
cisco Live!
Let's go

The bridge to possible

Architecting Hybrid Multi-Cloud Infrastructures

David Jansen, Distinguished Solutions Engineer
CCIE 5952

cisco Live!

BRKDCN-2916

A little bit about David...

Cisco role: Distinguished, Solutions Engineer; Global Solutions Engineering.

Experience: My career @Cisco has spanned half of my life.

Fun fact 1: An awesome husband; Father of a daughter and twin boys

Fun fact 2: Written / published 4 books; 4 video series and working on my last one.

Fun fact 3: Enjoy the outdoors, music, working out, running, etc.

Session focus areas

Abstract

Cloud continues to be a major disruptor; customers have been on the Cloud Journey for several years.

- Architectural trends impacting hybrid and multicloud infrastructures.
- Why private data centers and applications remain critical, even as applications and data move from traditional private data centers to public clouds.
- How to…
  - Address growing demands of users and applications being everywhere in a multicloud world.
  - Help you optimize and understand traffic flows, operational efficiencies, and visibility in a multicloud world.

Goals of today's session

- Explore top challenges and discuss relevant use cases
- Provide a deeper understanding of current Cloud Networking deployment options
- Offer practical guidance on how best to deliver:
  - A consistent Cloud Networking/Segmentation/Security solution
  - Visibility across entire solutions – including in-region, intra-region and inter-cloud connectivity options visibility

Agenda

- Current Challenges, Issues, Administration and Options
- Multi-Cloud Networking: Cisco Cloud Network Controller
  - Single region connectivity
  - Multi-region connectivity
  - Multi-Cloud connectivity
  - Private Data Center connectivity
- Branch to Cloud
- Cisco Cloud Network Controller Configuration
- New capabilities / features
- Sneak Peek
- Visibility
- Summary

Taxonomy

- Hybrid-Cloud: Public-Cloud and Private-Cloud
- Multi-Cloud: Hybrid-Cloud + 2 or more CSP(s)
- East / West Traffic Flows within (intra) across Cloud Regions/Branches
- North / South Traffic Flows across (inter) Clouds and on-prem Data Center as well as user to app flows
- I will be using AWS terminology but applies to other Cloud Providers (CSP):
  - IGW: Internet Gateway
  - DX-Gateway: Direct Connect Gateway
  - TGW: Transit Gateway
  - vPC: Virtual Private Cloud

Cloud Terminology Matrix

For Your Reference

Area	AWS service	Azure service	GCP Service
Cloud virtual networking	Virtual Private Cloud (VPC)	Virtual Network (VNet)	Virtual Private Cloud (VPC)
NAT gateways	NAT Gateways	Virtual Network NAT	Cloud NAT
Cross-premises connectivity	VPN Gateway	VPN Gateway	Cloud VPN Gateway
DNS management	Route 53	DNS	Cloud DNS
DNS-based routing	Route 53	Traffic Manager	Cloud DNS
Dedicated network	Direct Connect	ExpressRoute	Cloud Interconnect
Load balancing	Network Load Balancer	Load Balancer	Network Load Balancing
Application-level load balancing	Application Load Balancer	Application Gateway	Global Load Balancing
Route table	Custom Route Tables	User Defined Routes	Routes
Private link	PrivateLink	Azure Private Link	Private Service Connect
Private PaaS connectivity	VPC endpoints	Private Endpoint	Private Service Connect
Virtual network peering	VPC Peering Transit Gateway	VNet Peering	VPC Network Peering
Content delivery networks	Cloud Front	Azure CDN	Cloud CDN
Network Monitoring	VPC Flow Logs	Azure Network Watcher	Network Intelligence Center

Current Challenges, Issues, Administration and Options

Challenges to solve

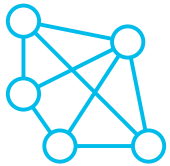

Network connectivity

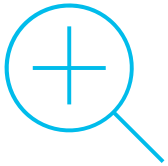

Operations and visibility

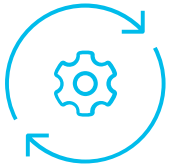

Services integration

Segmentation and security

Need For Homogenous Experience Across Heterogenous Cloud Environments

Creating two VPC(s) in one Region

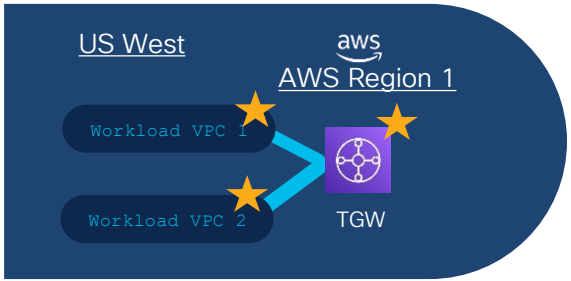

1
- Create VPC 1
- Create IP CIDR blocks

2
- Create VPC 2
- Create IP CIDR blocks

3
- Create TGW in Region 1

4
- Set up VPC to TGW attachment
- Create and associate routing table for VRFs VPC 1 and VPC2 in TGW
- Create all the relevant routes for VRFs in each of the component routing tables
- VPC routing table, TGW VPC attachment

★ Administrative touch points, routing controls and visibility

```
"TransitGatewayId": "tgw-0262a0e521EXAMPLE”
"TransitGatewayArn": "arn:aws:ec2:us-east-
2:111122223333:transit-gateway/tgw-0262a0e521EXAMPLE",
"AssociationDefaultRouteTableId": "tgw-rtb-018774adf3EXAMPLE”
"PropagationDefaultRouteTableId": "tgw-rtb-018774adf3EXAMPLE"
```

Adding a third VPC in a Second Region

1 - Create VPC 3 in region 2
- Create IP CIDR blocks

2 - Create TGW in Region 2

3 - Set up VPC to TGW attachment

4 - Tell VPC 3 about VPC 1 and VPC 2
- Creates and associates routing table for VRFs VPC 3 in the relevant TGWs
- Set up all the relevant routes for VRFs in each of the component routing tables
- VPC routing table, TGW VPC attachment

5 - Go to VPC 1 and VPC 2 and tell about VPC 3

6 - Set up TGW to TGW peering across regions
- VPC routing table, TGW VPC attachment, TGW peering attachment & routing tables

★ Administrative touch points, routing controls and visibility

Cloud Native Routing

- **Intent:** Reach 10.1.1.0/24 over all possible paths (ECMP) and Path Failure(s)
- Think of each Cloud as a distributed router
  - How do we know about path failures?
  - How do we have a backup path?
  - ECMP?

Inter-Region Connectivity Options

Connectivity Options

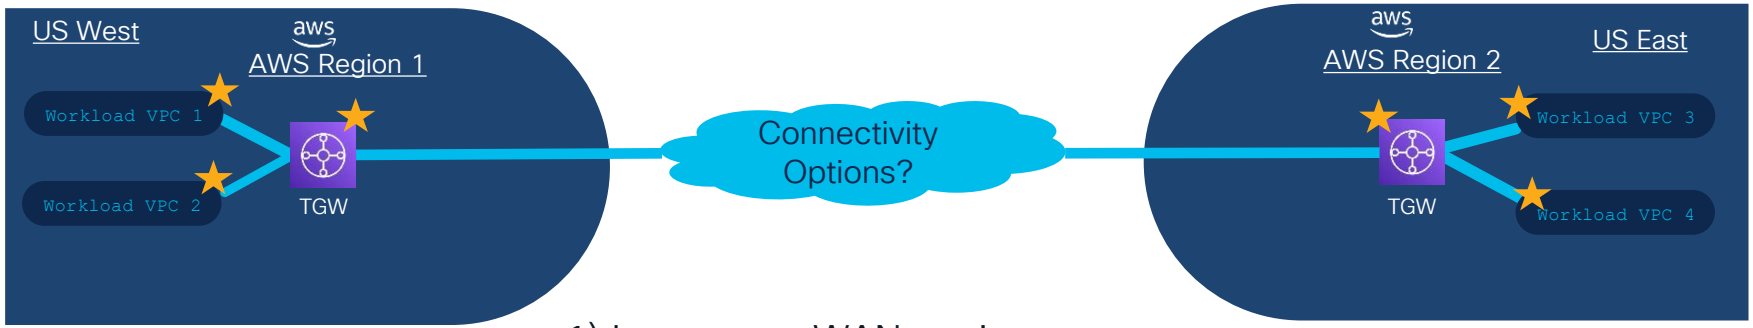

1) Leverage a WAN carrier
2) Leverage Colo (Equinix Fabric) + MACsec
3) Leverage Cloud Native TGW peering
4) Middle Mile
5) AWS Cloud WAN
6) Azure Virtual WAN

Customers are looking to build and create an environment they can switch out of easy (options) based on cost / pricing / charges.

Cloud Connectivity between same CSP

CSP Backbone – Different Regions

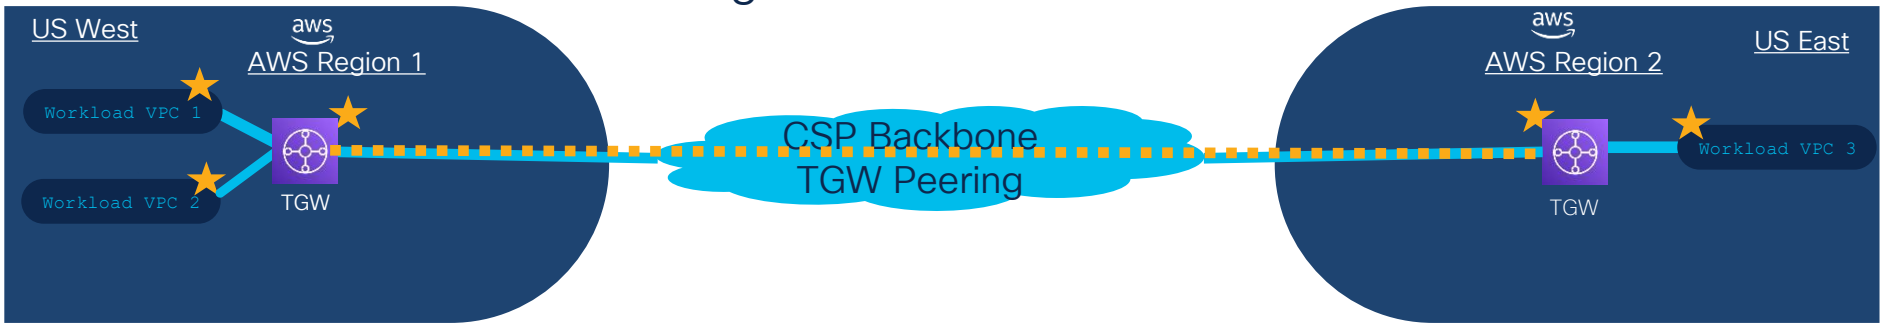

- No Routing protocol
- Leverage CSP Backbone
  - Each region is few hundred miles away. If Internet has latency of 500ms between the region; AWS backbone will offer less than 255ms
- No ECMP Path
- No Feedback mechanism for path-failures
- Misconfigurations

★ Administrative touch points, routing controls and visibility

Cloud Connectivity between same CSP

Customer Backbone/Private Peering – Different Regions

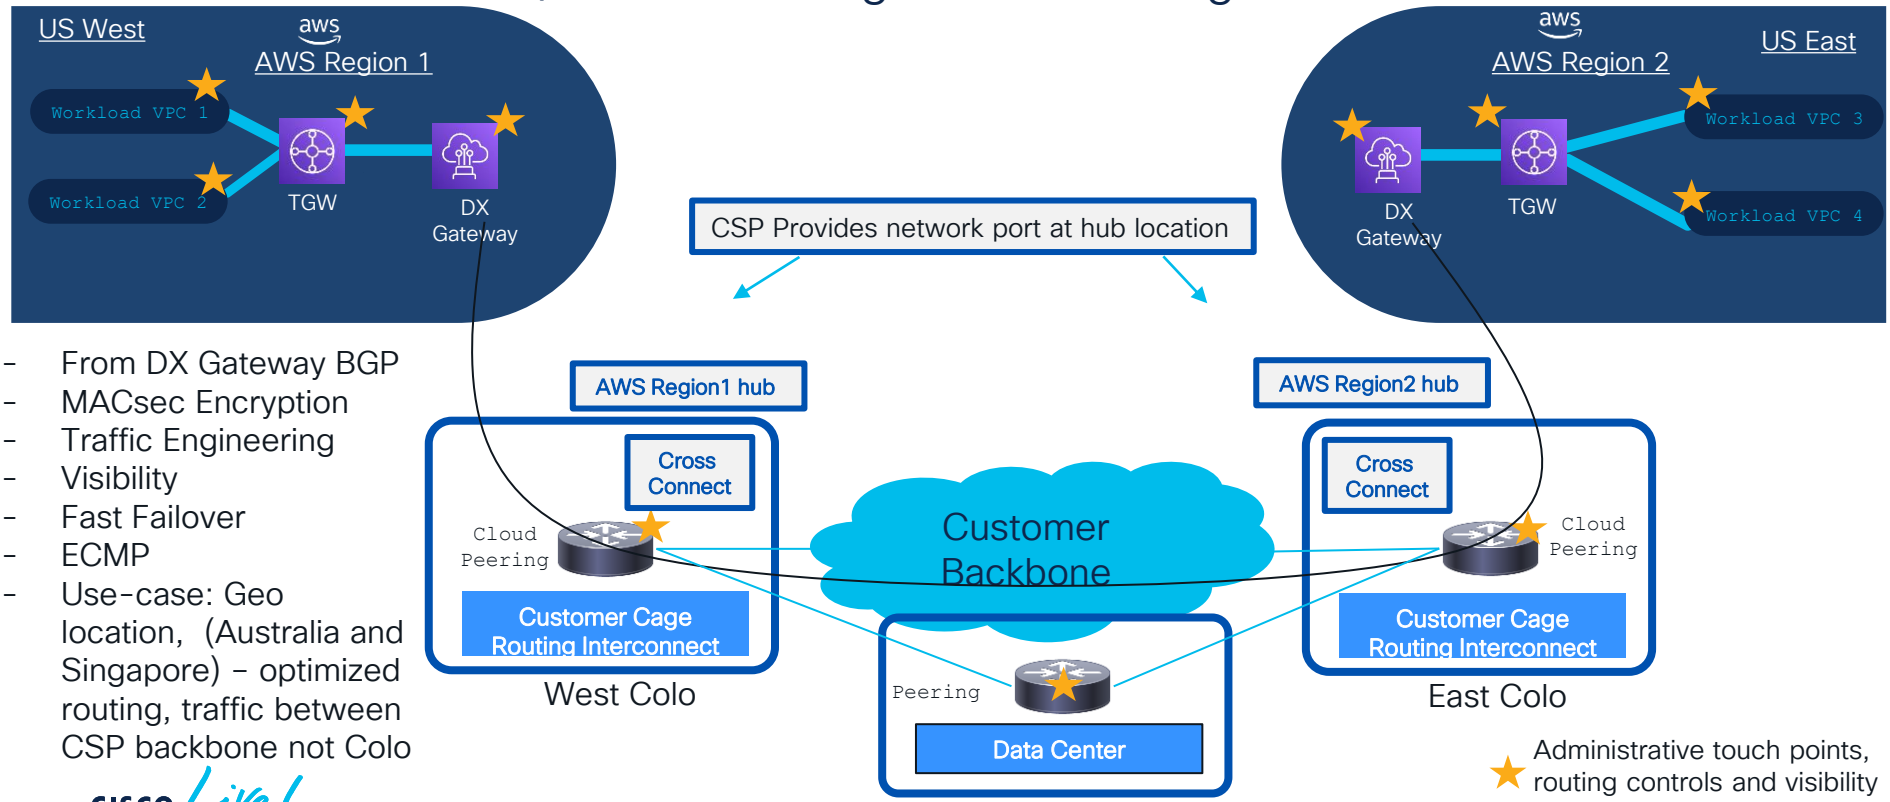

- From DX Gateway BGP
- MACsec Encryption
- Traffic Engineering
- Visibility
- Fast Failover
- ECMP
- Use-case: Geo location, (Australia and Singapore) – optimized routing, traffic between CSP backbone not Colo

★ Administrative touch points, routing controls and visibility

Multi-Cloud Connectivity and Routing

Multiple Ingress/Egress

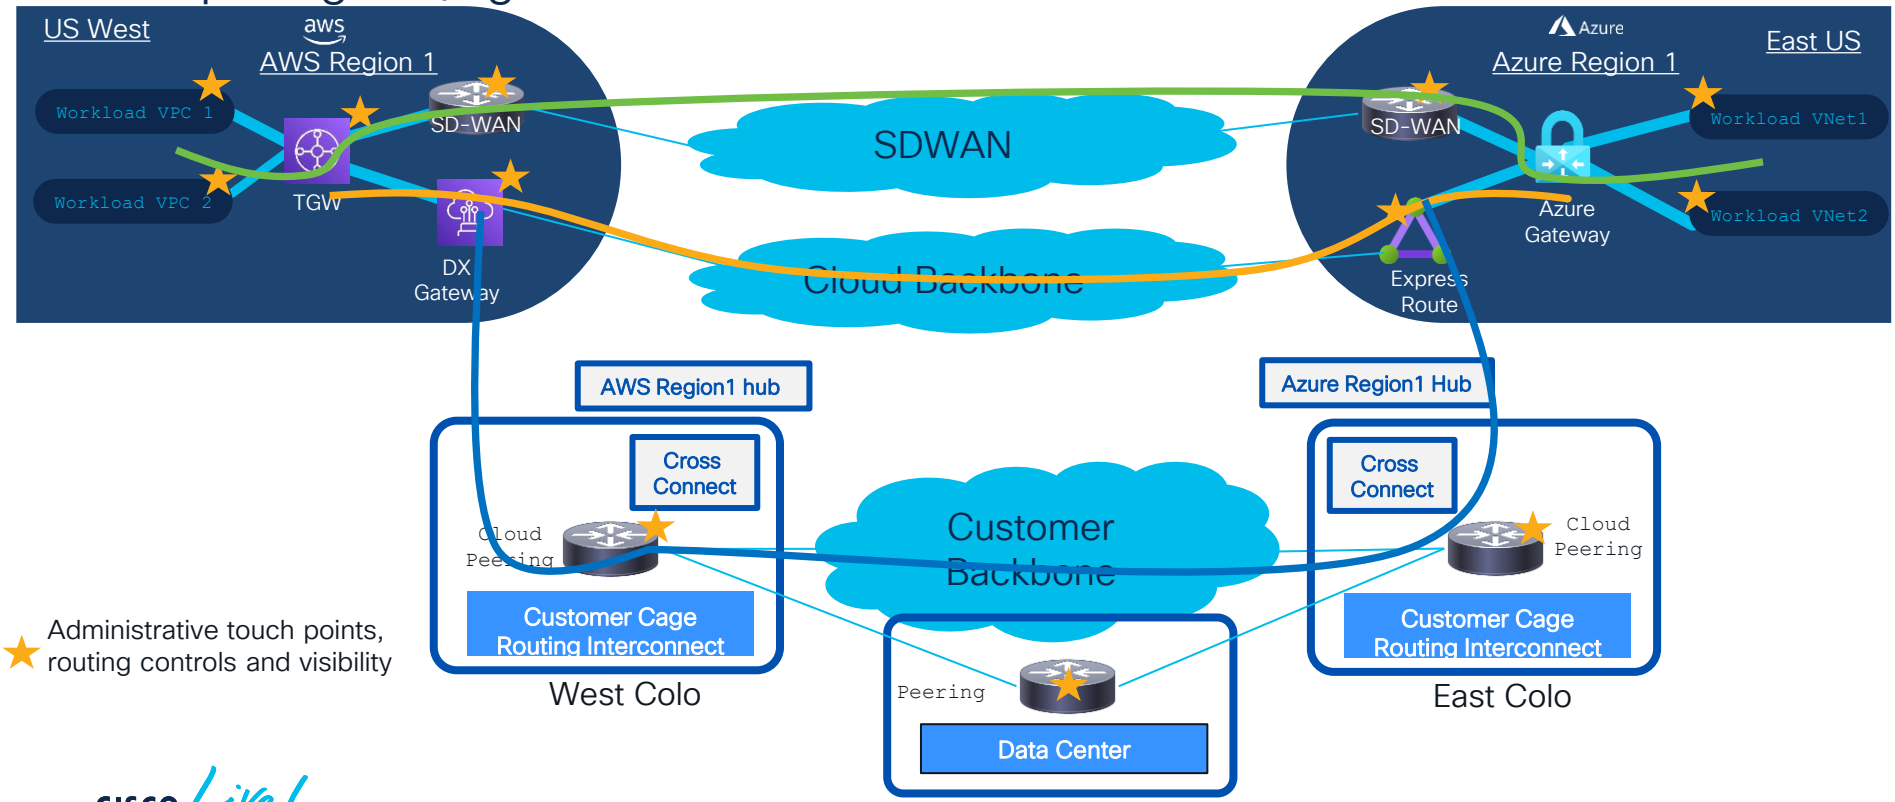

★ Administrative touch points, routing controls and visibility

Questions to solve for:

- Which path with the traffic flow?
- Is ECMP possible?
- Failover?
- Path failure detection and re-routing to a new / backup path?
- What about visibility?

Let’s take a closer look…

Multi-Cloud Networking: Cisco Cloud Network Controller

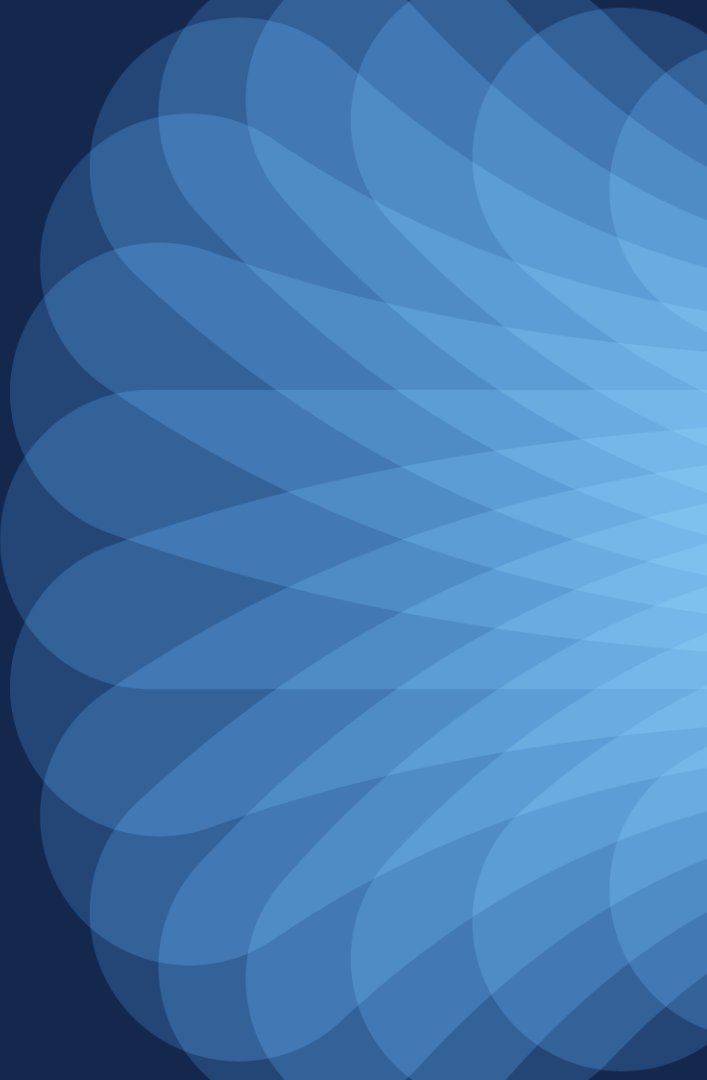

Cisco Cloud Network Controller

De-Coupled:
Security Policy and Network Connectivity

Specify Security Policy
Using Contracts

Network Connectivity
Enabled Per VRF, Route Maps

Multi-cloud Networking

Automate connectivity & segmentation, with full visibility across hybrid cloud

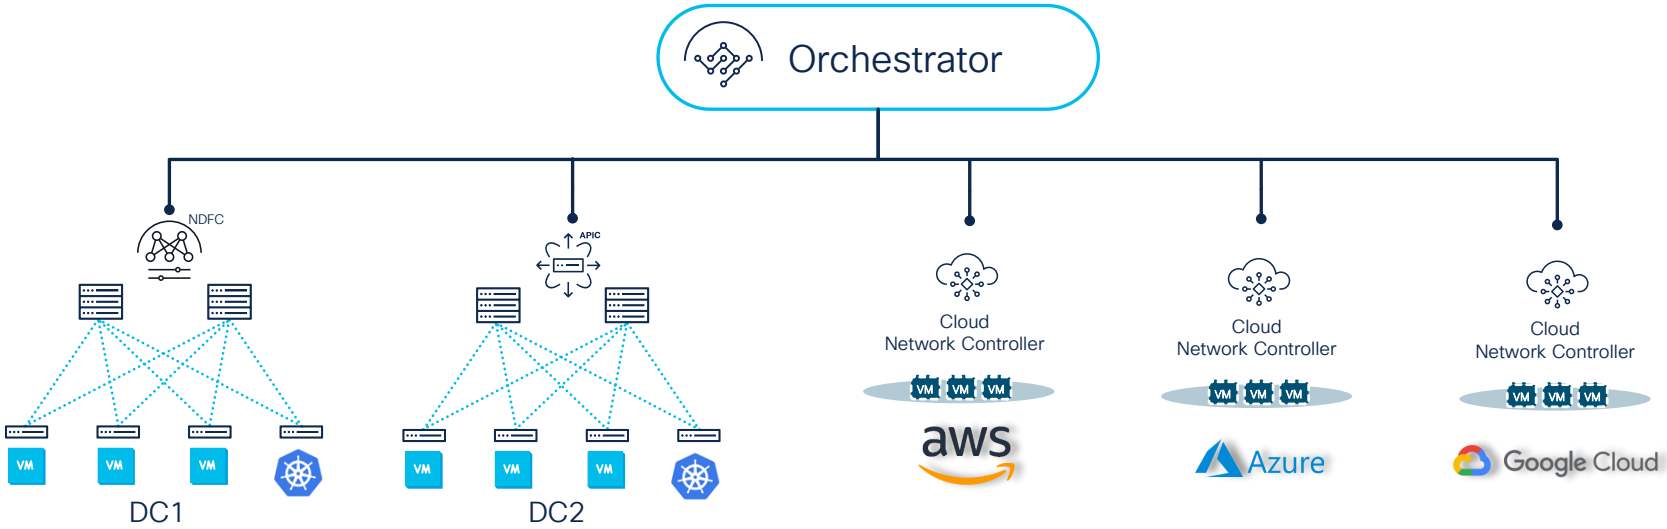

Automated connectivity and routing

Consistent security and segmentation

Single Point of Orchestration Visibility & Troubleshooting

Automated insertion of L4-7 services

Cisco Hybrid Multi-Cloud Networking Design Guide

Solution Building Blocks

Cloud Network Controller

Region Management
Network and App Policy
Cloud Resources: Inventory, Visibility, Topology View

Catalyst 8000v
Or Cloud Native Router

Intra-Cloud Connectivity
Inter-Cloud Connectivity
External Network Connectivity
BGP Routing (IPv4, EVPN*)

Nexus Dashboard Orchestrator

Secure Cloud Inter-Connect
Multi-Tenancy
Networking & Segmentation
L4-L7 Service Insertion

*Cat8kv only

Multi-Cloud: Flexible Deployment Models

Cloud only

- Intra-cloud single region
- Intra-cloud multi-region
- Inter-clouds

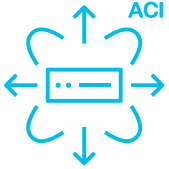

Hybrid with on-premises ACI

Hybrid with on-premises NDFC

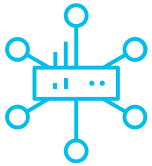

Connect to external networks

- SD-WAN router
- Branch router
- Data Center edge router

Distributed Cloud Networking

Cloud Network Controller / One per CSP

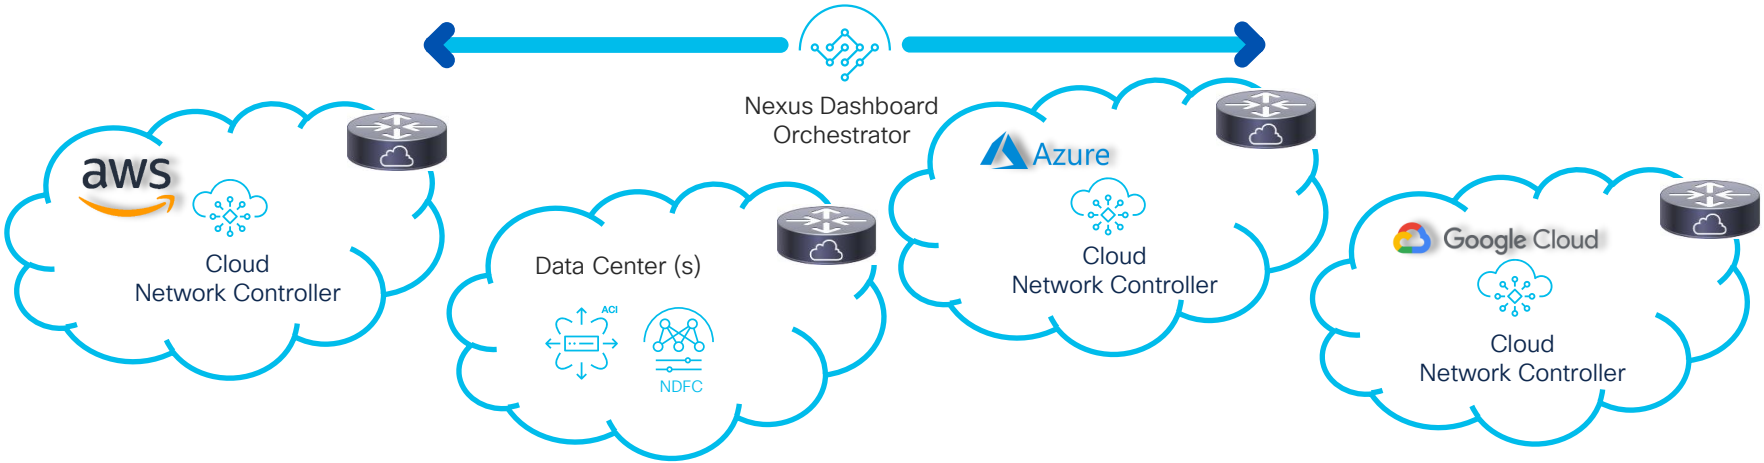

- The Cloud Network Controller per CSP is the Supervisor Module
  - Intent is stored in the Cloud Network Controller
- Each Cloud is a distributed router
- Leverage SLA probes to monitor Cloud Objects / Path Selection

Cloud Networking with Cloud Network Controller

- **Intent:** Reach 10.1.1.0/24 over all possible paths (ECMP) and Path Failure(s)
- Administrator expresses the Routing Intent (NetOps / Cloud NetOps)
- Validation of reachability
- Define Intent of Latency/Cost (future) not just AD

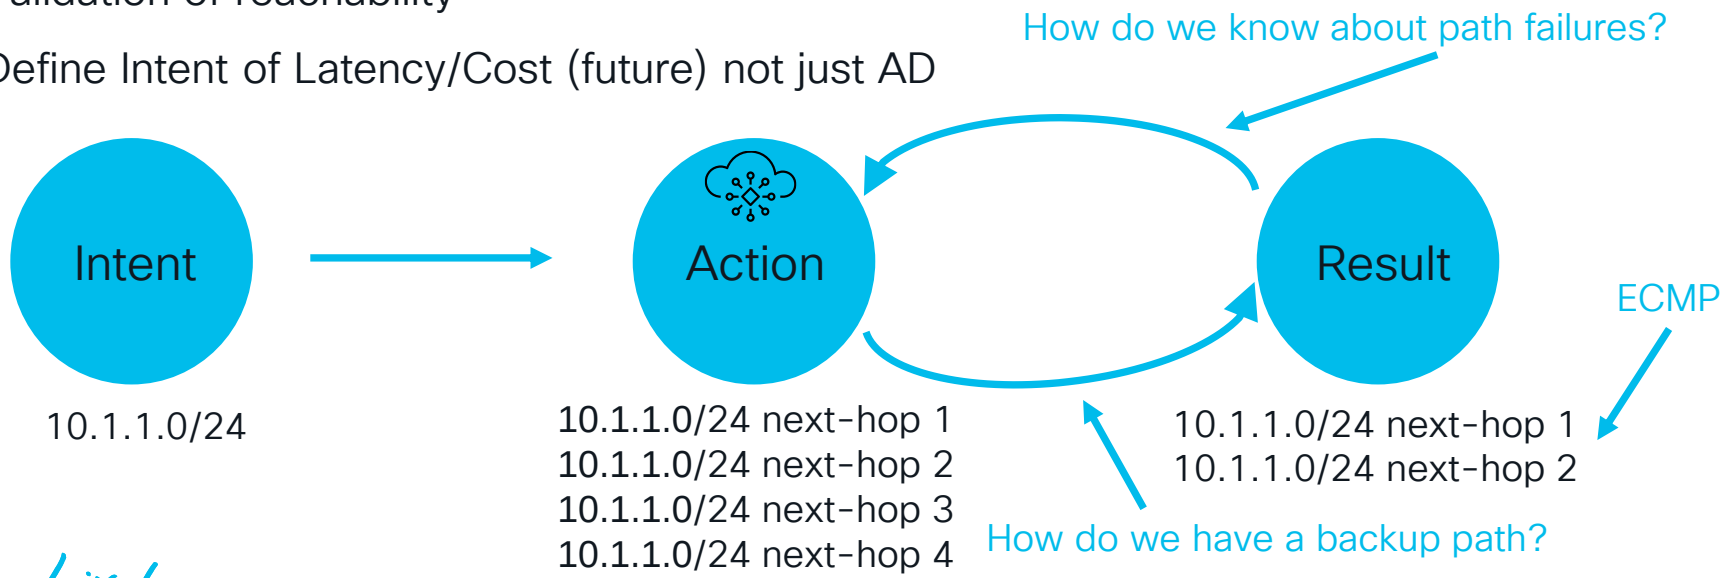

Connectivity between same CSP: Different Regions

Native TGW Peering

NetOps Operator: David

Intent:

- Create two regions
- Four VPC(s)
- Subnets per VPC / CIDR(s)
- VRF per VPC
- VPC attachment to TGW
- Life cycle of TGW
- TGW Routing
- TGW VRF Leaking
- TGW Peering automation

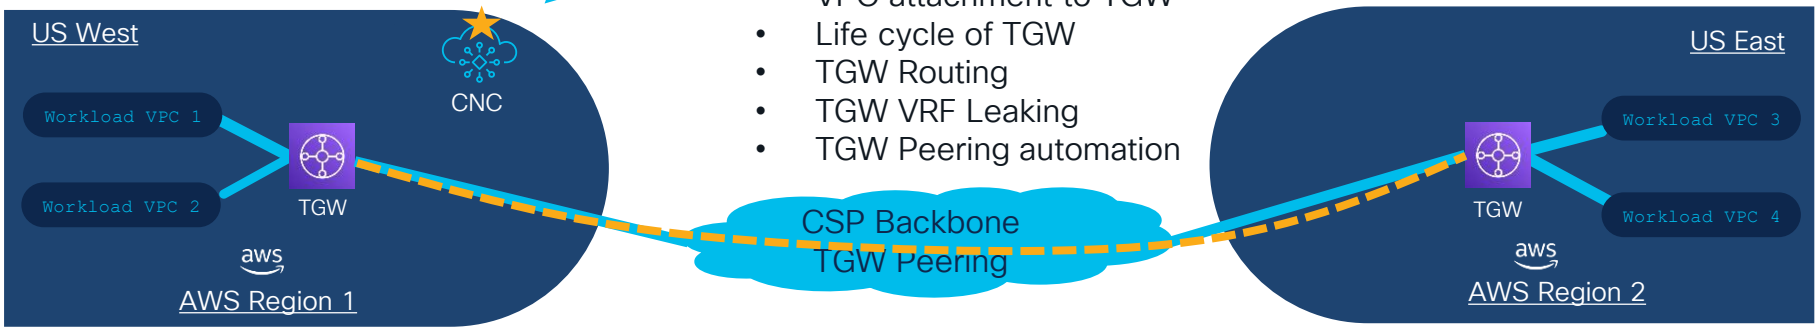

★ Administrative touch points, routing controls and visibility

- Keeps up with any further changes automatically
- Additions/Deletions of Subnets, CIDRs, VPCs, Regions.
- Updates to the network policy
- Selective Route Exchange between different routing domains VRF (i.e.: Shared Services)

Connectivity between same CSP: Different Regions

IPSec or Backbone Peering

NetOps Operator: David

Intent:

- Require Encryption
- Require ECMP
- Insert Cat8kv
- TGW Connect to Cat8kv
- GRE BGP IPv4
  - TGW to Cat8kv
  - Cat8kv to TGW
- IPSec or Backbone Peering (vpc-peering)
  - Cat8kv to Cat8kv

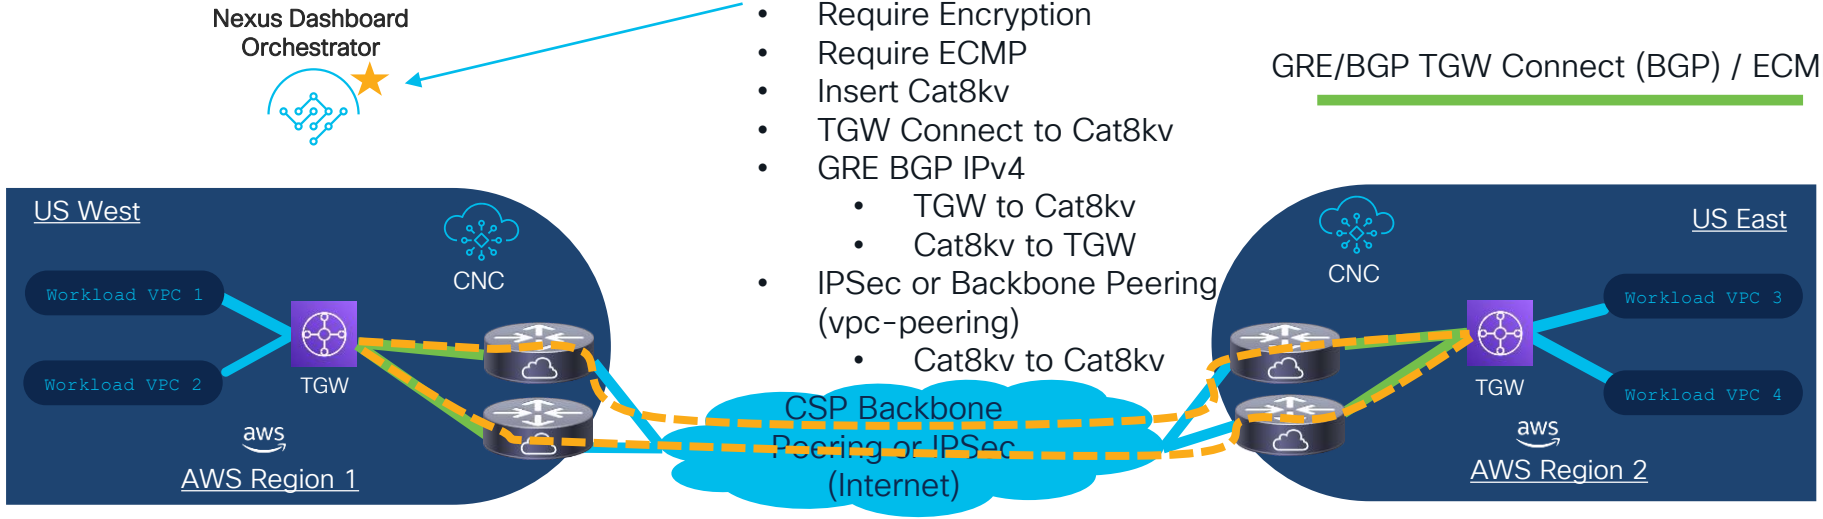

★ Administrative touch points, routing controls and visibility

- Keeps up with any further changes automatically
- Additions/Deletions of Subnets, CIDRs, VPCs, Regions.
- Updates to the network policy
- Selective Route Exchange between different routing domains VRF (i.e.: Shared Services)

Connectivity between same CSP: Different Regions

Private Peering

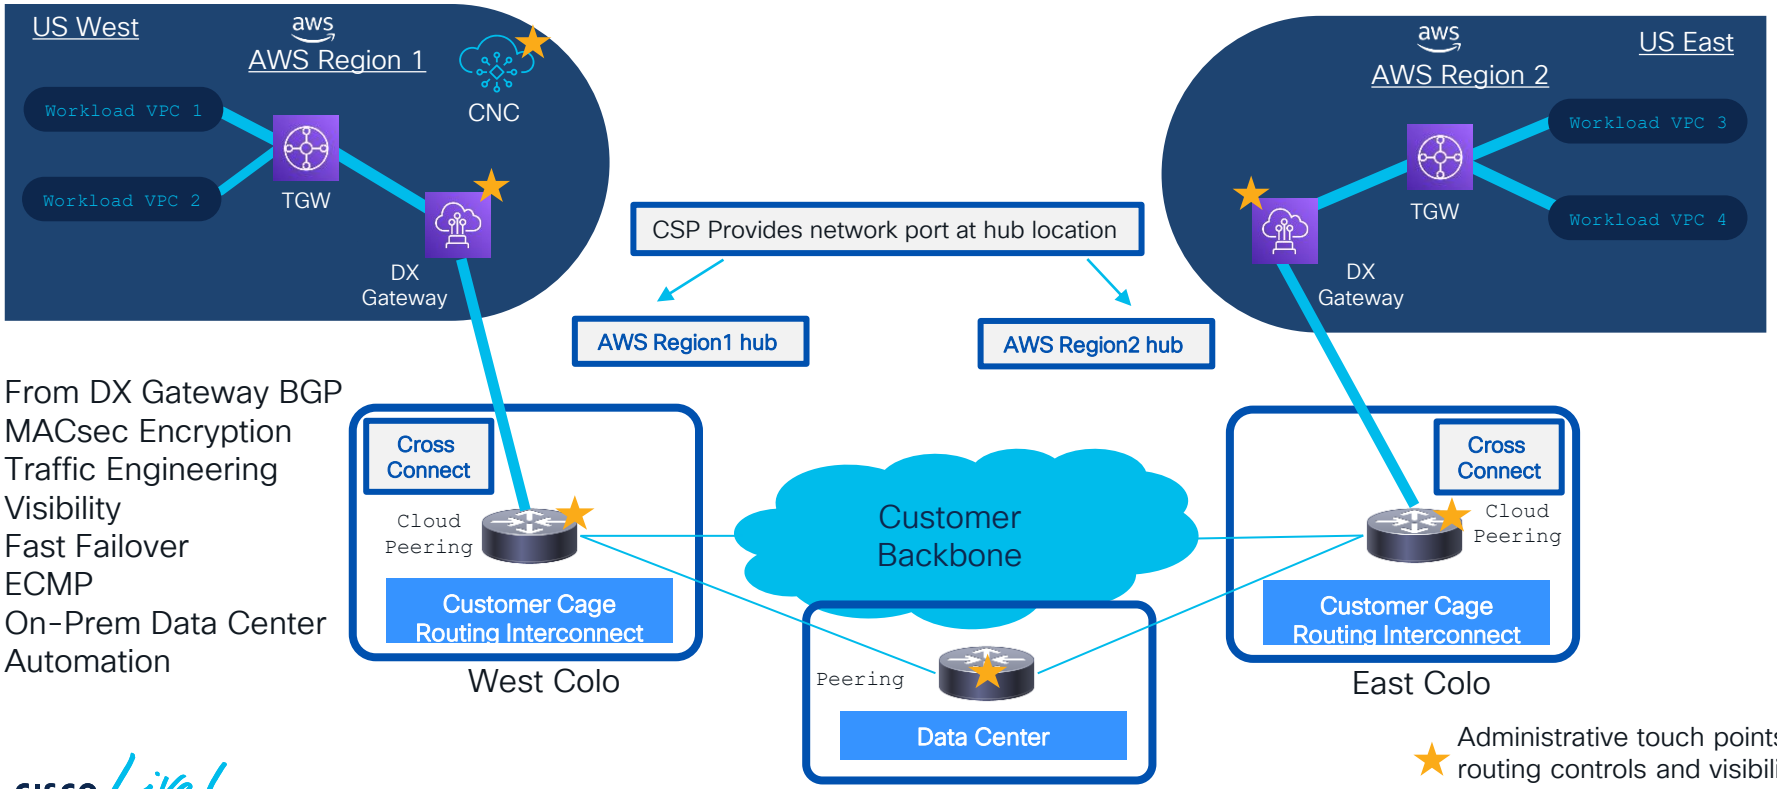

- From DX Gateway BGP
- MACsec Encryption
- Traffic Engineering
- Visibility
- Fast Failover
- ECMP
- On-Prem Data Center Automation

★ Administrative touch points, routing controls and visibility

L4-L7 Service Insertion

Azure Example

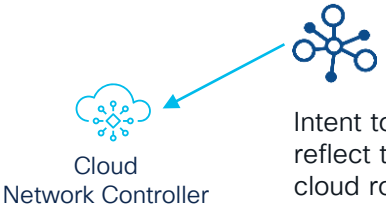

NetOps Operator: David

Intent to Cloud Network Controller; that change will reflect to route filtering on the Cat8kv or natively on cloud routing depending on your connectivity model.

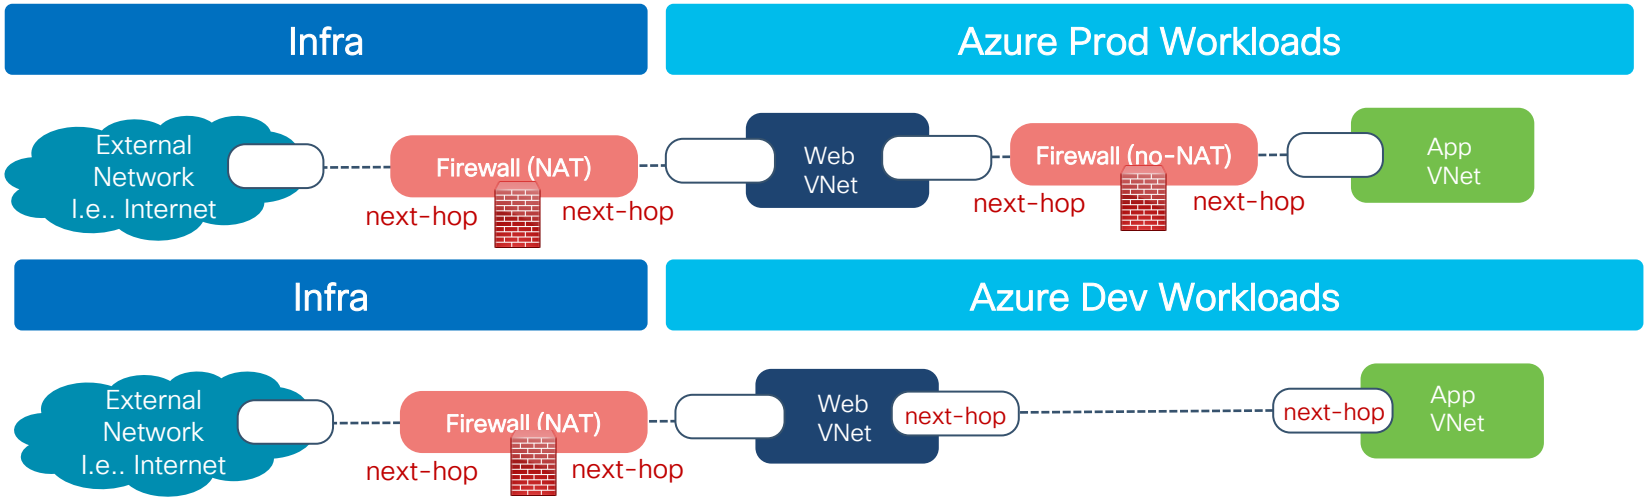

- **Workload(s):**
  - Prod Workloads use firewall as next-hop
  - Dev Workloads Azure "selective" use firewall services/inspection
- **External-Network:**
  - Based on routing policy, i.e. route community
  - route-map and BGP community configuration automatically (route-filtering)

When do you require Cloud Orchestrator?

For Your Reference

Orchestrator	aws	Azure	Google Cloud	On-Prem Data Center
aws				
Azure				
Google Cloud				
On-Prem Data Center				

Multi-Cloud Connectivity between different CSPs

Private Peering

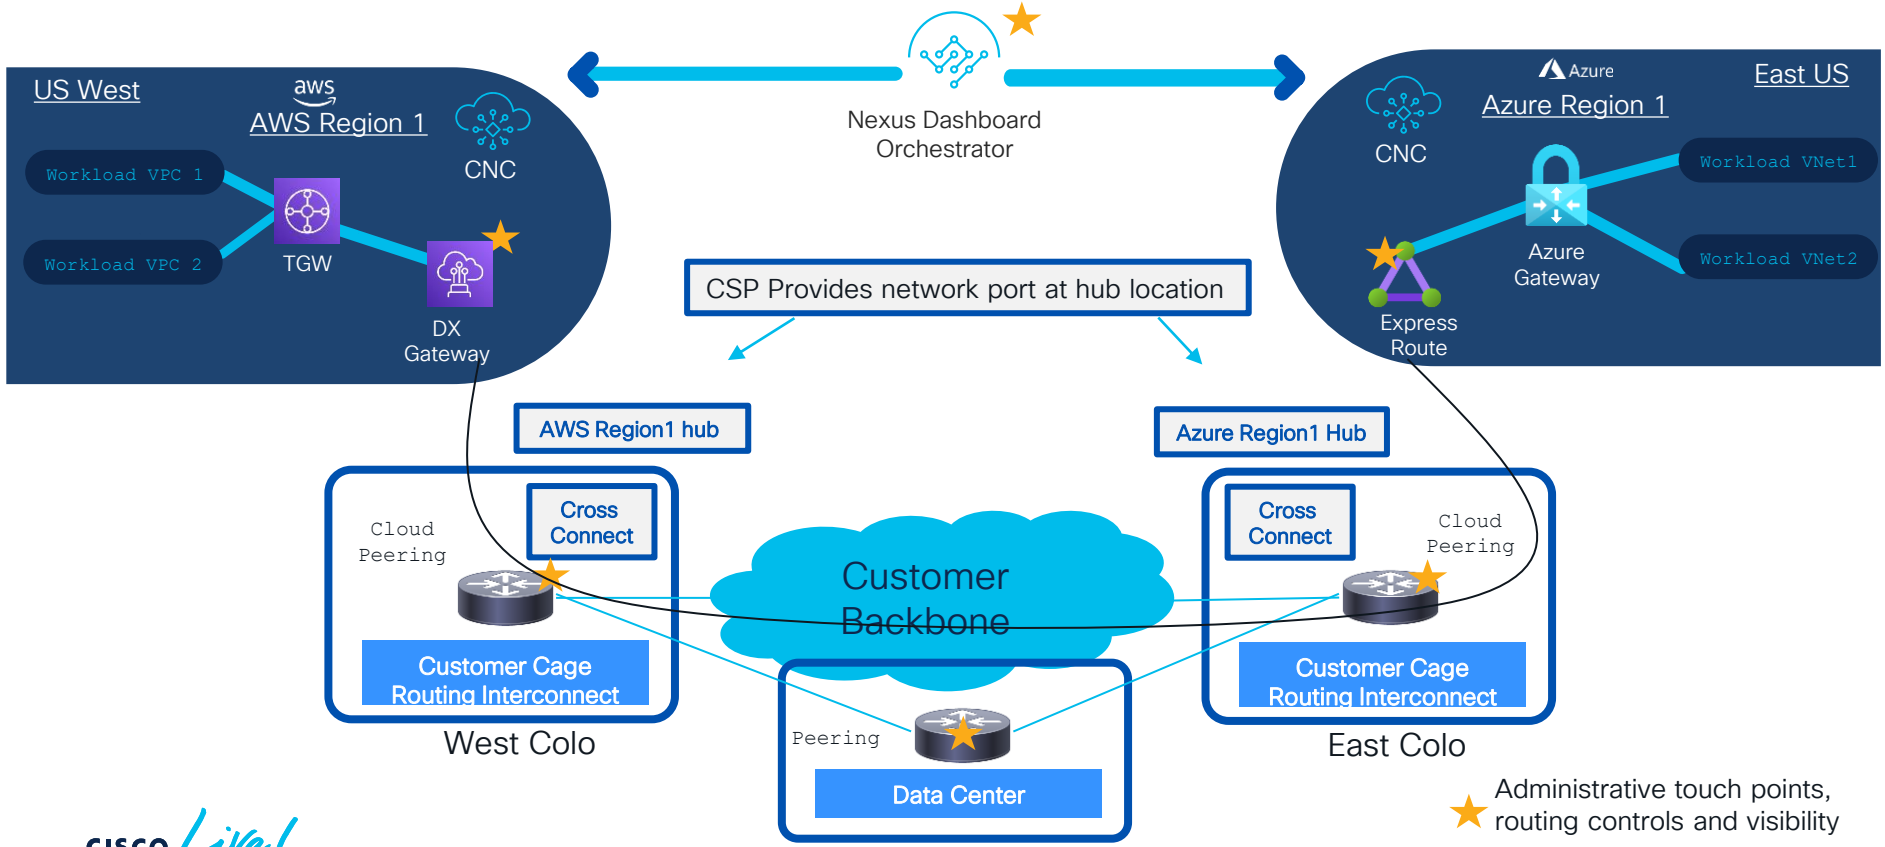

Multi-Cloud Connectivity between different CSPs

Private Peering with On-prem DC

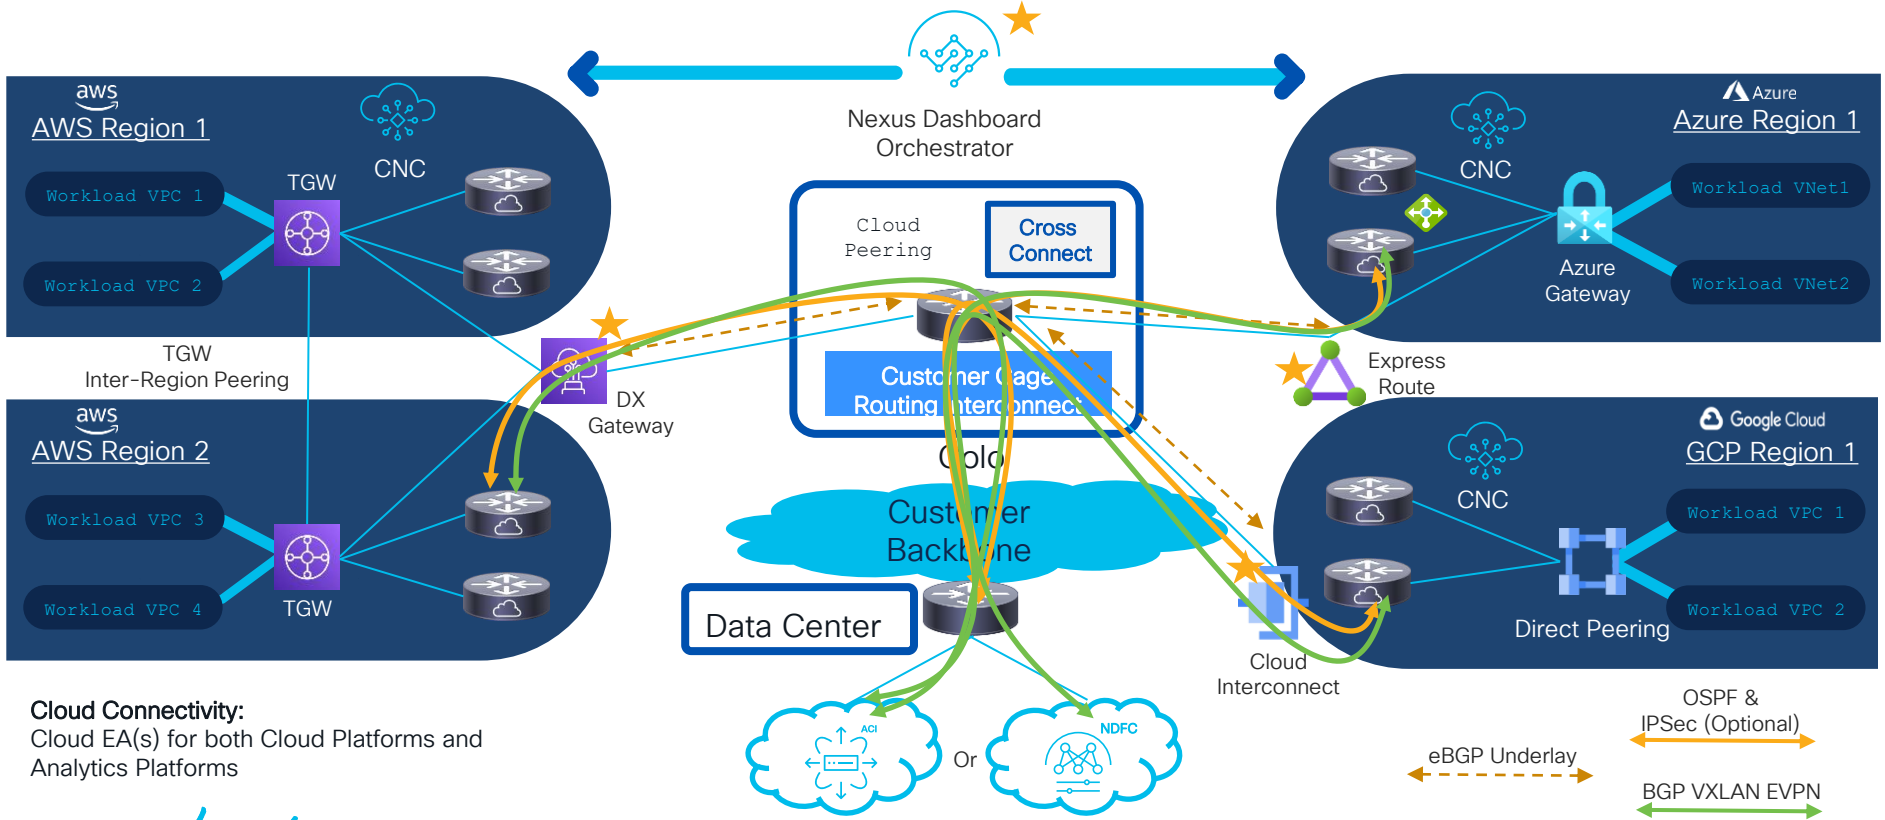

Cloud Connectivity:
Cloud EA(s) for both Cloud Platforms and Analytics Platforms

Hybrid Connectivity with AWS Direct Connect

BFD Example

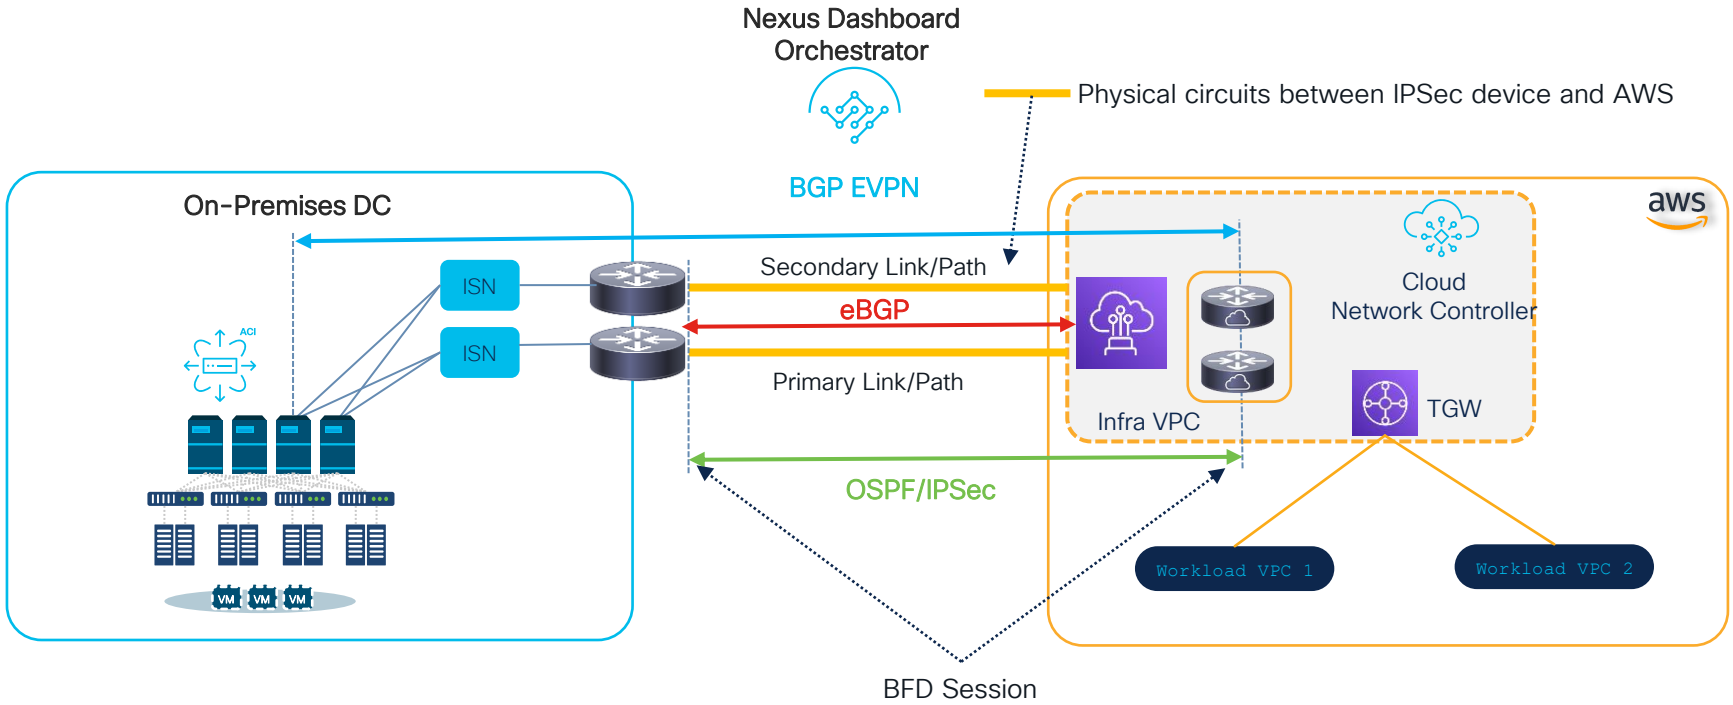

[Improving Convergence Between On-Premises ACI Sites and Cloud Sites]

Multi-Cloud Networking

AWS to Azure to GCP

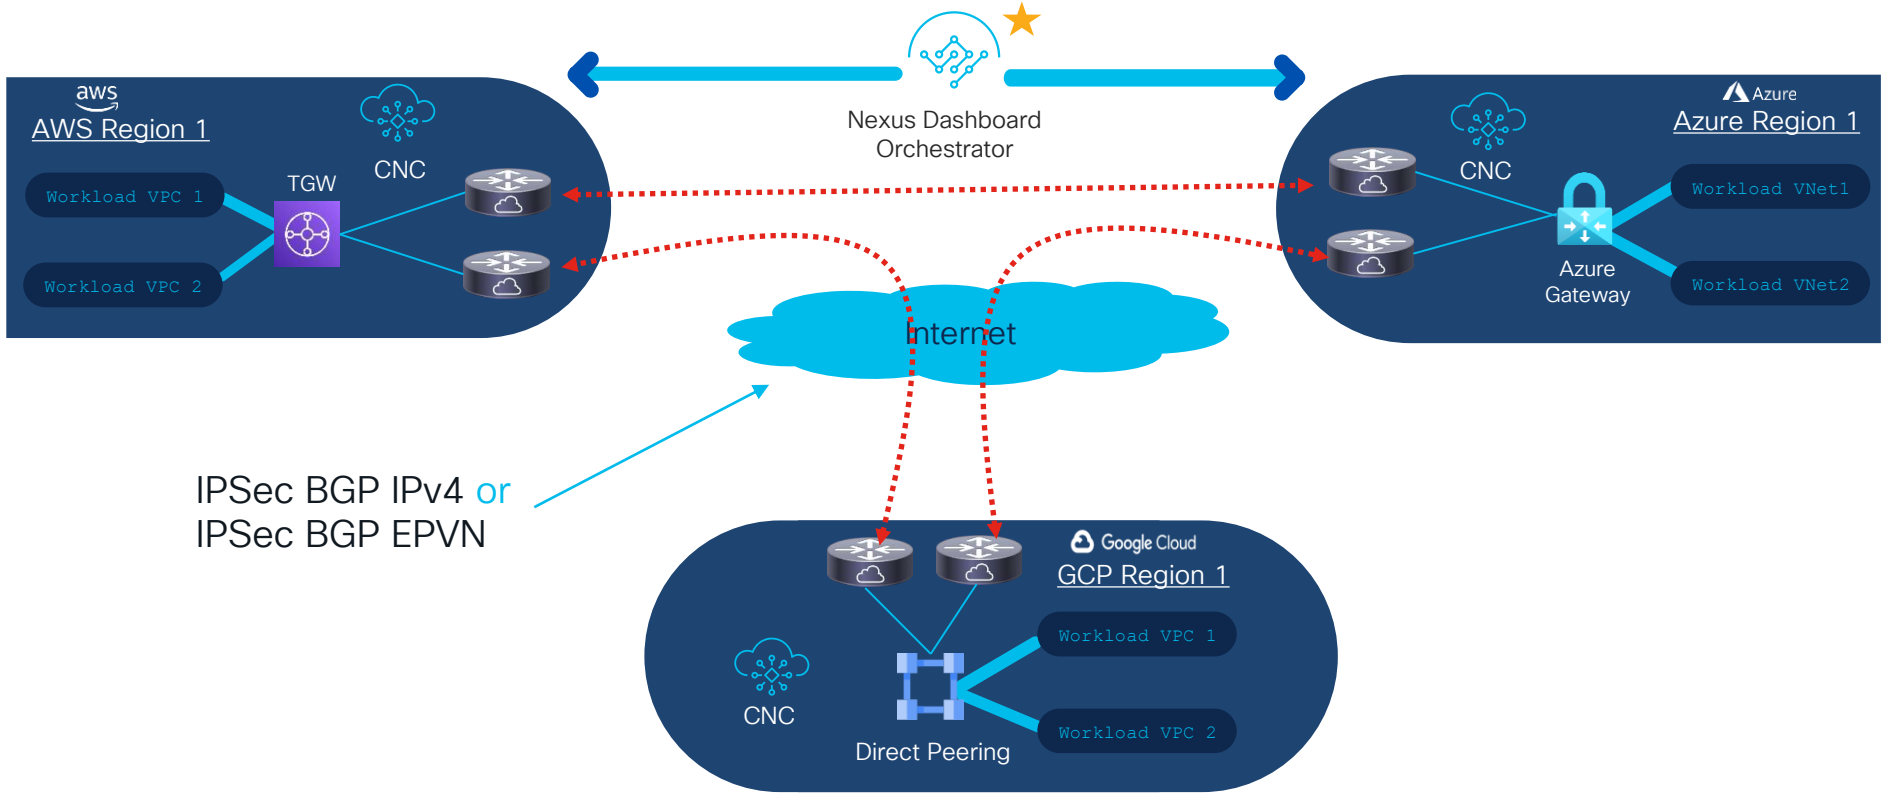

What About Brownfield?

- Connectivity to Azure brownfield VNets
- Connectivity to AWS brownfield VPCs
- Workload VPC / VNet
- VPC / VNet Route Table

- Attachment to new TGWs
- Import of existing VPC/Subnet
- VPCs/Subnets associated to the VRF Routing Domains

Branch to Cloud

Branch to Cloud Connectivity: Using Cat8kv

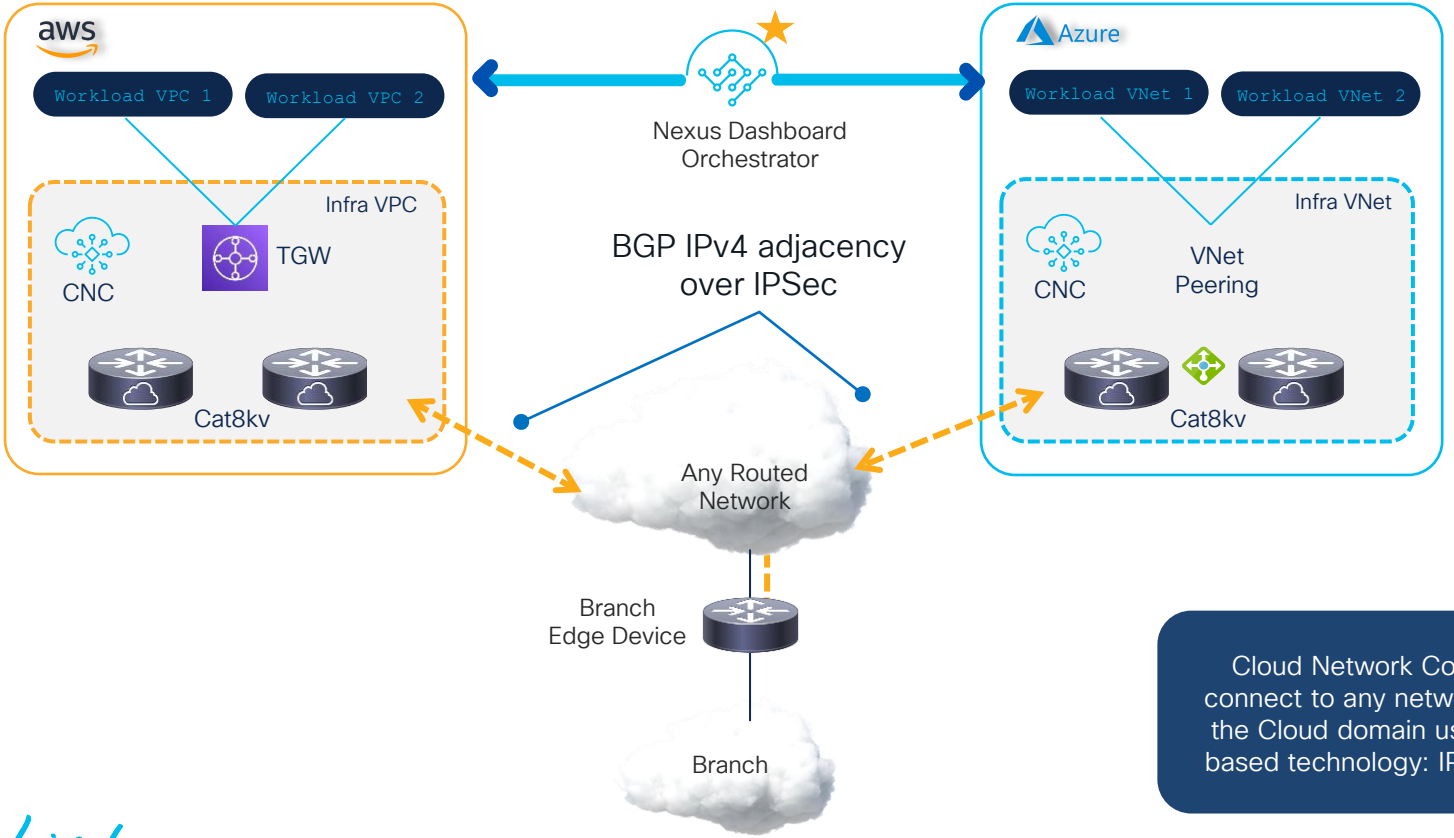

Cloud Network Controller can connect to any network outside of the Cloud domain using standard based technology: IPsec and BGP

Branch to Cloud Connectivity: Using SD-WAN

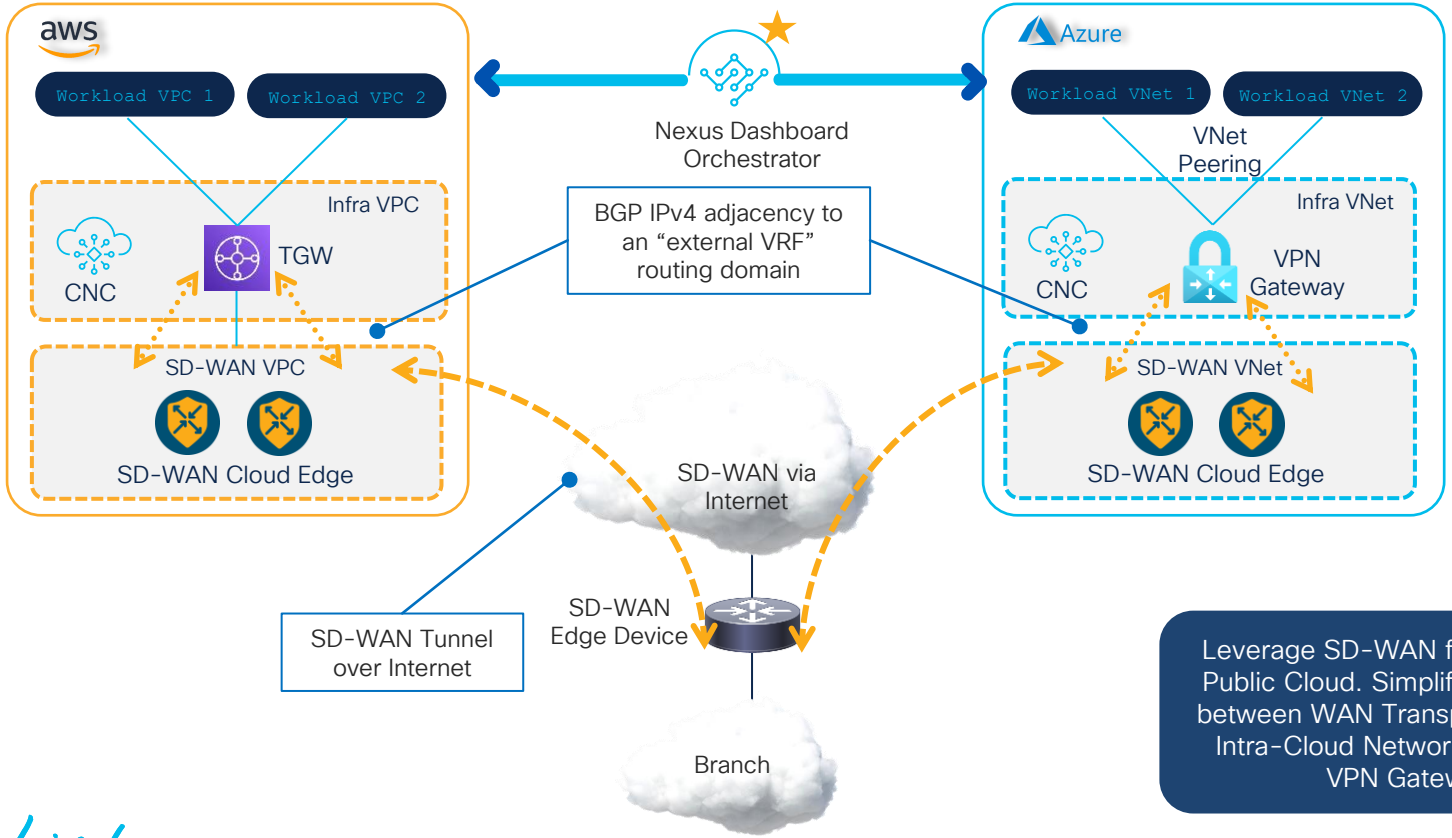

Leverage SD-WAN from Branch to Public Cloud. Simplified integration between WAN Transport and native Intra-Cloud Networking (TGW or VPN Gateway)

Interworking with SD-WAN

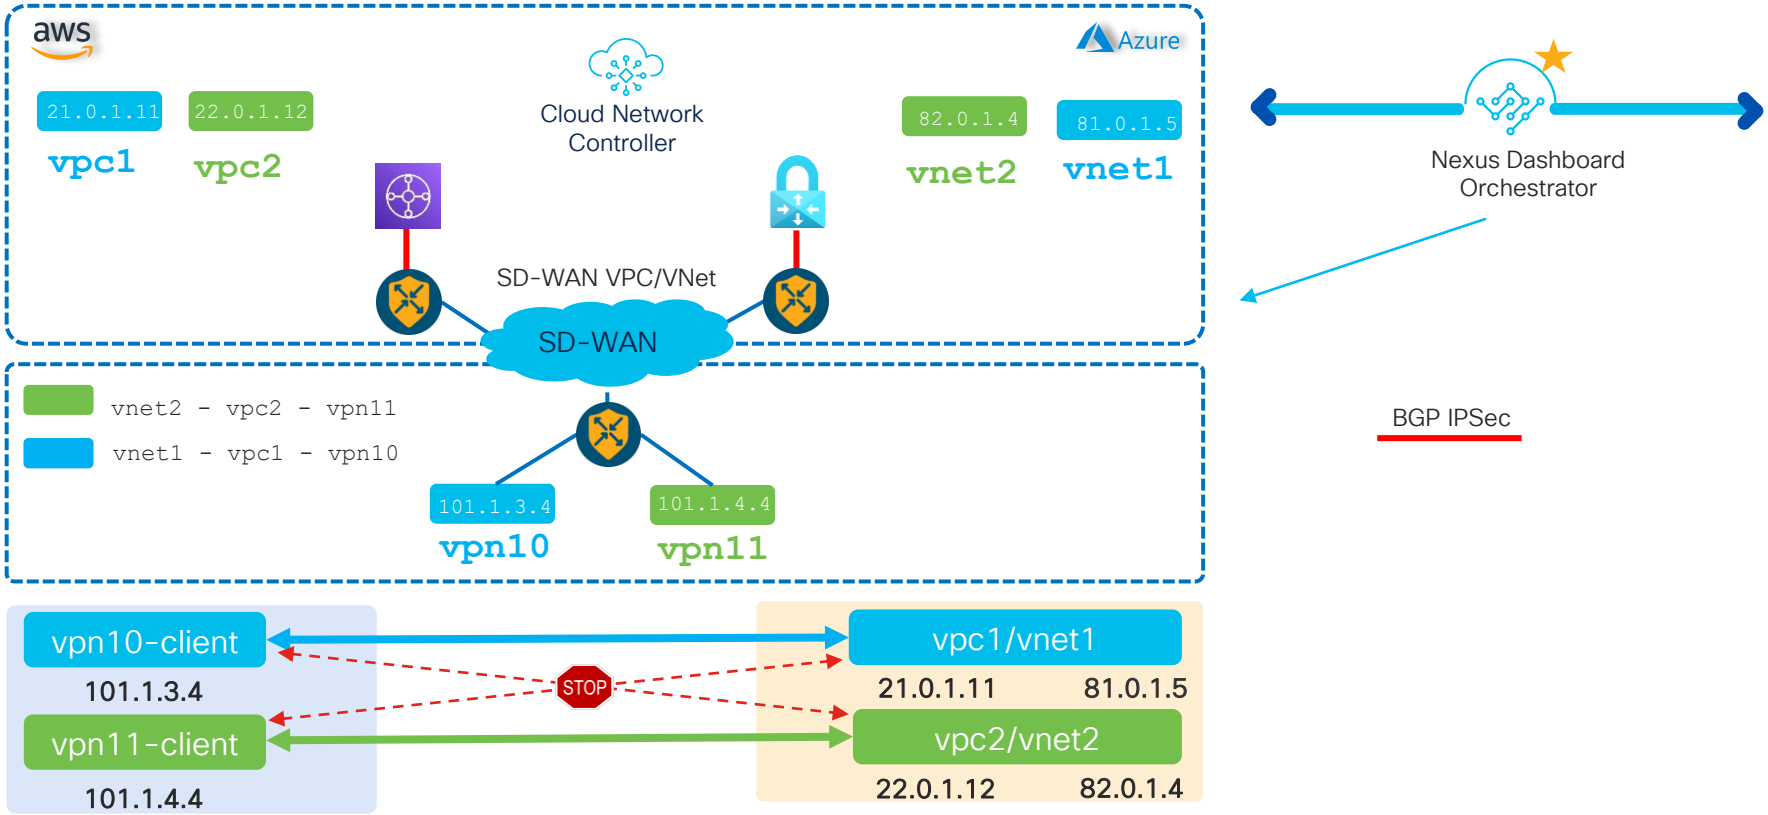

Contract-Based Routing or Routing Only with VRF leaking

Cloud Network Controller Implementation / Configuration

Enabling Contract-Based Routing

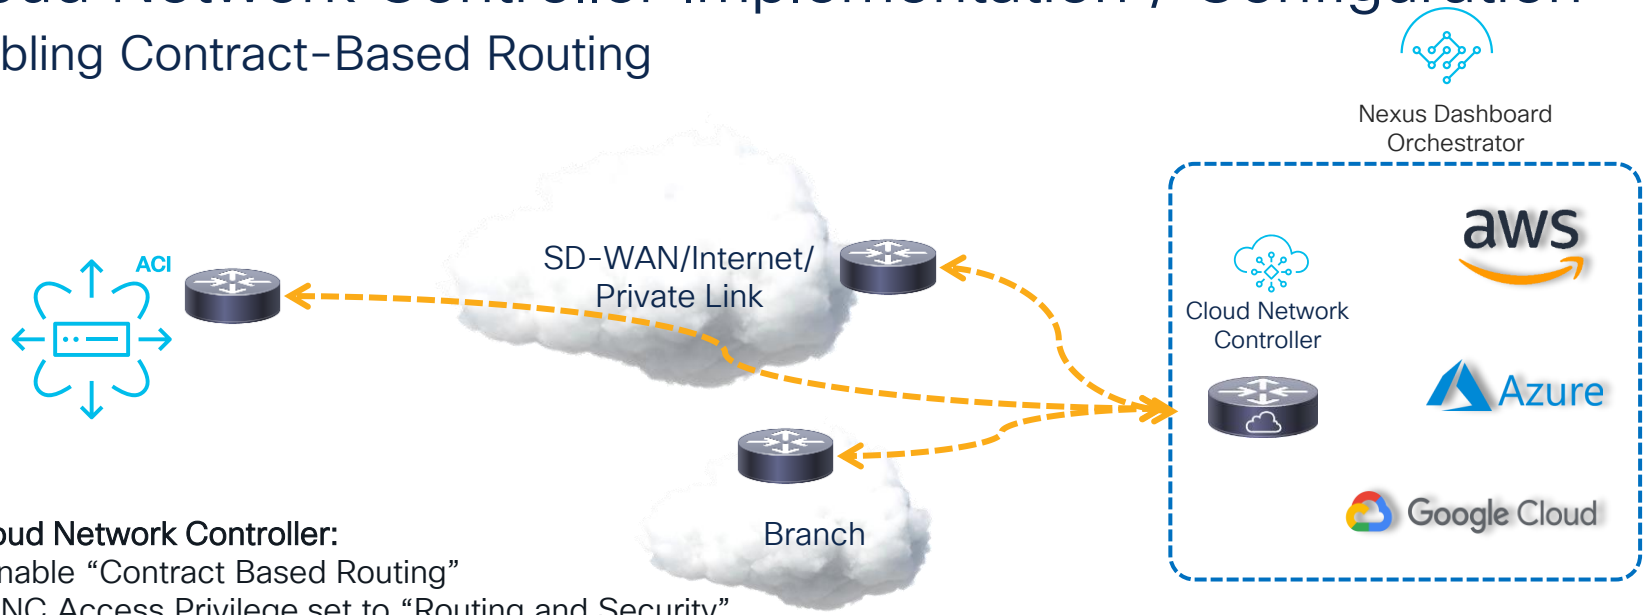

On Cloud Network Controller:

1) Enable “Contract Based Routing”
2) CNC Access Privilege set to “Routing and Security”

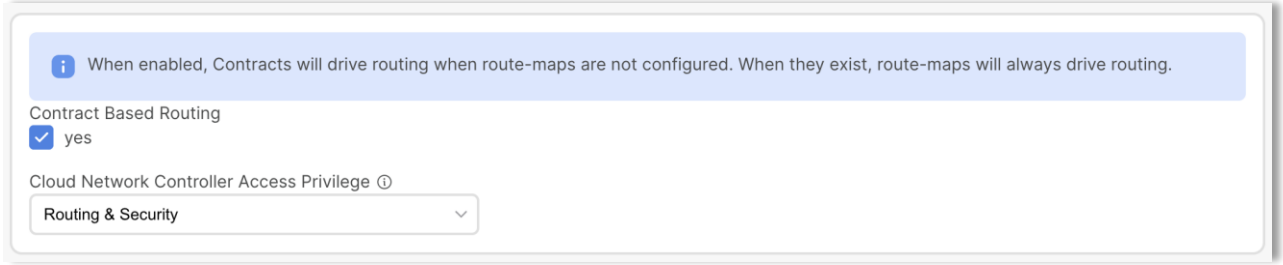

Cloud Network Controller Implementation/Configuration

Routing Only

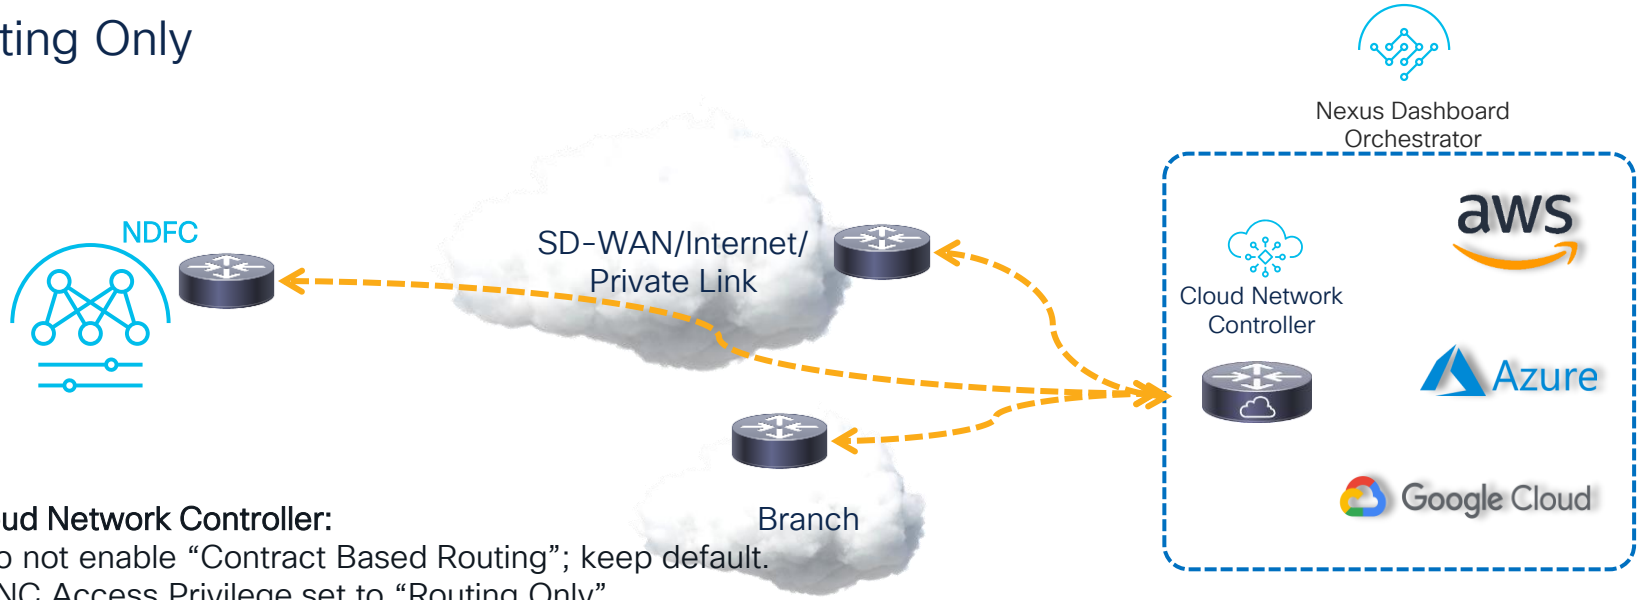

On Cloud Network Controller:

1) Do not enable “Contract Based Routing”; keep default.
2) CNC Access Privilege set to “Routing Only”

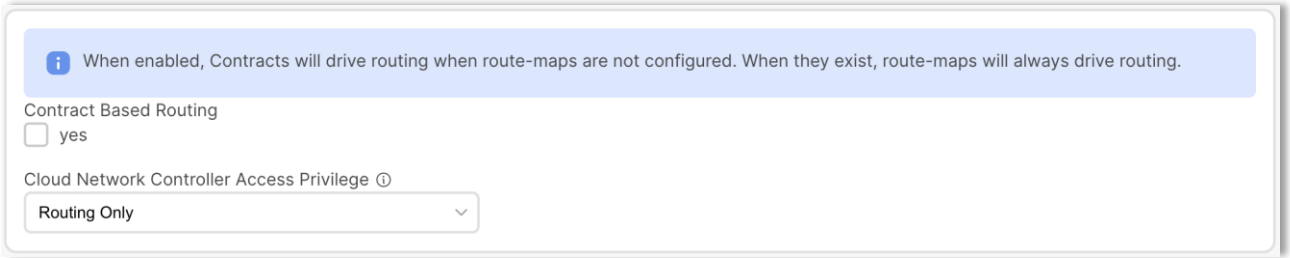

VRF Leak Example

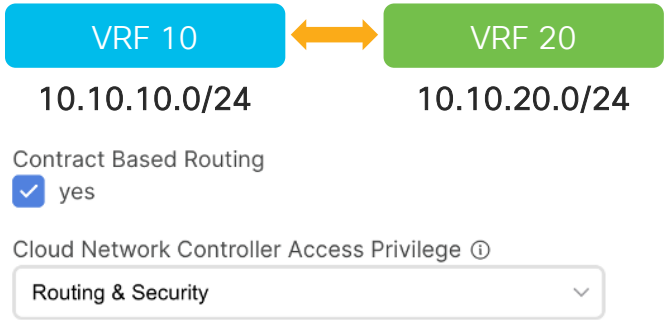

Contract is needed; enable contract-based routing. The contract will do the leaking implicitly.

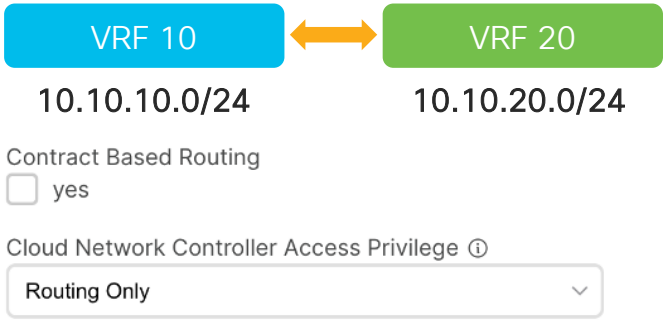

VRF leaking is still done by networking team but explicitly from CNC GUI; no contract needed/supported. Leaking can be done by FW/Routing Protocol.

CCNC 26.0.2
New Features

CCNC 26.0.2 New Features

- Multiple VPCs/VNets in a single VRF in a region AWS/Azure
- Per Subnet Route Table – Azure. (AWS next release)
- Multi-Account per tenant for AWS/Azure

Multiple VPCs/VNets in a single VRF for AWS and Azure in a region

Feature overview

- Prior to CCNC 26.0.2 release, ability have multiple VPCs/VNets map to one VRF in different cloud regions.
- Customers are asking for a smaller number of VRFs for a large number of VPCs/VNets in a single cloud region.
- From CCNC 26.0.2 release, allows customers to create multiple VPCs/VNets in same cloud region associated to one VRF.
- Resulting in reduced VRF management complex and building routing domain with a group of VPCs/VNets

Prior to CCNC release 26.0.2

One VRF maps to multiple VPCs/VNets in different cloud regions.

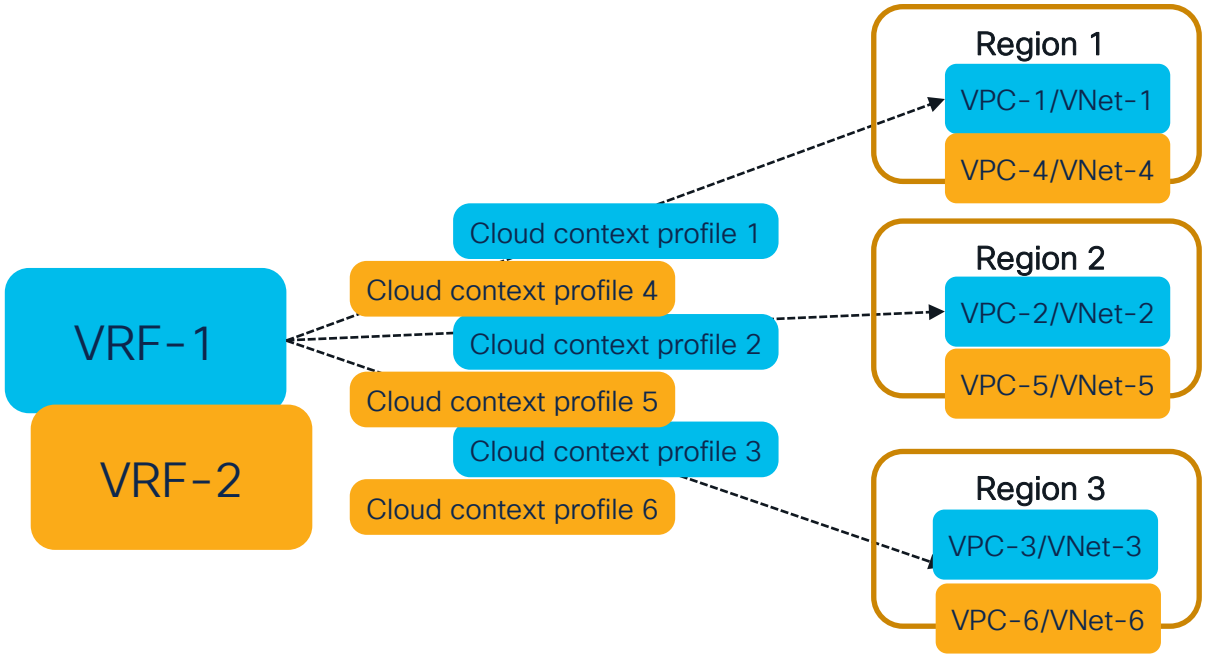

From CCNC release 26.0.2

One VRF maps to multiple VPCs/VNets in “Region 1” and “Region 2”

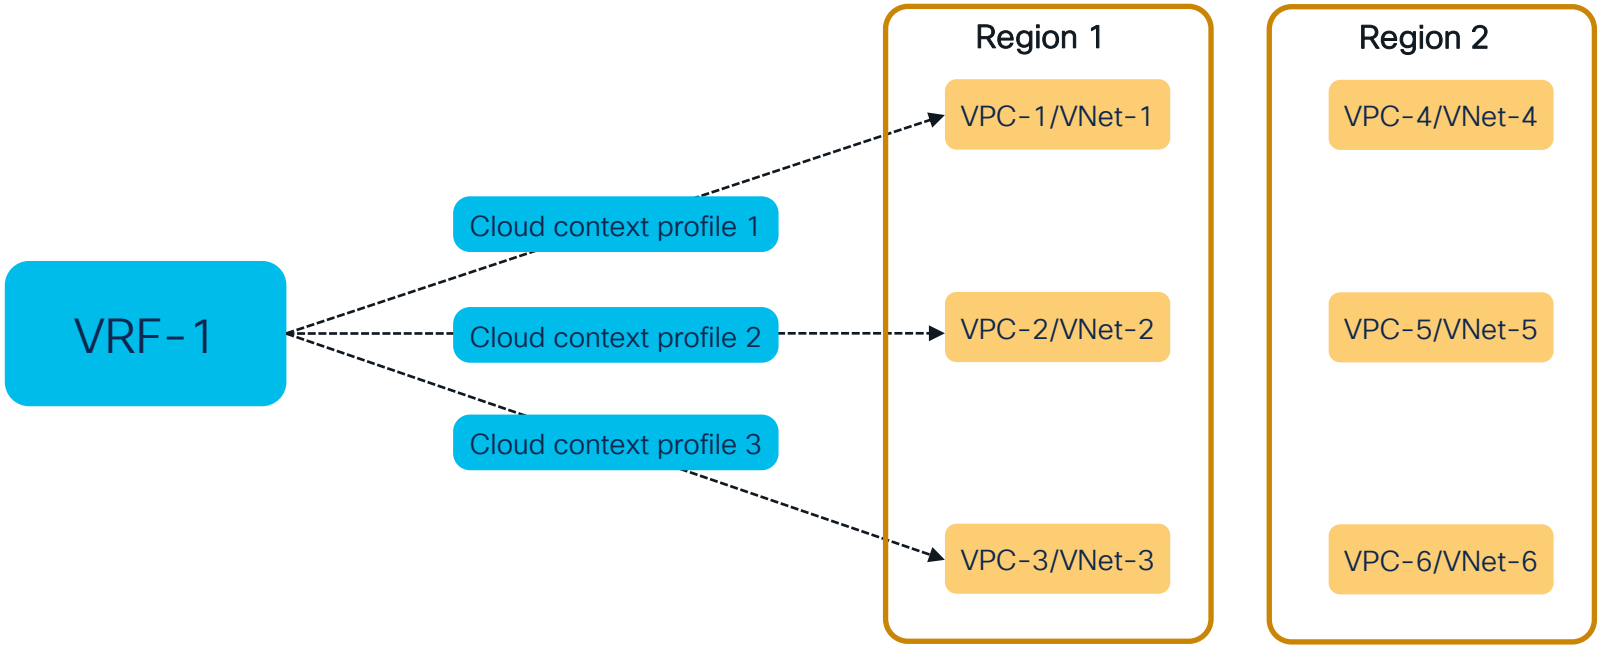

From CCNC release 26.0.2

One VRF maps to multiple VPCs/VNets in “Region 3”

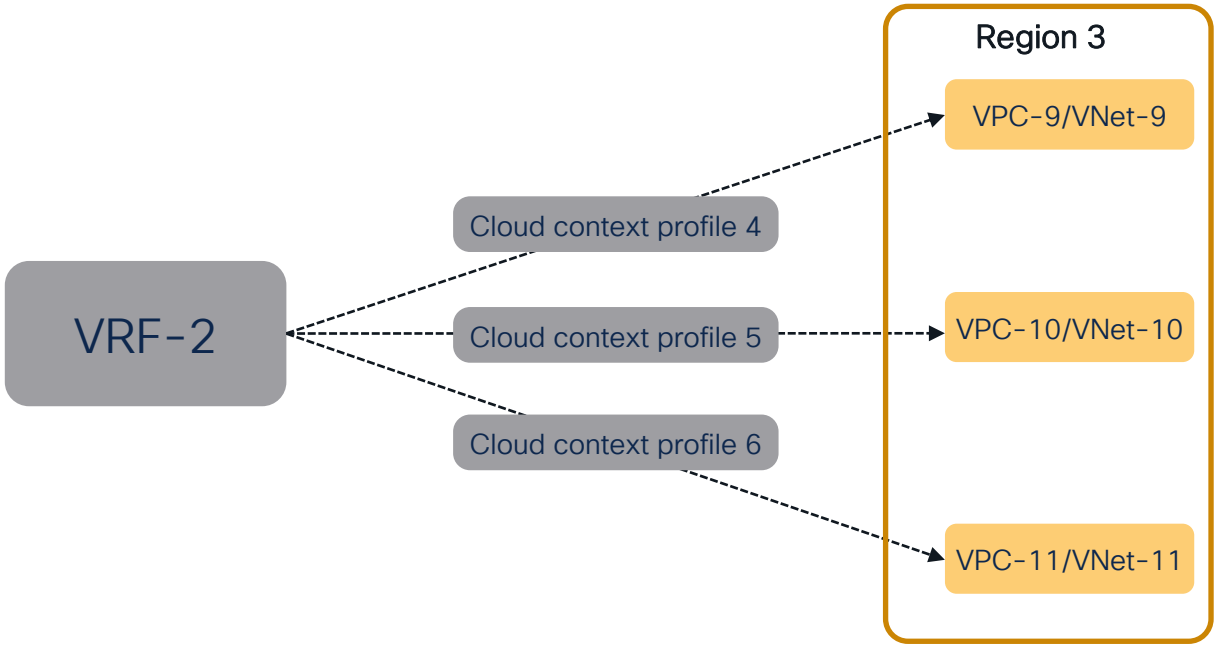

CCNC Routing Automation

NDO/CNC Allows ”VRF-1” and ”VRF-2” Communicate

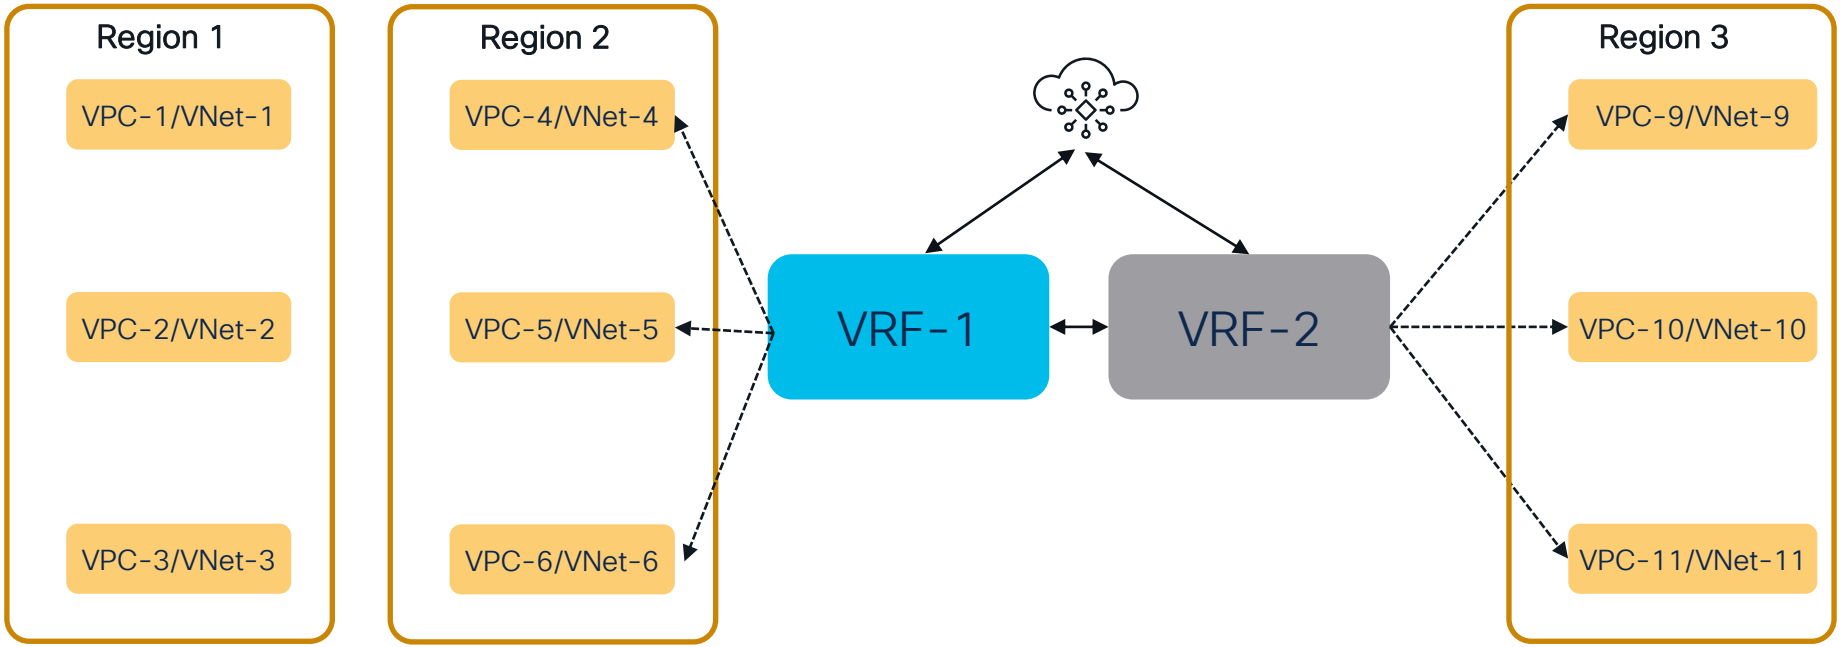

Per Subnet Route Table Azure

(AWS next release)

Feature Overview

- Per Subnet Route Table allows us to create a VNet with multiple subnets, where each subnet can map to its own Route Table, or we can group some subnets to map to the same Route Table.
- The benefit of this solution is allowing different routing treatments for different VNet's subnets. It also paves the way for other useful features in the future. For example, the insertion of a Firewall between 2 subnets.

Previously, all subnets were associated to one route table

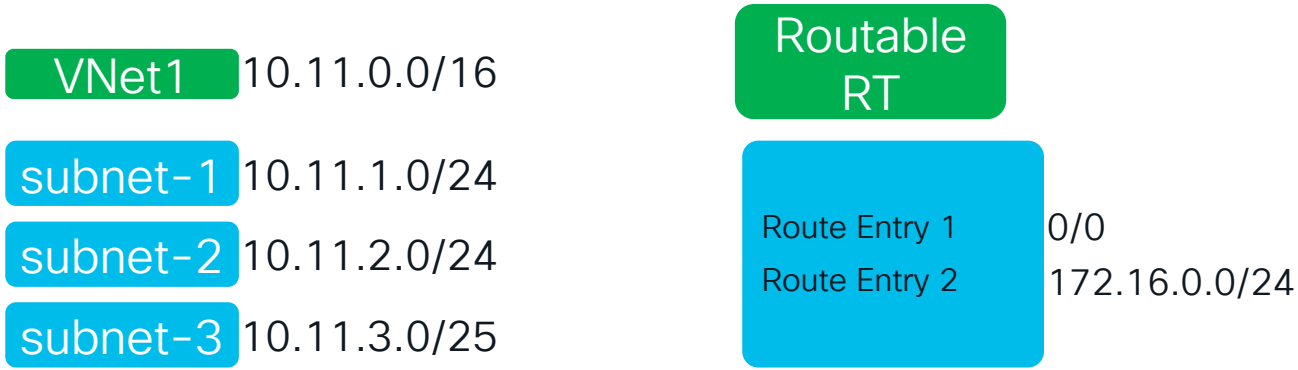

Per Subnet Route Table: Flexibility

Azure UDR

In the below example:

1) Subnets belong to the same VNet1
2) Different routing policy based its own route table
3) Subnet-1 cannot go to Internet; just Branch
4) Subnet-2 can reach Internet directly
5) Only subnet-3, subnet-4 can reach on-premises DC

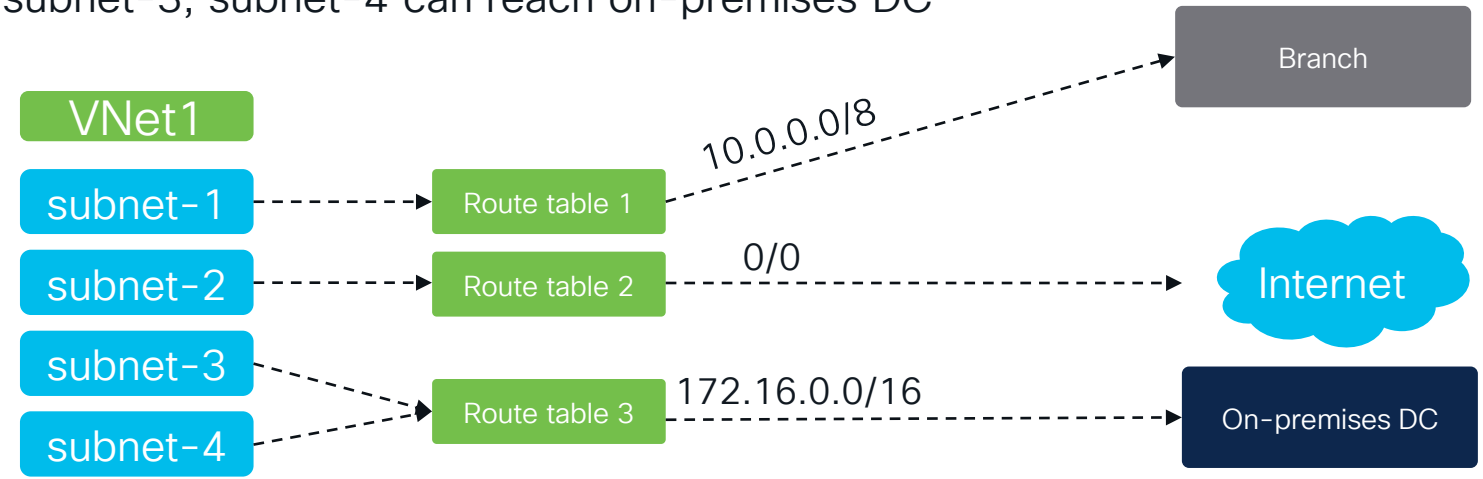

Per Subnet Route Table (Per Subnet UDR)

- Subnet-1 to Subnet-2 must go through a Firewall

- You need to have different route tables; otherwise, subnet-1 would just talk to subnet-2 w/o Firewall

- Granularity at the subnet level

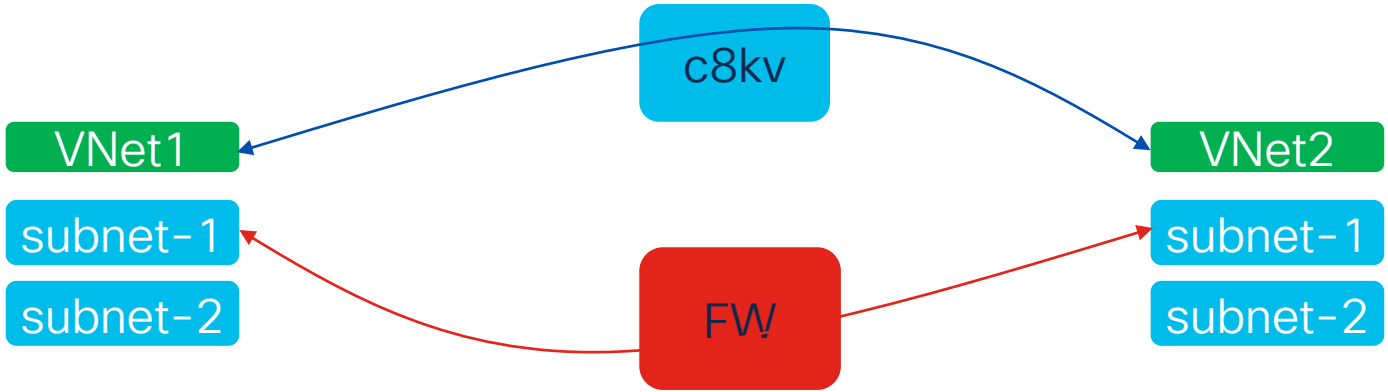

Per Subnet Route Table (Per Subnet UDR)

VNet1 -> VNet2:
1) c8Kv
2) FW

Within VNet1 subnet1 -> subnet 2:
1) cloud native

VNet1 subnet1-> VNet2 subnet1
1) c8KV
2) FW

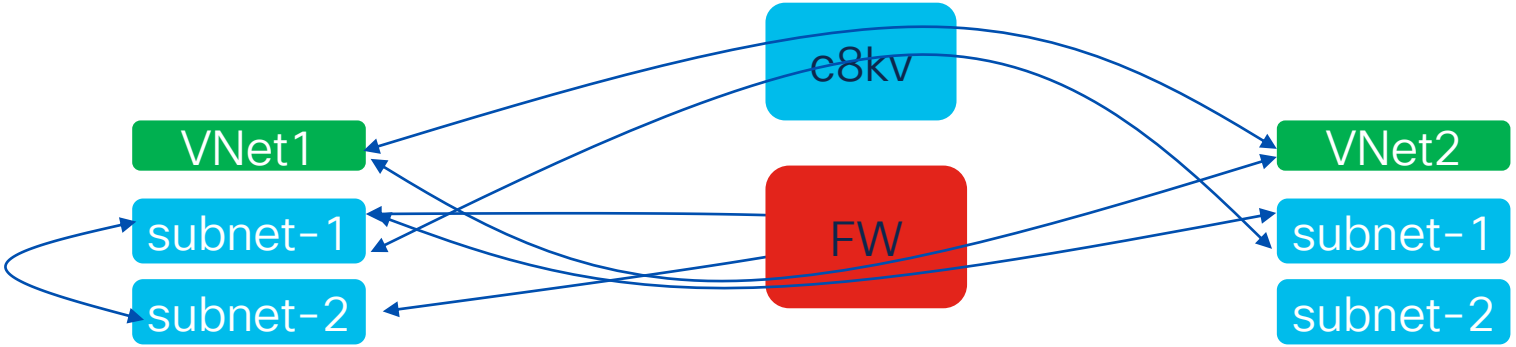

within VNet1 Subnet1 - subnet 2:
1) FW

Vnet1 -> VNet2:
1) cloud native

Multi-Account for AWS/Azure

Multi-Account

From CCNC 26.0.2 release, it is possible to create one tenant and associate it with multiple cloud accounts:

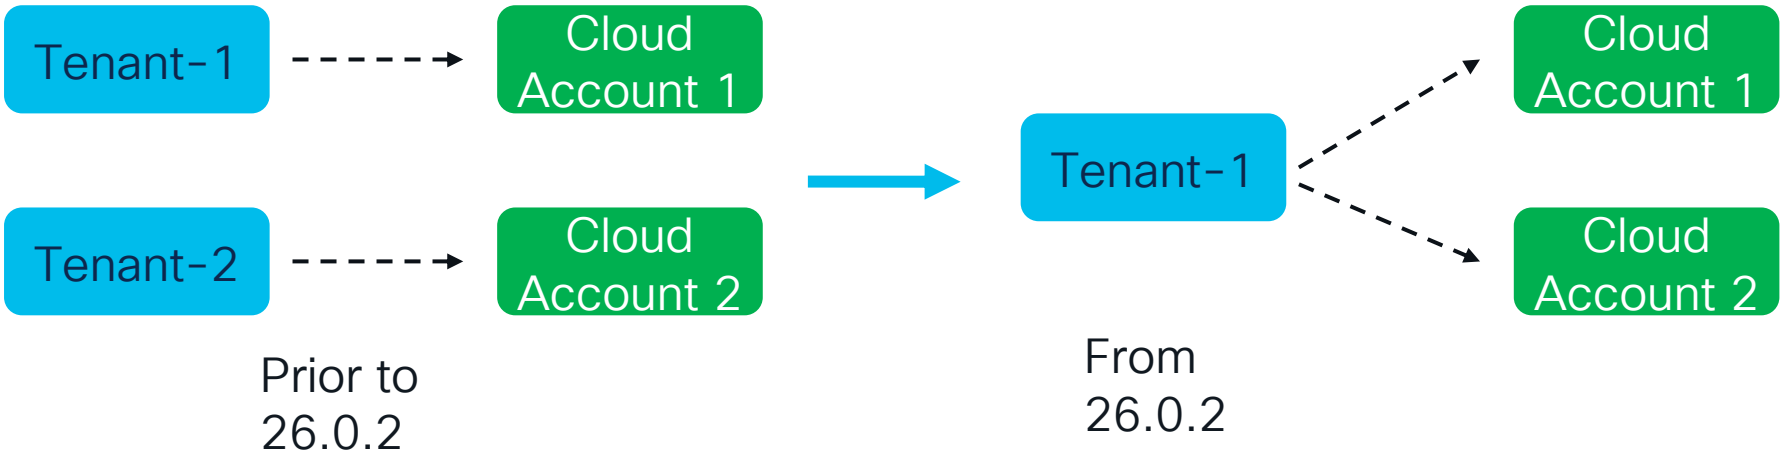

**Cloud account refers to AWS account ID or Azure subscription or GCP project.*

Sneak Peek

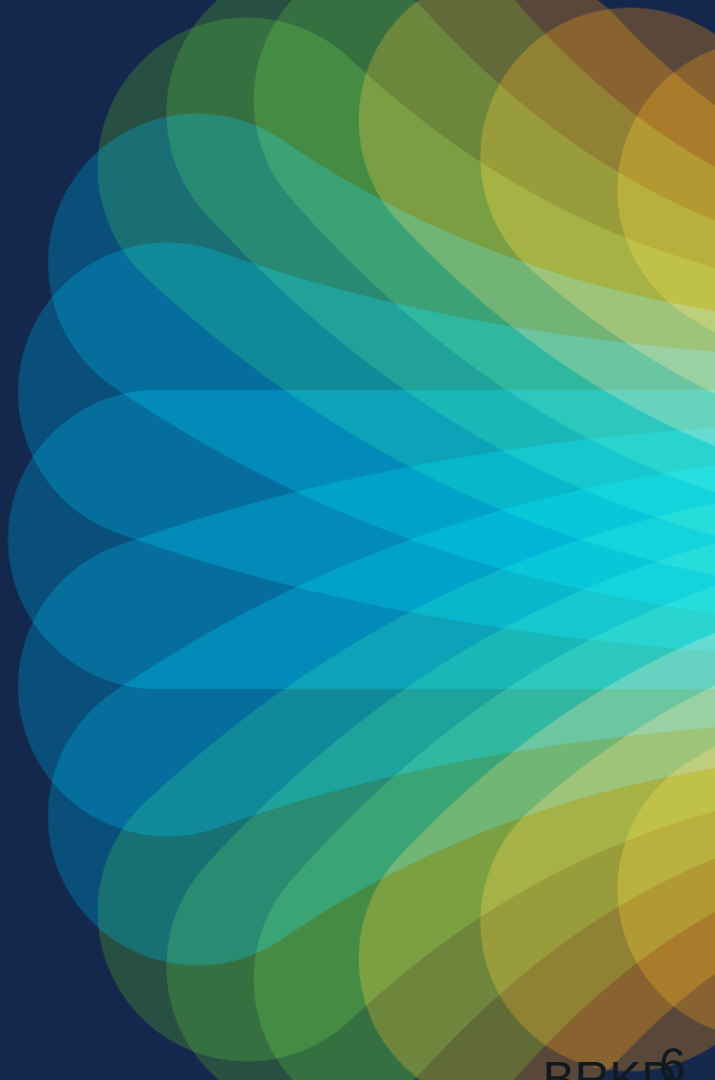

Common tenant hybrid cloud

26.0.3 - CMR-2 release

26.0.3

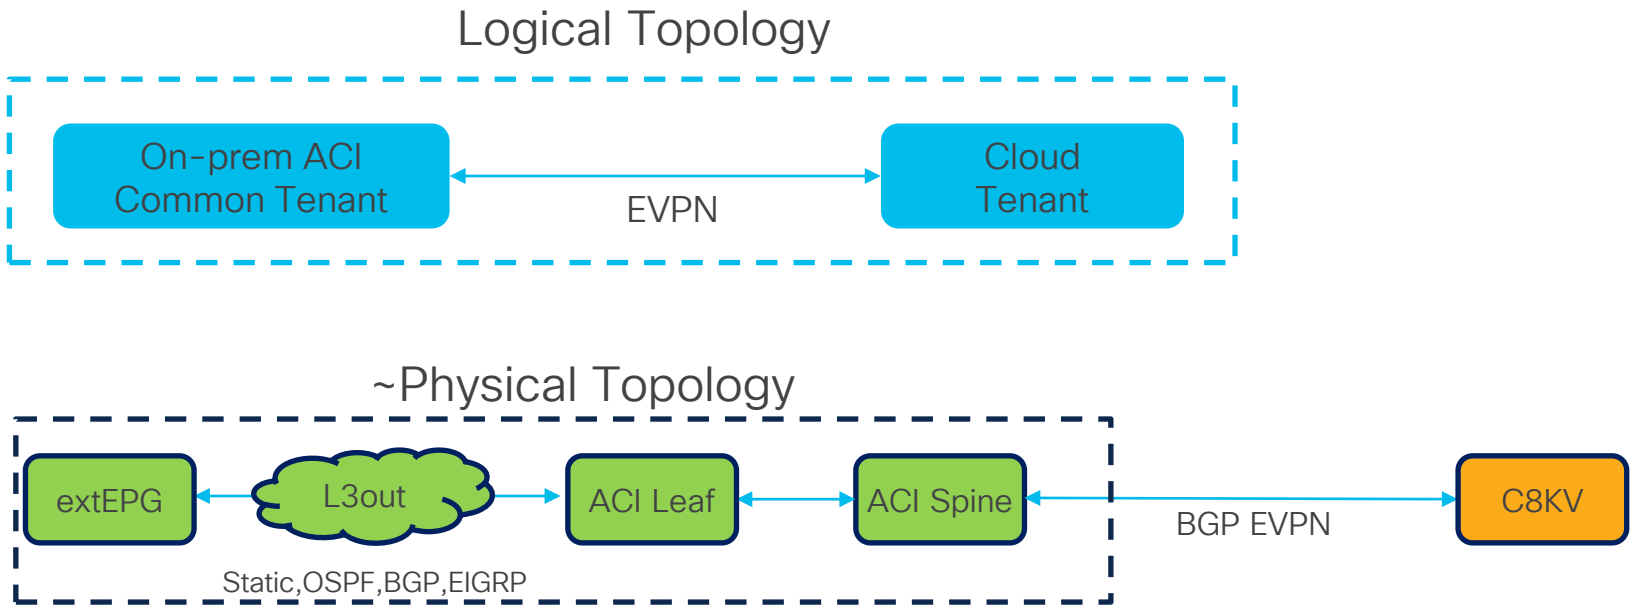

Hybrid Cloud

Use case Tenant Common

26.0.3

Tenant: Common
1) subnet under BD
2) subnet under EPG

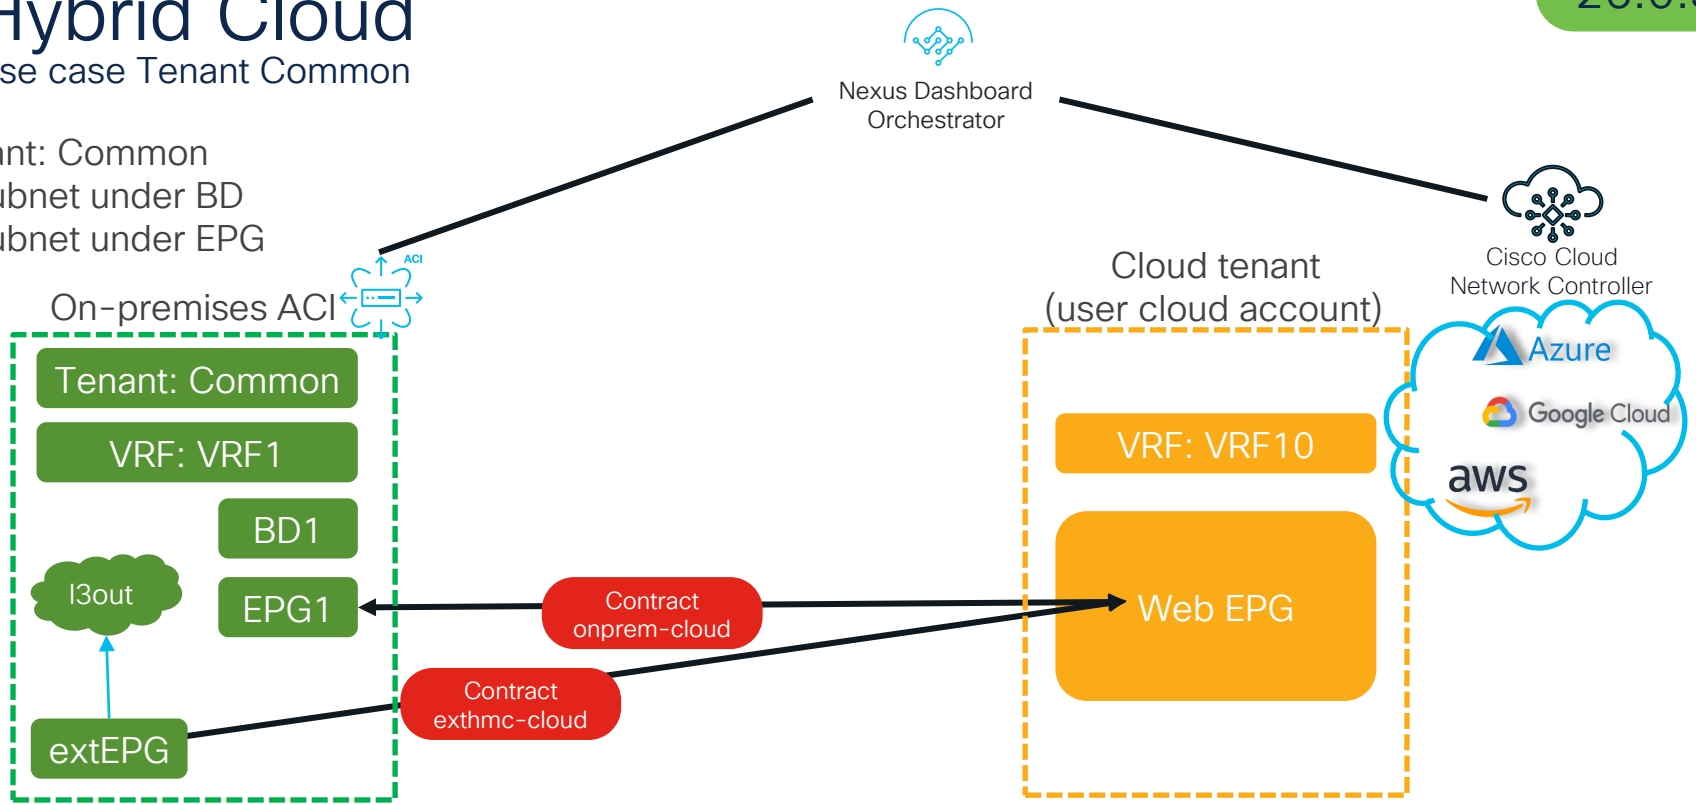

Hybrid Cloud

Use case Tenant Common

26.0.3

Tenant: Common
1) subnet under BD
2) subnet under EPG

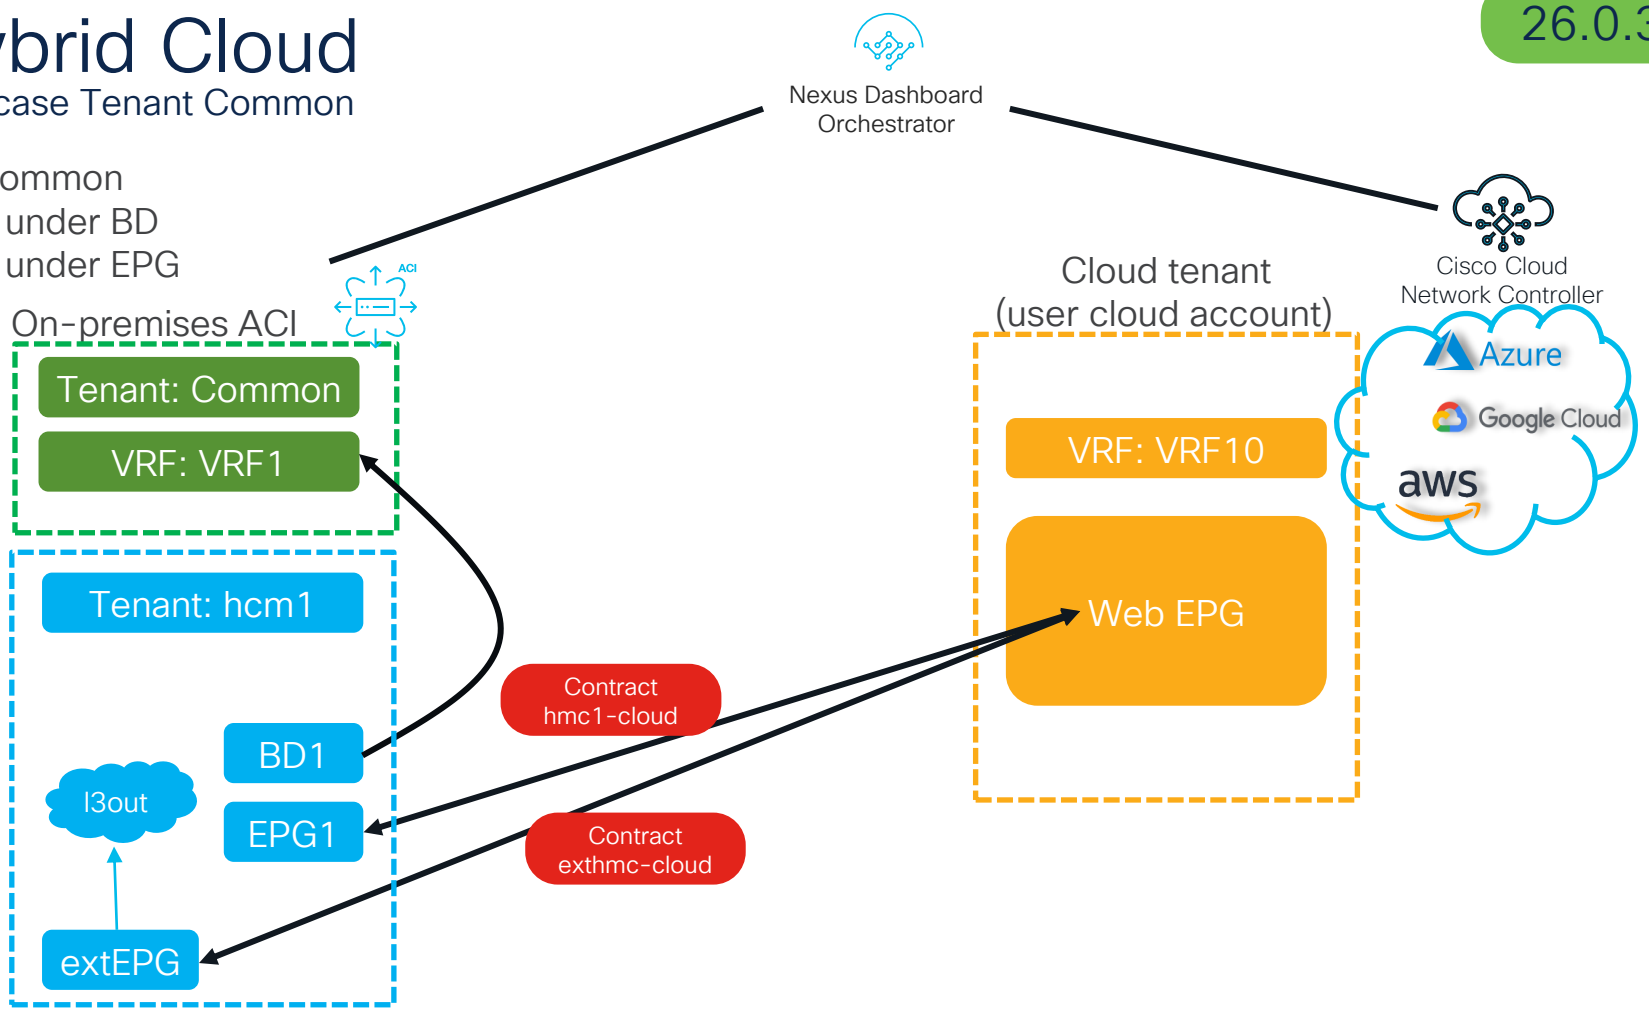

Hybrid Cloud

Use case Tenant Common

26.0.3

Tenant: Common
1) subnet under BD
2) subnet under EPG

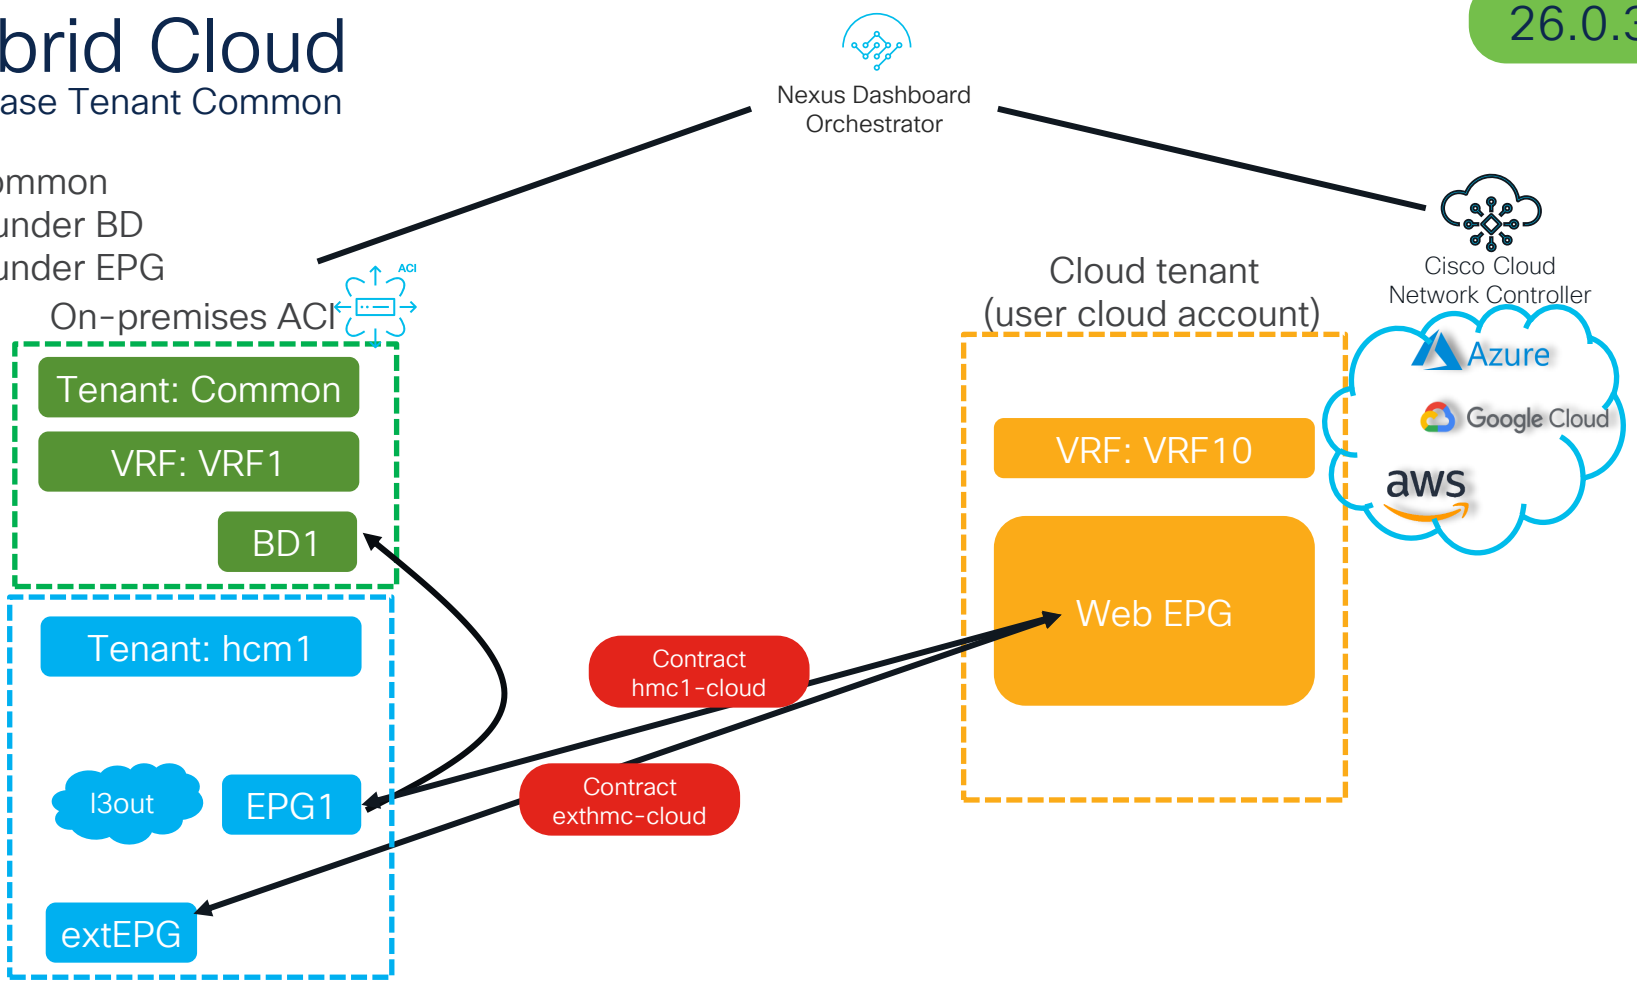

Hybrid Cloud

Use case Tenant Common + Stretched Tenant

26.0.3

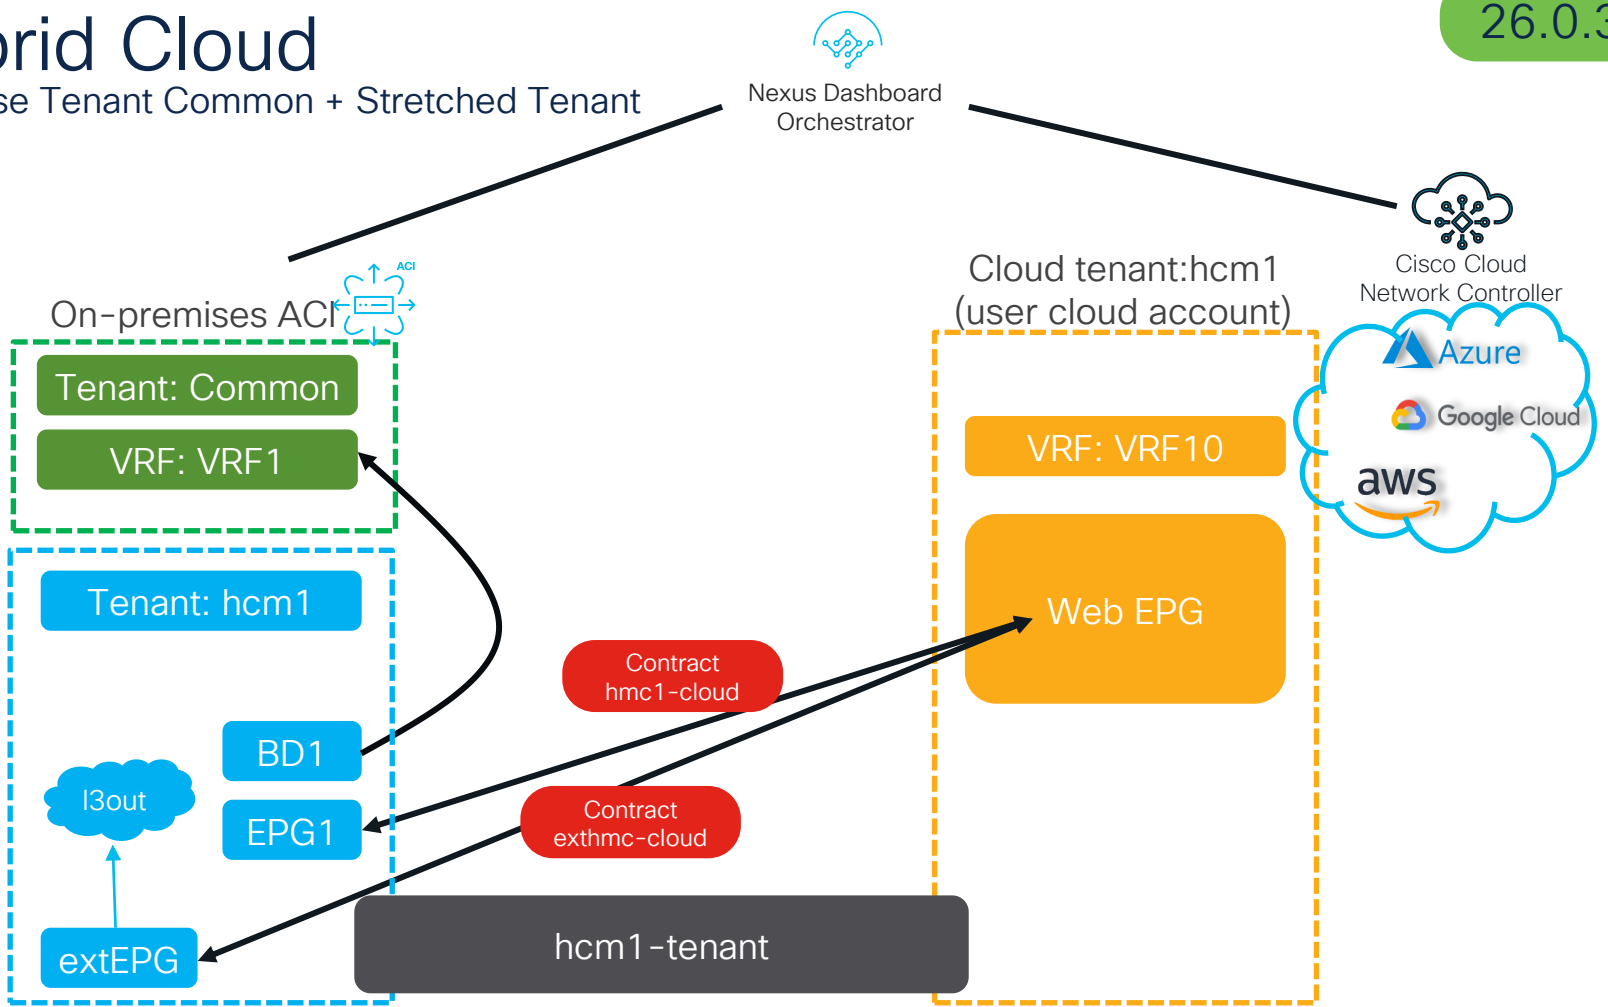

Service Chaining – Internet-Cloud Use case

27.0.1

27.0.1

Existing Deployment

Easier model going forward

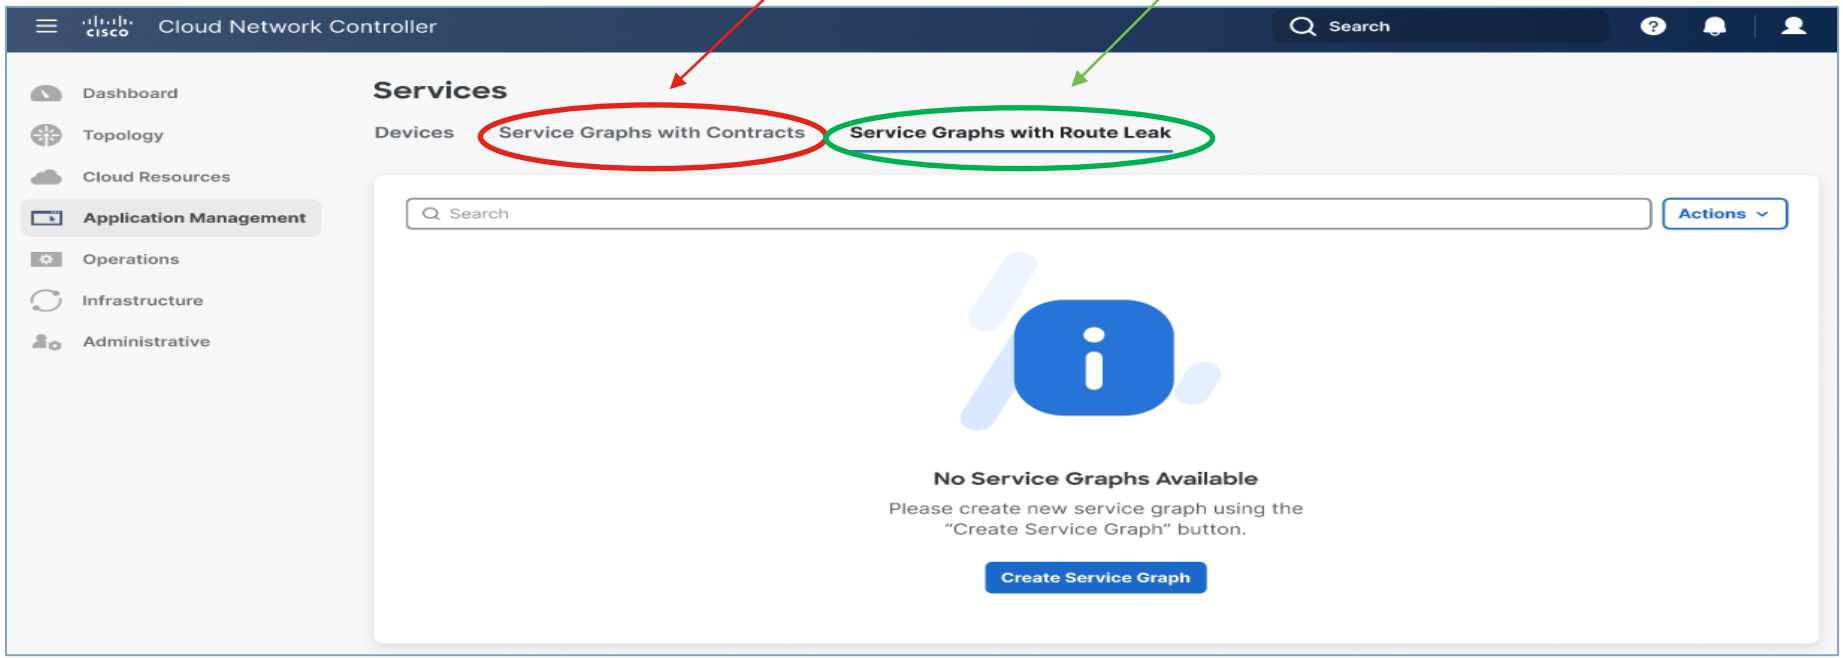

27.0.1

Service Graph with Route-leaking

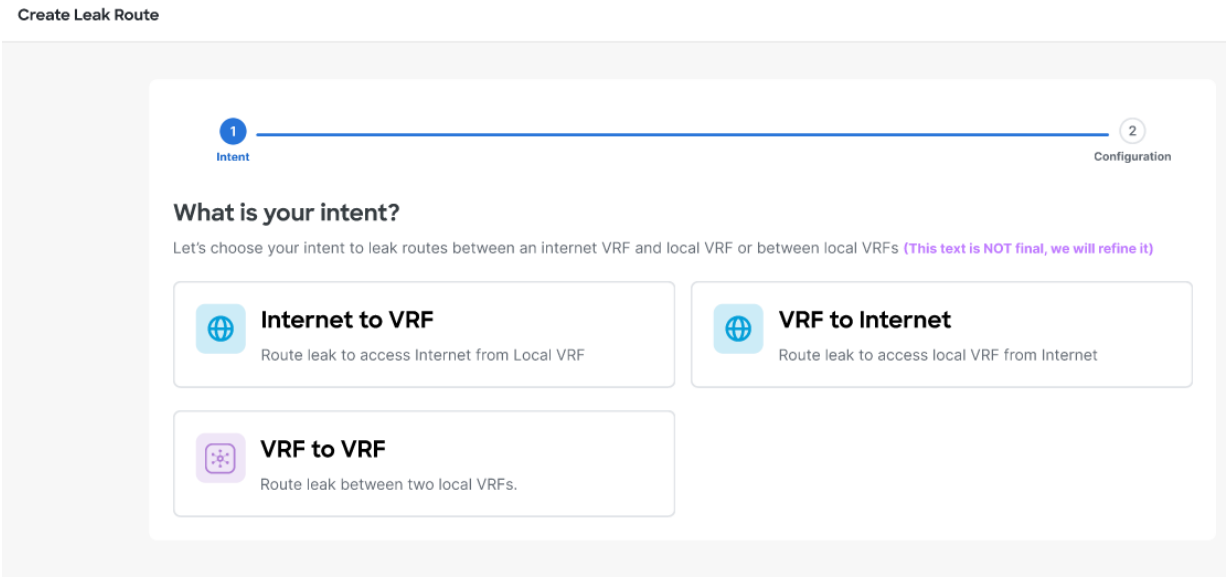

Service Graph with Route-leaking

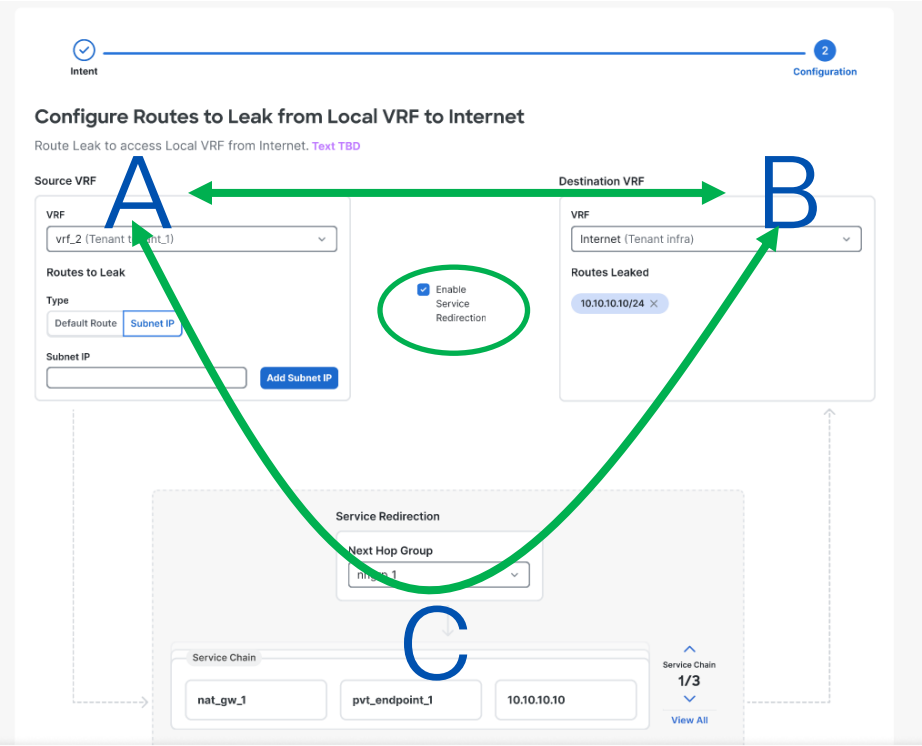

Hybrid Cloud

Use case Tenant Common

26.0.3

Tenant: Common
1) subnet under BD
2) subnet under EPG

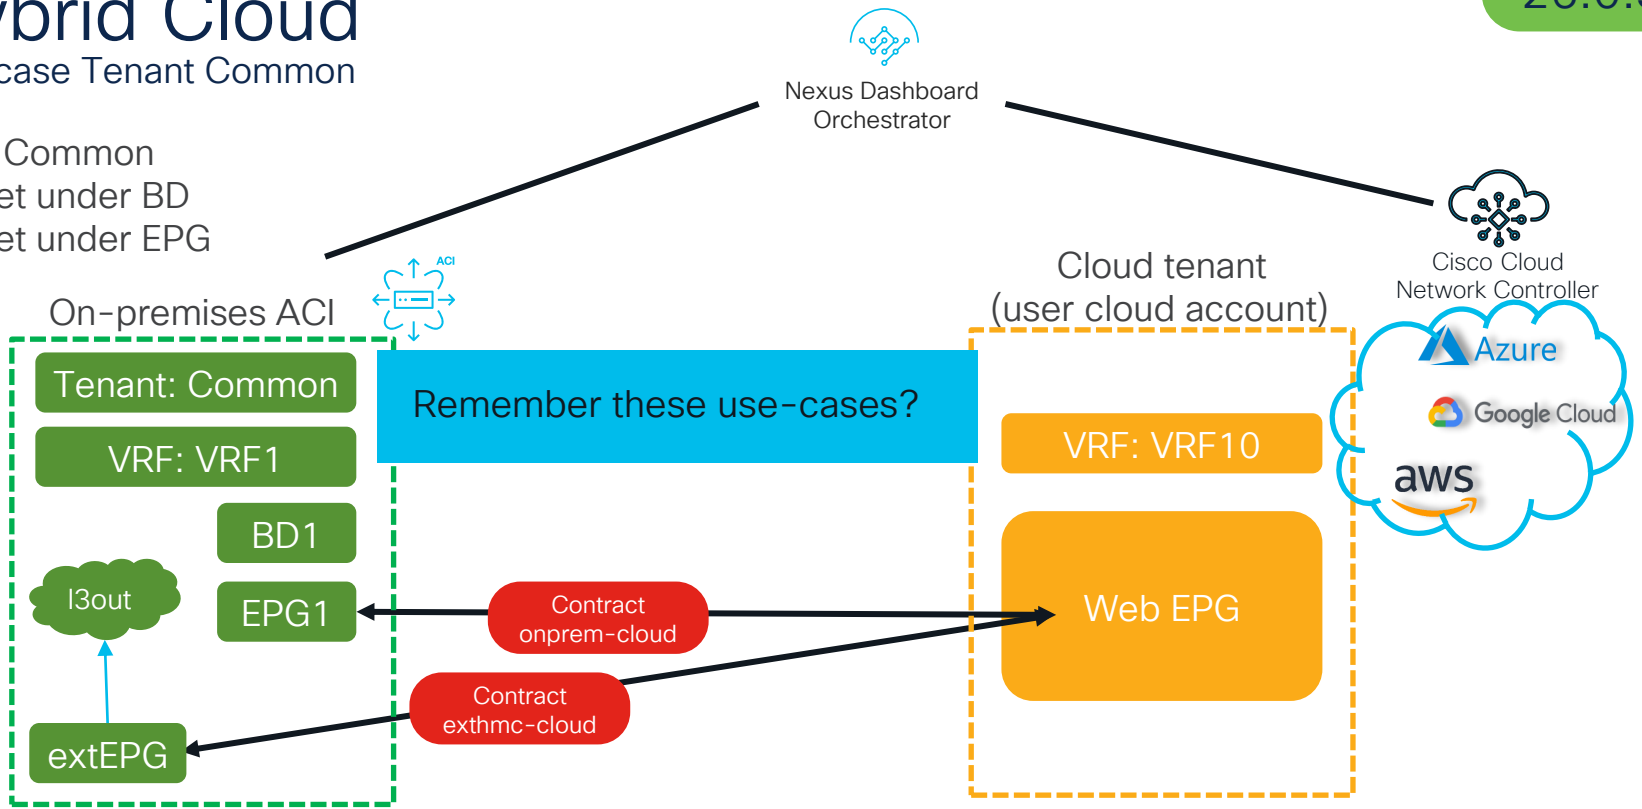

Per Subnet Route Table benefit

Amazon VPC Support

In the below example:

1) Subnets belong to the same VPC1
2) Different routing policy based its own route table
3) Subnet-1 cannot go to Internet; just Branch
4) Subnet-2 can reach Internet directly
5) Only subnet-3, subnet-4 can reach on-premises DC

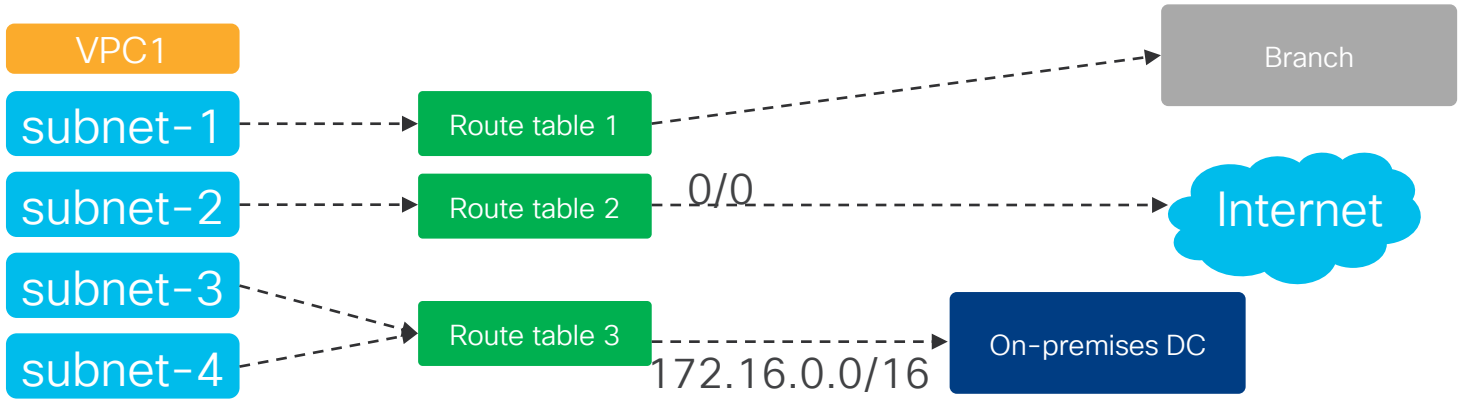

Hybrid Cloud network automation and segmentation

CCNC common use cases

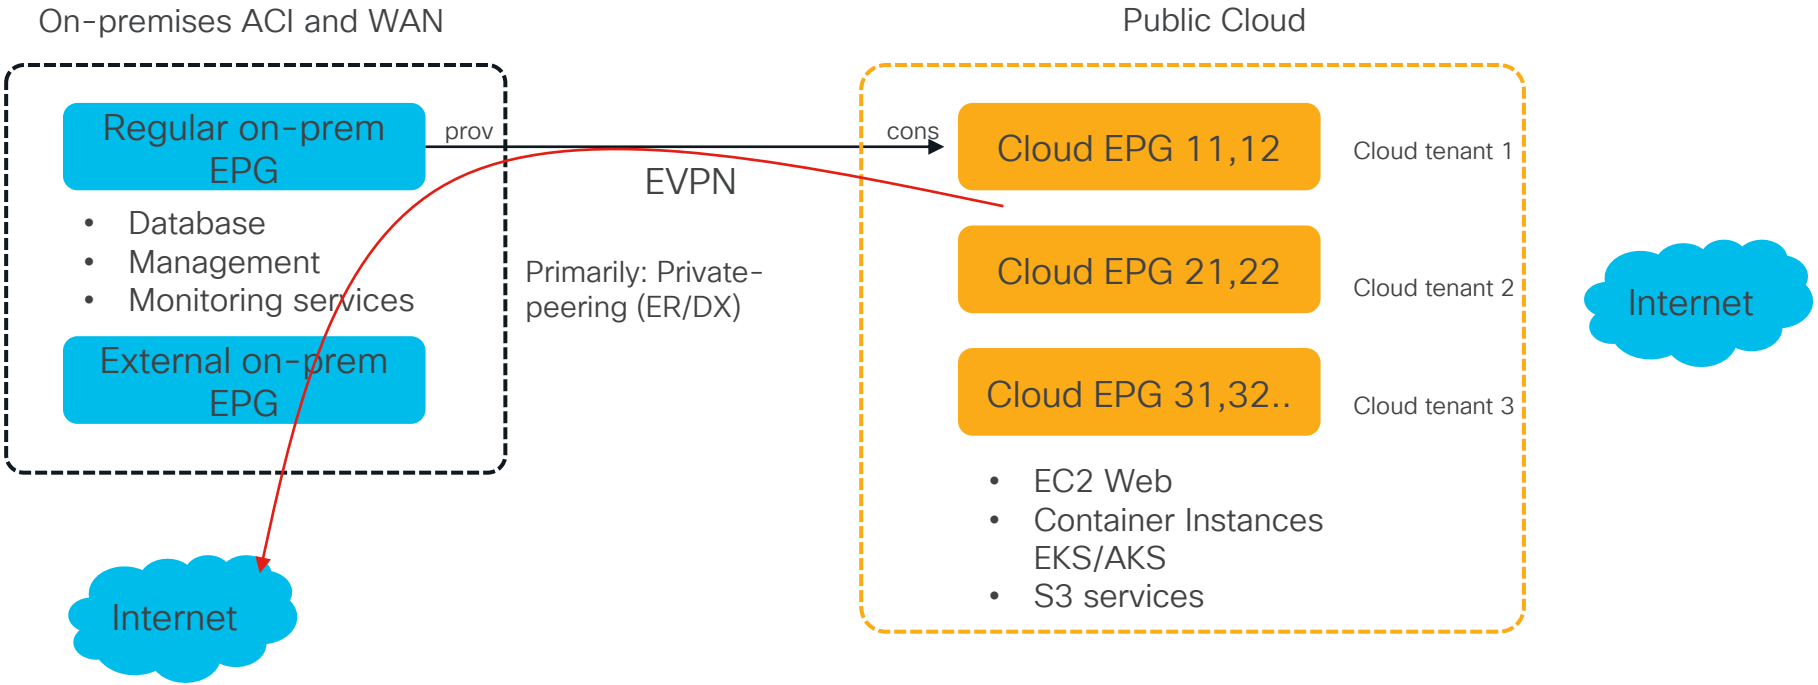

Visibility

Visibility Capabilities


Cloud
Network Controller


Nexus Dashboard
Orchestrator

- **Nexus Dashboard Orchestrator and CNC:**
  - Inventory
  - Topology (Underlay / Overlay)
  - Statistics and Bandwidth
  - L4-L7 Statistics
  - Troubleshooting tools
  - Raise faults (API failures, delays, resource constraints, state changes)
  - Provides visibility to cloud resources and the correlation with the configuration.
  - Drift Detection – Reconcile with cloud state across reboots, software upgrades
  - Event Analytics Tab in CNC (account permissions, licensing subscription, protocols down)
  - IP Connectivity Analysis
  - Site Details

Cloud Network Controller: Topology

Inventory

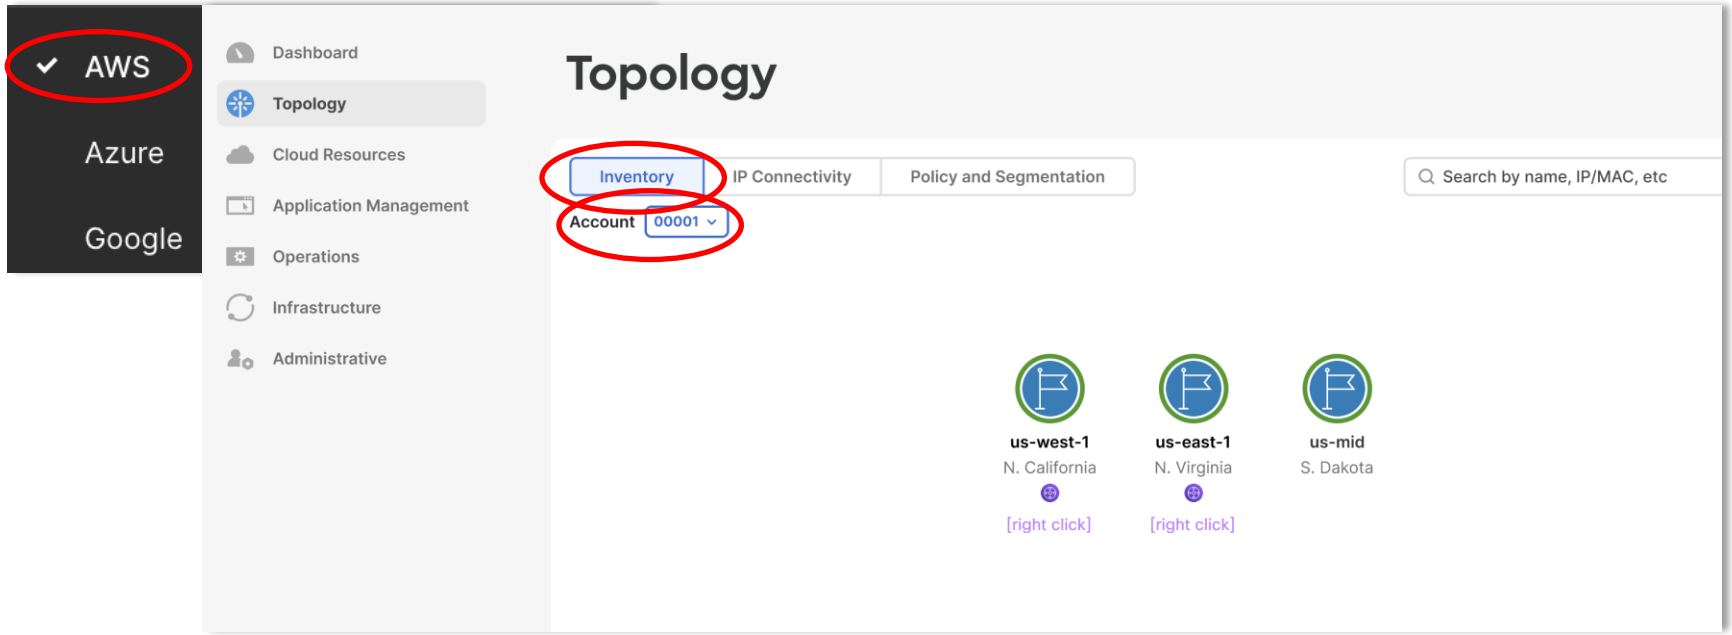

Cloud Network Controller: Topology

IP Connectivity

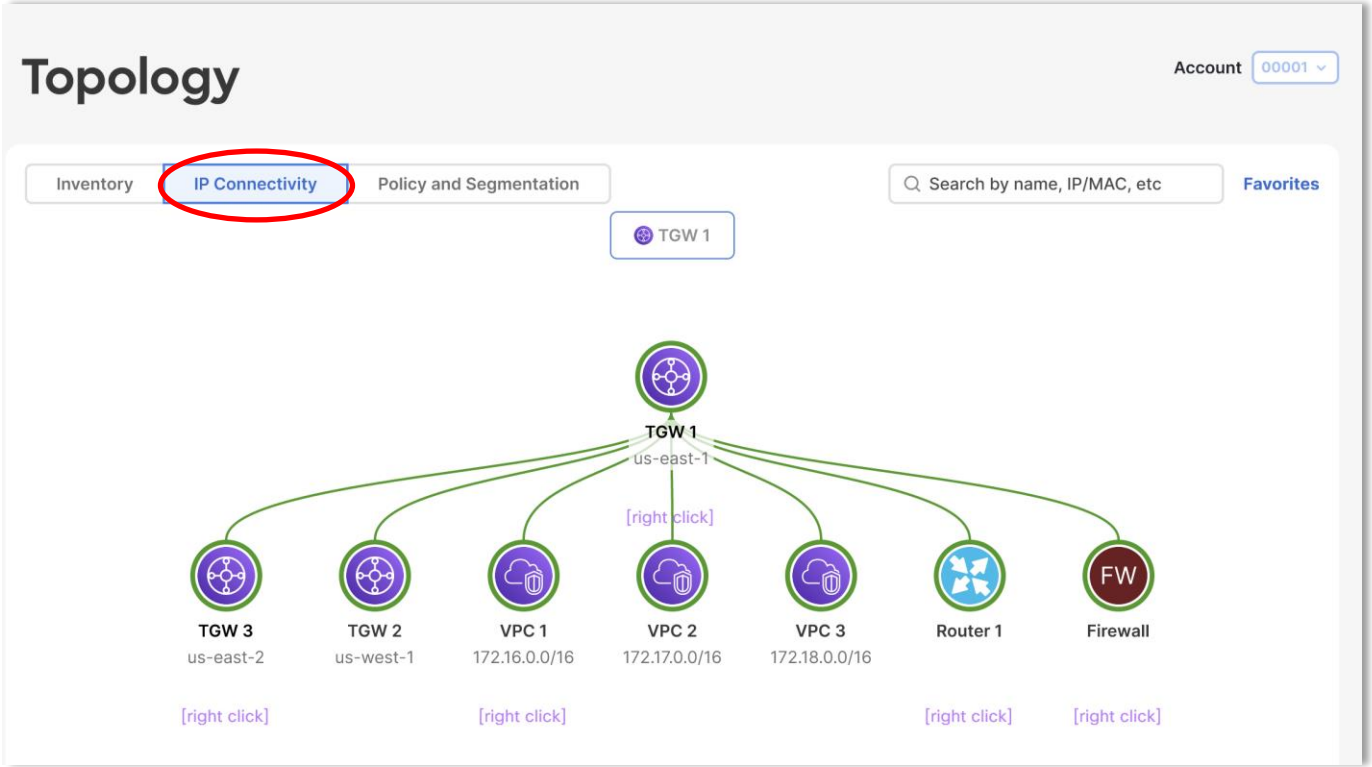

Cloud Network Controller: Topology

IP Connectivity

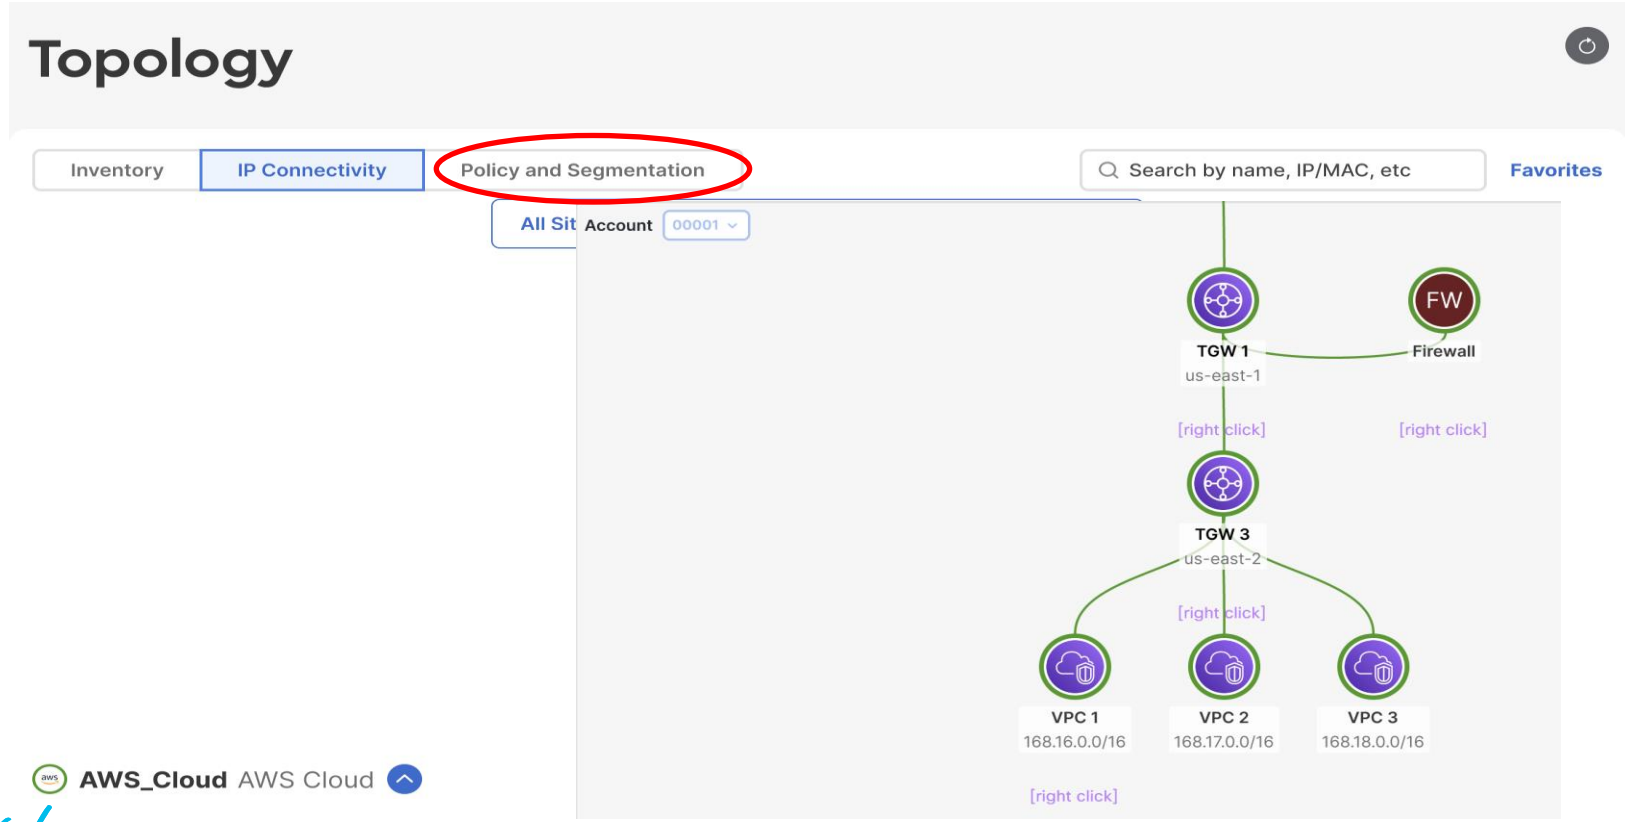

VNET / Route table detail

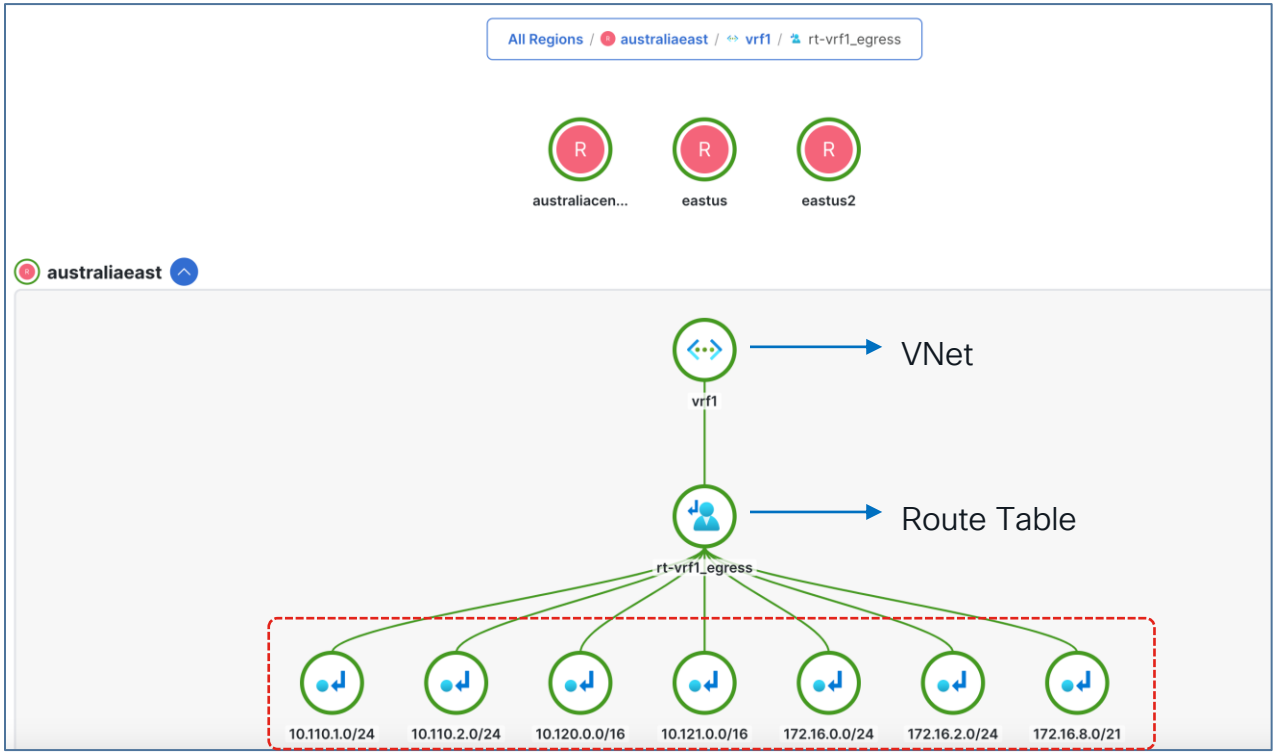

Route Table entries

Search

Nexus Dashboard Orchestrator

Overlay & Underlay Visibility: IPSec BGP IPv4 or BGP EVPN

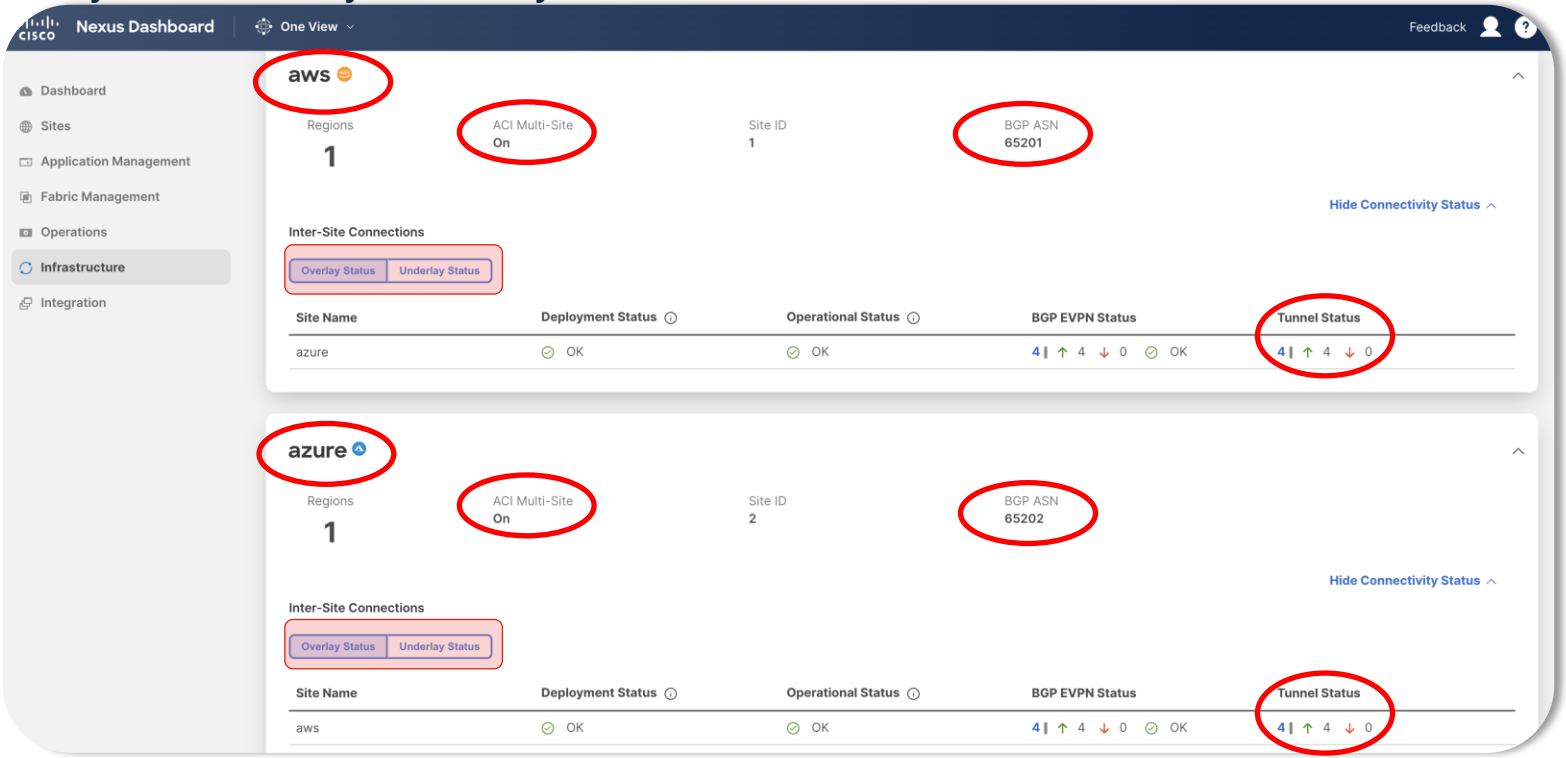

Visibility: AWS Inter Region

Thousand Eyes: Dallas, TX to Ashburn, VA

Visibility: AWS/Azure Inter Cloud

Thousand Eyes – Frankfurt, Germany AWS to Richmond, VA US Azure

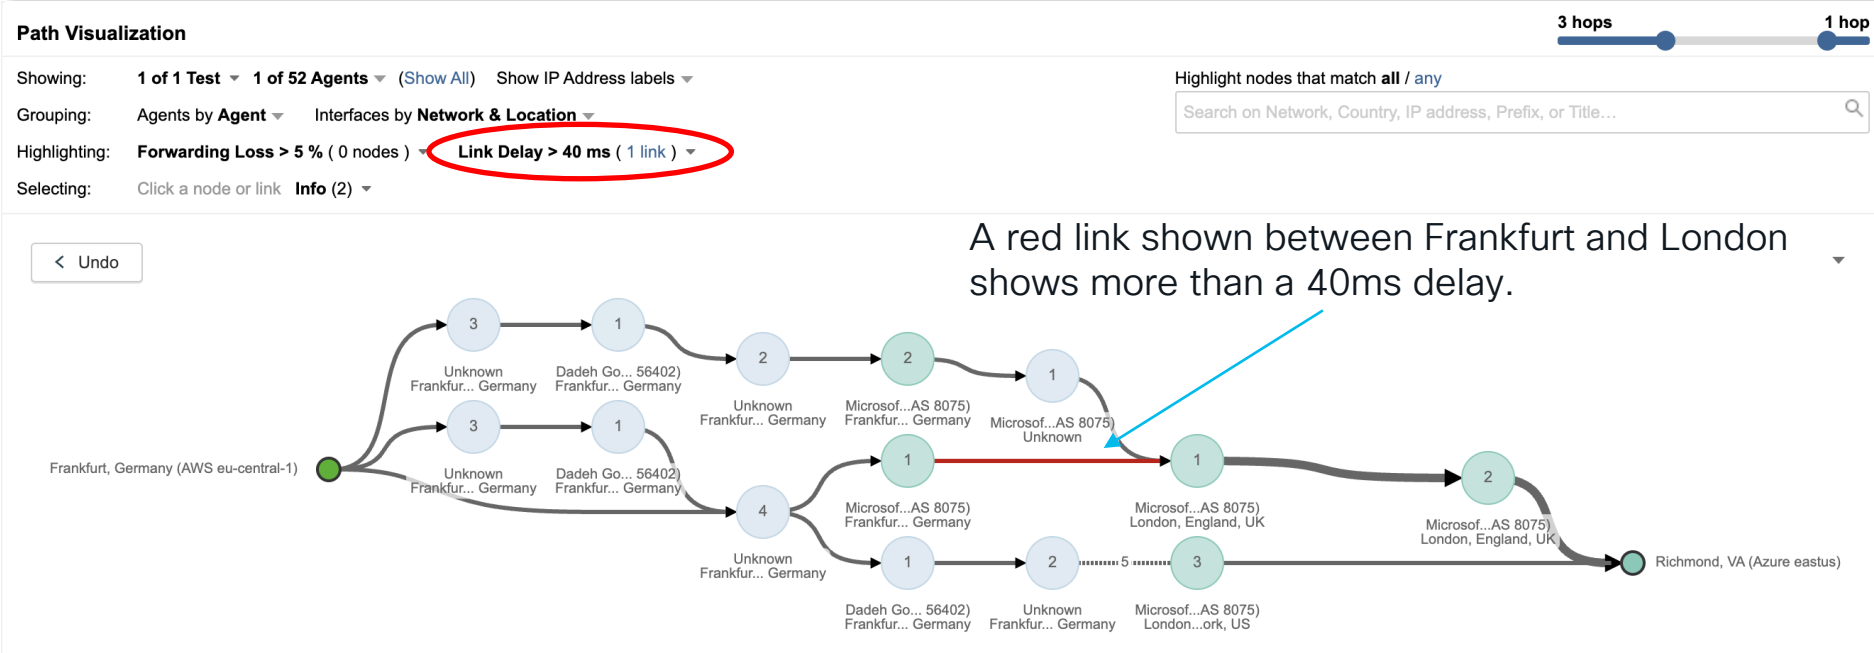

- Provider-to-provider path analysis w/ direct peering relationships.
- Traffic from AWS through an intermediate provider but still gets on the Azure backbone before leaving Frankfurt

Visibility: AWS/Azure Inter Cloud

Thousand Eyes – End user perspective from New Delhi to S3 Service

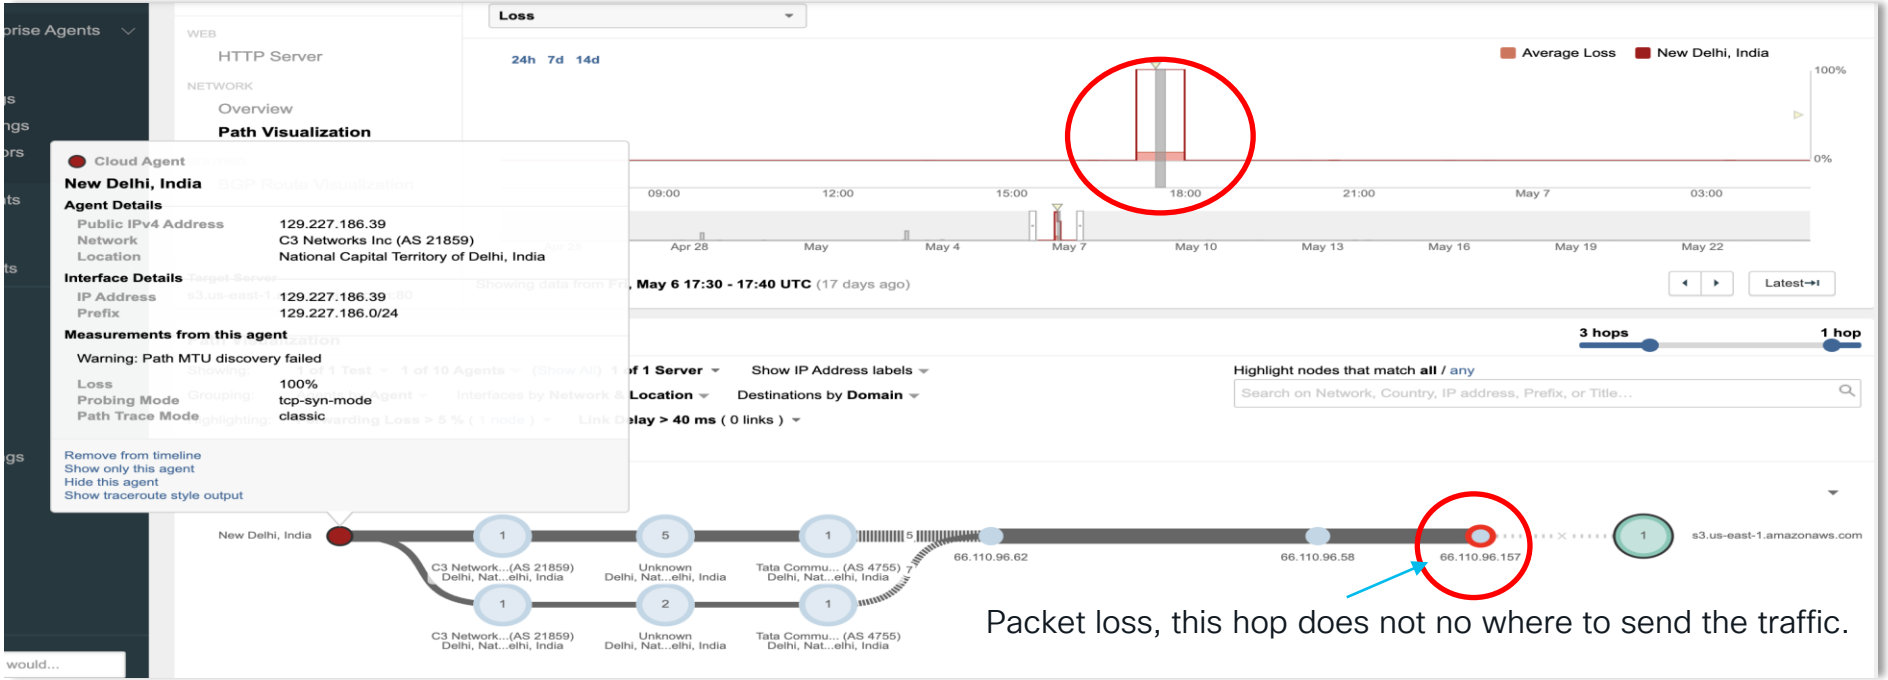

Depicted above is the loss of service, source of loss and time of outage.

Finally, what About Infrastructure as Code?

https://registry.terraform.io/namespaces/CiscoDevNet

https://galaxy.ansible.com/cisco

Summary

Cloud Network Controller

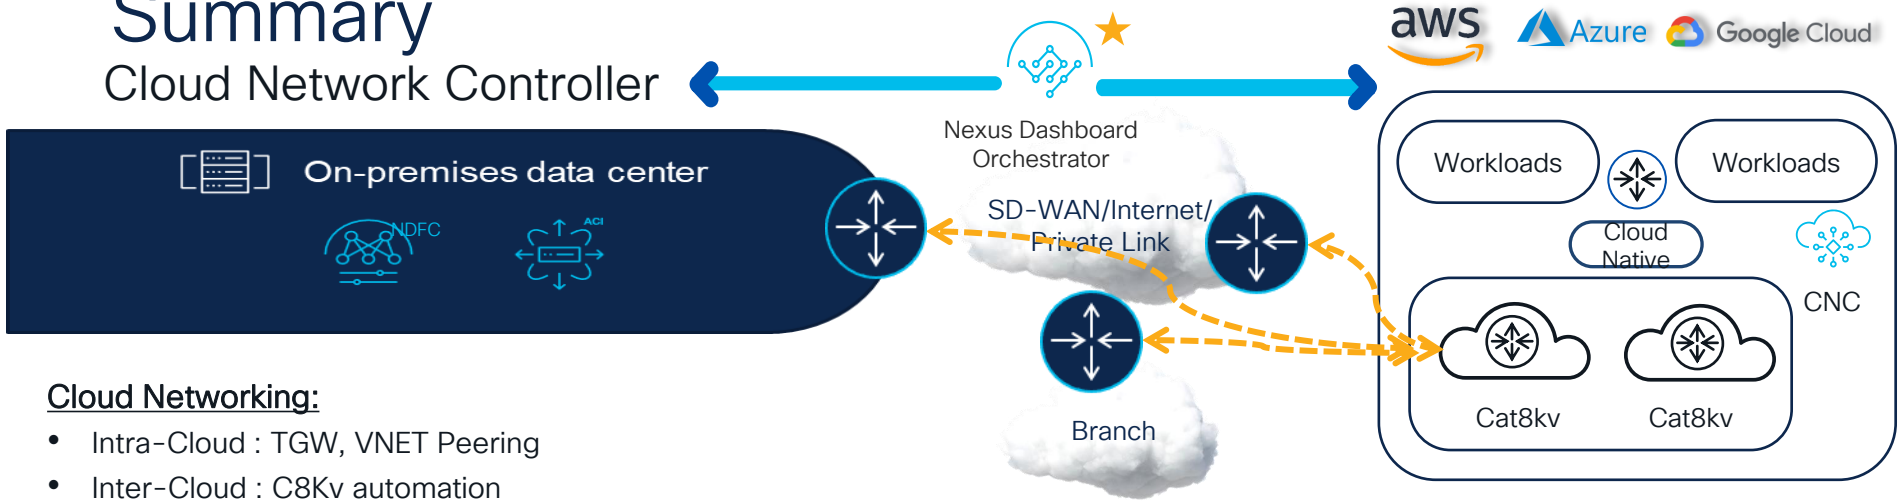

Cloud Networking:

- Intra-Cloud : TGW, VNET Peering
- Inter-Cloud : C8Kv automation
- Connectivity: IPSEC, Direct Connect, Express Rt
- End Point Discovery

Visibility:

- View and connect to brownfield VPC networks
- Inventory and topology view

L4-L7 Services:

- Automate service insertion and service chaining (Load balancers, Firewalls, …)

Segmentation:

- Extend segments from On-Premises to cloud
- Extend segments from cloud to cloud
- Security Group rule management

Public Cloud:

- AWS, Azure, Google Cloud

Open APIs:

- Enable automation using Terraform and Ansible

cisco Live!

Additional Information

- Equinix and Cisco are Enabling Cybersecurity on a Global Scale

  https://blog.equinix.com/blog/2022/06/06/equinix-and-cisco-are-enabling-cybersecurity-on-a-global-scale/

- Cisco Hybrid Multi-Cloud Networking Design Guide

  https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-cloud-aci-hybrid-multicloud-design-guide.html

- Network Service Mesh: Linking multi-cloud workloads

  https://www.networkworld.com/article/3662750/network-service-mesh-linking-multicloud-workloads.html

Please fill out the survey

Drop your email in the comments – I WILL respond!

The bridge to possible

Thank you

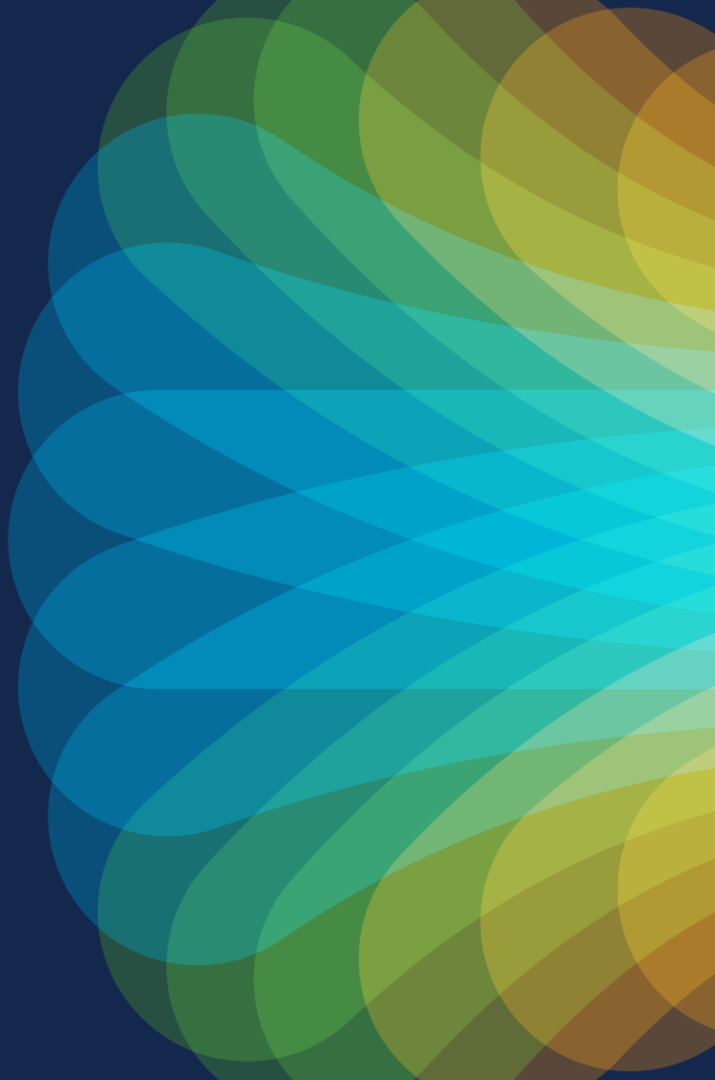

cisco Live!


cisco Live!
Let's go