

CISCO *Live!*



#CiscoLive



The bridge to possible

# Multi-Region Fabric

(Formerly, Hierarchical SD-WAN)

## Overview and Principles

Hamzah Kardame

Leader, Product Management, SD-WAN

BRKENT-2292



#CiscoLive

# 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 June 17, 2022.



<https://ciscolive.ciscoevents.com/ciscolivebot/#BRKENT-2292>



# Agenda

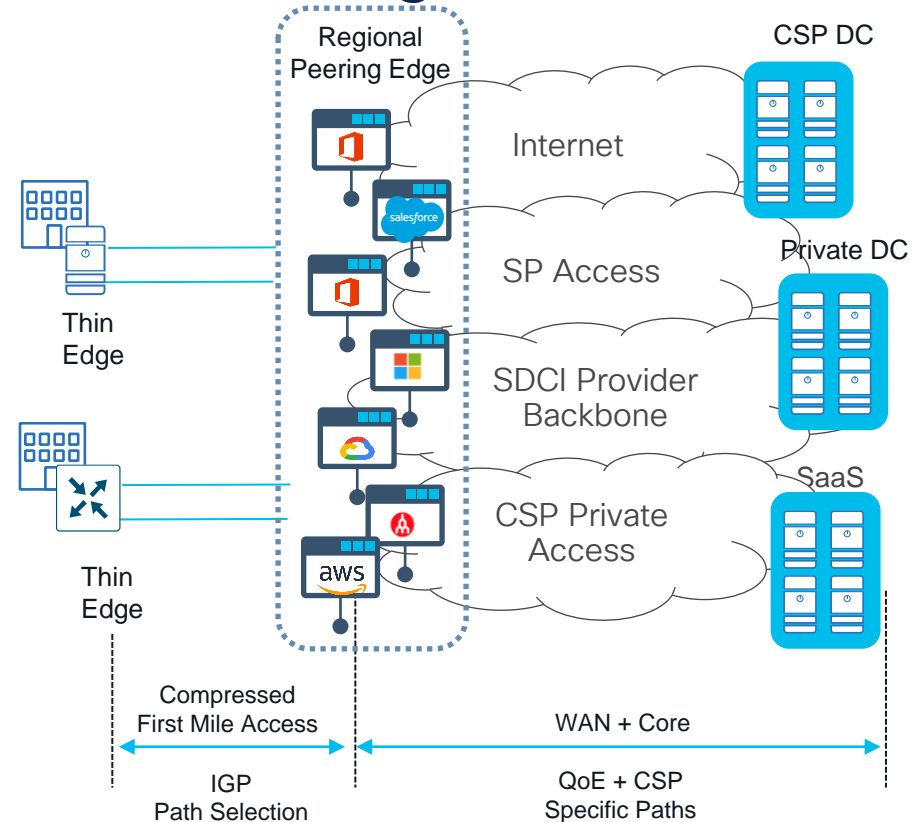
- Introduction
- Multi-Region Fabric – A Quick Look
- What's under the hood?
- Configuration
- Conclusion

# Introduction

# WAN is Evolving to a Service Exchange

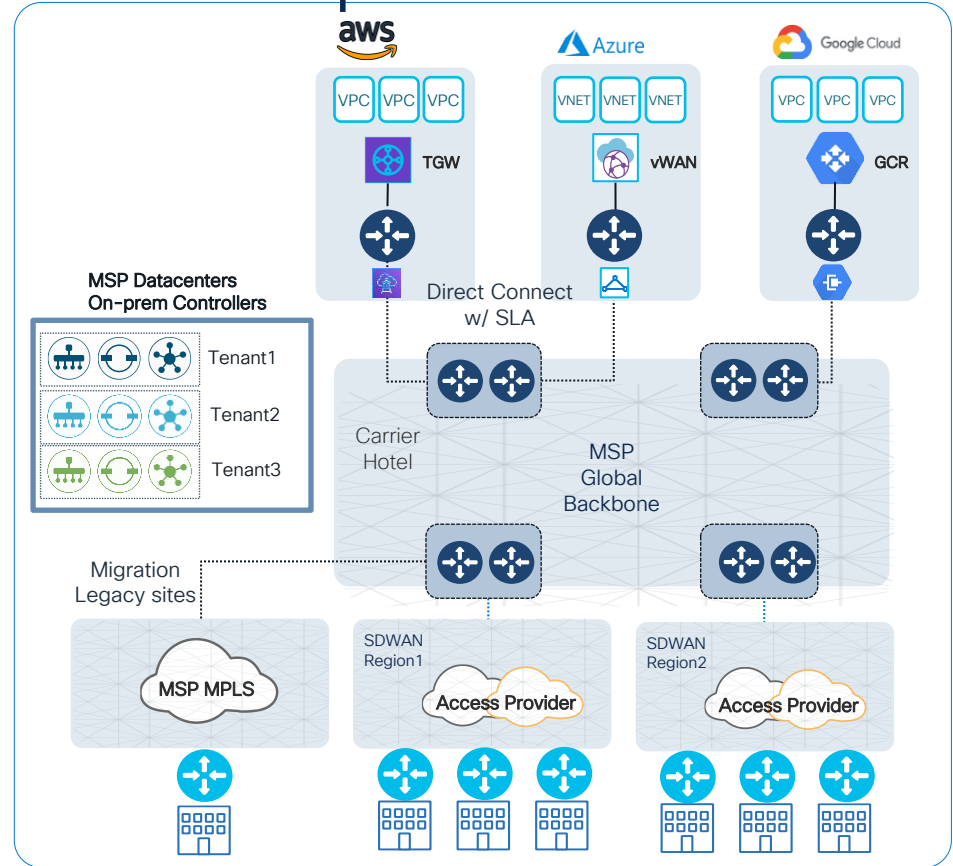
- The internet is changing from a network-of-networks to a network of data centers
- SDCI\* & Multiple Provider Backbones
- Large POP and Colo footprint
- Short-term Contracts, Usage-based
- Trending toward single ISP first-mile access
- On Demand

\*SDCI - Software Defined Cloud Interconnect

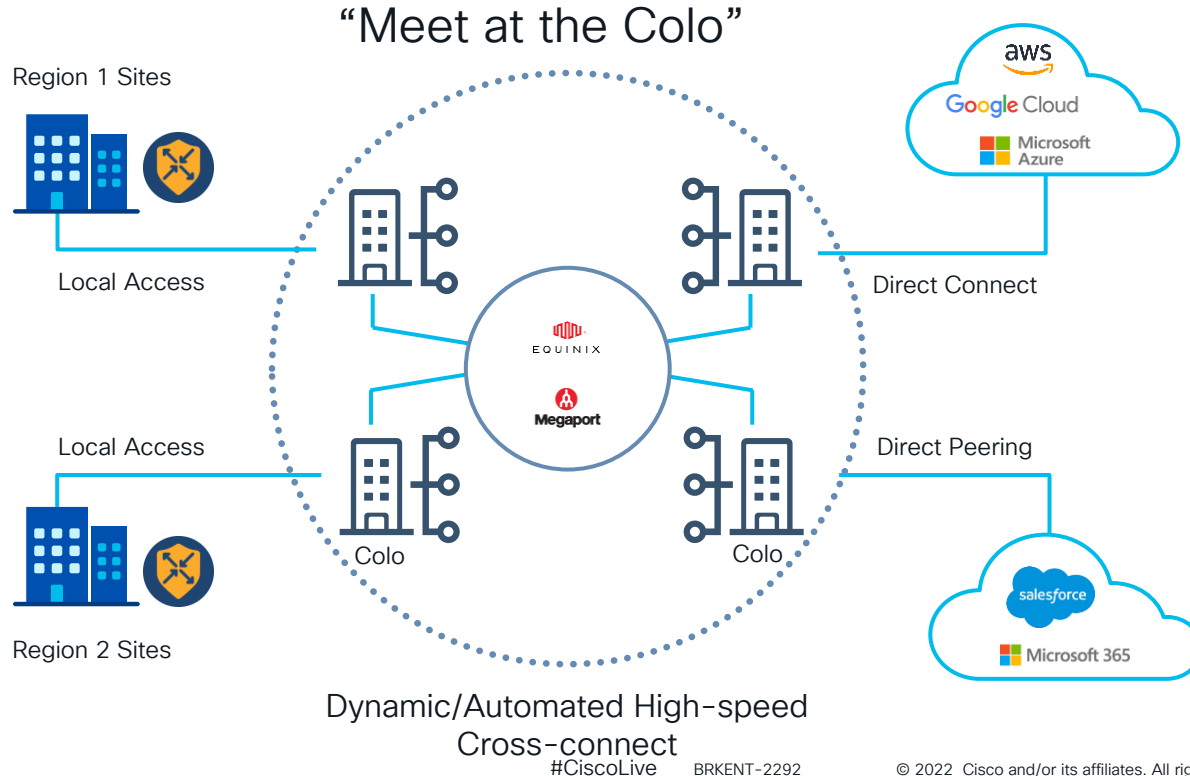


# MSP SD-WAN and middle-mile optimization

- SD-WAN as a Managed Service, evolving to SD-WAN as a Service
  - High Speed Backbone
  - SDN-POPs – Next Generation POPs with Compute Platforms running Openstack
- Cisco SD-WAN Overlay Networks
  - Gateways per customer (on Openstack)
- Access to Multi Cloud services – Private peerings with SLA
- Hybrid access connectivity

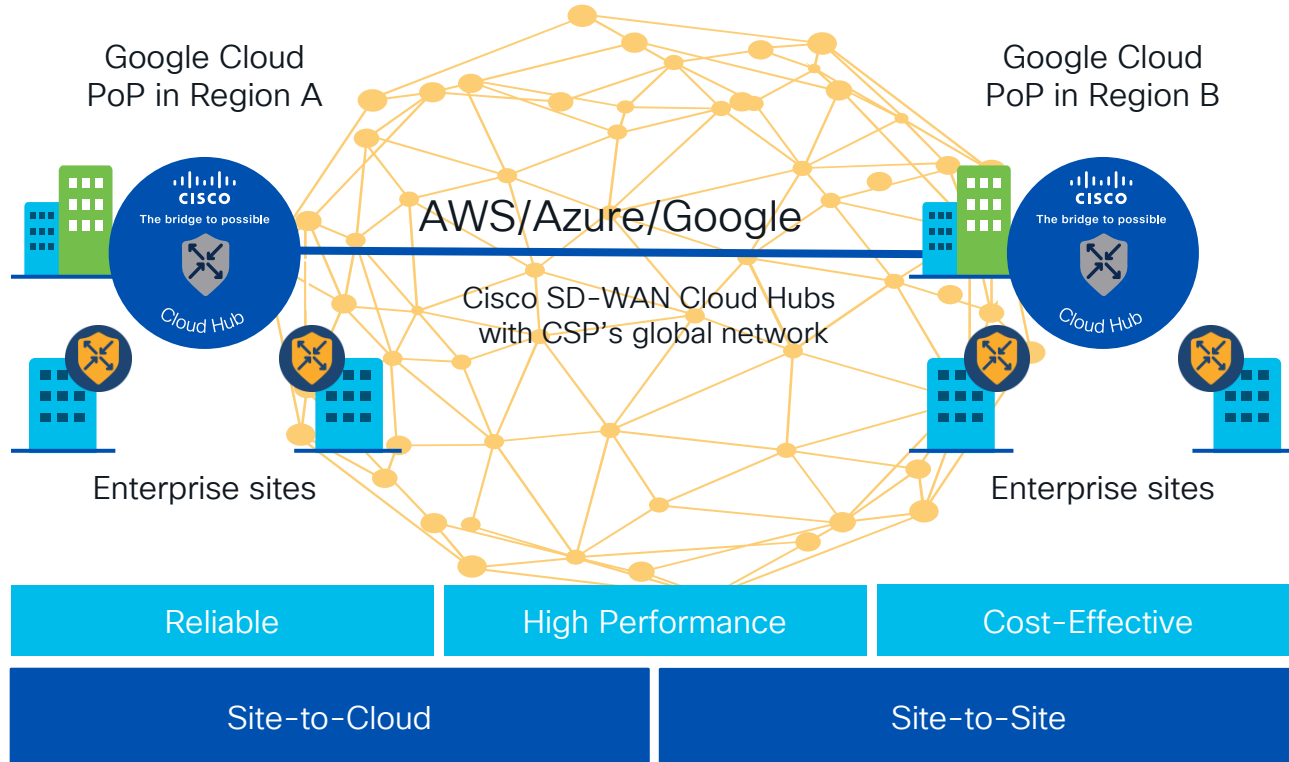


# SDCI providers and middle-mile optimization

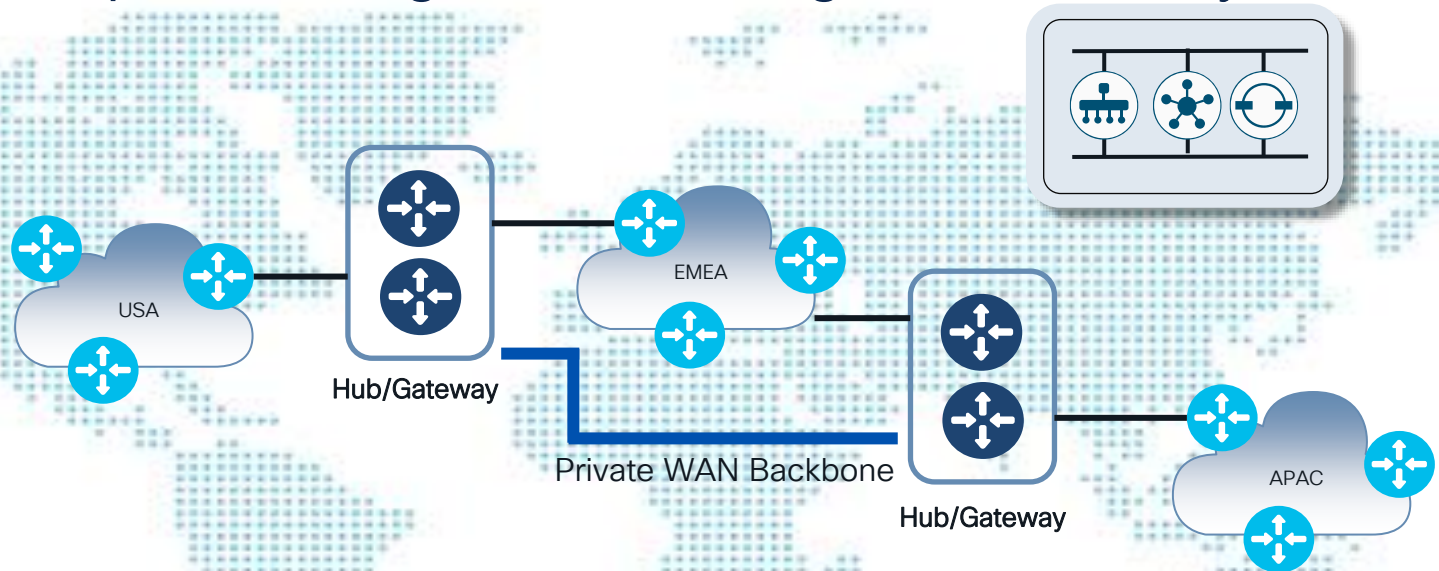




# Cloud Service Provider (CSP) SD-WAN Architecture



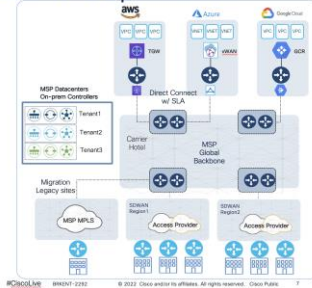
# Large Enterprise – Regional Meshing and Gateways



# Typical WAN Blueprint

## MSP SD-WAN and middle-mile optimization


- SD-WAN as a Managed Service, evolving to SD-WAN as a Service
- High Speed Backbone
- SDN-POPs – Next Generation POPs with Compute Platforms running Openstack
- Cisco SD-WAN Overlay Networks
  - Gateways per customer (on Openstack)
- Access to Multi Cloud services – Private peerings with SLA
- Hybrid access connectivity



CISCO Live!

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public


## SDCI providers and middle-mile optimization



CISCO Live!

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public


## Cloud Service Provider (CSP) SD-WAN Architecture



CISCO Live!

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

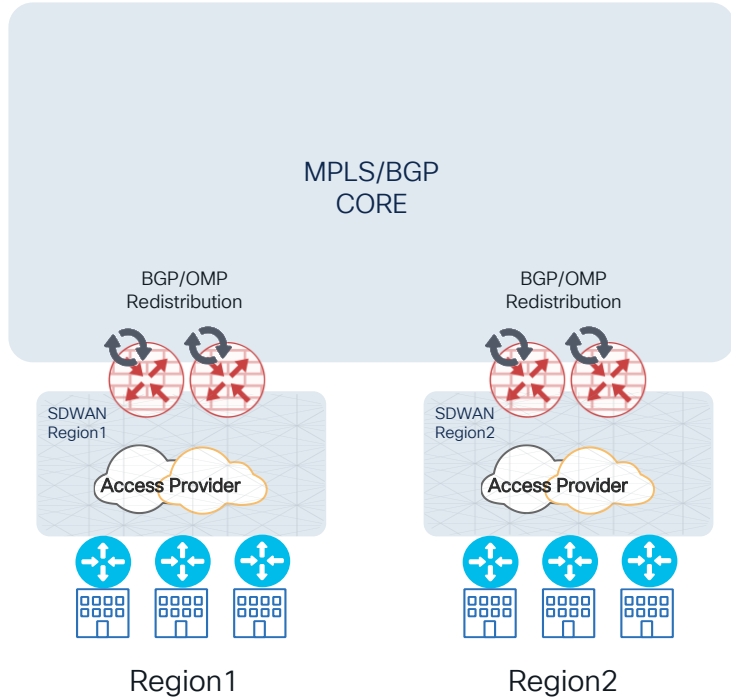
## Large Enterprise – Regional Meshing and Gateways



CISCO Live!

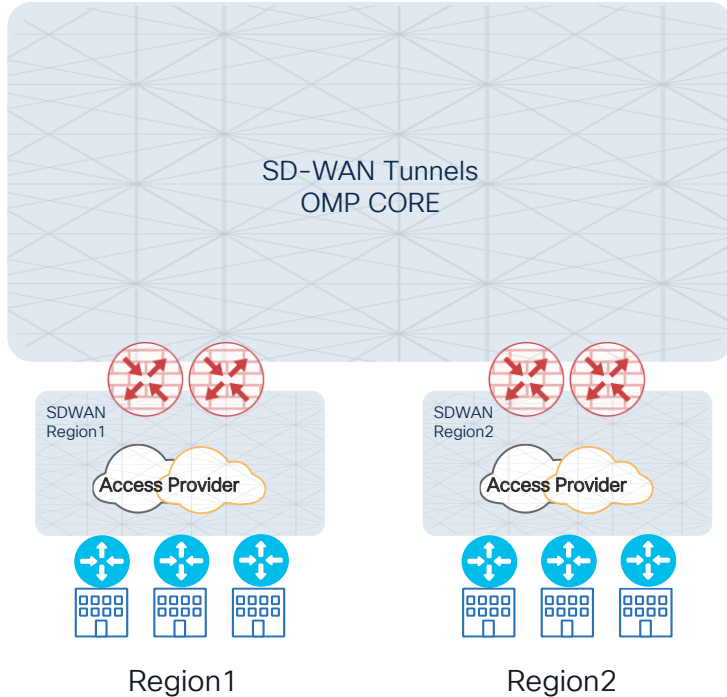
© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

# BGP Core Complexity



- Needs one MPLS-VPN VRF per customer overlay in the core
- OMP/BGP redistribution on each GTW
- Tagging to avoid routing loops
- Full routing must be advertised to vSmart by each GTW, to avoid blackhole
- No end-to-end path performance monitoring
- Only available on MPLS (no support for Public Internet Core)

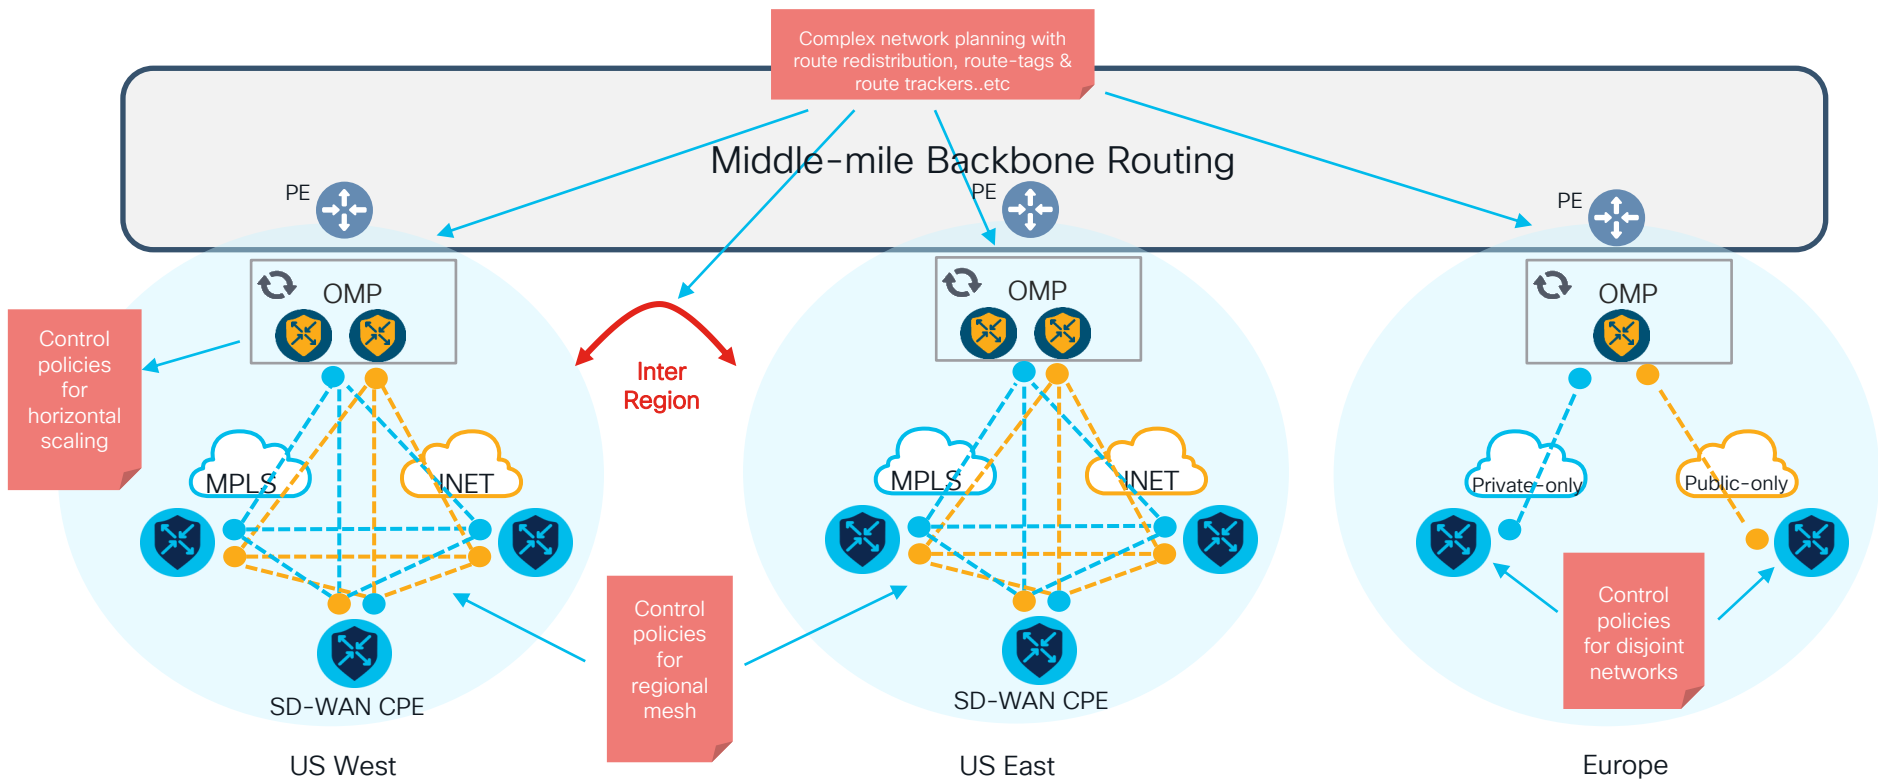

# OMP Multi Region Limitations



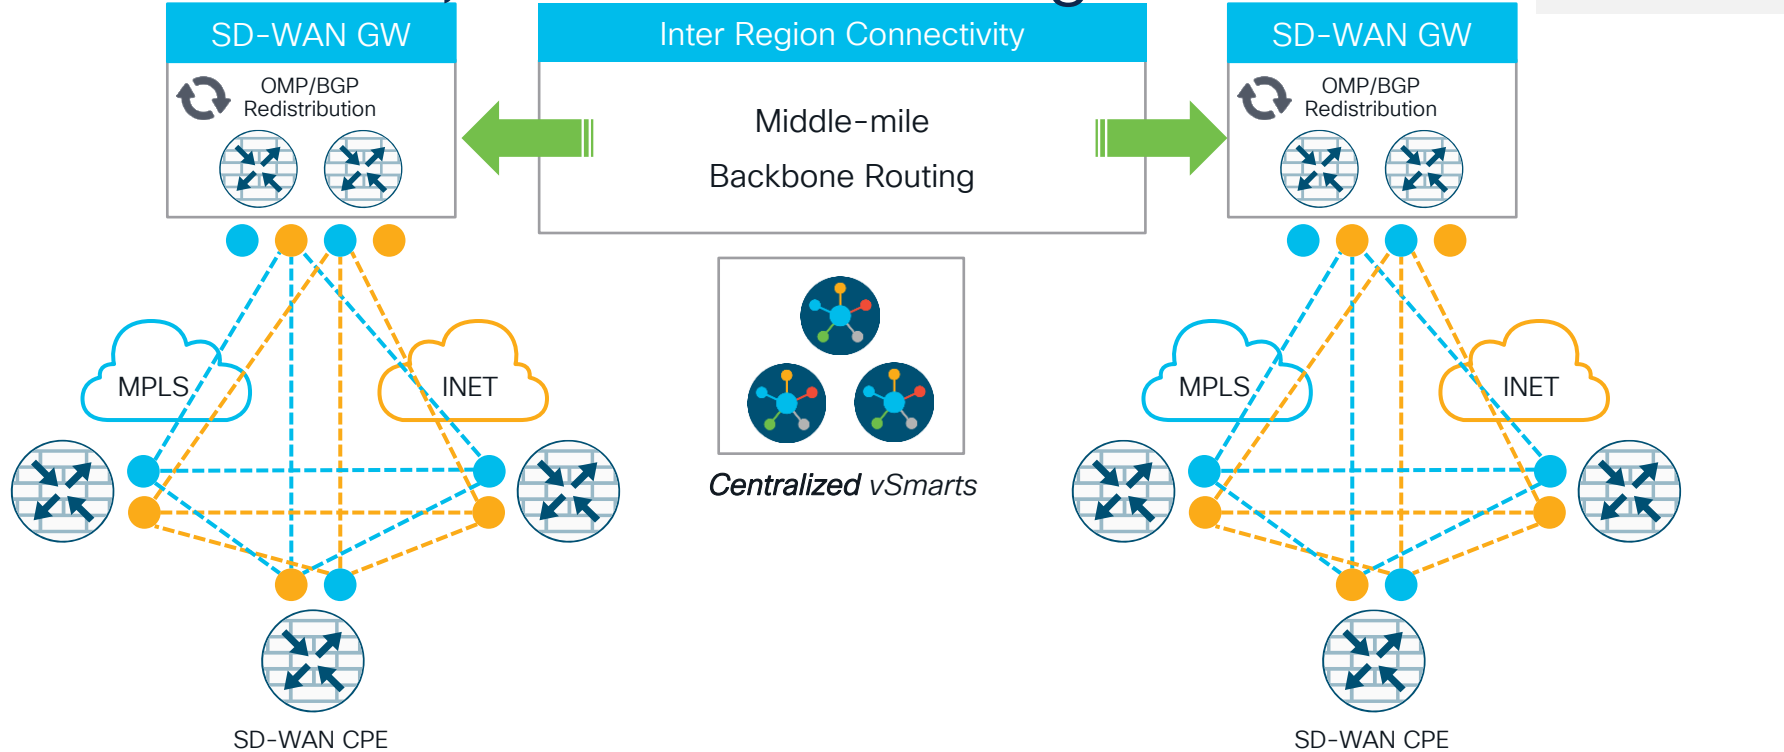
- Current SDWAN architecture supports a flat overlay model, where site to site tunnels are directly established between them
- To define regions, use Control Policies
  - Restrict direct tunnels to devices in the same region
  - Change next-hop for routes outside of the region. Replace with GWs
  - Static definition of routing across regions (aka PBR like)
  - Complex Control Plane Policies
- Potential Traffic blackholing for access and core failures
- No route summarization
- No end-to-end path forwarding decision

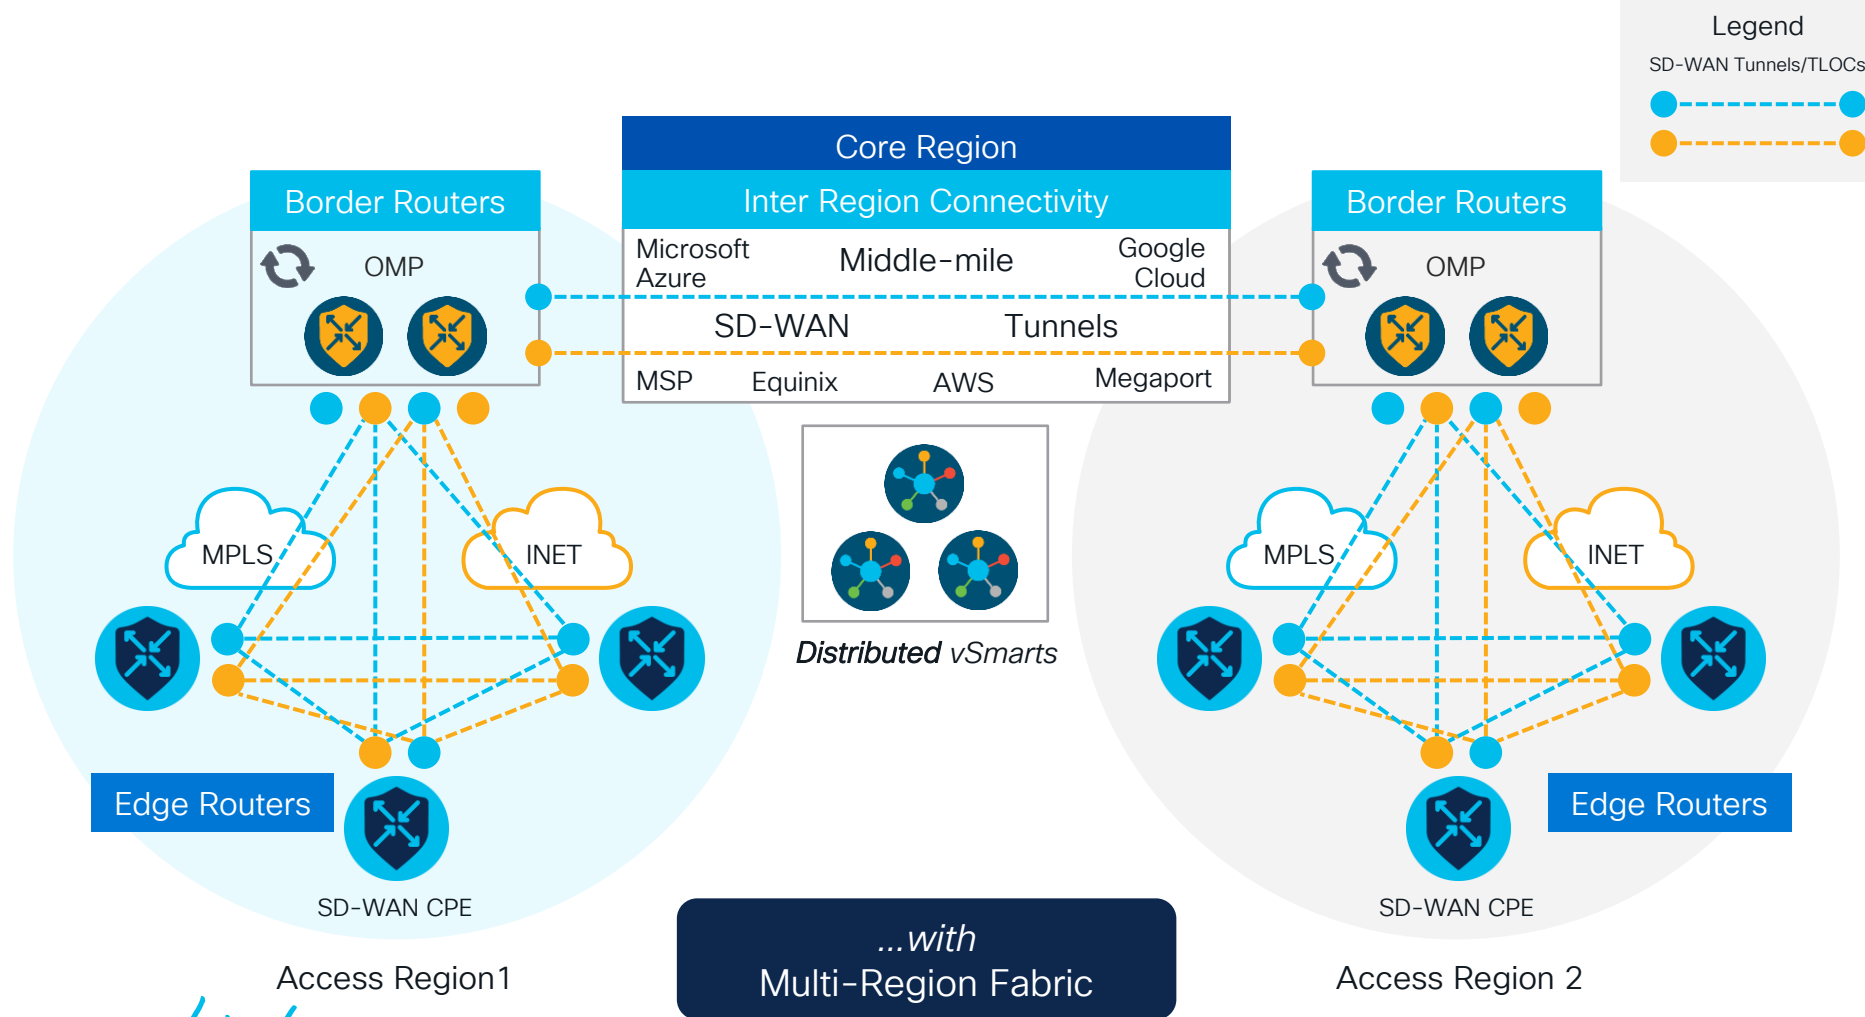
# Without Multi-Region Fabric

## Summary



# The Network, *without* Multi-Region Fabric



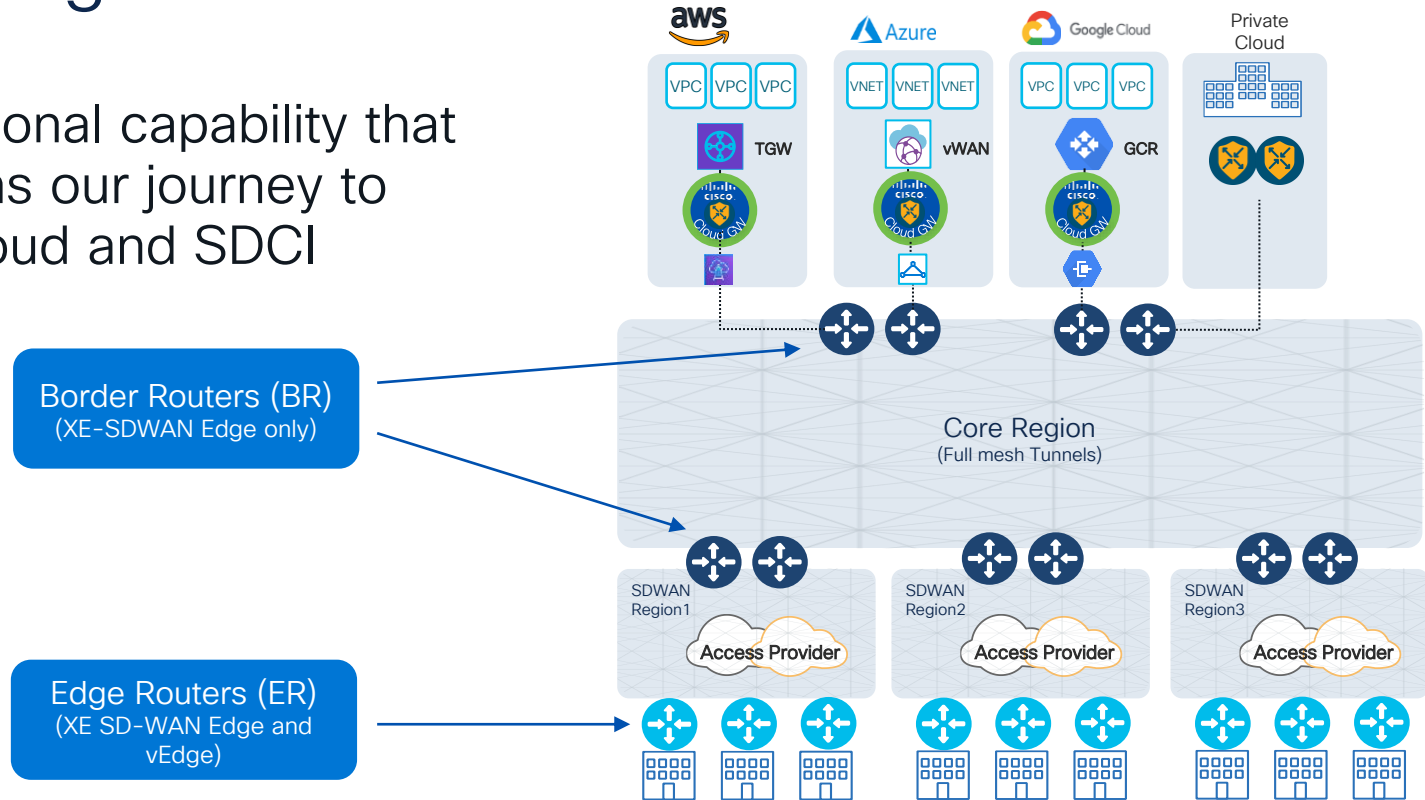




# Multi-Region Fabric *A Quick Look*

# Multi Region Fabric – New Roles

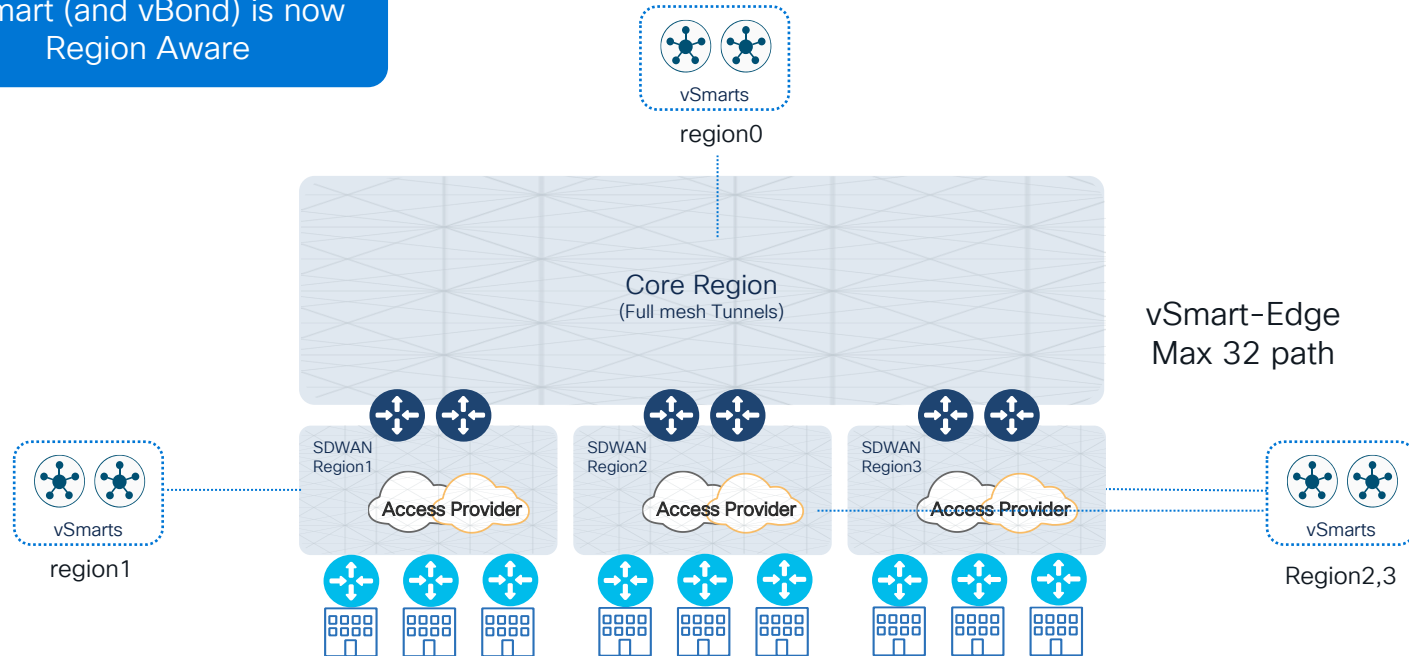
Foundational capability that underpins our journey to multi-cloud and SDCI



# Distributed vSmart Controllers

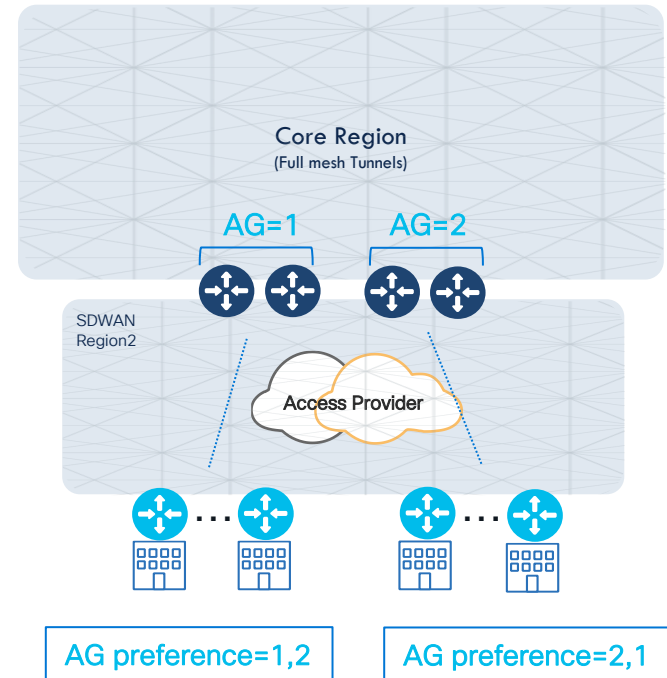
vSmart-vSmart  
Max 128 path

vSmart (and vBond) is now  
Region Aware

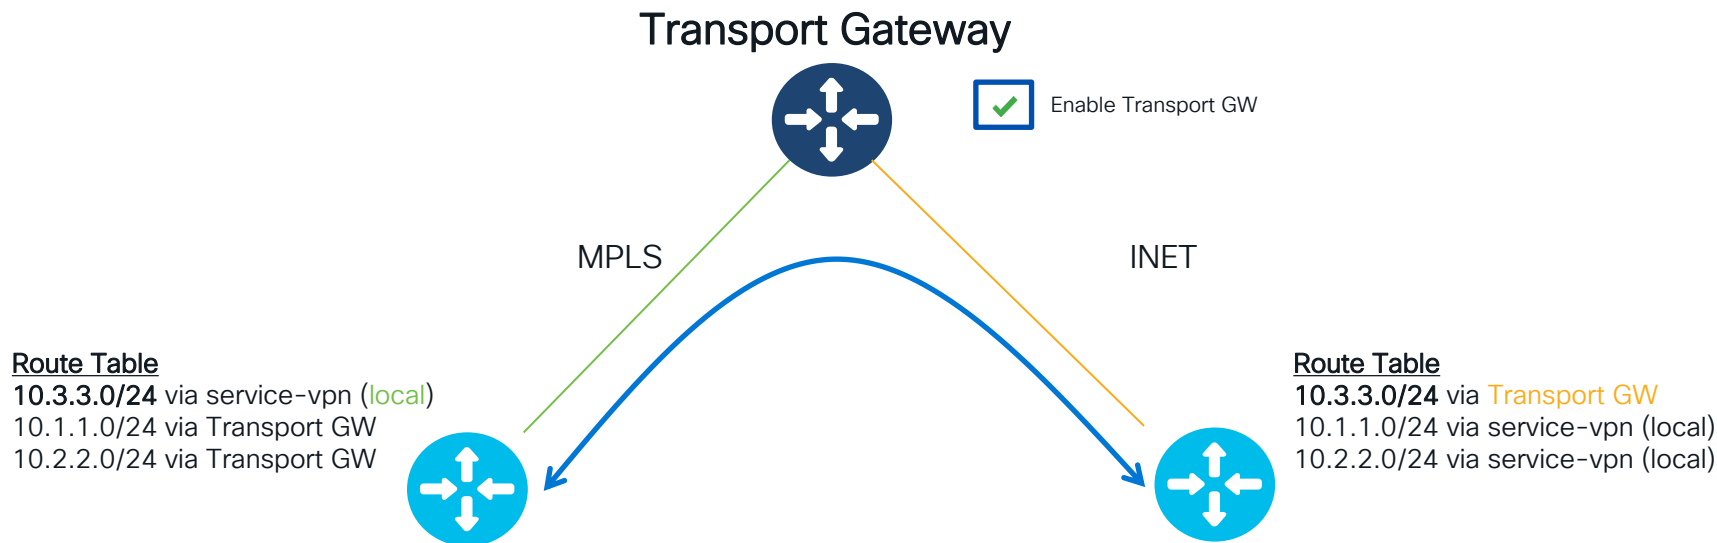


# Horizontal Scaling using Affinity Groups

- A simpler way to achieve horizontal scaling
- Affinity groups (AG) configured under the system settings on the Edge. E.g.
  - Edge Routers with AG preference=1,2 will build tunnels to all BRs but prefer to forward traffic to BRs with AG=1,2
  - If BRs serving AG=1 go down, then branches fallback to BRs serving AG=2
  - Works bi-directionally as well



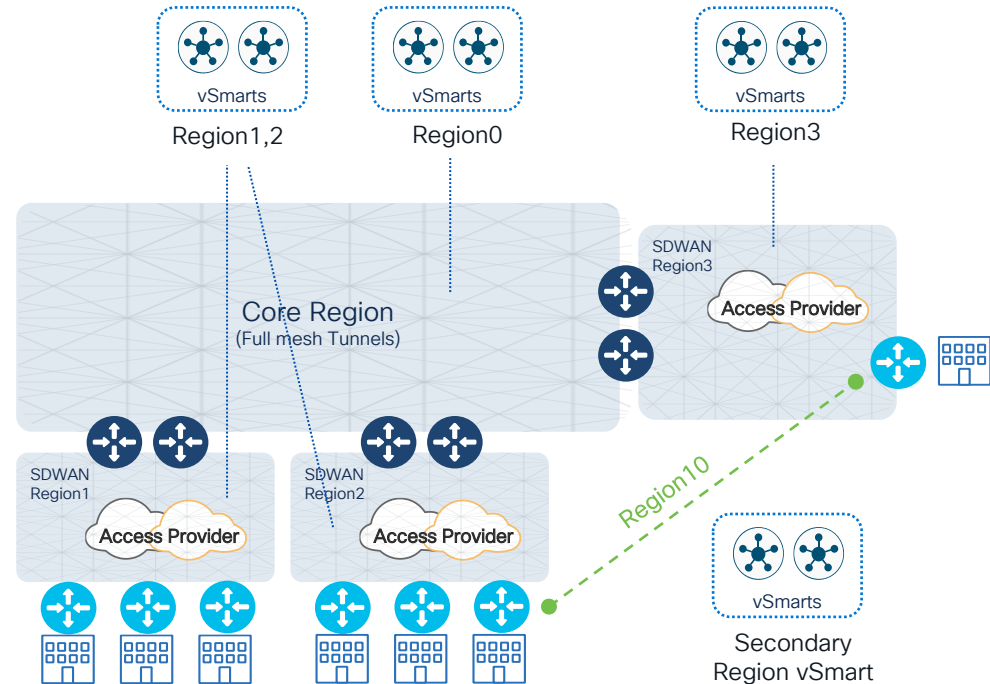
# Disjoint Underlays – Transport Gateway (TR)



- Enable transport-gateway (TR)
  - XE SD-WAN Edge will re-originate all vRoutes it has learned, but with its own TLOCs as the next-hop.
- Simple way to connect sites with underlay WANs that cannot directly communicate with each other

# Direct Tunnels – Secondary Region

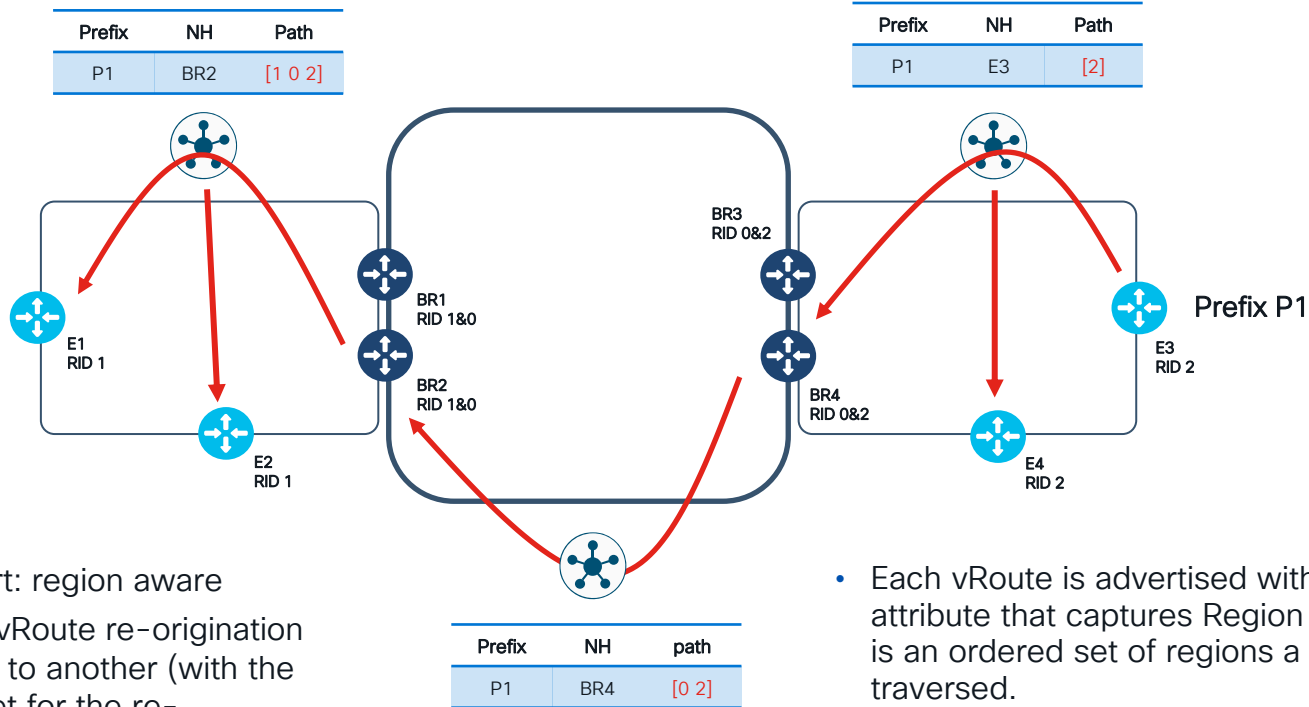
- Typically, multi-region or hierarchical network designs restrict *direct* tunnel instantiation and route exchanges across regions
- However, some customer use-cases cited as an exception to the above:
  - Send **non-critical traffic using cheaper WAN links** rather than using the optimized middle-mile WAN or PAYG links
  - Control/limit scalability requirements at BRs, to **reduce Edge costs at PoP/COLO**
  - Connect to common central site(s) from all regions or specific regions



# Multi-Region Fabric

## *What's under the hood?*

# Routing in Multi-Region Fabric




- OMP and vSmart: region aware
- Border routers: vRoute re-origination from one region to another (with the correct TLOC set for the re-originated route)

- Each vRoute is advertised with a new attribute that captures Region path- which is an ordered set of regions a route has traversed.
- Re-originated routes are withdrawn if the connectivity goes down. This helps prevent blackholing scenarios.



# Region and Tunnel



## Policy

```
policy
  no need for control policy
!
```


- vSmart advertises only intra-region TLOCs to WAN Edge
  - Spoke has only TLOCs from the same region
  - Border Node has TLOCs from edge region and core
- Region-id used to restrict tunnels between WAN Edge devices in the same region
- Full mesh within region

# Region and Routes

Policy

```
policy
no need for control policy
!
```


Prefix	Edge	Path
P1	T4	1 0 2
	T5	
P2	T3	1



- vSmart advertises **intra-region** routes unchanged to Edge and Border Routers
- Border Routers check **inter-region** prefixes reachability
  - > Intra-region prefix reachability using direct tunnels
  - > Inter-region prefix reachability via Border Nodes – Default load Balancing

# TLOCs

show sdwan omp tlocs



```
-----  
tloc entries for 10.0.0.109  
      mpls  
      ipsec  
-----  
RECEIVED FROM:  
tenant-id      0  
peer          10.0.0.21  
status        C,I,R  
loss-reason    not set  
lost-to-peer   not set  
lost-to-path-id not set  
Attributes:  
  attribute-type installed  
  encap-key     not set  
  encap-proto   0  
  encap-spi     261  
  encap-auth    sha1-hmac,ah-sha1-hmac  
  encap-encrypt aes256  
  public-ip     10.1.1.25  
  public-port   12346  
  private-ip    10.1.1.25  
  private-port  12346  
  public-ip     ::  
  public-port   0  
  private-ip    ::  
  private-port  0  
  bfd-status    up  
  domain-id     not set  
  site-id       101  
  overlay-id    not set  
  preference     0  
  region-id     1  
  tag           not set  
  stale         not set  
  weight        1  
  version       3  
  gen-id        0x8000000d  
  carrier       default  
  restrict      0  
  on-demand     0  
  groups        [ 11 ]  
  bandwidth     0  
  bandwidth-dmin 0  
  bandwidth-down 0  
  bandwidth-dmax 0  
  adapt-qos-period 0  
  adapt-qos-up    0  
  qos-group     default-group  
  border        not set  
  extended-ipsec-anti-replay not set  
  unknown-attr-len not set
```

# OMP Routes

## New attributes



Edge in region 2 show sdwan omp routes

```
C3#sh sdwan omp routes
Generating output, this might take time, please wait ...
```


Code:

```
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA -> On-demand inactive
U -> TLOC unresolved
Reo -> reoriginated
```

TENANT	VPN	PREFIX	FROM PEER	ID	LABEL	STATUS	TYPE	TLOC IP	COLOR	ENCAP	PREFERENCE	REGION ID	REGION PATH
0	10	10.10.10.109/32	10.0.0.21	19	1003	C,I,R	installed	10.0.0.63	mpls	ipsec	-	2	2 0 1
			10.0.0.21	20	1003	C,I,R	installed	10.0.0.63	biz-internet	ipsec	-	2	2 0 1
			10.0.0.21	21	1003	C,I,R	installed	10.0.0.64	mpls	ipsec	-	2	2 0 1
			10.0.0.21	22	1003	C,I,R	installed	10.0.0.64	biz-internet	ipsec	-	2	2 0 1
			10.0.0.22	9	1003	C,R	installed	10.0.0.63	mpls	ipsec	-	2	2 0 1
			10.0.0.22	10	1003	C,R	installed	10.0.0.63	biz-internet	ipsec	-	2	2 0 1
			10.0.0.22	29	1003	C,R	installed	10.0.0.64	mpls	ipsec	-	2	2 0 1
			10.0.0.22	30	1003	C,R	installed	10.0.0.64	biz-internet	ipsec	-	2	2 0 1



# Routes



```
BR1#sh sdwan omp routes 10.10.10.103/32
```

```
Code:
```

```
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA -> On-demand inactive
U -> TLOC unresolved
Reo -> reoriginated
```

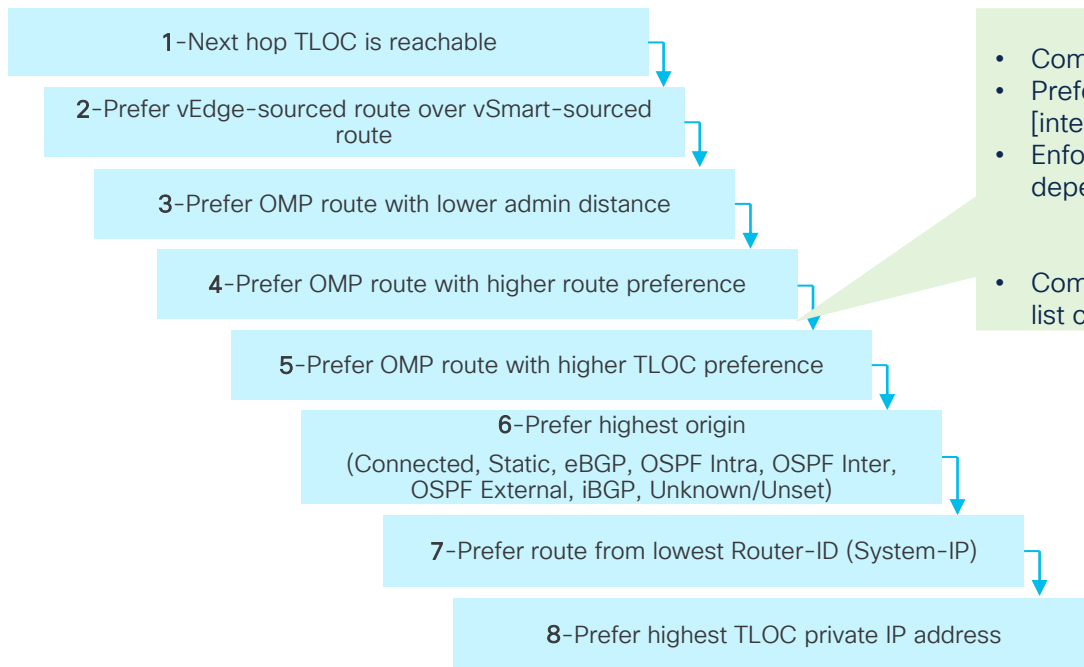
TENANT	VPN	PREFIX	FROM PEER	PATH ID	LABEL	STATUS	ATTRIBUTE TYPE	TLOC IP	COLOR	ENCAP	PREFERENCE	REGION ID	REGION PATH
0	10	10.10.10.103/32	0.0.0.0	21474	1003	C,Red,R,Reo	installed	10.0.0.61	mpls	ipsec	-	1	1 0 2
			0.0.0.0	83714									
			0.0.0.0	21474	1003	C,Red,R,Reo	installed	10.0.0.61	biz-internet	ipsec	-	1	1 0 2
				83716									
			10.0.0.21	9	1003	C,I,R	installed	10.0.0.64	private6	ipsec	-	0	0 2
			10.0.0.21	14	1003	C,I,R	installed	10.0.0.63	private6	ipsec	-	0	0 2
			10.0.0.22	13	1003	C,R	installed	10.0.0.63	private6	ipsec	-	0	0 2
			10.0.0.22	14	1003	C,R	installed	10.0.0.64	private6	ipsec	-	0	0 2

```
BR1#
```



# OMP Best Path Selection Algorithm

## Updated following MRF



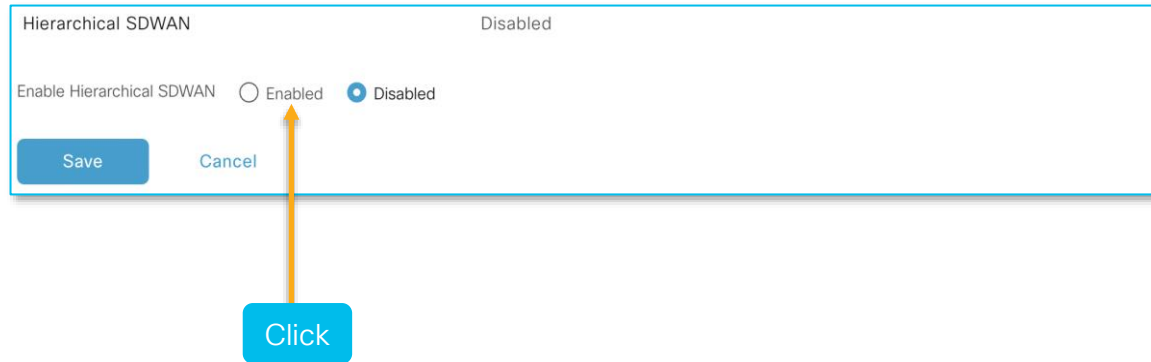
*Between step 4 and 5*

- Compare region-path-length
- Prefer access-region paths [intra-region] over core-region paths [inter-region]
- Enforce the transport-gateway router (TR) path check (user-config dependent):
  - Prefer TR-sourced paths, or, prefer 'direct' paths' (if available) or ECMP between all paths
- Compare the affinity in the paths based on the affinity preference list configuration

# Configuration

# vManage – Global Settings

Step 1: Enable Multi-Region Fabric under Administration > Settings



The screenshot shows a configuration dialog box titled "Hierarchical SDWAN" with a status indicator "Disabled" in the top right corner. Inside the dialog, the text "Enable Hierarchical SDWAN" is followed by two radio buttons: "Enabled" (which is unselected) and "Disabled" (which is selected, indicated by a blue dot). Below the radio buttons are two buttons: "Save" and "Cancel". An orange arrow originates from a blue button labeled "Click" positioned below the dialog and points directly to the "Enabled" radio button, indicating the action to be taken.



# vManage – Region Information

## Step 2: Assign MRF region ID to vSmarts and WAN edges

The screenshot displays the Cisco vManage interface for configuring a template. The breadcrumb navigation shows 'Feature Template > System > vsmart-FT'. The 'Configuration Groups' section includes tabs for 'Configuration Groups', 'Device Templates', and 'Feature Templates'. The 'Device Type' is set to 'vSmart'. The 'Template Name' and 'Description' fields both contain 'vsmart-FT'. The 'Basic Configuration' tab is selected, showing fields for 'Site ID', 'System IP', 'Hostname', and 'Location'. The 'Region ID List' field is highlighted with a dashed orange box and contains the value '012'.

Cisco vManage Select Resource Group Configuration - Templates

Configuration Groups Device Templates Feature Templates

Feature Template > System > vsmart-FT

Device Type vSmart

Template Name\* vsmart-FT

Description\* vsmart-FT

Basic Configuration GPS Advanced

▼ BASIC CONFIGURATION

Site ID [system\_site\_id]

System IP [system\_system\_ip]

Hostname [system\_host\_name]

Location

Region ID List 012

# vManage – Device Role

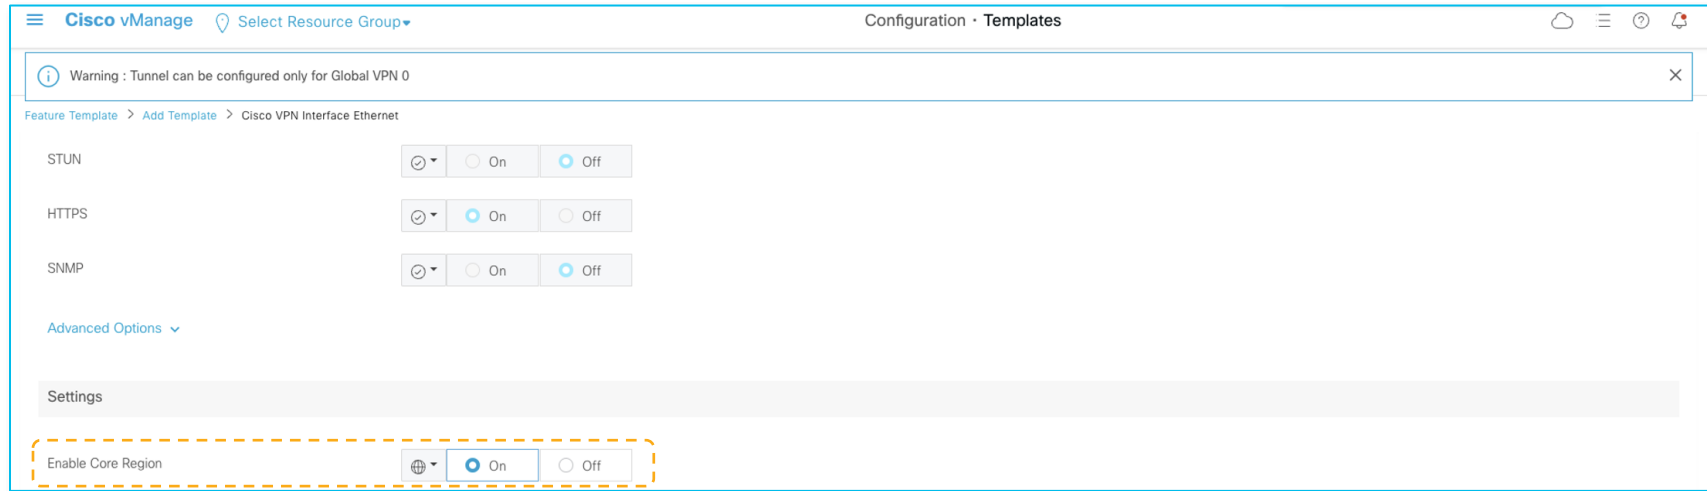
## Step 3: Assign device ‘role’ to the WAN edges

The screenshot shows the Cisco vManage interface for configuring a device template. The breadcrumb trail is: Feature Template > Add Template > Cisco System. The form includes the following fields:

- Urenary IU: 1
- Timezone: UTC
- Hostname: [system\_host\_name]
- Location:
- Device Groups:
- Controller Groups:
- Description:
- Console Baud Rate (bps): -- Choose --
- Maximum OMP Sessions:
- Region ID: 1
- Role: (highlighted with a dashed orange box, dropdown menu is open showing 'Edge Router' and 'Border Router')

# vManage – Border Router Core Interface

Step 4: On BRs, assign interfaces to the ‘core’ region



# Key Takeaways

# Try it today!

- Multi-Region Fabric is the core enabler for WAN architectures involving a middle-mile
  - For Managed Services SD-WAN
  - Large Enterprise deployments using MSP/Cloud/SDCI backbone
- Brownfield migration capability available this August (20.9/17.9 release)!



Eliminates need  
for lengthy global  
network policies



Automatic  
hop-by-hop  
inter-region  
routing



Architected to  
scale



Simpler  
redundancy  
planning



Flexible  
architecture to  
cater to dynamic  
network needs



Operationally  
easier to  
deploy and  
manage

# Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.



# Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. [www.cisco.com/go/certs](https://www.cisco.com/go/certs)

## Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.



## Learn



### Cisco U.

IT learning hub that guides teams and learners toward their goals

### Cisco Digital Learning

Subscription-based product, technology, and certification training

### Cisco Modeling Labs

Network simulation platform for design, testing, and troubleshooting

### Cisco Learning Network

Resource community portal for certifications and learning



## Train



### Cisco Training Bootcamps

Intensive team & individual automation and technology training programs

### Cisco Learning Partner Program

Authorized training partners supporting Cisco technology and career certifications

### Cisco Instructor-led and Virtual Instructor-led training

Accelerated curriculum of product, technology, and certification courses



## Certify



### Cisco Certifications and Specialist Certifications

Award-winning certification program empowers students and IT Professionals to advance their technical careers

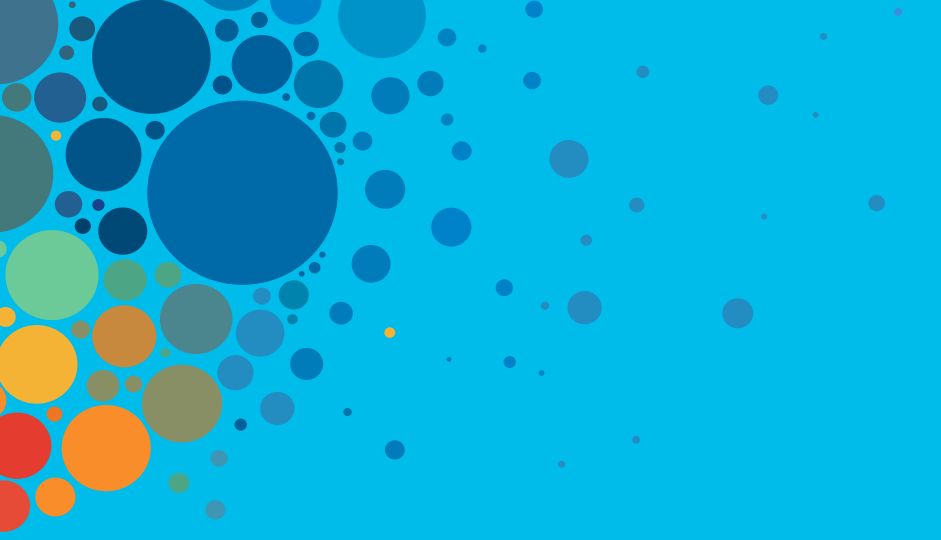
### Cisco Guided Study Groups

180-day certification prep program with learning and support

### Cisco Continuing Education Program

Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**



# Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer 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](https://www.CiscoLive.com/on-demand)





The bridge to possible

# Thank you

CISCO *Live!*



#CiscoLive