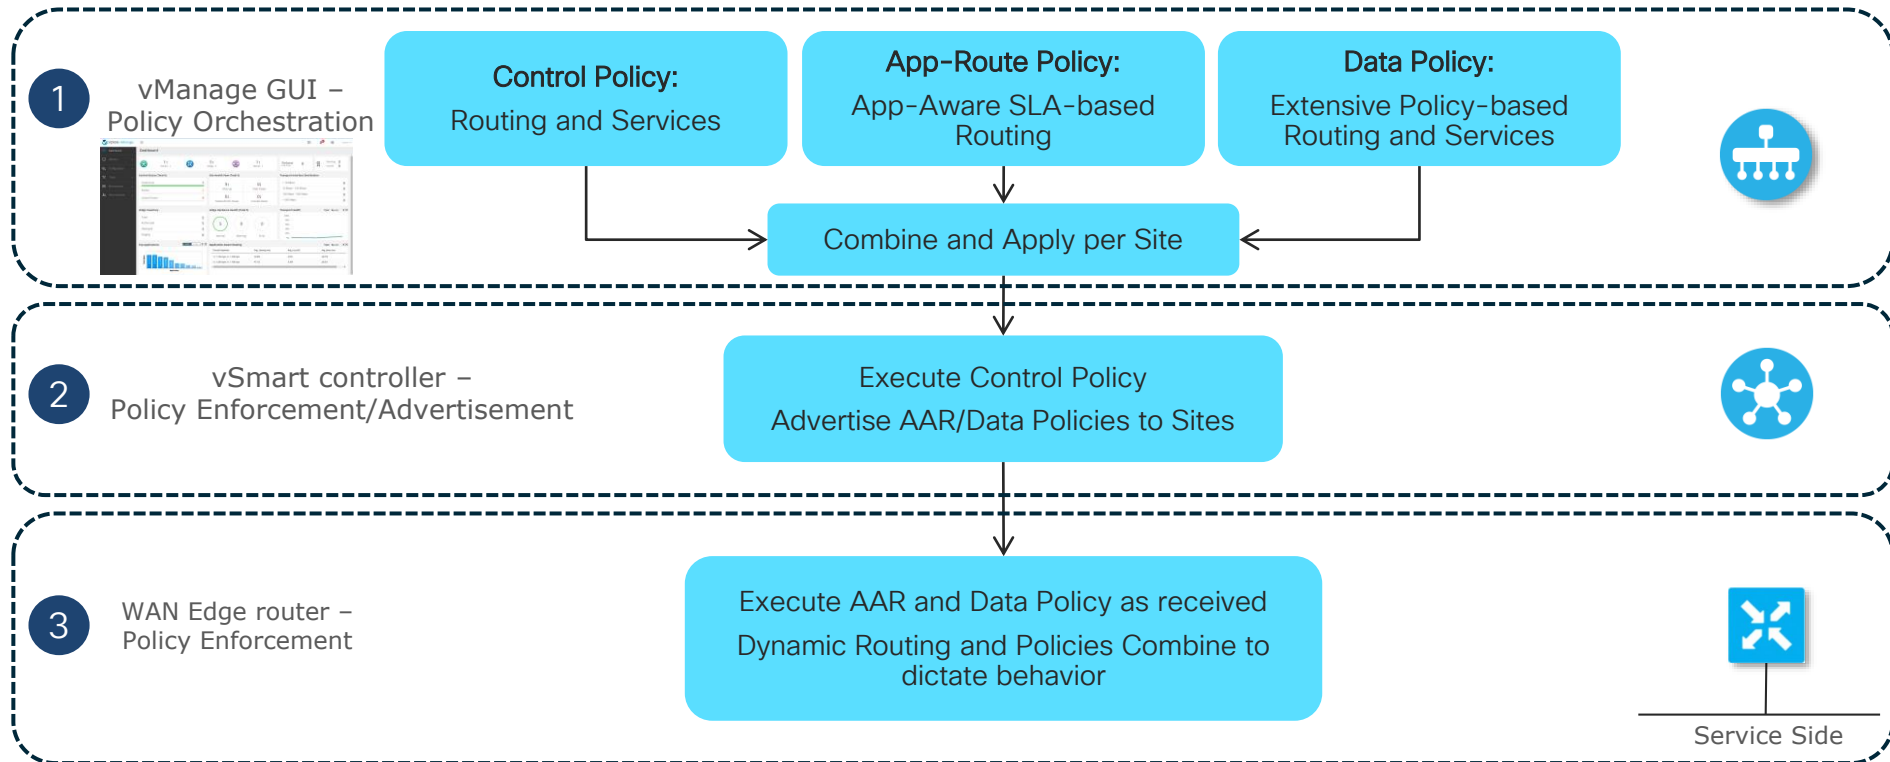
- App-route policies:
 - Applied on vSmart
 - Advertised to and executed on WAN Edge
- Monitors SLAs for active overlay paths to direct Applications along qualified paths
- Allows for the use of L3/L4 keys or DPI Signatures for application identification
- Delivers a fully distributed SLA-driven routing mechanism

App-Aware Routing Policies

SLA-Driven Routing / Performance Routing

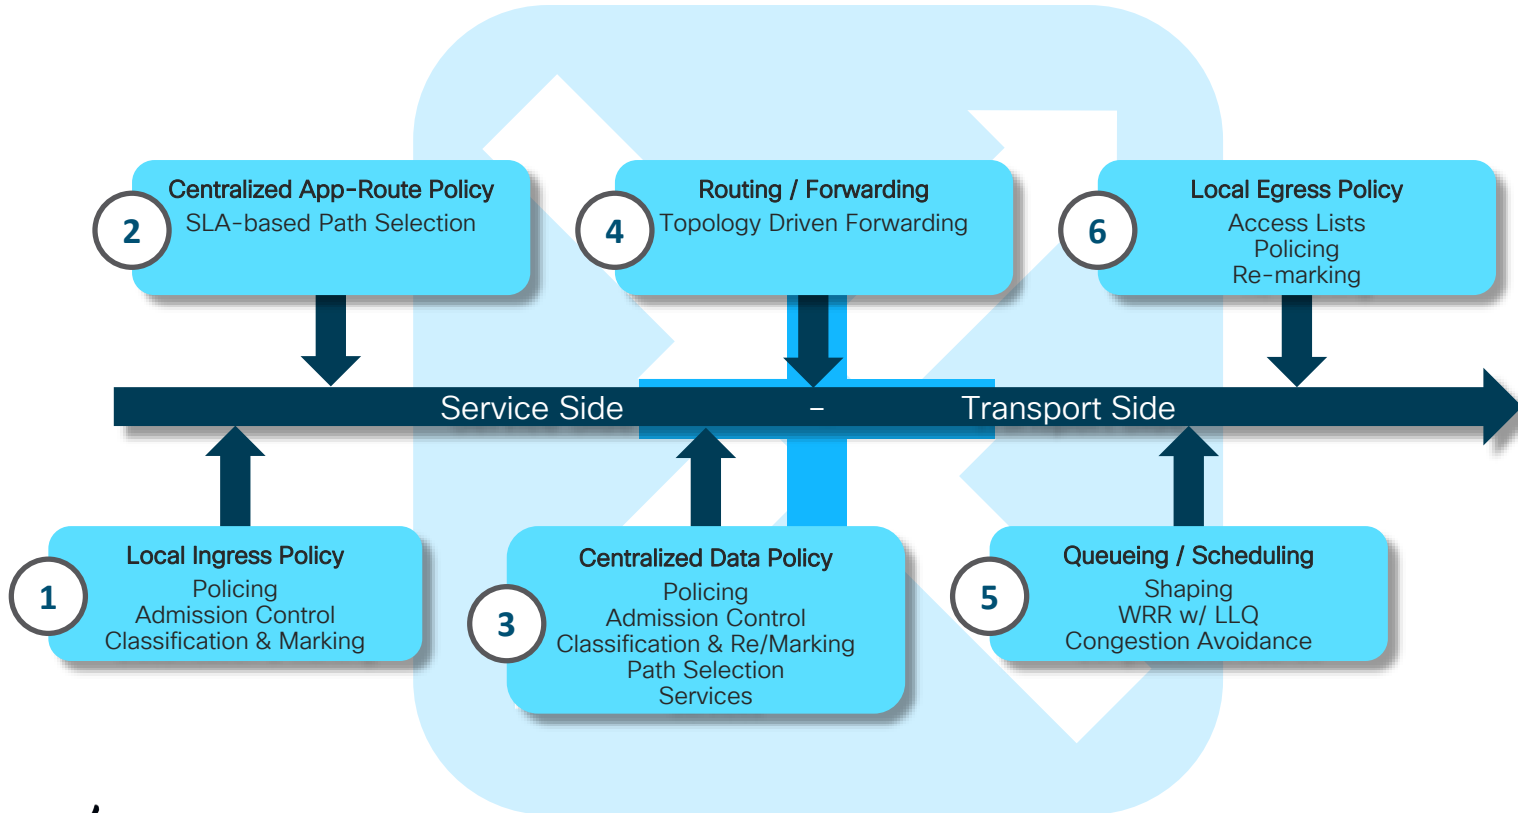


Cisco SD-WAN Policy Orchestration Process



Cisco SD-WAN Policy Execution

Topology-driven routing and Policy execution chain




Cisco SD-WAN Common Overlay Design Options

Common Designs and Services

- Network Topologies and Connectivity Models
- Multiple Underlays with direct/indirect attach
- Primary and Backup path/resource definition
- Dis-contiguous Data Plane / Underlay Transport
- Multi-Domain Overlay with Regional meshing
- Mobile / LTE Attach

Network Topologies

Network Topologies - Review




- A fully meshed Fabric Data plane and Service (VPN) plane is established by default
- This is done on the basis of TLOCs present in the OMP TLOC Table
- Every branch now has every other branch 1 hop away
- VPNs are advertised as an active service from every node
- vSmart applies route distribution constraints based on VPN service advertisements


Constructing Topologies – Data or VPN Plane

Capabilities

- Fabric Data Plane or Individual VPNs subject to specific topologies / connectivity models



- Fully meshed fabric data plane
- Individual VPNs can use any topology

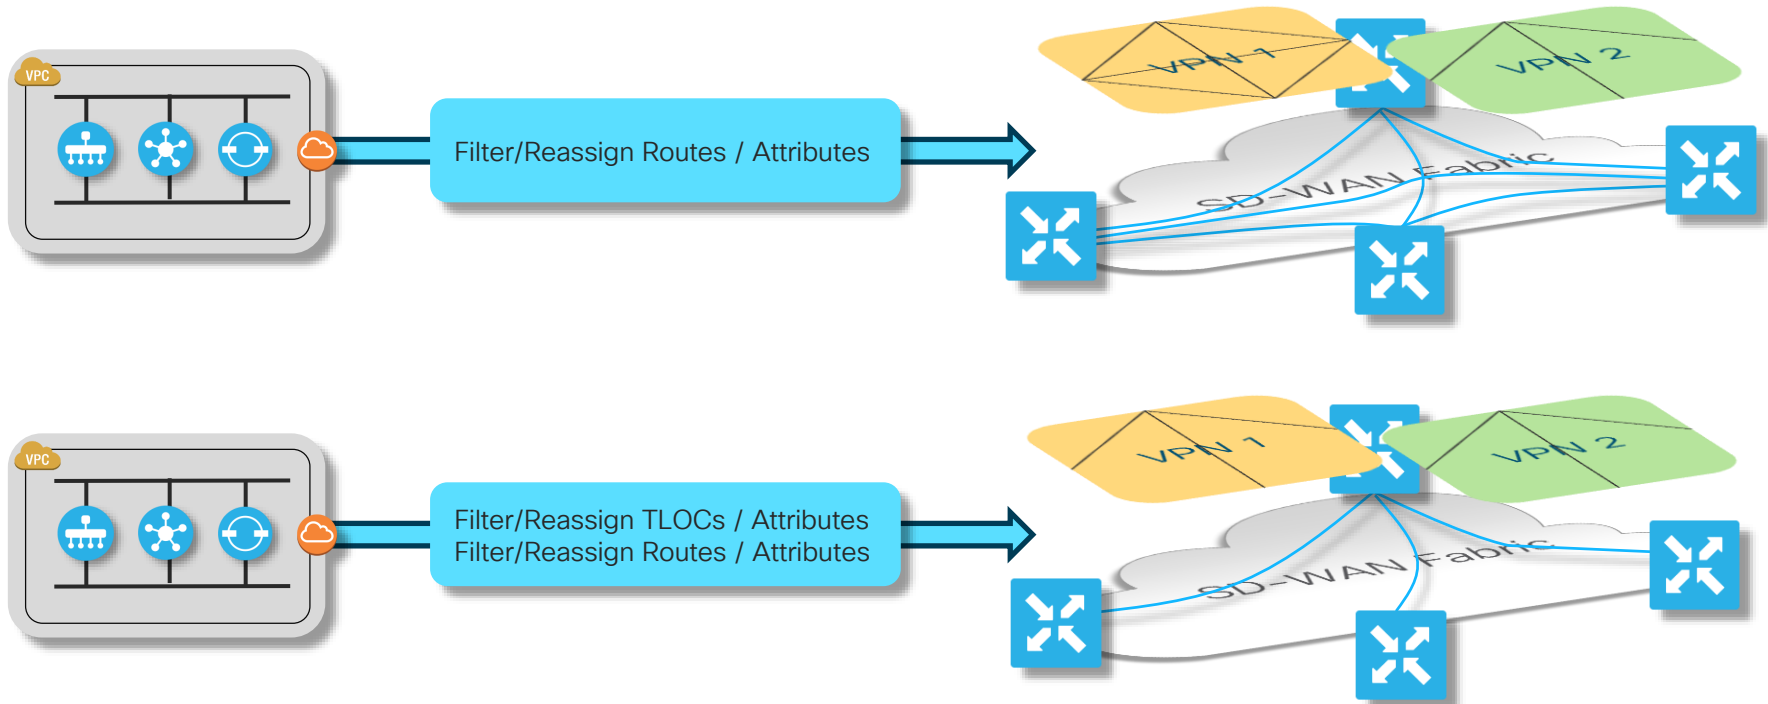


- Restricted fabric data plane
- Individual VPNs restricted to connectivity model used by underlying fabric

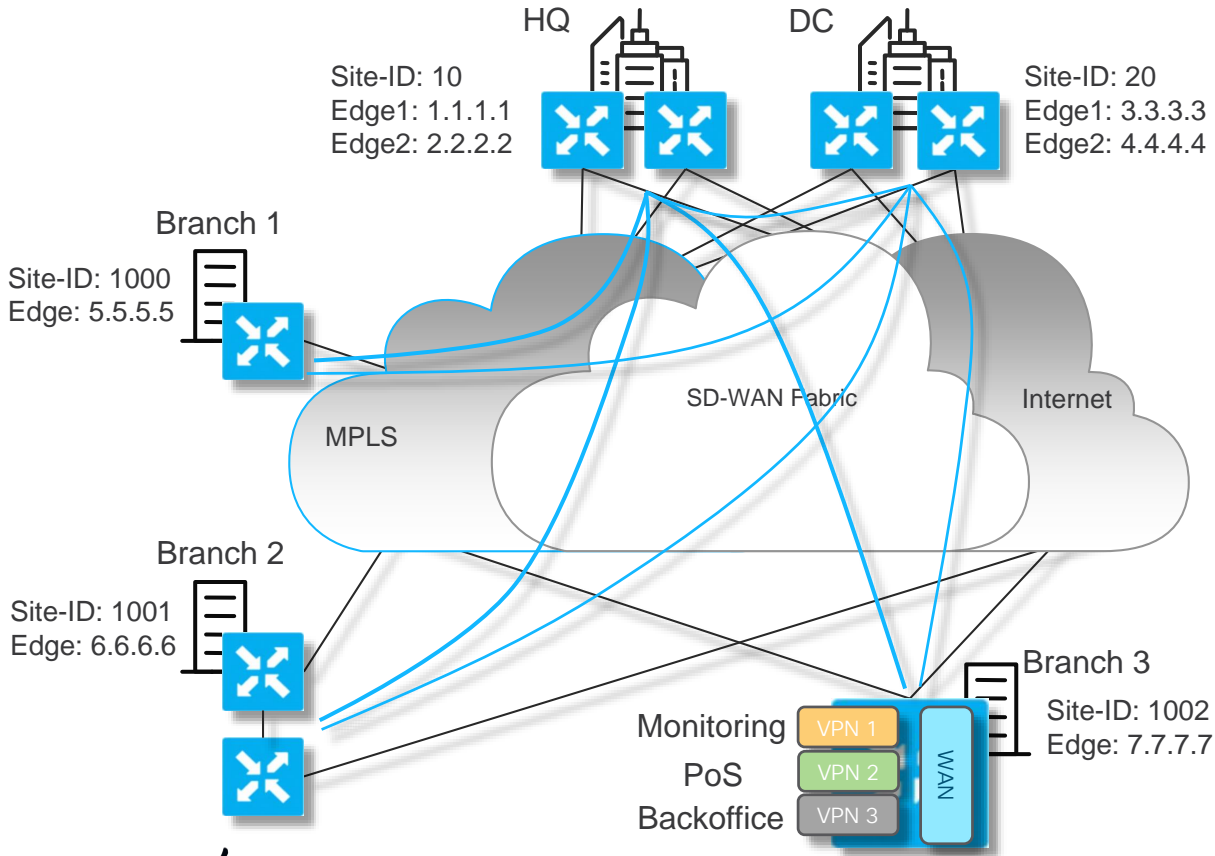
Constructing Topologies

Fabric Data Plane or VPN Plane Enforcement

- Data Plane or Individual VPNs subject to specific topologies / connectivity models



Hub-and-Spoke Design



- Topology can be applied to the Data Plane or individual VPNs as described
- Loose and Strict options are possible as well
 - Loose: Branch to Branch via hub/DC
 - Strict: Branch to hub only

Hub-and-Spoke Topologies

Fabric Data Plane and VPN Plane Control Policy

Policy

```
lists
  tloc-list hub-site tlocs
    tloc 1.1.1.1 color biz-internet encap ipsec preference 100
    tloc 1.1.1.1 color mpls encap ipsec preference 50
    tloc 3.3.3.3 color biz-internet encap ipsec preference 100
    tloc 3.3.3.3 color mpls encap ipsec preference 50
  !
  site-list branch sites
    site-id 1000-2000
  !
  site-list hub sites
    site-id 1-100
  !
  !
```

1 Define Hub Site TLOC-list

2 Declare Branches

3 Declare Hubs

```
apply-policy
  site-list branch sites
  control-policy restricted data plane out
  !
  !
```

5 Apply Policy to the target site-list

Policy

```
control-policy restricted data plane
  sequence 10
    match tloc
      site-list hub sites
    !
    action accept
  !
  !
  sequence 20
    match route
      site-list branch sites
    !
    action accept
    set
      tloc-list hub site tlocs
    !
    !
  !
  !
  sequence 30
    match tloc
    !
    action reject
  !
  !
  default-action accept
```

4 Define the Control Policy

Advertise Hub TLOCs

Branch Prefixes

Drop Branch TLOCs

Hub-and-Spoke Topologies

VPN 1 Full Mesh and VPN 2 Hub-and-Spoke Topologies

Loose Hub-and-Spoke
Spokes communicate via hub(s)

```

Policy
 lists
  vpn-list VPN2
  vpn 2
 !
 site-list branch_sites
  site-id 1000-2000
 !
 !
 control-policy vpn_multi-topology
  sequence 10
  match route
   site-list branch_sites
   vpn-list VPN2
  !
  action accept
  set
   tloc 1.1.1.1 color mpls
  !
 !
 default-action accept
  
```

Branch Prefixes

Hub site TLOC

Strict Hub-and-Spoke
No spoke to spoke communication

```

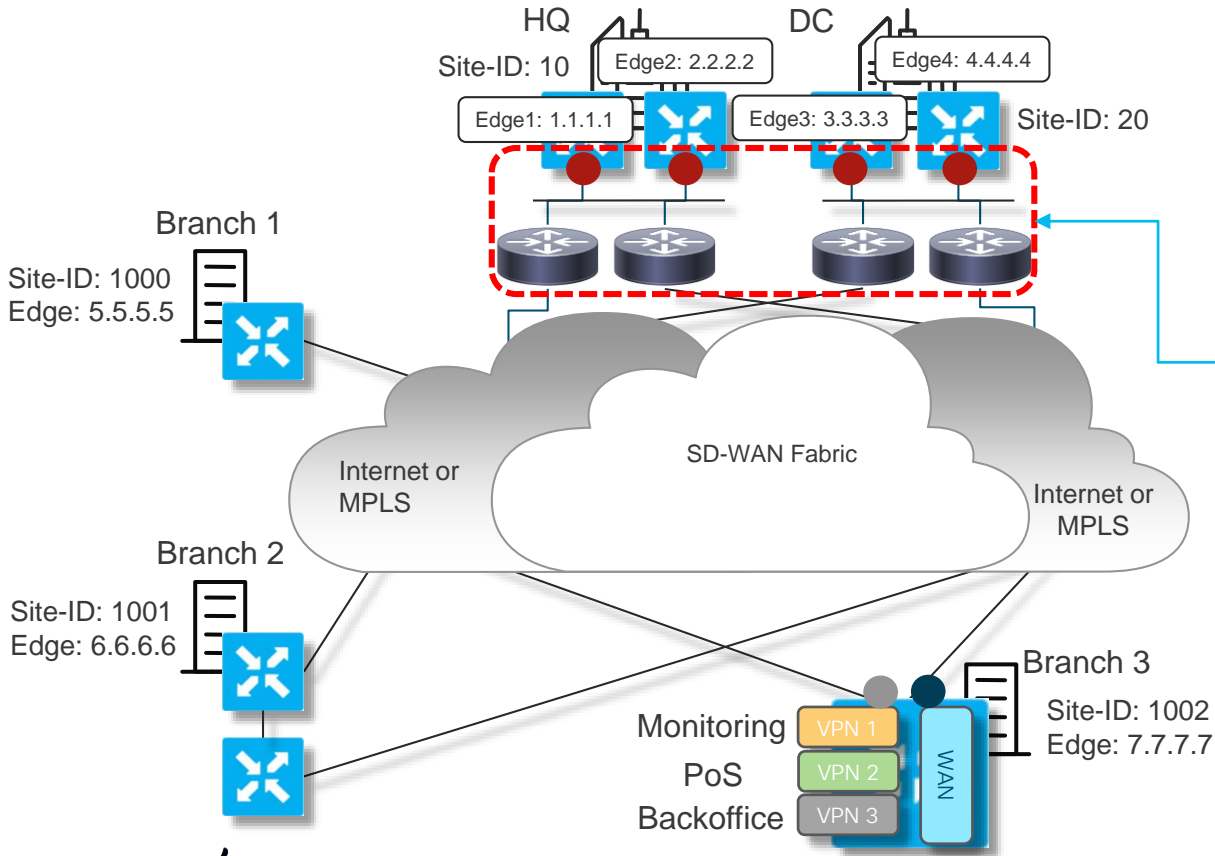
Policy
 lists
  vpn-list VPN2
  vpn 2
 !
 site-list hub_sites
  site-id 1-100
 !
 !
 control-policy vpn_multi-topology
  sequence 10
  match route
   site-list hub_sites
   vpn-list VPN2
  !
  action accept
  !
  sequence 20
  match route
  !
  action reject
  !
 default-action accept
  
```

Advertise Hub Prefixes

Drop Branch Prefixes

Multiple Underlays with direct/indirect attach

Multiple Underlays with direct/indirect attach



- Branches have multiple direct physical attachments to multiple underlays

This is represented by multiple colors at the branch

- DC/HQ nodes aren't directly attached to the transport but is provided an internal link

Hence, only a single color is normally used at DC/HQ

- TLOC: MPLS
- TLOC: Public-Internet
- TLOC: X


Multiple Underlays with direct/indirect attach

Challenges in Uniformly Enabling Capabilities Across the Network

- Branches typically have direct physical attach – standard operation
 - Underlay/Transport Routing and Path Preference
 - Application-Aware Routing and SLA measurements
 - Nothing changes from a standard design
- Central locations are challenged by lack of direct connectivity
 - Routing traffic per underlay
 - Path Preference – Using policies or static assignments
 - Application-aware routing measurements and switch-over


Multiple Underlays with direct/indirect attach

- Use Loopback interfaces to represent each underlay Network
- Assign IP-addressing that allows for routing to specific underlays



- BGP or OSPF can be used to share Loopback IPs with rest of network
- In case of disparate underlays, VPN 0 routing must be properly setup

Multiple Underlays with direct/indirect attach



```
vpn 0
 interface ge0/0
   ip address 192.168.1.2/24
 !
 interface loopback0
   ip address 192.169.1.1/32
   tunnel-interface
   color mpls
 !
 interface loopback1
   ip address 192.169.1.2/32
   tunnel-interface
   color public-internet
```

VPN 0 - Static Routing

```
vpn 0
 ip route 10.0.0.0/8 >> mpls-gw
 ip route 0.0.0.0/0 >> internet-gw
```

Transport - Static Routing


```
ip route 192.168.1.1/32 >> WAN Edge
ip route 192.168.1.2/32 >> WAN Edge
```

VPN 0 - OSPF

```
vpn 1
 router
  ospf
   area 0
    interface loopback0
    interface loopback1
  exit
 exit
 !
 !
```

Primary and Backup path/resource definition

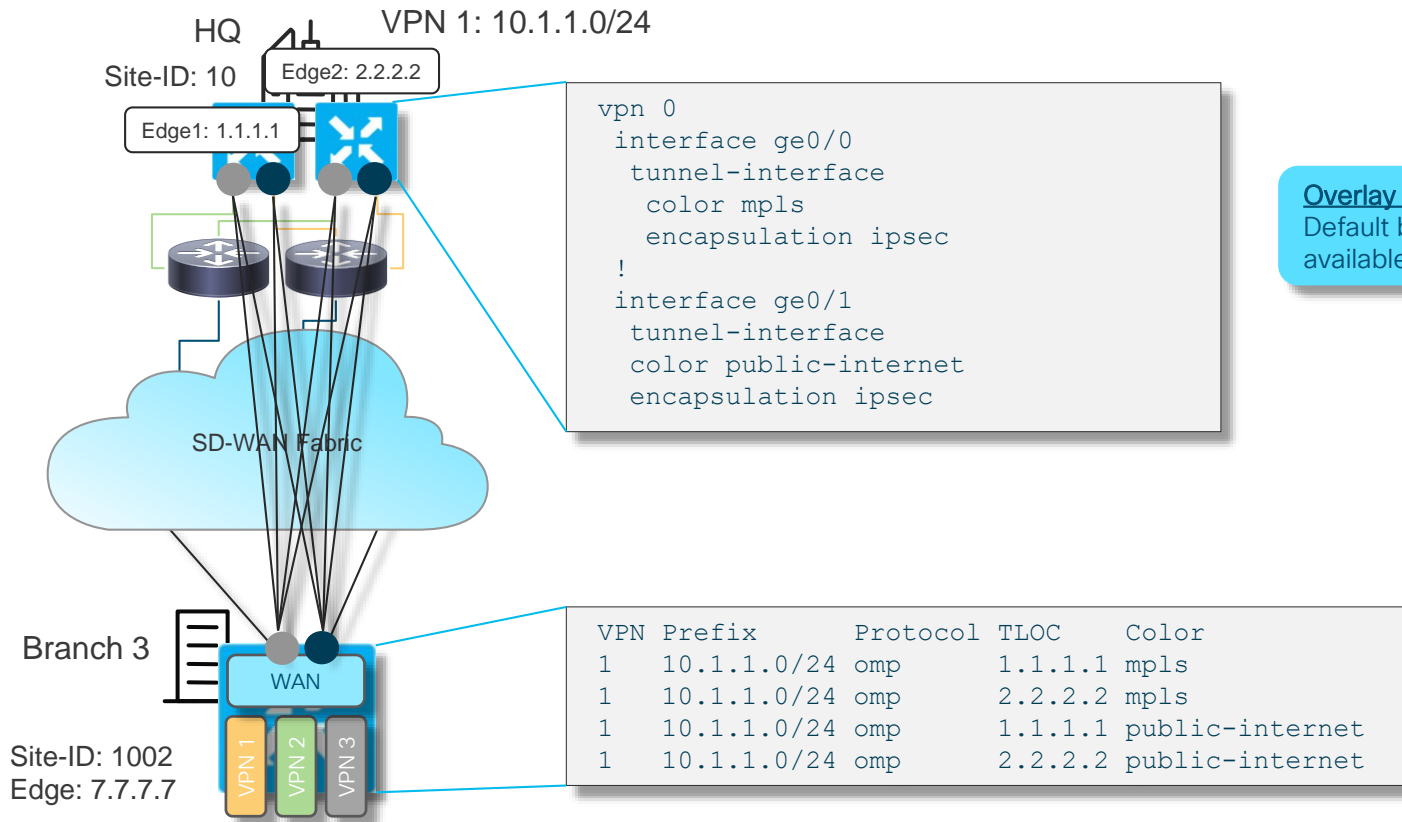
Primary and Backup path/resource definition



- Default behavior is to advertise branch prefixes with every TLOC as a valid NH for ECMP
- This can cause asymmetric distribution of flows
- Several techniques can be used to manage this
- In some cases, having all transports active is a higher priority so then default is ok

- TLOC: MPLS
- TLOC: Public-Internet

Primary and Backup path/resource definition

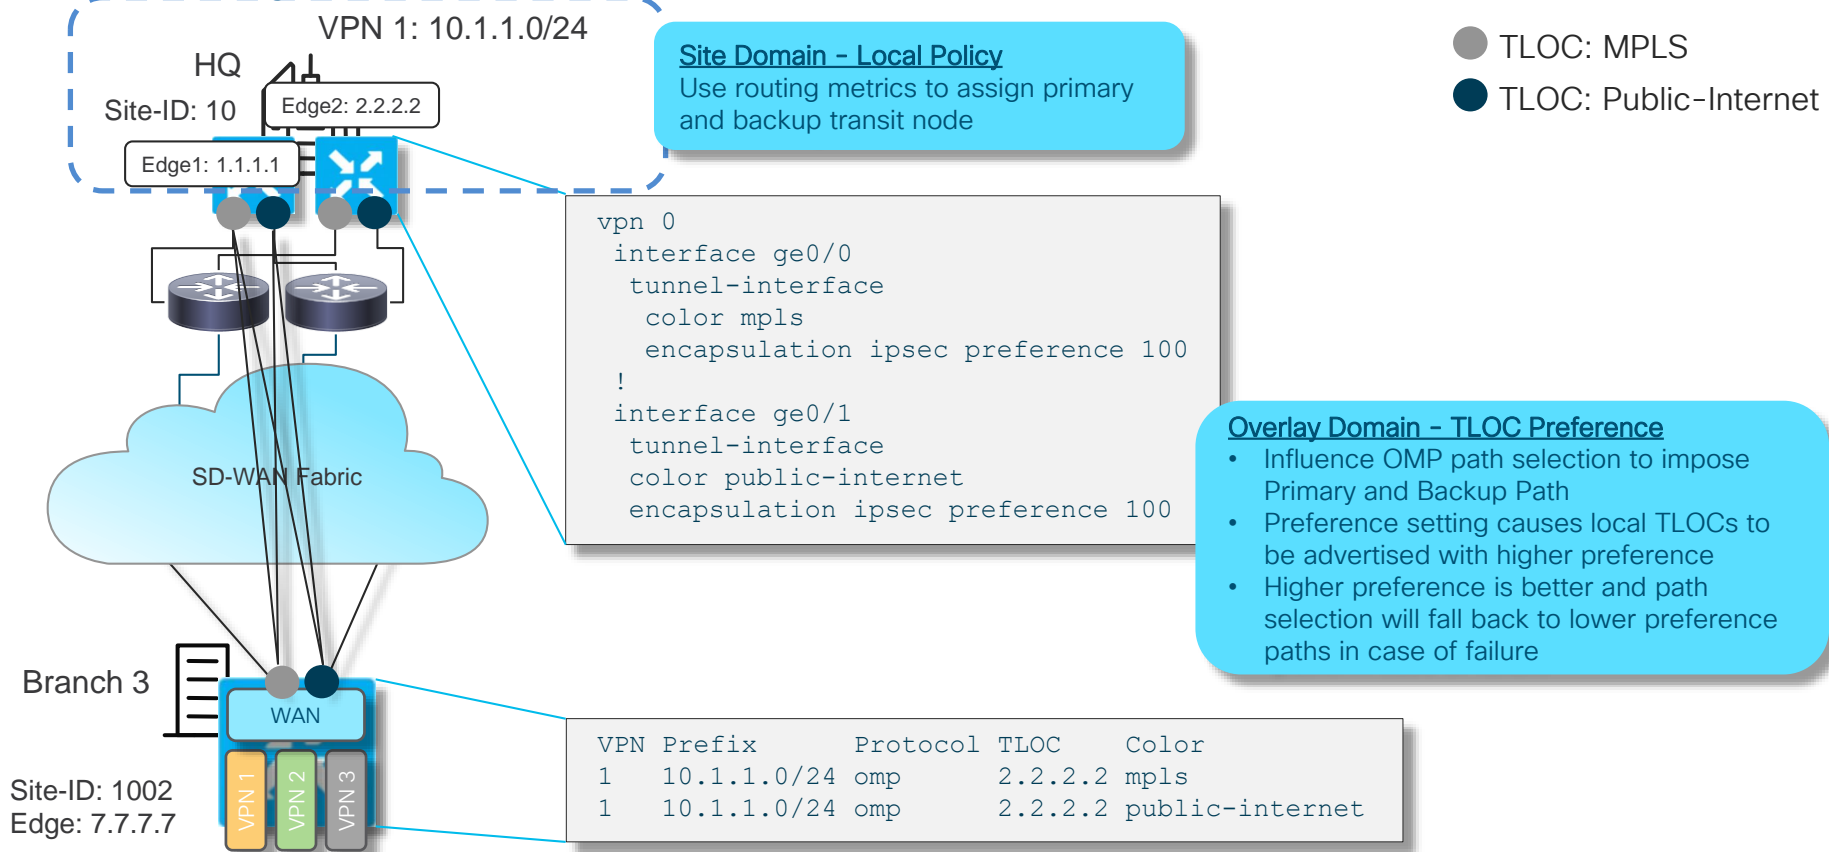


- TLOC: MPLS
- TLOC: Public-Internet

Overlay Domain - Path Preference
 Default behavior is to use every available path for ECMP




Primary and Backup path/resource definition



Primary and Backup path/resource definition

Network Resource (e.g. Data Center) Preference or Active/Backup



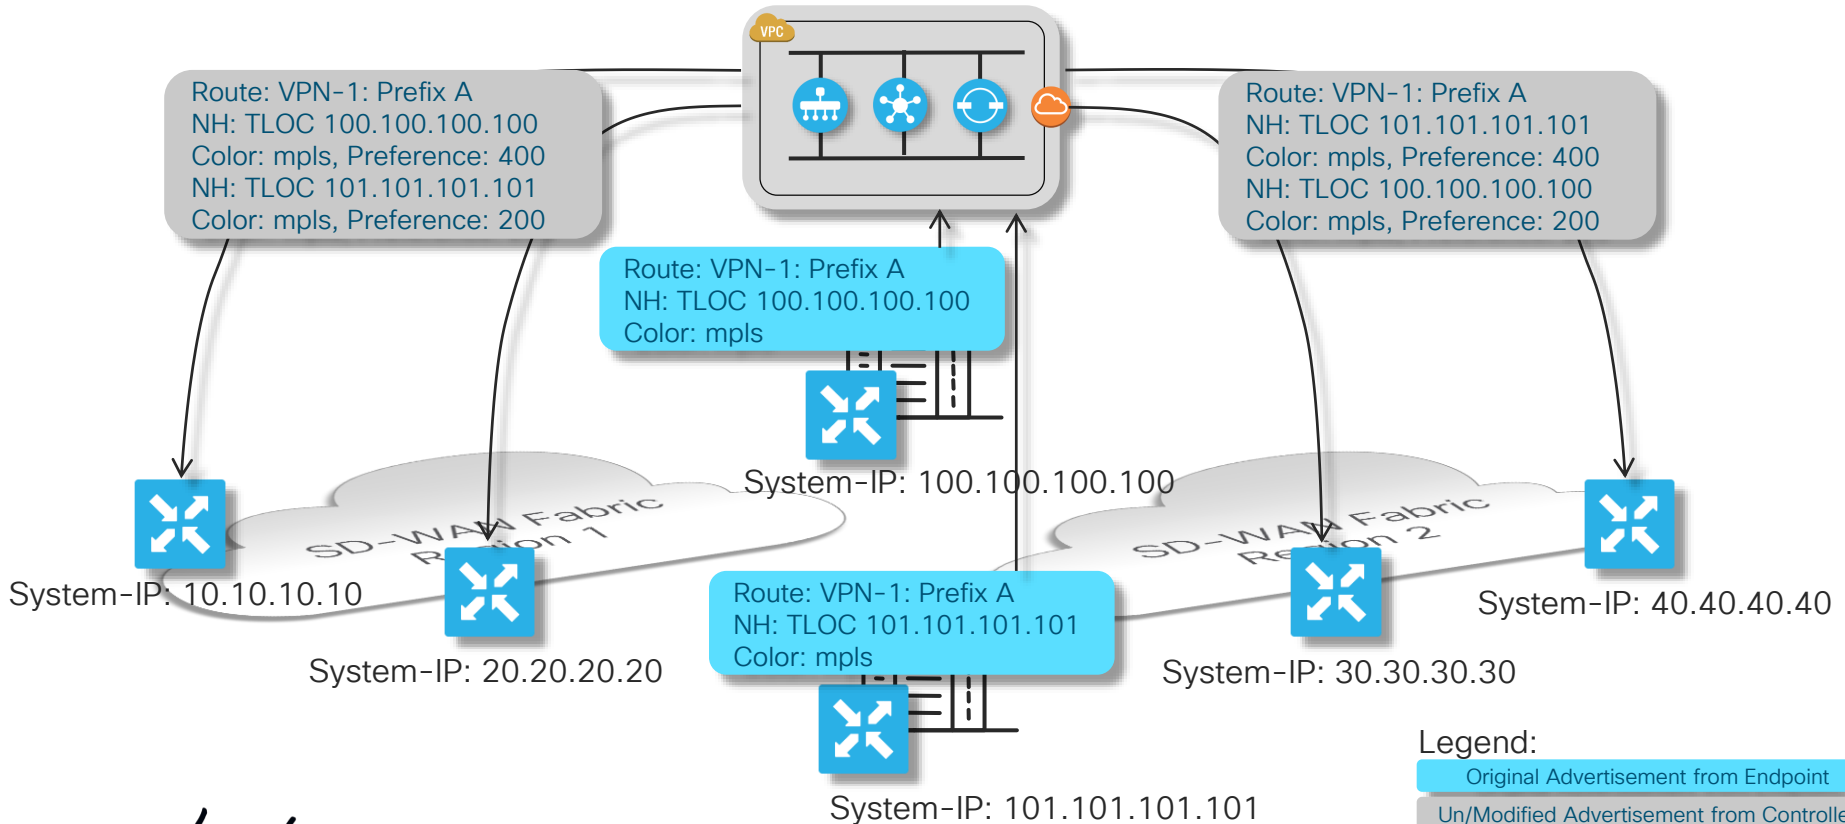
- Solution:

Identify regions by Site-Id and associate Primary and Backup DC locations with the regions

A control policy is used to make the associations and defining DC preference

Primary and Backup path/resource

Control Policy Operation



Primary and Backup path/resource definition

Control Policy Configuration

1 Define Data Center TLOC-lists

```
policy
lists
  tloc-list dc-preference-west
    tloc 100.100.100.100 color mpls encap ipsec preference 400
    tloc 101.101.101.101 color mpls encap ipsec preference 200
  !
  tloc-list dc-preference-east
    tloc 100.100.100.100 color mpls encap ipsec preference 200
    tloc 101.101.101.101 color mpls encap ipsec preference 400
  !
  site-list sites-region-west
    site-id 1-20
  !
  site-list sites-region-east
    site-id 21-40
  !
  site-list dc-sites
    site-id 100-101
```

2 Declare Regions

3 Declare Data Centers

```
apply-policy
  site-list sites-region-west
    control-policy adv-dc-preference-west out
  !
  site-list sites-region-east
    control-policy adv-dc-preference-east out
  !
  !
```

5 Apply Policies to the target site-lists


```
control-policy adv-dc-preference-west
  sequence 10
  match route
    site-list dc-sites
  !
  action accept
  set
    tloc-list dc-preference-west
  !
  !
  !
  default-action accept
  !
control-policy adv-dc-preference-east
  sequence 10
  match route
    site-list dc-sites
  !
  action accept
  set
    tloc-list dc-preference-east
  !
  !
  !
  default-action accept
  !
  !
```

4 Define the Control Policies

Interconnecting Dis- contiguous Data Planes

Interconnecting Dis-contiguous Data Planes

Interconnecting nodes single-homed to different underlays

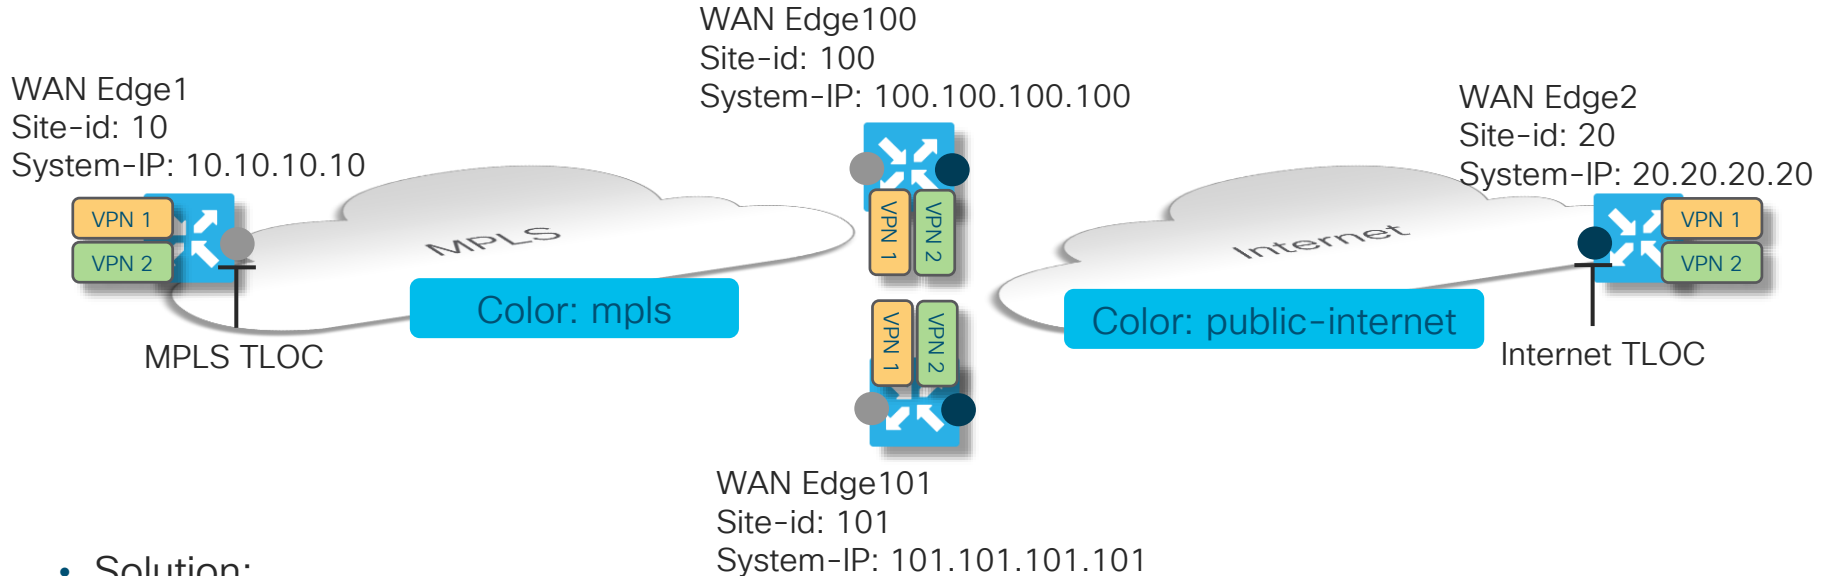


- Problem:

Overlay with a dis-contiguous data plane and endpoints need to communicate end-to-end

Interconnecting Dis-contiguous Data Planes

Interconnecting nodes single-homed to different underlays



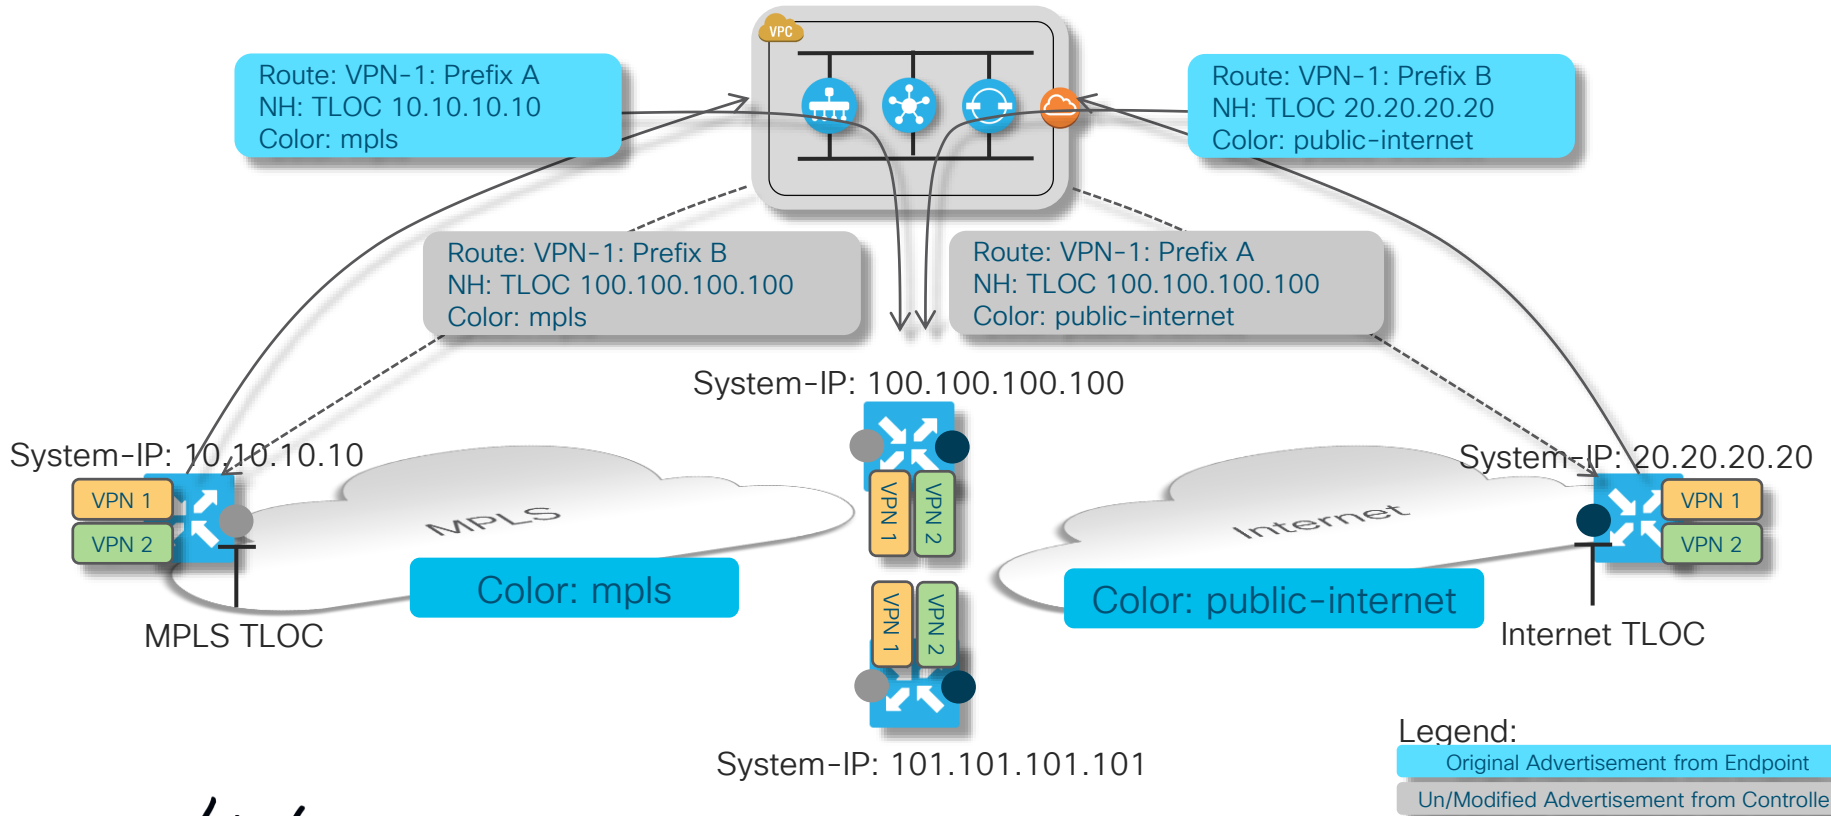
- Solution:

Identify one or more multi-homed sites to bridge the data plane gap and act as gateways

Use a control policy to enable distribution of routing information between domains enabling gateway-supported paths

Interconnecting Dis-contiguous Data Planes

Control Policy Operation



Interconnecting Dis-contiguous Data Planes

Control Policy Configuration

1 Define Gateway TLOC-lists

```
policy
lists
  tloc-list internet-gateways
  tloc 100.100.100.100 color mpls encap ipsec
  tloc 101.101.101.101 color mpls encap ipsec
  !
  tloc-list mpls-gateways
  tloc 100.100.100.100 color public-internet encap ipsec
  tloc 101.101.101.101 color public-internet encap ipsec
  !
  site-list internet-sites
  site-id 20
  !
  site-list mpls-sites
  site-id 10
```

2 Declare Target Sites

```
apply-policy
  site-list internet-sites
  control-policy announce-mpls-sites out
  !
  site-list mpls-sites
  control-policy announce-internet-sites out
  !
```

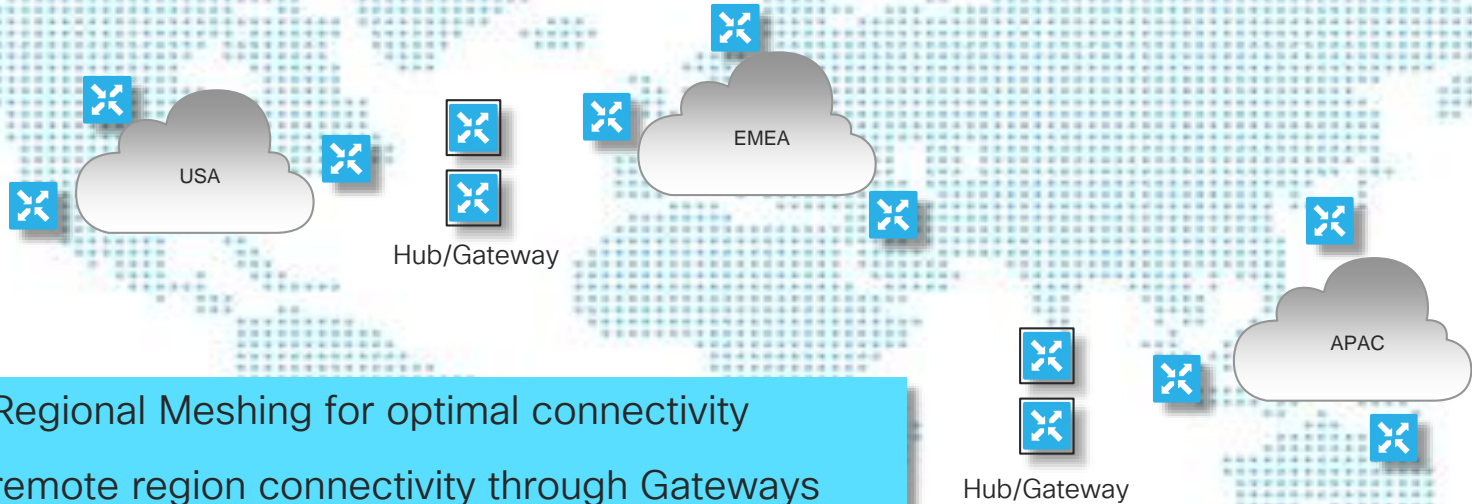

4 Apply Policies to the target site-lists

3 Define the Control Policies

```
control-policy announce-internet-sites
sequence 10
  match route
    site-list internet-sites
  !
  action accept
  set
    tloc-list internet-gateways
  !
  !
  !
  default-action accept
  !
control-policy announce-mpls-sites
sequence 10
  match route
    site-list mpls-sites
  !
  action accept
  set
    tloc-list mpls-gateways
  !
  !
  !
  default-action accept
  !
  !
```

Multi-Region Overlay

Multi-Region Overlay Requirements



- Support Regional Meshing for optimal connectivity
- Support remote region connectivity through Gateways
- Provide Redundant Gateway Connectivity

Multi-Region Overlay

Definitions and Dependencies


- Site-ID assignment allowing for Site identification – 32 bits

	Continent	Country	Site number
	X	YYY	ZZZZ
	1-7	1-999	1-9999
Example {	Europe	Sweden	Site
	5	046	1000

- TLOC Colors illustrating how sites are attached
- System-IP identifying individual nodes

Multi-Region Overlay

Site Assignments



Control Policy Case Study

Reachability Information Distribution Requirements

US

Inbound TLOC Advertisement
US Region - All Colors
US Gateways - All Colors
EMEA Gateways - All Colors
APAC Gateway - All Colors

Outbound TLOC Advertisement
US Gateways - All Colors

Inbound vRoute Advertisement
US Region - Original NH
EMEA Region - EU GW NH
APAC Region - APAC GW NH

Outbound vRoute Advertisement
US Region - US GW NH

EMEA

Inbound TLOC Advertisement
EMEA Region - All Colors
EMEA Gateways - All Colors
US Gateways - All Colors
APAC Gateways - All Colors

Outbound TLOC Advertisements
EMEA Gateways - All Colors

Inbound vRoute Advertisement
EMEA Region - Original NH
US Region - US GW NH
APAC Region - APAC GW NH

Outbound vRoute Advertisement
EMEA Region - EU GW NH

APAC

Inbound TLOC Advertisement
APAC Region - All Colors
APAC Gateways - All Colors
EMEA Gateways - All Colors
US Gateways - All Colors

Outbound TLOC Advertisement
APAC Gateways - All Colors

Inbound vRoute Advertisement
APAC Region - Original NH
EMEA Region - EU GW NH
US Regions - US GW NH

Outbound vRoute Advertisement
APAC Region - APAC GW NH

Control Policy Case Study

Policy Definition - Lists

```
policy
 lists
  site-list US branch sites
  site-id 60010000-60018999
  !
  site-list US gateway sites
  site-id 60019000-60019999
  !
  site-list EMEA branch sites
  site-id 50010000-50338999
  site-id 50340000-59999999
  !
  site-list EMEA gateway sites
  site-id 50339000-50339999
  !
  site-list APAC branch sites
  site-id 30010000-30668999
  site-id 30670000-39999999
  !
  site-list APAC gateway sites
  site-id 30669000-30669999
  !
  !
  !
```

```
policy
 lists
  tloc-list US gateway tlocs
  tloc 1.1.1.1 color mpls encap ipsec preference 100
  tloc 1.1.1.1 color biz-internet encap ipsec preference 100
  tloc 2.2.2.2 color mpls encap ipsec preference 50
  tloc 2.2.2.2 color biz-internet encap ipsec preference 50
  !
  tloc-list EMEA gateway tlocs
  tloc 3.3.3.3 color mpls encap ipsec preference 100
  tloc 3.3.3.3 color biz-internet encap ipsec preference 100
  tloc 4.4.4.4 color mpls encap ipsec preference 50
  tloc 4.4.4.4 color biz-internet encap ipsec preference 50
  !
  tloc-list APAC gateway tlocs
  tloc 5.5.5.5 color mpls encap ipsec preference 100
  tloc 5.5.5.5 color biz-internet encap ipsec preference 100
  tloc 6.6.6.6 color mpls encap ipsec preference 50
  tloc 6.6.6.6 color biz-internet encap ipsec preference 50
  !
  !
  !
```

Control Policy Case Study

Policy Definition Cont'd – Control Policy – Applied to US Sites

```
policy
  control-policy us_domain
    sequence 10
      match tloc
        site-list US_branch_sites
      !
      action accept
      !
    !
    sequence 20
      match tloc
        site-list US_gateway_sites
        SNIP ... (accept)
    sequence 30
      match tloc
        site-list EMEA_gateway_sites
        SNIP ... (action accept)
    sequence 40
      match tloc
        site-list APAC_gateway_sites
      !
      SNIP ... (action accept)
```

```
sequence 50
  match route
    site-list US_branch_sites
  !
  action accept
  !
sequence 60
  match route
    site-list US_gateway_sites
    SNIP ... (action accept)
sequence 70
  match route
    site-list EMEA_branch_sites
  !
  action accept
  set
    tloc-list EMEA_gateway_tlocs
  !
  !
sequence 80
  match route
    site-list EMEA_gateway_sites
    SNIP ... (action accept)
```


Control Policy Case Study

Policy Definition Cont'd – Control Policy – Applied to US Sites

```
sequence 90
  match route
    site-list APAC branch sites
  !
  action accept
  set
    tloc-list APAC gateway tlocs
  !
  !
sequence 100
  match route
    site-list APAC gateway sites
  !
  action accept
  !
  !
default-action accept
```

```
apply-policy
  site-list US branch sites
  control-policy us domain out
  !
  site-list US gateway sites
  control-policy us domain out
  !
  !
```

- Policy Logic

Sequence 10: Advertise US Branch TLOCs

Sequence 20: Advertise US GW TLOCs

Sequence 30: Advertise EMEA GW TLOCs

Sequence 40: Advertise APAC GW TLOCs

Sequence 50: Advertise US Branch routes

Sequence 60: Advertise US GW routes

Sequence 70: Advertise EMEA Branch routes w/ NH of EMEA GW

Sequence 80: Advertise EMEA GW routes

Sequence 90: Advertise APAC Branch routes w/ NH of APAC GW

Sequence 100: Advertise APAC GW Routes

Cisco SD-WAN Common Overlay Services


Common Designs and Services

- Extranet Service
- Quality of Service
- Application Pinning
- Internet Breakout
- SLA-driven Path Selection

Extranet Service

Providing Extranet Services

Scenario



- Problem: Shared Services to be consumed from Extranet VPN hosted location
- Solution: Provision Extranet Access from other overlay VPNs

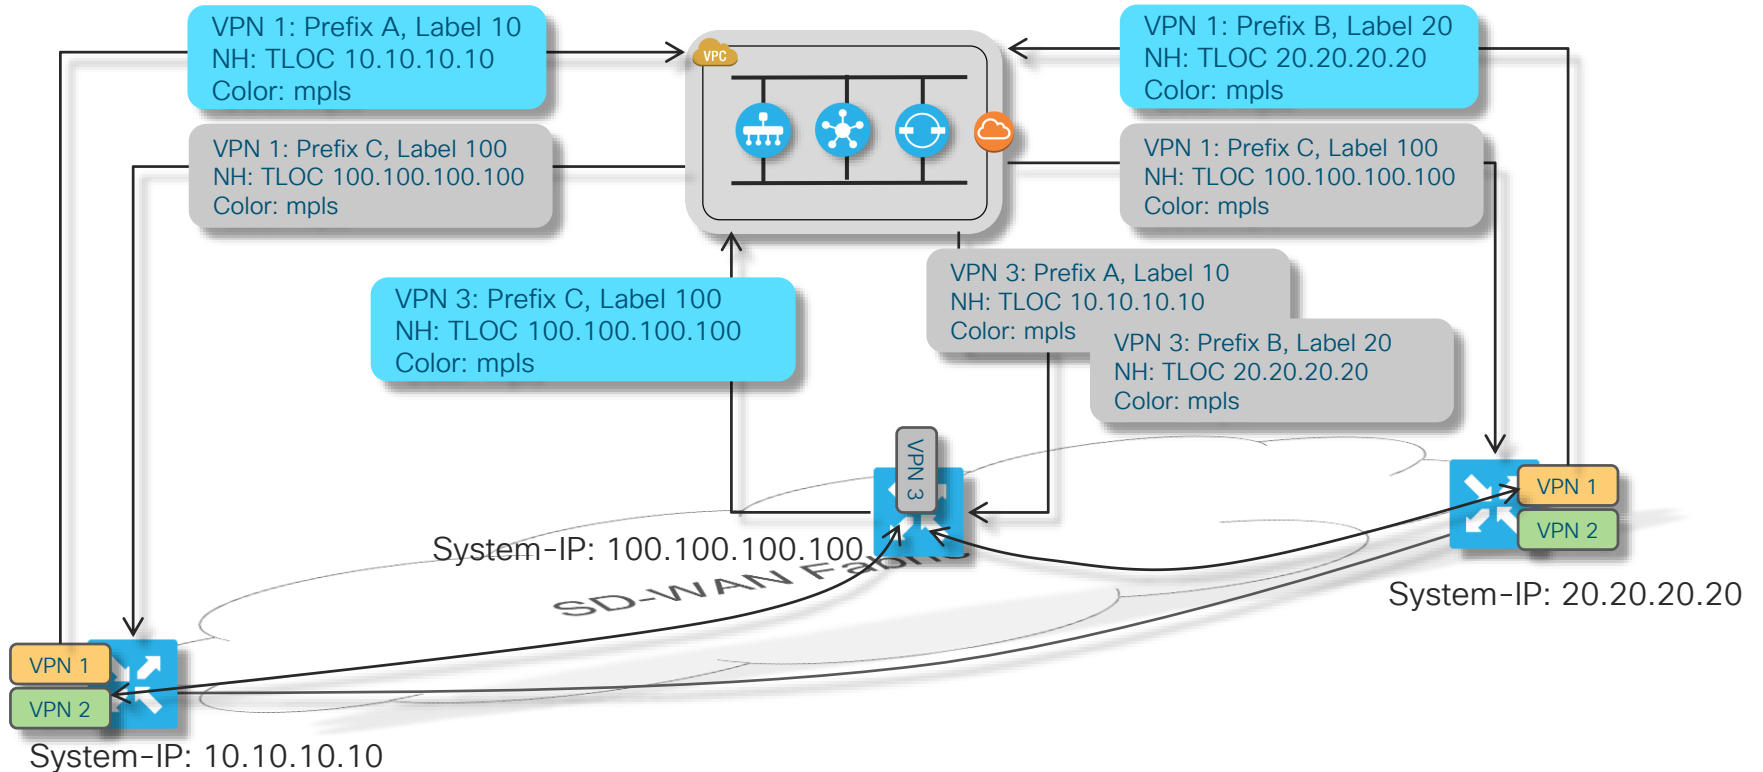
Providing Extranet Services

Control Policy Operation

Legend:

Original Advertisement from Endpoint

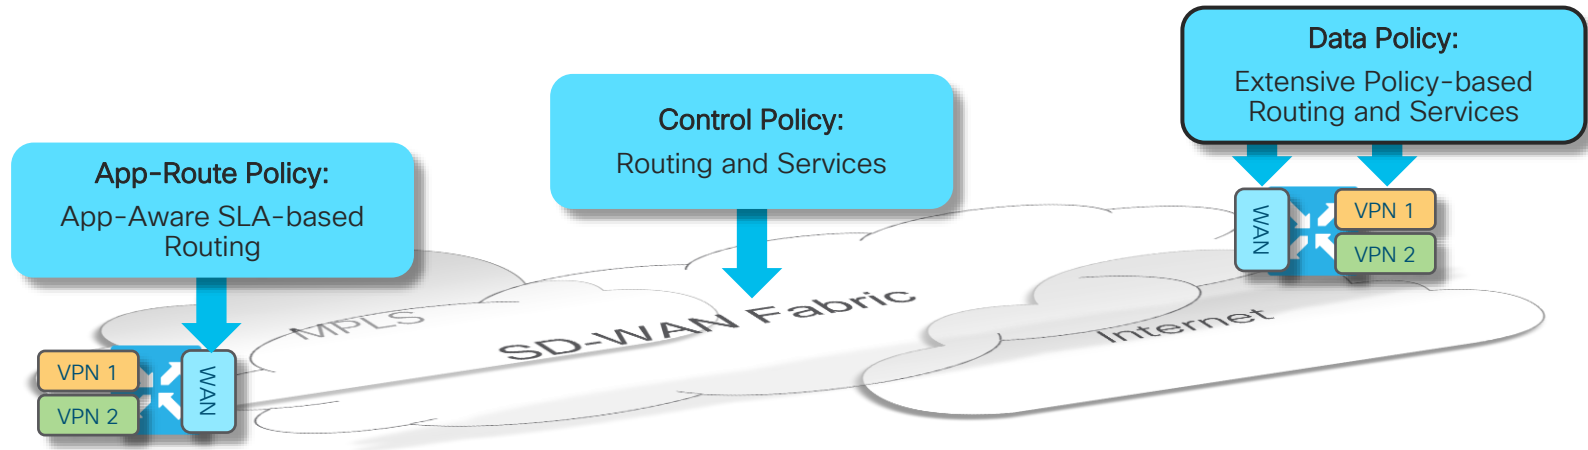
Un/Modified Advertisement from Controller



Introduction to Data Policies

Cisco SD-WAN Policy Architecture

Suite of Policies to address different functional domains



- Control Policies are applied at vSmart: Tailors routing information advertised to vEdge endpoints
- App-Route Policies are applied at vEdge: SLA-driven path selection for applications
- Data Policies are applied at vEdge: Extensive Policy driven routing


Data Policy Application

Direction of Processing

- A Data Policy can be applied in three modes:
 - From-service (Upstream)
 - From-tunnel (Downstream)
 - All (Up and Downstream)
- Different Data-policies can be applied to the same site if they apply to different directions

```
apply-policy site-list <name>  
  data-policy <name> all | from-service | from-tunnel
```

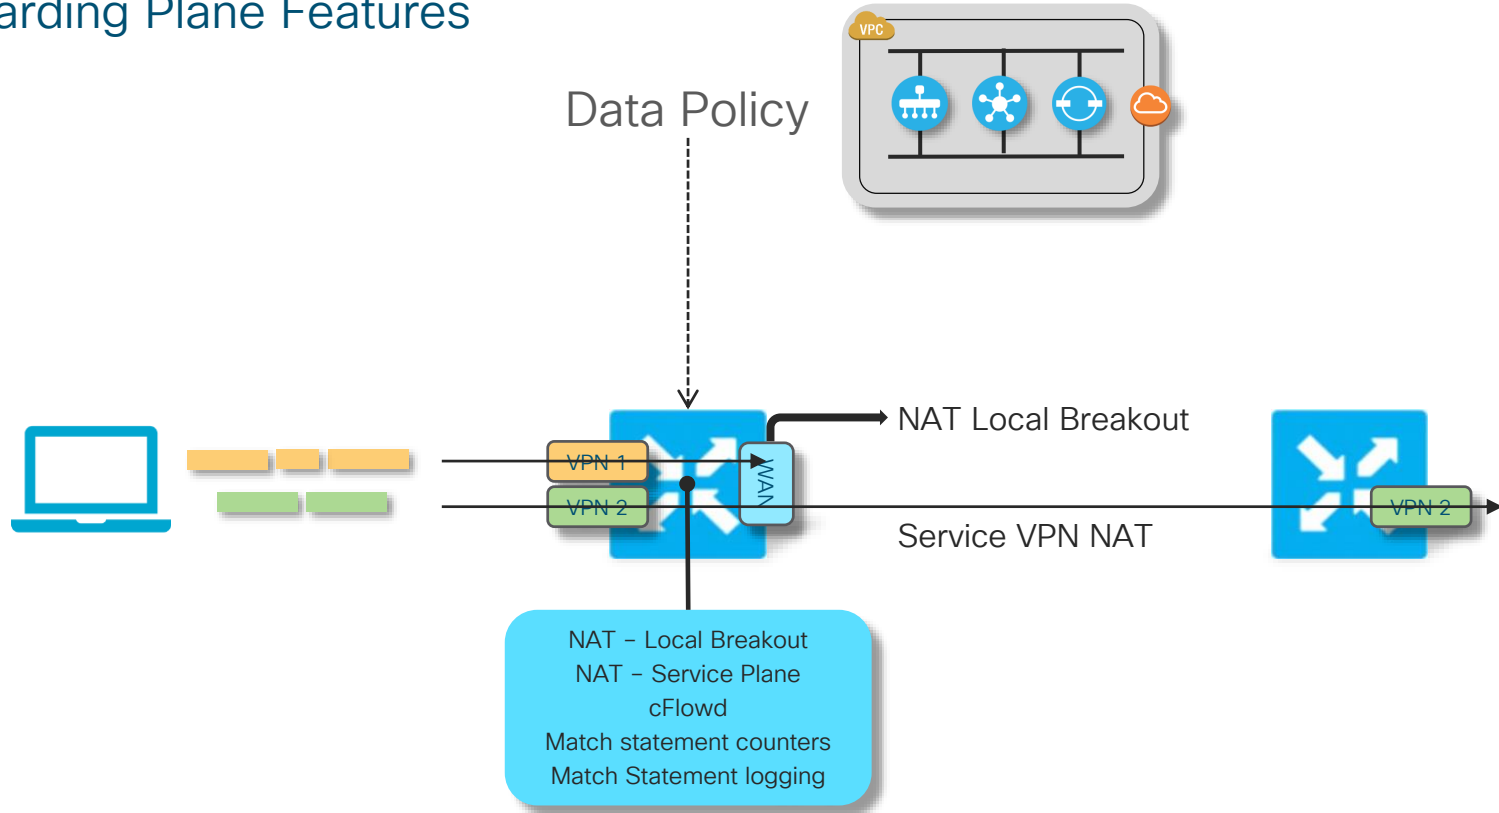
Upstream Traffic matched by Data-policy



Downstream Traffic matched by Data-policy

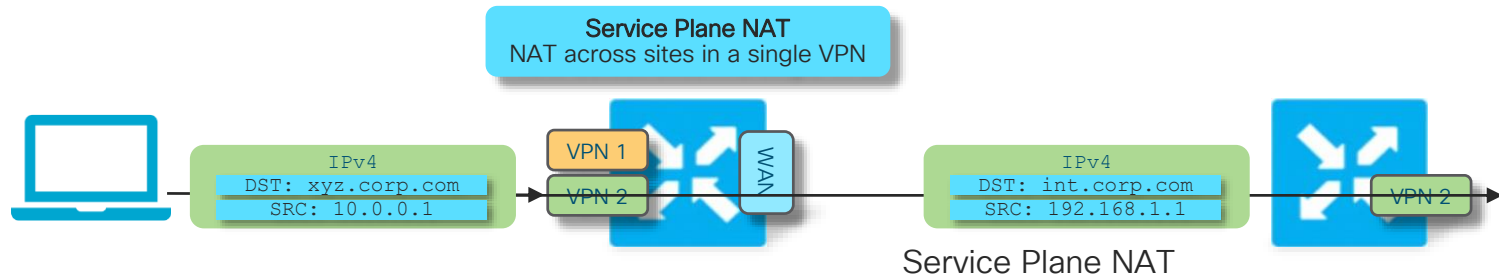
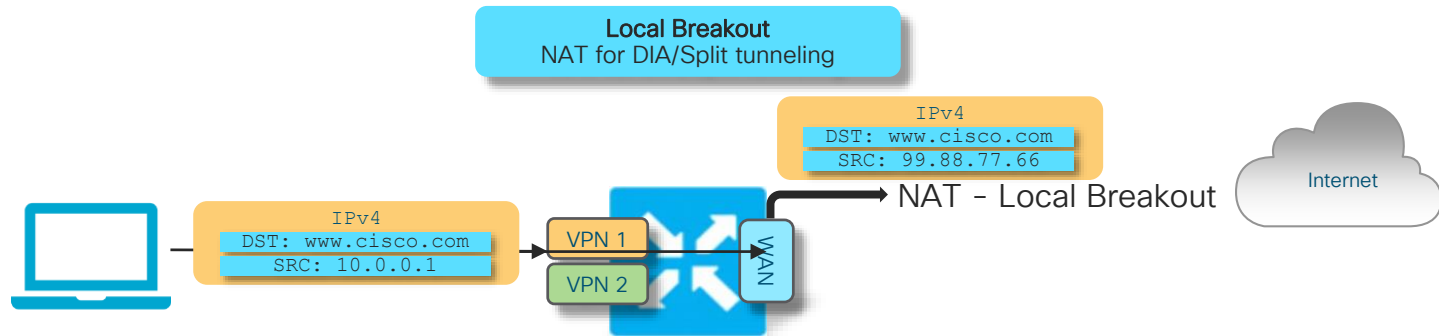
Data Policy Capabilities

Forwarding Plane Features



Data Policy Case #1

Forwarding Plane Features – NAT for DIA and Service VPN



Data Policy Capabilities

Forwarding Plane Feature Enablement – Policy Structure

Service Plane NAT NAT across sites in a single VPN

```
policy data-policy Srvc Plane NAT
  vpn-list VPN1
  sequence 10
  match source-ip 10.0.0.1/32
  !
  action accept
  nat pool 1
  !
  !
  default-action accept
  !
```

```
vEdge
vpn 1
  interface natpool1
  ip address 192.168.1.1/32
  no shutdown
  !
```

Local Breakout NAT for DIA/Split tunneling

```
policy data-policy DIA NAT
  vpn-list VPN1
  sequence 10
  match source-ip 10.0.0.1/32
  !
  action accept
  nat use-vpn 0
  !
  !
  default-action accept
  !
```

```
vEdge
vpn 0
  interface ge0/0
  ip address 192.168.1.1/32
  no shutdown
  nat
  !
```

Data Policy Capabilities

Forwarding Plane Feature Enablement – Policy Structure

Local Breakout cFlowd and Counting

```
policy data-policy DIA NAT
  vpn-list VPN1
  sequence 10
  match source-ip 10.0.0.1/32
  !
  action accept
  cflowd
  count local-breakout-traffic
  nat use-vpn 0
  !
  !
  default-action accept
  !
```

- Counters visible using GUI/Realtime or via CLI

```
show policy data-policy-filter
```

- Use cflowd template for export-destination configuration

Local Breakout Logging breakout traffic

```
policy data-policy DIA NAT
  vpn-list VPN1
  sequence 10
  match source-ip 10.0.0.1/32
  !
  action accept
  log
  nat use-vpn 0
  !
  !
  default-action accept
  !
```

```
vEdge
System
logging
server syslog.company.com
vpn 1
source-interface loopback1
exit
!
```

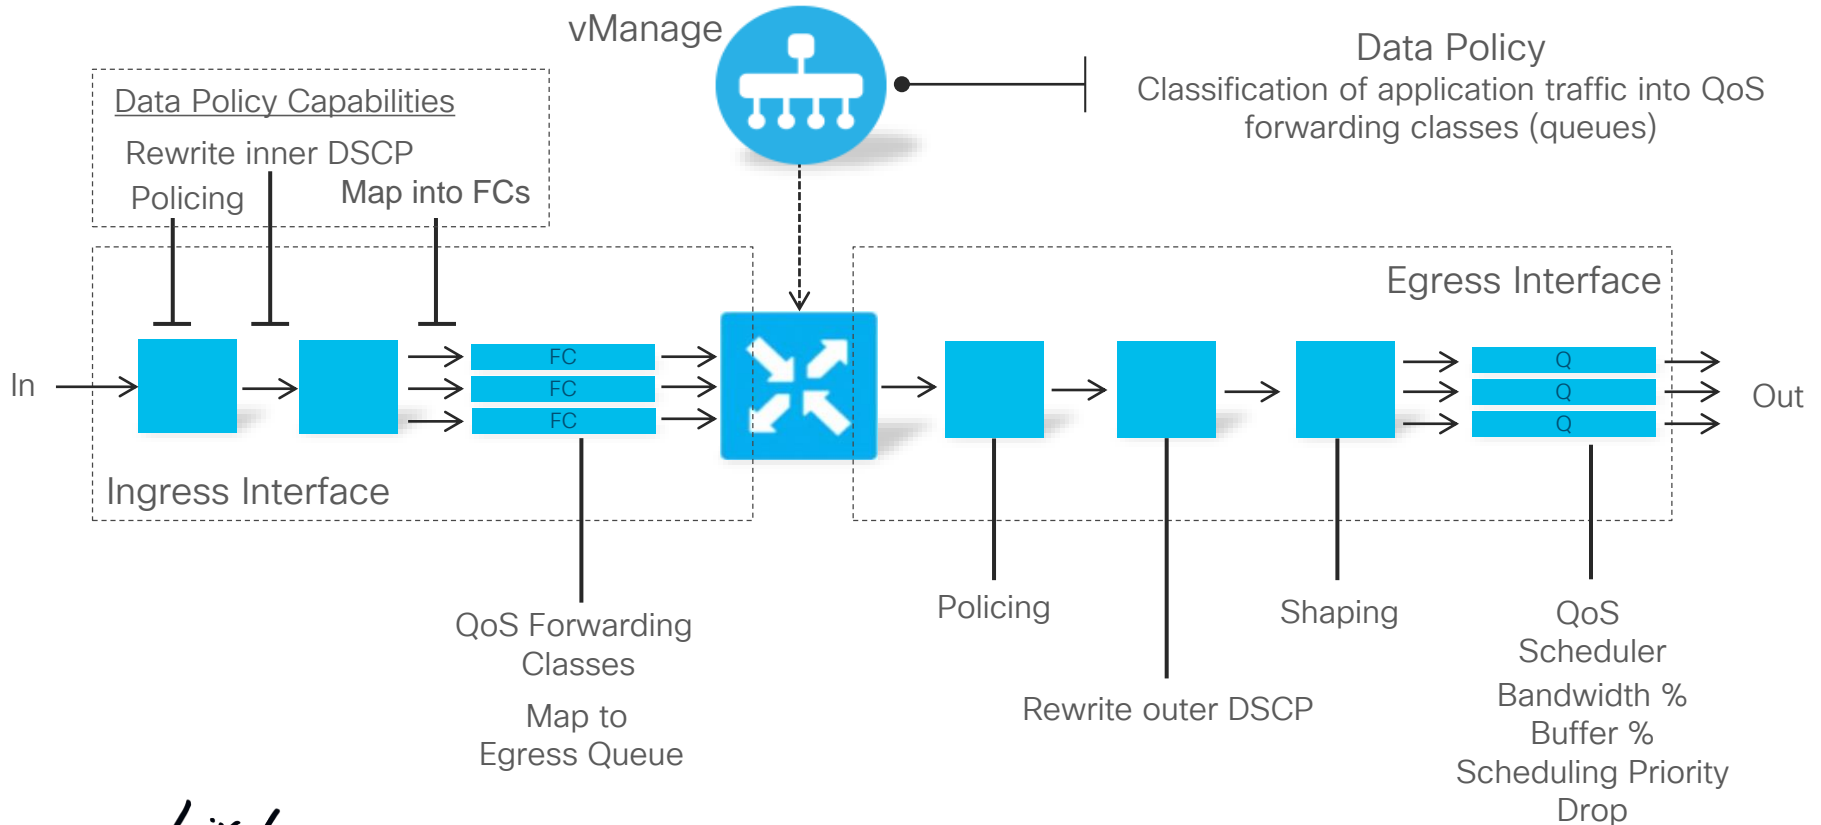
```
vEdge
policy
log-frequency <number>*
```

* Default is every 1000 packets

Quality of Service

WAN Edge Router Device QoS Overview

WAN Edge Router



Data Policy for QoS

Quality of Service – Policy Structure

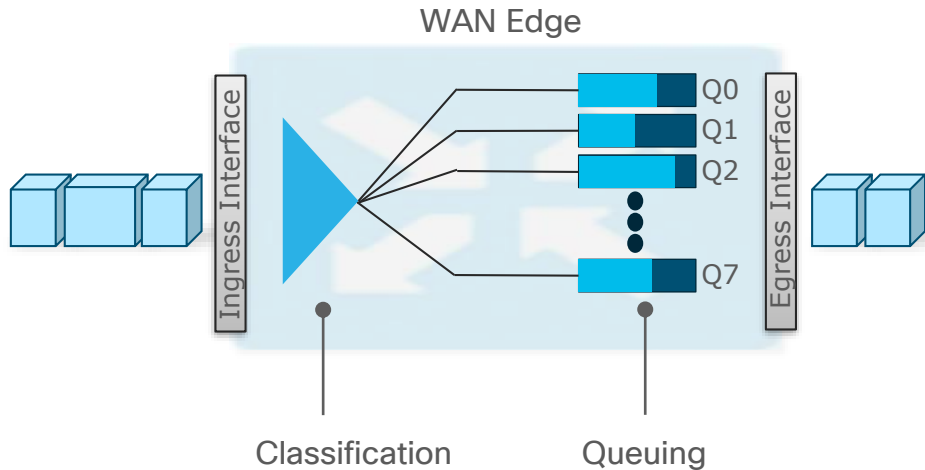
```
policy
data-policy enterprise traffic
  vpn-list VPN1
  sequence 10
  match app-list audio-video
  !
  action accept
  set
  dscp 46
  forwarding-class EF-class
  !
  !
  !
data-policy DIA
  vpn-list VPN10
  sequence 10
  match source-ip 10.0.0.0/8
  !
  action accept
  set
  policer police_DIA
  !
  !
  !
  default-action accept
  !
```

- App-list consists of DPI signature references
- Forwarding-class referring to configured QoS-class (Ref: qos-group in Cisco IOS)

```
policy
  policer police_DIA
  rate 10000000
  burst 1000000
  exceed drop
  !
  !
```

Policer configured as part of Policy

WAN Edge Router Qos Capabilities



- Classification
 - Flow match on 6-tuple (ACL, Data Policy)
 - Application match on DPI (Data Policy)
- Per-Egress Interface Queuing
 - Q0 is LLQ
 - vEdge control traffic (DTLS/TLS, BFD, routing protocols) goes into Q0
 - Assign a some small value for control (5%)
- Scheduling for Q1-Q7 is WRR*
 - Bandwidth percent determines queue weight
 - Unused Q0 bandwidth is distributed between other queues
- Queue drop is RED** or tail-drop
 - Linear drop probability, i.e. X% queue depth results in X% drop probability

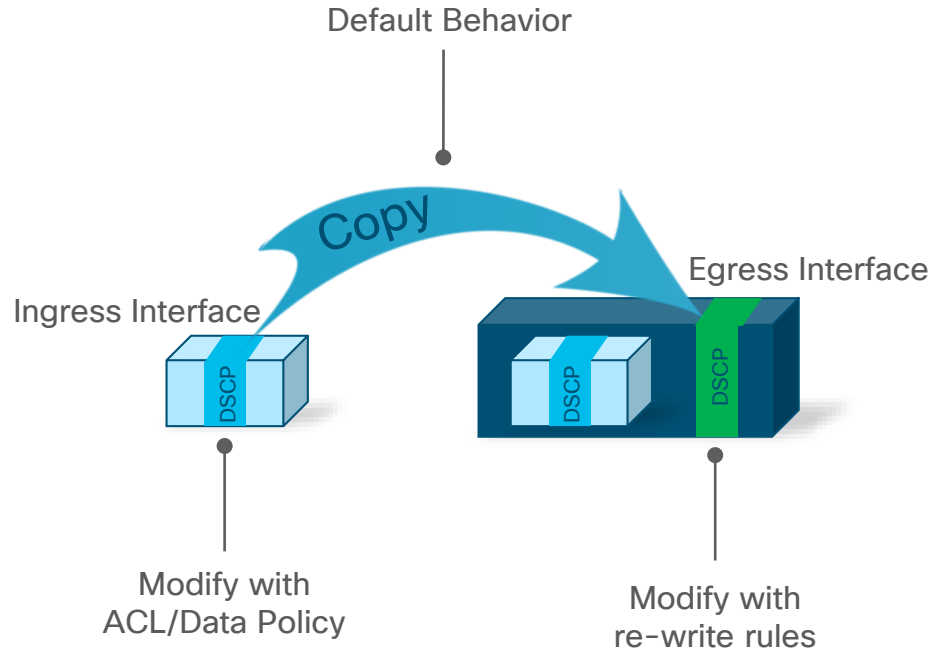
* Weighted Round-Robin

** Random Early Discard

CISCO *Live!*

Marking and Remarking

Supporting Enterprise and Provider DSCP schemes concurrently



- Comply with service providers provisioned classes of service
- Ingress Classification
 - DPI or 6 tuple matching using centralized or localized data policy
- Ingress interface marks/re-marks inner DSCP bits
- Inner DSCP bits are copied to the outer DSCP bits
- Egress interface re-write rules remark outer DSCP bits

Re-Marking the BFD to match Application Traffic

Supporting Control and Critical traffic E2E

access-list LAN-Classification

```
...
sequence 50
match
  dscp 48
!
action accept
  class NetworkControl
!
!
default-action accept
!
```

access-list MarkBFDPackets

```
sequence 10
match
  class NetworkControl
!
action accept
!
!
sequence 20
match
  dscp 48
  protocol 17
!
action accept
set
  dscp 46
!
!
!
default-action accept
!
```

- WAN Edge QoS Default Behavior
 - All user traffic get mapped to Q2
 - All control traffic mapped to Q0 LLQ
 - Controller traffic
 - BFD Packets
 - Marked with DSCP of 48
- 100ms of buffer per port – Buffer Allocation can be configured per queue
- Recommendation
 - Always reserve minimum of 5% of BW and buffer for LLQ

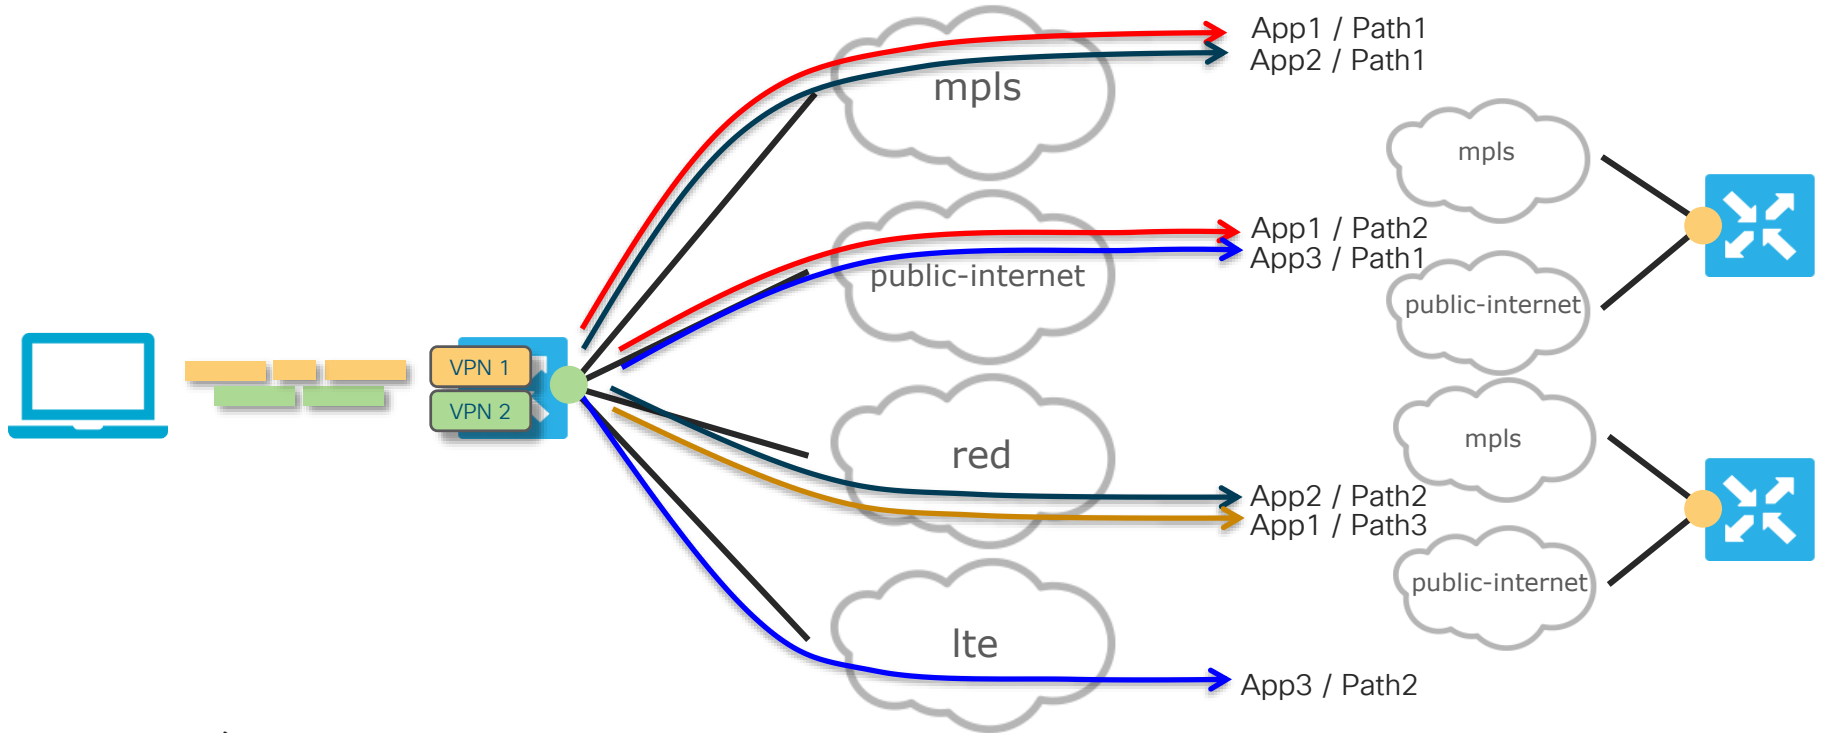


Application Pinning

Application Pinning

Transport selection per Application

- Local TLOC Selection: Loose preference, falls back to routing upon failure
- Remote TLOC Selection: Strict preference, traffic dropped upon failure



Application Pinning

Data Policy Configuration

Local TLOC
Prefer Local Underlay Path

```
vSmart
policy
data-policy local-tloc-preference
vpn-list VPN1
sequence 10
  match source-ip 10.0.0.0/8
  !
  action accept
  local-tloc red blue
```

- local-tloc - Loose match that will fall back to routing if all TLOCs in list are down
- tloc-list refers to specific remote TLOCs and will not fall back to routing

(Remote) TLOC
Prefer a remote Node/TLOC

```
vSmart
policy
data-policy local-tloc-preference
vpn-list VPN1
sequence 10
  match source-ip 10.0.0.0/8
  !
  action accept
  set
    tloc 1.1.1.1 color biz-internet

Or
  action accept
  set
    tloc-list remote-node
```

```
policy
lists
tloc-list remote-node
  tloc 1.1.1.1 color mpls encaps ipsec preference 100
  tloc 1.1.1.1 color biz-internet encaps ipsec preference 50
```

Internet Breakout – DIA / DCA

Internet Breakout / DIA

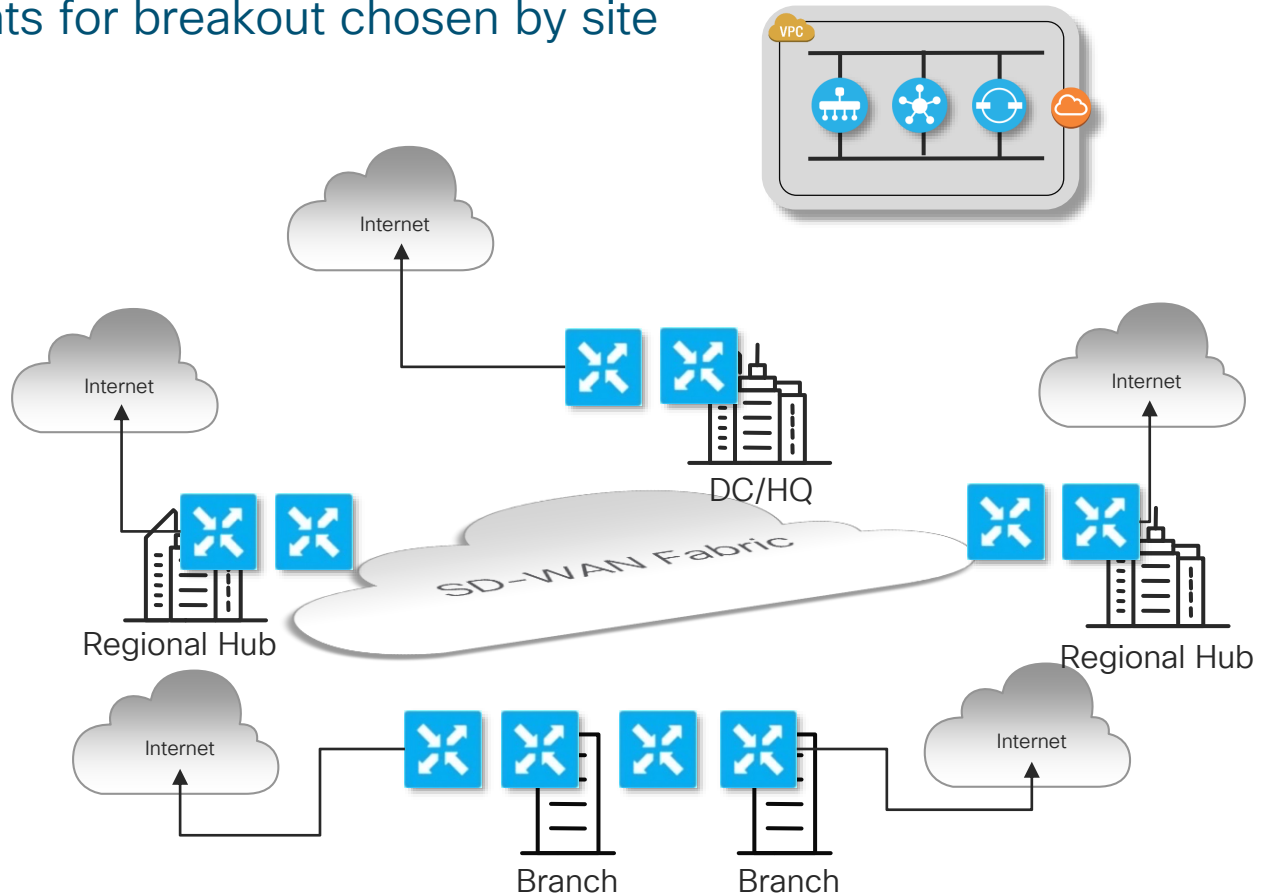
Routing and/or Policy-driven Capabilities

- The Cisco SD-WAN Architecture provides a lot of flexibility in enabling DIA
- Breakouts can be presented via:
 - Routing
 - Policy
 - In combination, with Preference and Backup options
 - Cloud-based Security as a Local Service using a Policy
- NAT is a required feature when providing a local breakout
- Service-side breakouts can be provided in case NAT is not needed or special care is needed for public addressing
- Can be deployed in combination with Service Chaining for monitoring/security/processing requirements

Internet Breakout Leverage


Most appropriate points for breakout chosen by site

- Enterprises can gradually progress from centralized to distributed breakouts
- Routing plane enables primary/backup as needed
- Policies further enhance selection and breakout granularity
- Align well with deployment of Cloud-based Security solutions



SD-WAN Internet Breakout Options

Local Breakout using a Default Route




```
vpn 0
 interface ge0/0
   nat
   tracker my_tracker
 !
vpn 1
 ip route 0.0.0.0/0 vpn 0
```

```
System
 tracker my_tracker
 endpoint-ip 1.2.3.4
 Interval 5
 Multiplier 3
 Threshold 500
```

- Static route in Service VPN
 - Can be default or more granular
- Redirects traffic to interfaces in VPN 0:
 - Interfaces must have NAT enabled
 - Multiple interfaces enables per-flow load-sharing
 - Relies on VPN 0 routing table
- Can be complemented with a Tracker to monitor Internet availability beyond first hop gateway

SD-WAN Internet Breakout Options

Local Breakout using Data Policy



Branch

```
vEdge
vpn 0
interface ge0/0
nat
```

```
vSmart
policy
data-policy internet-breakout
vpn-list VPN1
sequence 10
match source-ip 10.0.0.0/8
!
action accept
nat use-vpn 0
local-tloc public-internet
```

- Policy now redirects instead of static route
 - In case local exit fails, lookup can fall back to local service VPN routing table
- Redirects traffic to interfaces in VPN 0:
 - Interfaces must have NAT enabled
 - Multiple interfaces enables per-flow load-sharing
 - Relies on VPN 0 routing table
- Can be complemented with a Tracker to monitor Internet availability beyond first hop gateway
- Local TLOC to be used can be specified

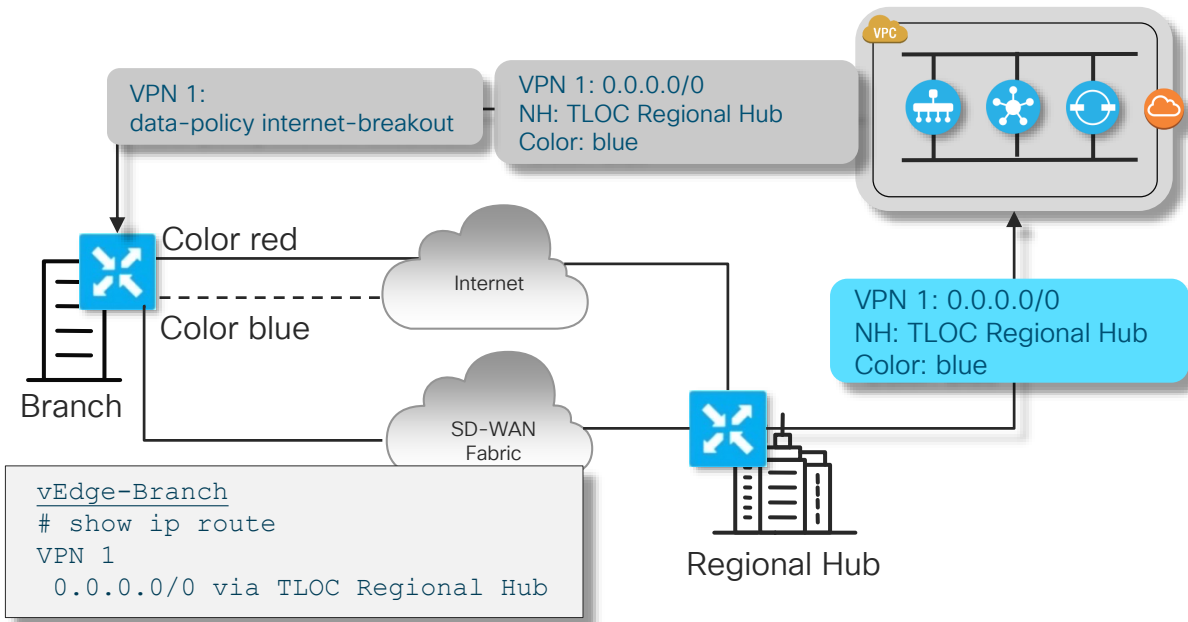
SD-WAN Internet Breakout Options

Joint Local and Regional Breakout using Data Policy + Routing

Legend:

Original Advertisement from Endpoint

Un/Modified Advertisement from Controller



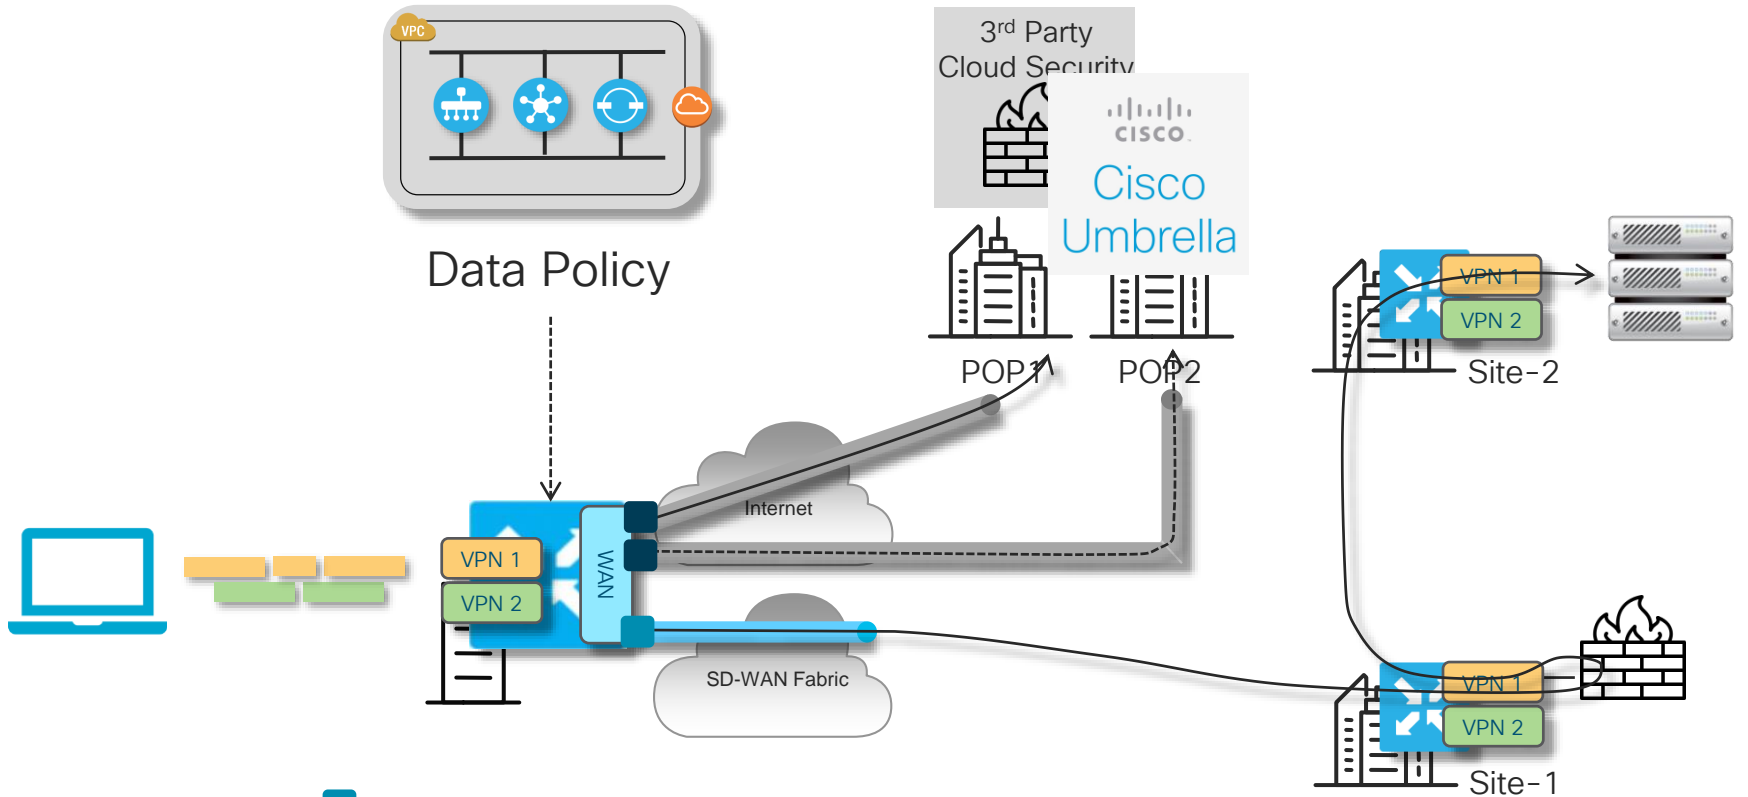
```
vSmart
policy
data-policy internet-breakout
vpn-list VPN1
sequence 10
match source-ip 10.0.0.0/8
!
action accept
nat use-vpn 0
local-tloc red blue
```

```
vEdge-Regional Hub
VPN 1
ip route 0.0.0.0/0 null0 or
default from OSPF/BGP
```

- Data Policy allows for granular breakout policy matching L3/L4/L7 information
 - Data Policy takes precedence
 - Default route from Regional Hub acts as backup in case TLOC Red & Blue are both down

SD-WAN Internet Breakout Options

Service Chaining – Cloud Security and Shared Services



SD-WAN Internet Breakout Options

Service Chaining – Local Services – Policy Configuration

```
vSmart
policy
  data-policy Cloud Security
  vpn-list vpn all
    sequence 10
    match protocol 6
    match destination-port 80 443
    !
    action accept
    set
      service FW local
    !
    !
    !
  default-action accept
```

2 Match Traffic

3 Apply Local Service

1 Define Local Service FW

```
WAN Edge
vpn 1
  service FW interface gre1 gre2
vpn 0
  interface ge0/0
    ip address 99.88.77.66/32
    no shutdown
  nat
  !
  interface gre1
    ip address 12.13.14.15/24
    tunnel-source-interface ge0/0
    tunnel-destination 123.123.123.123
    no shutdown
  !
  interface gre2
    ip address 16.17.18.19/24
    tunnel-source-interface ge0/0
    tunnel-destination 124.124.124.124
    no shutdown
```

Primary Tunnel

Backup Tunnel

- Data Policy redirection to locally configured service
- Service represented by local GRE or IPsec tunnel pre-configured on each WAN Edge

SD-WAN Internet Breakout Options

Service Chaining – Remote Services – Policy Configuration

```
vSmart
policy
  data-policy Central Security
  vpn-list vpn all
    sequence 10
    match protocol 6
    match destination-port 80 443
    !
    action accept
    set
      service FW vpn 1
    !
    !
    !
  default-action accept
```

2 Match Traffic

3 Apply OMP FW Service

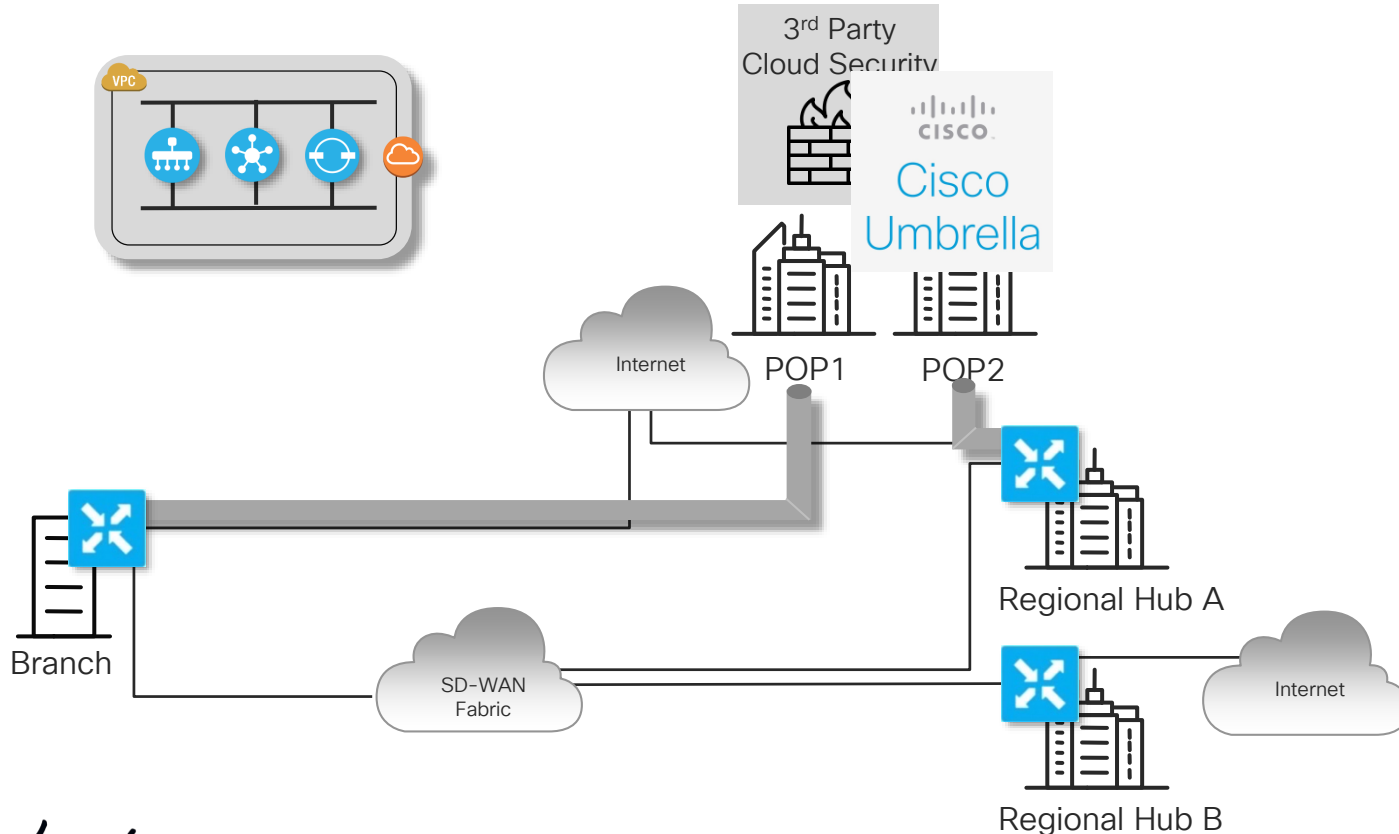
```
WAN Edge - Site1
vpn 1
  service FW address 12.13.14.100
  !
  interface ge0/0
    ip address 12.13.14.15/24
    no shutdown
```

1 Define Service FW for OMP Announcement

- Data Policy redirection to remotely configured service
- Service represented by OMP advertised service identifier
- Service association can be specified via TLOC or TLOC-list (with priorities) if needed

SD-WAN Internet Breakout Options

Joint Local and Regional Breakout using Data Policy and Cloud Security + Routing Preference



SD-WAN Internet Breakout Options

Joint Local and Regional Breakout using Data Policy and Cloud Security + Routing Preference

```
vSmart
policy
  data-policy Cloud Security
  vpn-list vpn all
    sequence 10
    match
      destination-data-prefix-list internal-prefixes
      !
    action accept
    !
  !
  sequence 20
  match
  !
  action accept
  count count_fw
  set
    service FW local [restrict]
  !
```

Exclude Internal Prefixes
from Internet Breakout

Any other traffic sent to
Internet Breakout

Drop Traffic if
Service Down

```
policy
lists
  data-prefix-list internal-prefixes
  ip-prefix 10.0.0.0/8
  ip-prefix 172.16.0.0/12
  ip-prefix 192.168.0.0/16
```

```
WAN-Edge-Branch
vpn 1
  service FW interface gre1
vpn 0
  interface gre1
  ip address 12.13.14.15/24
  tunnel-source-interface ge0/0
  tunnel-destination 123.123.123.123
  no shutdown
```

```
WAN-Edge-Regional Hub A
vpn 1
  service FW interface gre1
  ! ip route 0.0.0.0/0 null0 or
  ! default from OSPF/BGP
```

```
WAN-Edge-Regional Hub B
vpn 1
  ! ip route 0.0.0.0/0 null0 or
  ! default from OSPF/BGP
```

SD-WAN Internet Breakout Options

Joint Local and Regional Breakout using Data Policy and Cloud Security + Routing Preference

vSmart Control Policy

```
vSmart
Policy
lists
  prefix-list default_route
    ip-prefix 0.0.0.0/0
  !
!
control-policy default_priority
sequence 10
  match route
    prefix-list default_route
    site-id Regional Hub A
  !
  action accept
  set
    preference 100
  !
!
!
default-action accept
```

Default from Hub A gets
higher preference

WAN Edge Static TLOC preference

```
WAN-Edge-Regional Hub A
vpn 0
  interface ge0/0
    tunnel-interface
      encapsulation ipsec preference 100
  !
!
vpn 1
  ! ip route 0.0.0.0/0 null0 or
  ! default from OSPF/BGP
```

```
WAN-Edge-Regional Hub B
vpn 0
  interface ge0/0
    tunnel-interface
  vpn 1
  ! ip route 0.0.0.0/0 null0 or
  ! default from OSPF/BGP
```

SD-WAN Internet Breakout Options

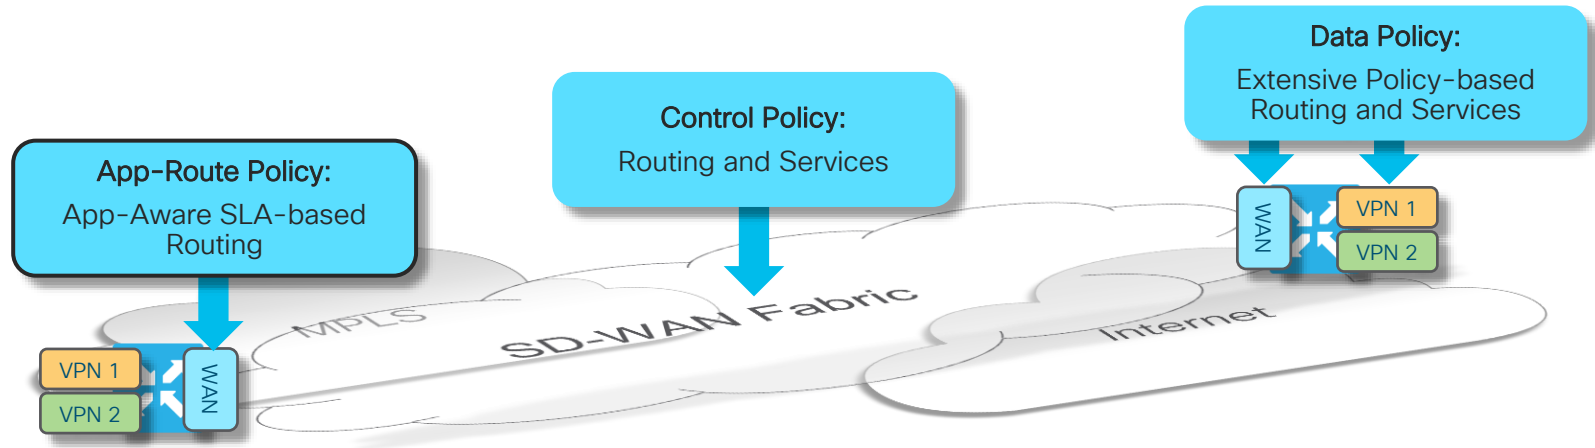
Application Specific Breakout

- The Data Policy construct can also be used to locally breakout specific applications with defined DPI signatures (e.g. O365, FaceBook, Youtube)
- Example:
 - Office365 to be locally broken out
 - All other Internet traffic via regional exit
- Arrangements required for supporting O365
 - Data Policy for breaking out locally
 - Default route from regional exit for two purposes:
 - Breakout for all non O365 traffic
 - O365 session establishment involves quite a few protocols beyond the core O365 protocols – A default route from somewhere is required to deal with those applications and allow for successful O365 operations
- SD-AVC support to provide Application Recognition from the first packet

SLA-Driven Path Selection using App-Route Policies

Cisco SD-WAN Policy Architecture

Suite of Policies to address different functional domains



- Control Policies are applied at vSmart: Tailors routing information advertised to Edge endpoints
- App-Route Policies are applied at WAN Edge: SLA-driven path selection for applications
- Data Policies are applied at WAN Edge: Extensive Policy driven routing

App-Route Policies

App-route Components and Dependencies

BFD Settings

BFD rx_interval and multiplier settings
(only rx_interval is relevant to AAR)

```
bfd
color <color>
hello-interval <msec>
multiplier <number>
```

App-route algorithm configuration

Define how SLA data is used to influence path selection

```
bfd
app-route
multiplier <number>
poll-interval <msec>
```

App-route Policy Definition

Define SLA-classes, Application associations, VPN applicability and Policy actions/preferences

```
SLA-classes
Policy Construct
match
action
```

DPI Engine Enablement

AAR relies on DPI for L7 signatures

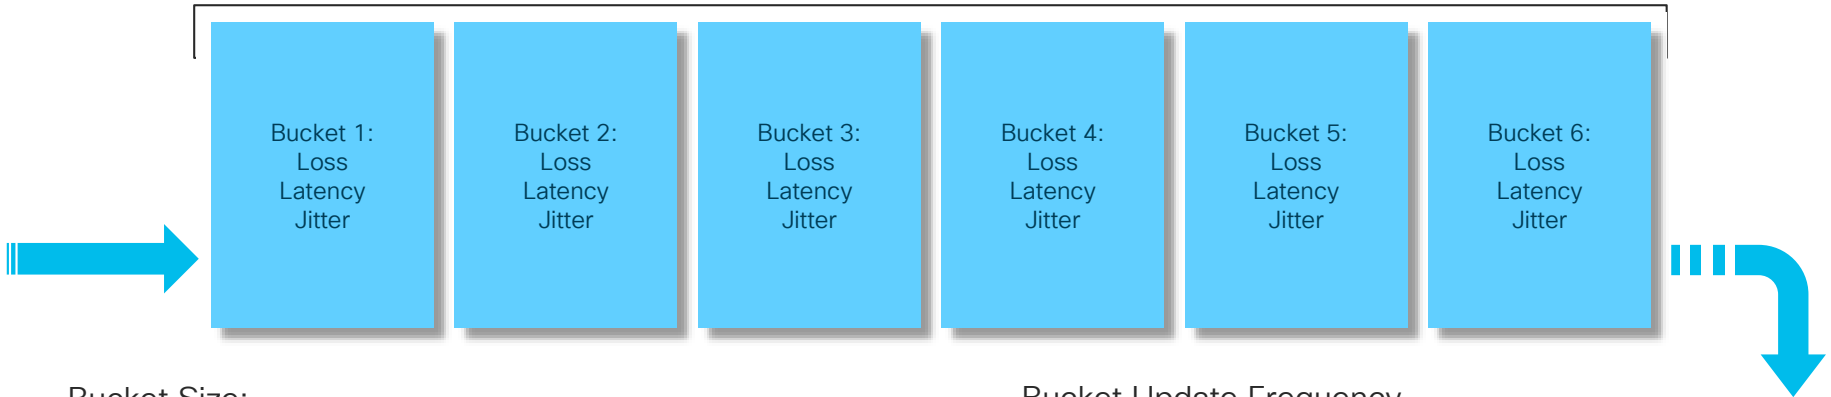
```
policy
app-visibility
```

https://sdwan-docs.cisco.com/Product_Documentation/Software_Features/Release_18.4/07Policy_Applications/01Application-Aware_Routing/01Configuring_Application-Aware_Routing

App-Route Policies

App-route Algorithm

Avg (B1 + B2 + B3 + B4 + B5 + B6) = Mean
Mean recalculated every Bucket completion cycle



Bucket Size:

`bfd`

`app-route poll-interval (default 600,000 ms)`

Bucket Update Frequency

`bfd`

`hello-interval (default 1000ms)`

of Buckets:

`bfd`

`app-route multiplier (default 6)`

App-Route Policies

App-route Policy Definition

SLA Classes
Loss, Latency, Jitter per Class

```
Policy
sla-class <name>
jitter <msec>
latency <msec>
loss <percentage>
```

App-list
Use L3/L4 or DPI Signatures

```
Policy
lists
app-list <name>
app <name> | app-family <family>
```

App-route Policy
VPN applicability and Policy actions/preferences

App-route Logging
Enable logging of packet headers

App-Route Policies

App-route Policy Definition

SLA Classes

Loss, Latency, Jitter per Class

App-list

Use L3/L4 or DPI Signatures

App-route Policy

VPN applicability and Policy actions/preferences

App-route Logging

Enable logging of packet headers

- 1 For traffic not explicitly matched in policy
- 2 For traffic with an SLA-class disqualified across all links
- 3 Drop traffic if SLA-class is disqualified
- 4 One or more preferred colors if multiple links qualify

```
Policy
app-route-policy <name>
vpn-list <vpn-list>
default-action sla-class <name> 1
sequence <number>
match
...
action
backup-sla-preferred-color [list] 2
count <name>
log
sla-class <name> [strict] [preferred-color [list]]
3 4
```

App-Route Policies


Policy Example

```
policy
 lists
  vpn-list VPN1
  vpn 1
  !
  site-list app-route-sites
  site-id 3003
  !
  app-list AVV
  app-family audio_video
  !
  app-list SFDC
  app salesforce
  !
```

```
apply-policy
 site-list app-route-sites
 app-route-policy SLA-Routing
```

```
Policy
 sla-class EF
  loss 1
  latency 100
  !
 sla-class Biz-apps
  loss 2
  latency 150
  !
 app-route-policy SLA-Routing
  vpn-list VPN1
  sequence 10
  match app-list AVV
  !
  action
  sla-class EF
  !
  !
  sequence 20
  match app-list SFDC
  !
  action
  sla-class Biz-apps
  !
  !
```

App-route Policy Path Convergence



Current Mean Latency is 20ms, when Latency jumps to 150ms as Bucket 1 collection starts

Recommended Settings and Operational Best Practices

vManage Statistics Collection

Configuration and Volumes

The screenshot shows the 'Statistics Setting' section in the vManage interface. It lists various categories with radio buttons for 'Enable All', 'Disable All', and 'Custom'. All categories are currently set to 'Enable All'.

Category	Enable All	Disable All	Custom
Approute	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bridge Interface	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
BridgeMac	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
CloudExpress	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Device System Status	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
DPI	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Flow Log	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Interface	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Wlan Client Info	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

- Configure collection per category and per device
- Custom allows to control collection of each category on a per device basis

The screenshot shows the 'Statistics Database Configuration' page with a table of storage usage. The total storage used is 0.3713 GB, and the maximum available space is 59.0586 GB. The table lists various statistics types and their current and maximum sizes.

Statistics Type	Current Size(GB)	Size(GB)
Audit Log	0.0053	5
Interface	0.0145	5
Device Configuration	0.0001	5
Device System Status	0.192	5
BridgeMac	0	5
DPI	0	5
Bridge Interface	0	5
Approute	0.1325	5
Total	0.3713 GB	70.0000 GB

- Storage can be assigned for individual categories to reflect:
 - Collection not being enabled
 - Storage assignments and data lifetime

Overlay and vEdge Recommended Settings

Useful Settings to get Right the First Time

- System-IP
 - Pick a range for the entire network that does not overlap with other addressing
 - Not routed but significant to anything present in VPN 0 / Transport
 - An incorrectly chosen range or System-IP setting can cause connectivity issues
- Site-ID
 - The target for policy application and identifier of routing sources (ref: BGP AS)
 - Several schemes documented and one is discussed later on
- Vmanage connection preference
 - Determines which TLOC is used for vManage traffic (statistics upload etc)
 - Advised to use the highest bandwidth link and avoid cellular interfaces
- Max-control-connections
 - Determines how many vSmart sessions are established per TLOC
 - For Transports without controller access, it must be set to Zero (0)

Template Creation Guidelines

Templates are Friends

- Plan for template creation and test out features to be deployed
 - Allows for the optimization of template structure and maintenance
- Use a simple "bootstrap" template for distributed devices that are not yet in production
 - The device is then in a known state and vManaged
 - Tracking events is easier if a logical name is applied
 - The local configuration of the device can't be changed
 - The device can be moved to production (or any other state) at will from vManage
- The template can be changed at any time from within vManage
- Template Variables can be managed in several different ways:
 - Entered manually at time of template attachment
 - Stored in a .csv file that is referenced at time of template application
 - Using the REST API (possibly in conjunction with other platforms such as Infoblox)

Template Creation

Feature Template Components and Sources

Device Template - Aggregate Configuration Template

CONFIGURATION | TEMPLATES

Device Model: C1311-IP13TEAR0

Template Name: C1000-Template

Description: C1000-Template

Basic Information

System: Factory_Default_cEdge_System_Template

Logging: Factory_Default_Logging_Template

AAA: Factory_Default_AAA_Template

OMP: Factory_Default_cEdge_OMP_ipv46_Template

BFD: Factory_Default_BFD_Template

Security: Factory_Default_cEdge_Security_Template

Additional System Templates: NTP

Additional Templates

AppQoS: Choose...

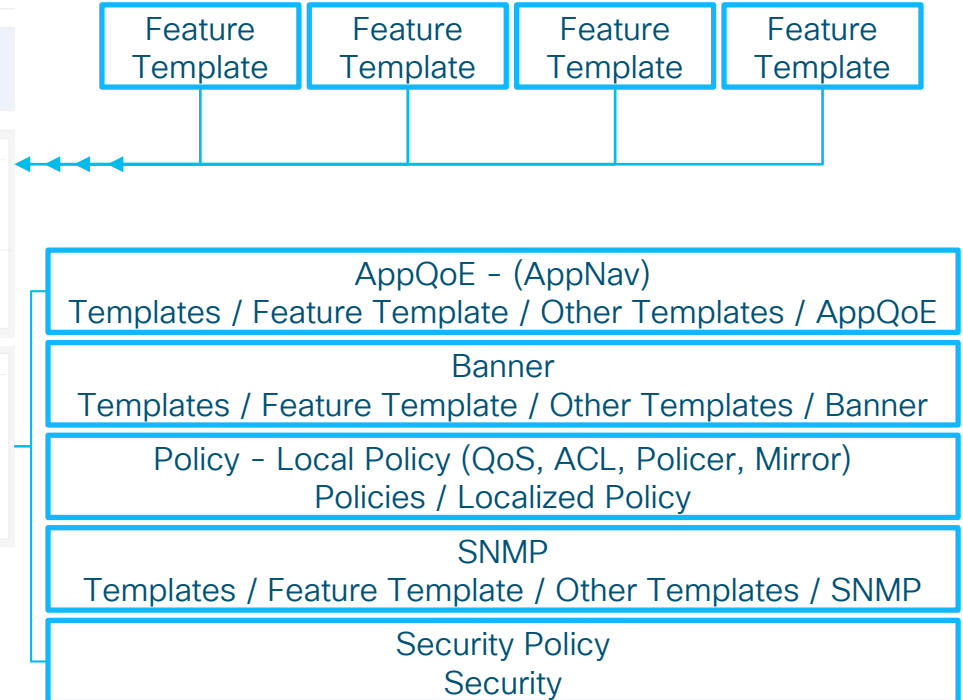
Banner: Choose...

Policy: Choose...

SNMP: Choose...

Security Policy: Choose...

Dedicated or Shared Feature Templates



Template Creation - Device Template

Optimizing object use in a Device Template - Optional Objects

The screenshot shows the configuration page for an IPv4 route. At the top, there is a header 'IPv4 ROUTE' and a button '+ New IPv4 Route'. Below this, there is a form with the following fields and options:

- Prefix:** A text input field with a globe icon and a dropdown arrow.
- Gateway:** Three radio button options: 'Next Hop' (selected), 'Null 0', and 'VPN'.
- Next Hop:** A text input field with a '+ Add Next Hop' button.
- Mark as Optional Row:** A checkbox with an information icon, highlighted by a dashed red box.
- Buttons:** 'Add' and 'Cancel' buttons at the bottom right.

- Using Device Templates, quite a few objects can be tagged as Optional
- Simply not assigning a value at template application leaves the object out of the created configuration
- This makes Device Templates flexible to support a variety of different configurations

Template Creation - CLI Template

Optimizing object use in CLI template by means of variables

The image consists of three screenshots from a CLI configuration editor, illustrating the process of creating a variable for a block of configuration text.

- Left Screenshot:** Shows a CLI configuration for a VPN. Lines 82-97 are highlighted in grey, representing the configuration for interface ge0/1. The configuration includes settings for description, ip dhcp-client, tunnel-interface, encapsulation ipsec, color mpls, and various allow-service commands.
- Middle Screenshot:** Shows the 'Create Variable' dialog box. A red dashed box highlights the configuration block from the left screenshot. The dialog box contains the text: 'Replacing Text: interface ge0/1 description "MPLS" ip dhcp-client tunnel-interface encapsulation ipsec color mpls 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 shutdown !'. Below this, the 'Variable Name:' field is set to 'Optional MPLS Intf'. A red arrow points from the highlighted configuration block to the 'Create Variable' button.
- Right Screenshot:** Shows the resulting CLI configuration. The configuration for interface ge0/1 is now wrapped in a variable definition: `!{{Optional MPLS Intf}}!`. A red dashed box highlights this variable definition. A red arrow points from the 'Create Variable' button in the middle screenshot to this variable definition in the right screenshot.

- In a CLI template, an arbitrary number of lines can be turned into a variable
- Assigning this variable a “;” at template application leaves the section out of the created configuration
- This makes CLI Templates flexible to support a variety of different configurations

CISCO *Live!*

Policy Creation and Management Guidelines

Really not different from standard operations

- Define Requirements up front
 - Important Applications
 - Segmentation and Connectivity Models
 - SLA and QoS Requirements
 - Application Pinning, Breakout, Hosting, Routing i.e. Application Management Requirements
- Use a sandbox for verification and testing
 - A separate domain where policies and requirements can be tested
 - Can be part of the production network, simply a separate Site-ID range
- Limit Policy Management to a few capable resources

Construction of SD-WAN Policies

Policy Building Blocks

Lists


Application
Color
Data Prefix
Policer
Prefix
Site
SLA Class
TLOC
VPN

Policy

Policy Type
Policy Sequence 1
Match <route tloc Application>
Action <Accept Reject set >
Policy Sequence 2
Match <route tloc Application>
Action <Accept Reject set >
Default Action
<Accept Reject>

Apply Policy

Site-List
Policy <type> <name>
Direction (if applicable)



Policy Management

Best Practices



- Create and Maintain separate Lists and Policies per network region (and a sandbox if possible)
- Make modifications to a copy of the original and swap the copy with the original when applying
- More complex policies can be large and updates should be tested before applied to the live network



TECCRS-2014
SD-WAN Technical Deep Dive 8 Hours



TECRST - 2191
SD-WAN design, deploy and best
practices 4 Hours



TECCRS-3006
ENFV Deep Dive and Hands on Lab 8 Hours

Cisco SD-WAN



CISCO *Live!*

Tectorials



SD-WAN

Breakouts

CISCO *Live!*

- Keynote 09:30
- BRKCRS-1579 SD-WAN Powered by Meraki 11:00
- BRKRST-2041 WAN Architecture and Design Principal 11:00
- BRKCRS-2110 Delivering Cisco Next gen SD-WAN with Viptela 14:00
- BRKCRS-2113 Cloud Ready WAN for IAAS and SAASA with Cisco SD-WAN 17:00

- BRKRST-2377 SD-WAN Security 08:00
- BRKRST-2095 SD-WAN Routing Migration 16:00
- BRKRST-3404 How to choose the correct branch device 16:00

- BRKRST-2791 Building and using Policies with Cisco SD-WAN 08:00
- BRKRST-2560 SD-Wan Machine Analytics, Machine Learnings and IA 08:00
- BRKRST-2096 SD-Wan Proof Of Concept 11:00
- BRKRST-2093 Deploy, monitor and troubleshoot 11:00
- BRKARC-2012 ENFV Architecture, Configuration and troubleshooting 11:00
- BRKRST-2559 3 Steps to design SD-WAN On Prem 14:00
- BRKRST-2097 Conquer the Cloud with SD-WAN 14:45
- BRKRST-2095 SD-WAN Routing Migrations 16:45
- Keynote 17:00
- Cisco Live Celebration 18:30

- BRKRST-2091 SD-WAN Datacenter and Branch Integration Design 09:00
- BRKOPS-2826 SD-WAN as Managed Services 11:00

Continue your education



Demos in the
Cisco campus



Walk-in labs



Meet the engineer
1:1 meetings



Related sessions



Thank you





You make **possible**