Cisco ACI Anywhere

Xiaozhen Lu
Technical Solution Specialist

Agenda

- ACI Foundations
  - What is ACI
  - ACI Fundamentals
- ACI Anywhere
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Host Based Networking (Forwarding)

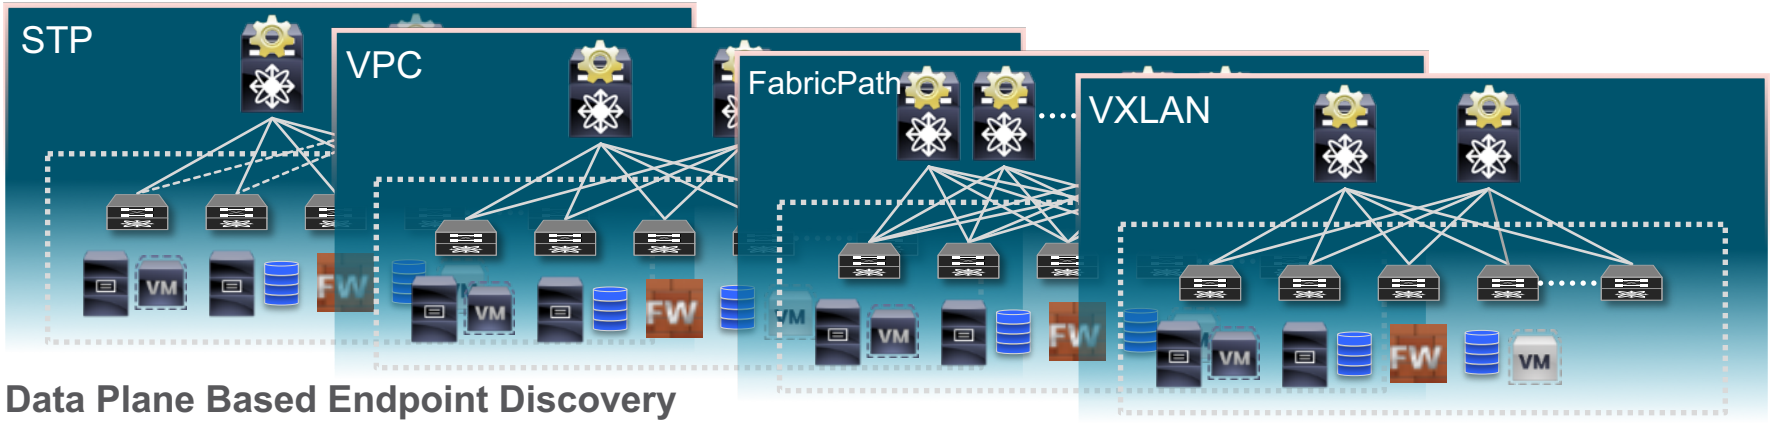

Data Plane Based Endpoint Discovery

Control Plane Based Endpoint Location Tracking

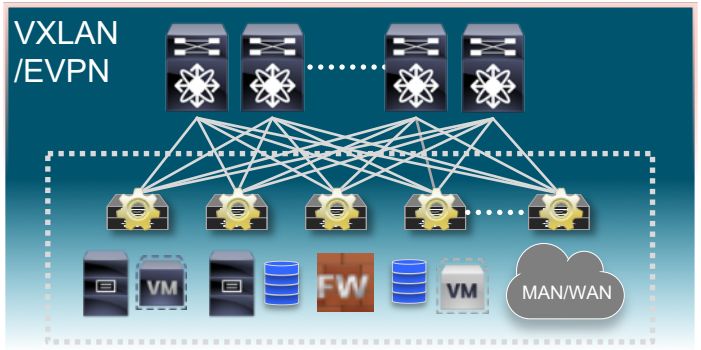

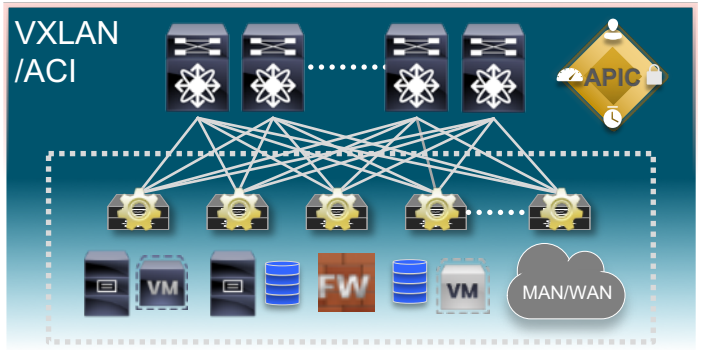

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Identity Based Networking

VTEP IP	Group Policy	VNID	Tenant Packet

Devices and users are authenticated and authorized into end-point groups (aka EPG's or SGT's)

End Point Group Tags (EPG's, SGT's) are encoded in a VXLAN header

Policies between scalable groups are established following the provider/consumer model

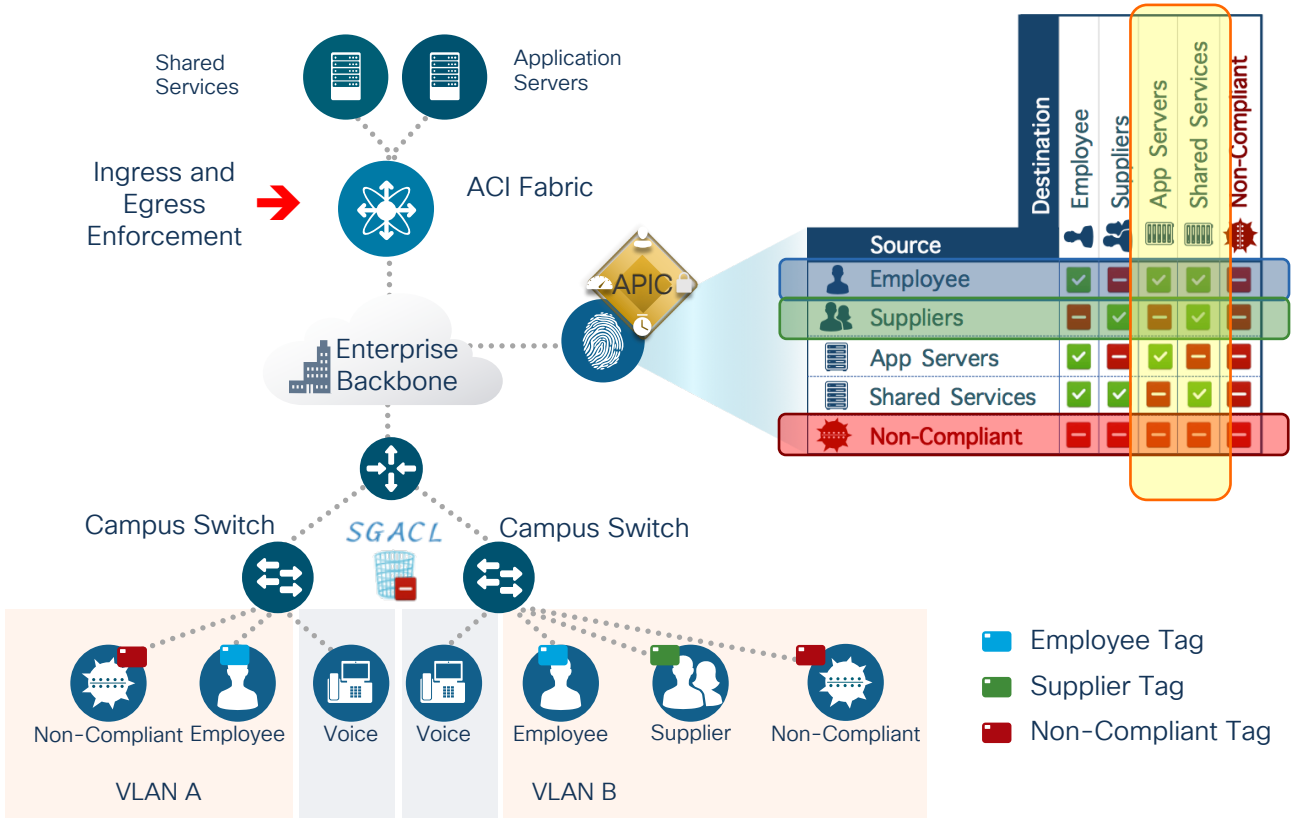

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- ACI Foundations
  - What is ACI
  - ACI Fundamentals
- ACI Anywhere
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Fabric – An IP network with an Integrated Overlay

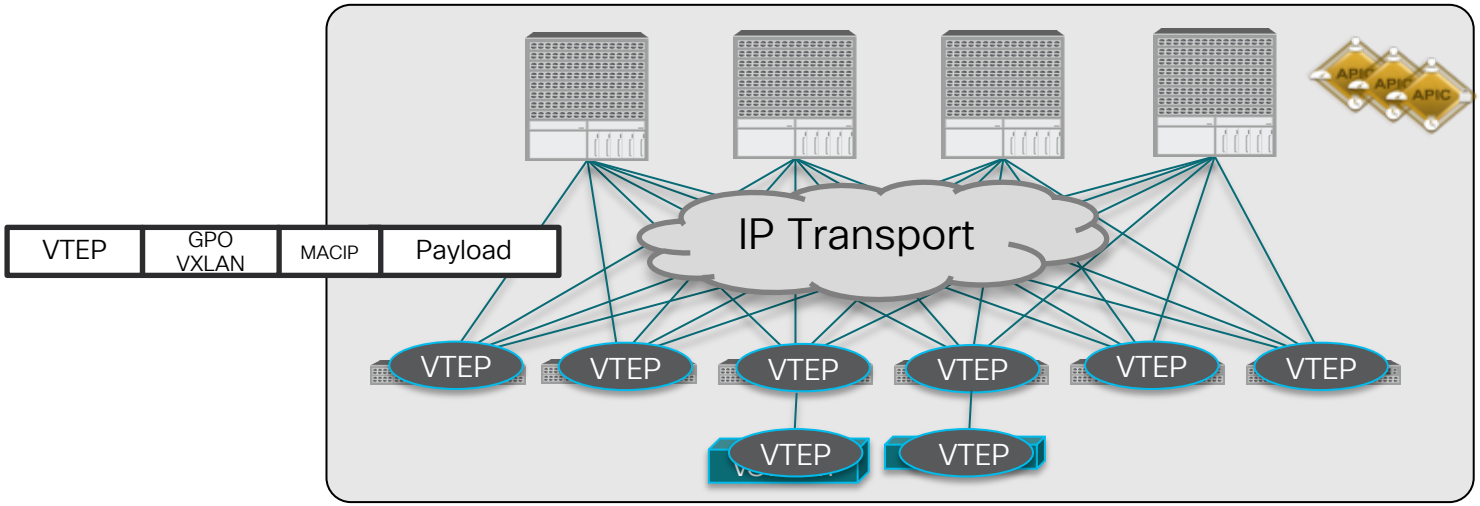

- Cisco ACI leverages an integrated VXLAN based overlay
  - IP Network for Transport
  - VXLAN based tunnel end points (VTEP)
  - VTEP discovery via infrastructure routing
  - Directory (Mapping) service for EID (host MAC and IP address) to VTEP lookup

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Removing the Classic L2/L3 Boundaries
Layer 2 and Layer 3 integrated forwarding

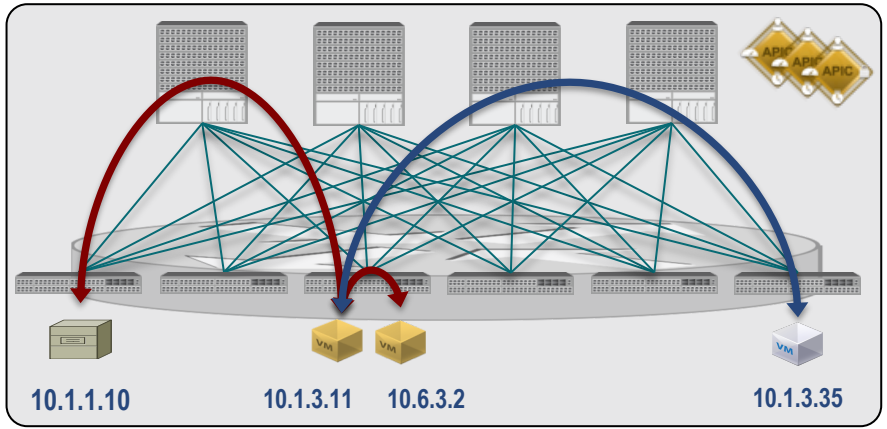

Distributed Default Gateway

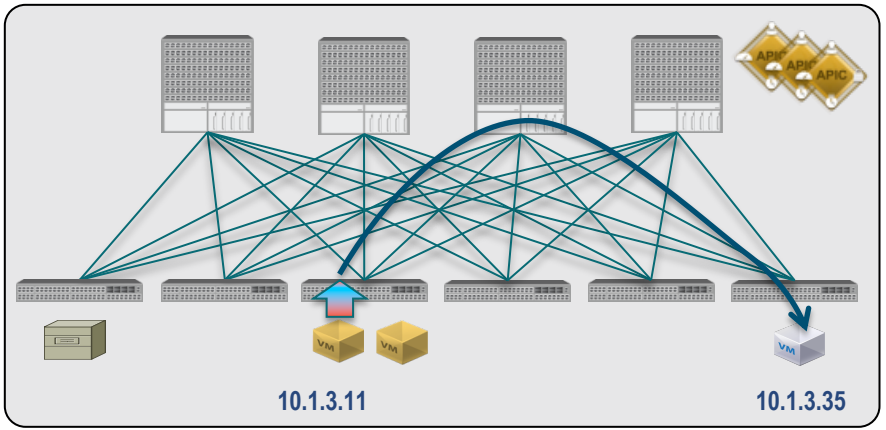

Directed ARP Forwarding and/or Proxy-ARP based default GW

- ACI Fabric supports full layer 2 and layer 3 forwarding semantics, no changes required to applications or end point IP stacks
- ACI Fabric provides optimal forwarding for layer 2 and layer 3
  - Fabric provides a pervasive SVI which allows for a distributed default gateway function

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

What is an EPG?
A Logical Group of Endpoints Attached to the Network

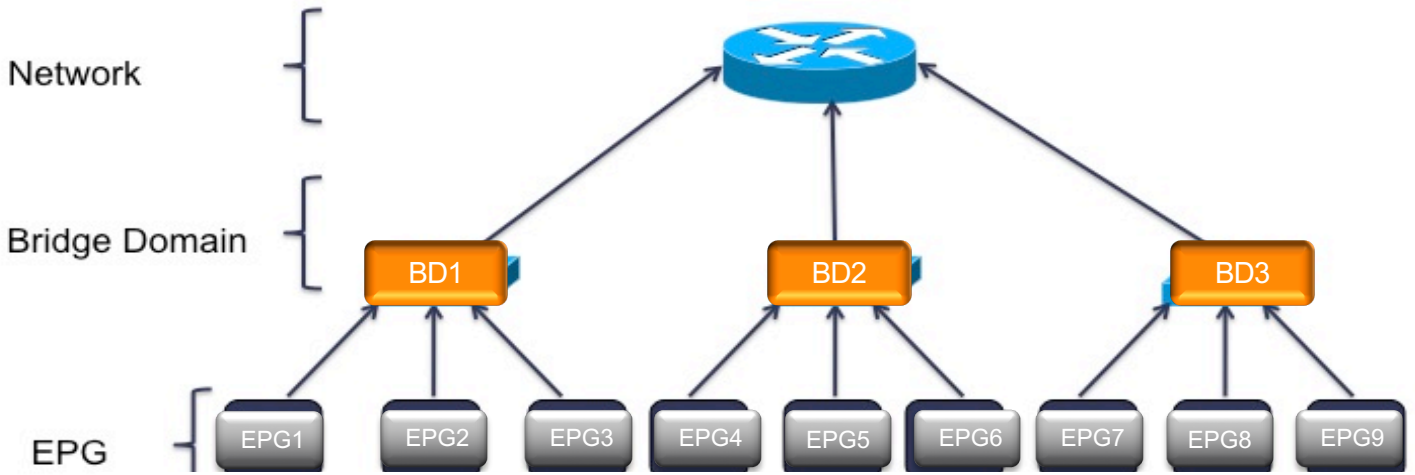

- All of the endpoints (things attached to the network) in the same EPG are treated to the same rules (policy)
  - A security group using the same access lists (similar to an SGT in TrustSec)
  - A services group using the same QoS rules, same L4-7 services, …
  - It could be as simple as all the servers on the same VLAN or subnet

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

What are Contracts and Filters?
The Network Rules

- Contracts are semantics to specify EPG to EPG communication in ACI
- Communication policy includes filters (ACLs), QoS, Route Leaking, L4-7 Service Graphs
- Contract filters are similar to Access Control Lists (e.g. match this TCP port)
- Contracts can be defined between EPGs or between L3Out External EPGs and regular EPGs

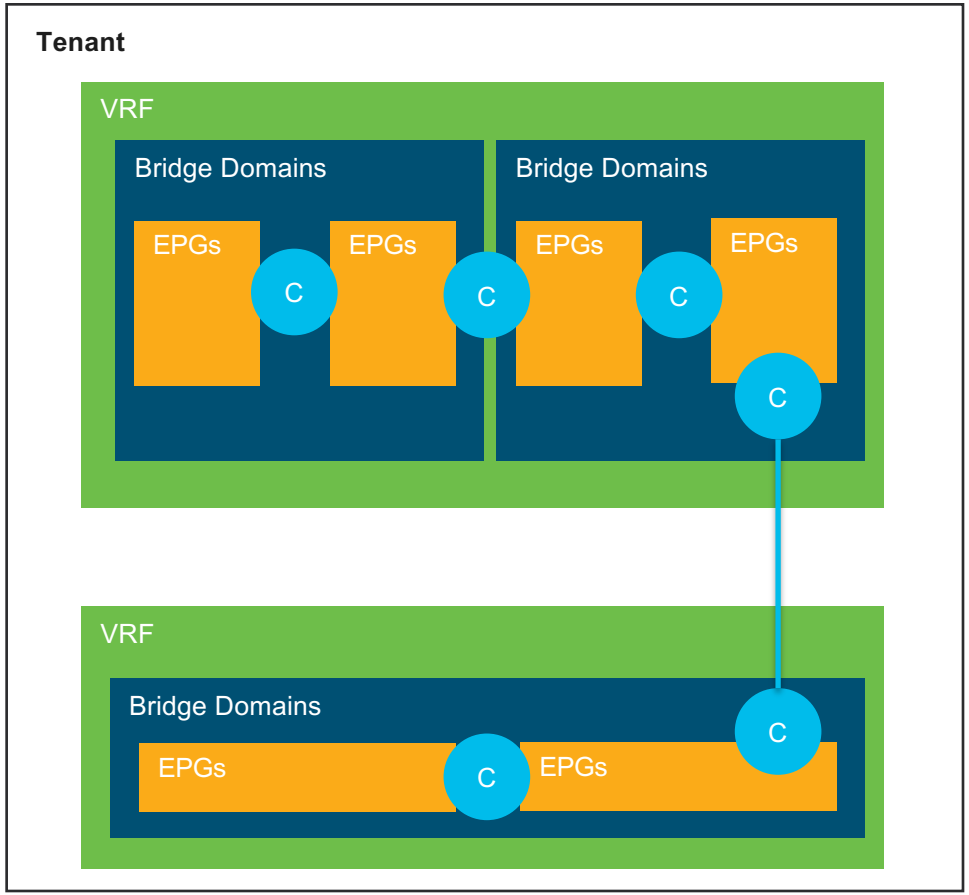

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

What is an Application Profile?

A Logical Group of EPG's and Associated Contracts

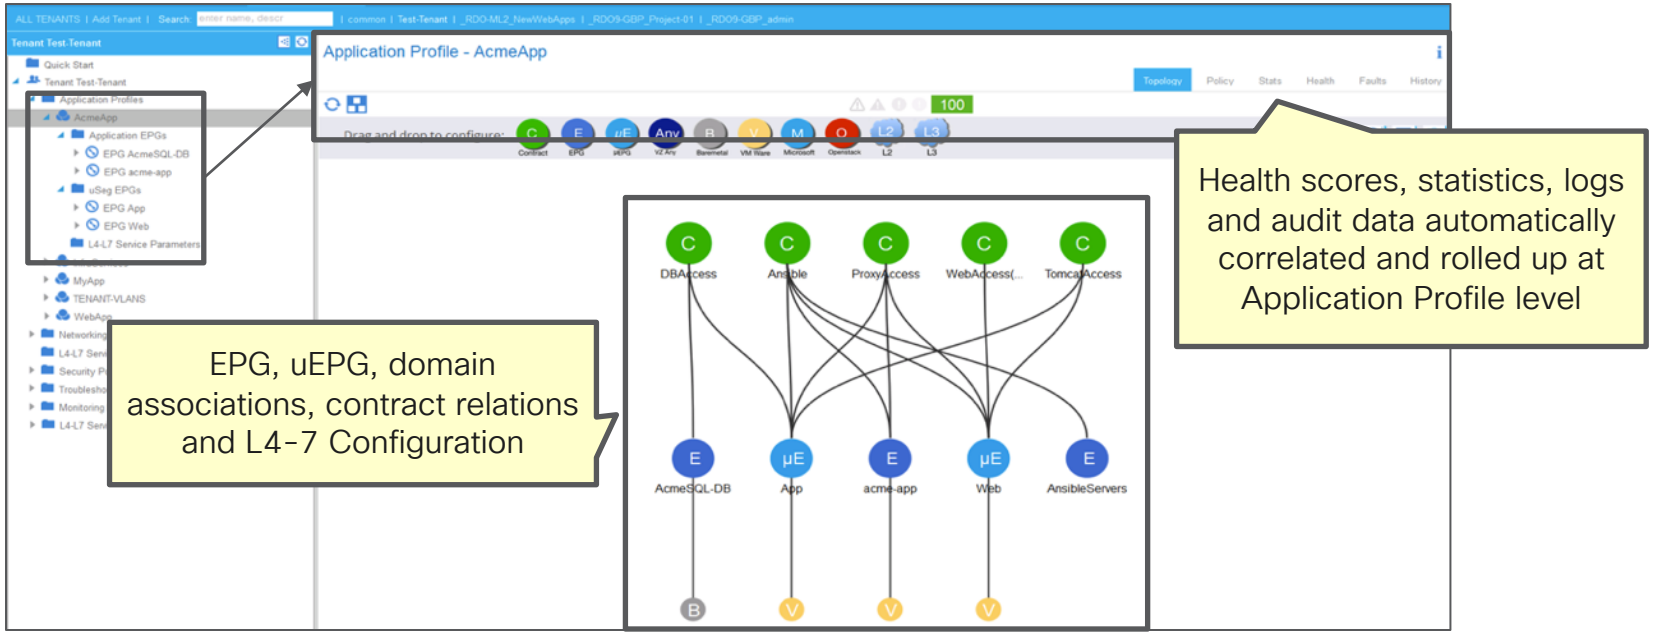

Developer: My application template NOC: What I monitor and troubleshoot

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

What is a Tenant?
A Virtual Private Cloud

Outside View: A Tenant is a group that owns a virtual private cloud instance

Inside View: The portion of the database that your login has access to

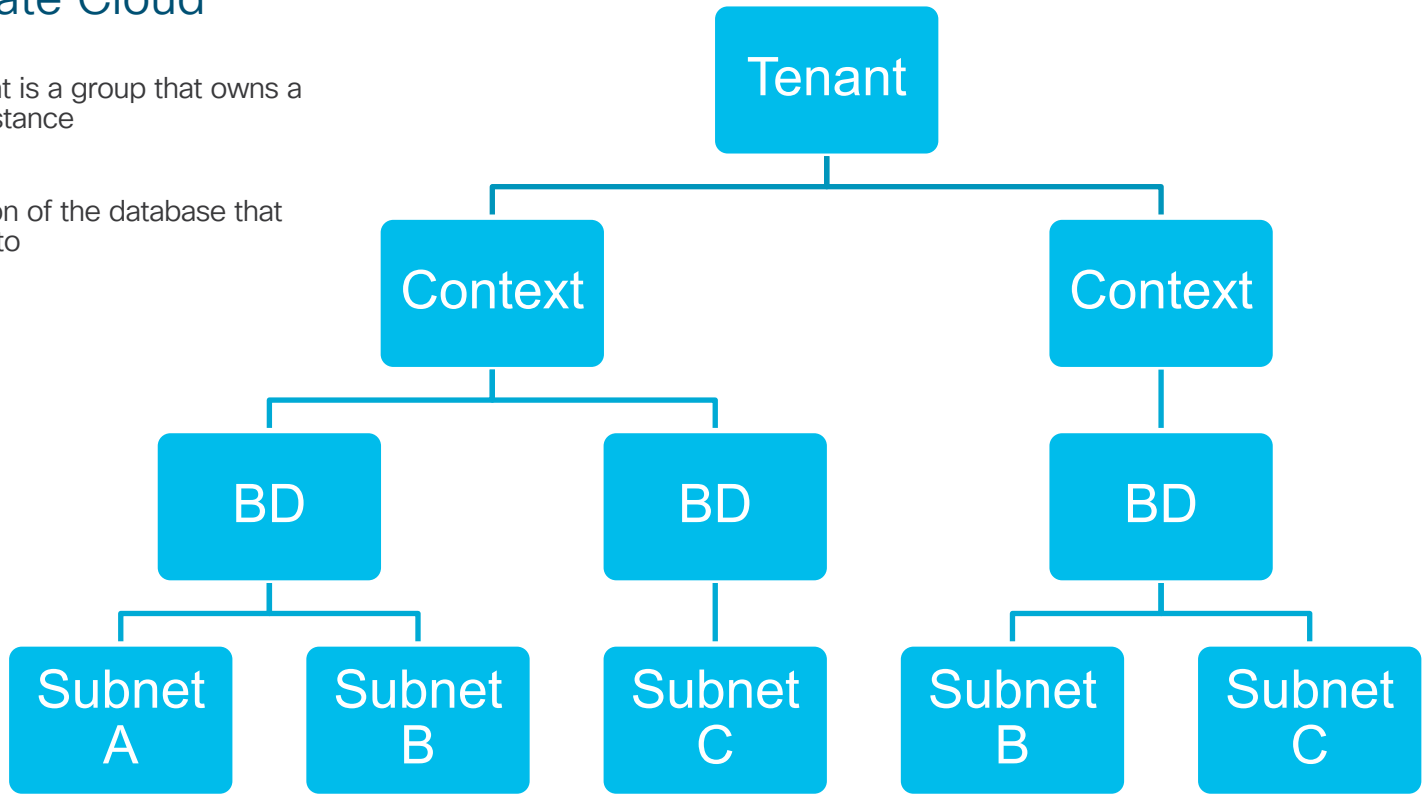

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

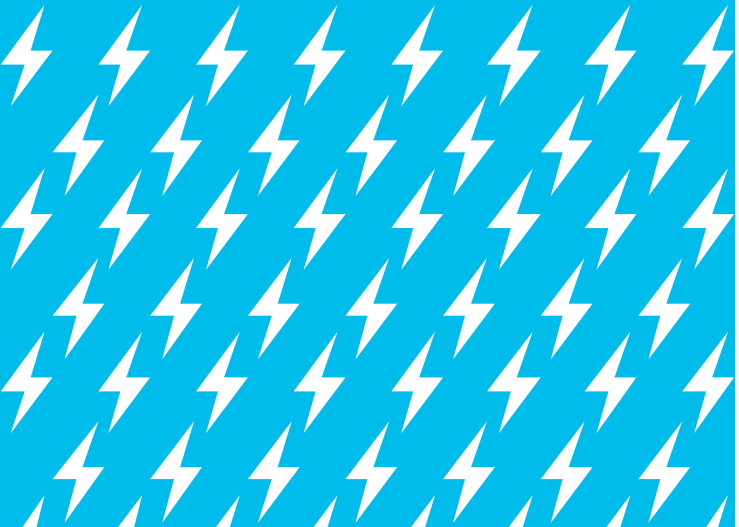

- ACI Foundations
- ACI Anywhere
  - ACI Multi-Pod
  - ACI Multi-Site
  - Physical Remote Leaf
  - Virtual Remote Leaf (vPod)
  - ACI with Public Cloud
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Application Centric Infrastructure
Fabric and Policy Domain Evolution

ACI Single Pod Fabric

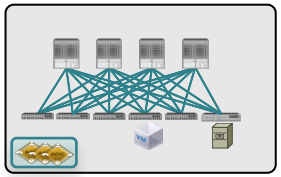

ACI 1.0 - Leaf/Spine Single Pod Fabric

ACI Multi-Pod Fabric

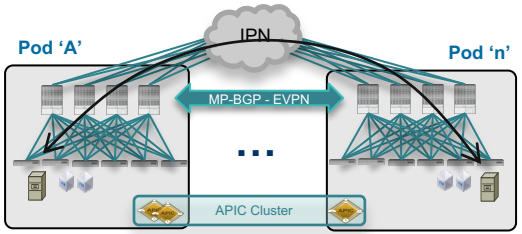

ACI 2.0 - Multiple Networks (Pods) in a single Availability Zone (Fabric)

ACI Multi-Site

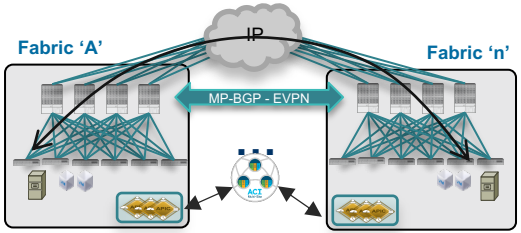

ACI 3.0 – Multiple Availability Zones (Fabrics) in a Single Region 'and' Multi-Region Policy Management

ACI Remote Leaf

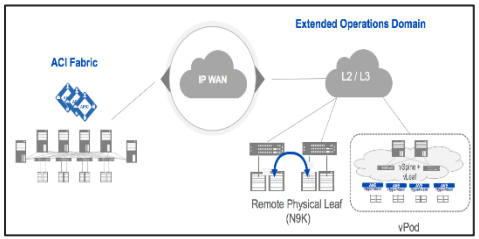

ACI 3.1/3.2 - Remote Leaf and vPod extends an Availability Zone (Fabric) to remote locations

ACI Multi-Cloud

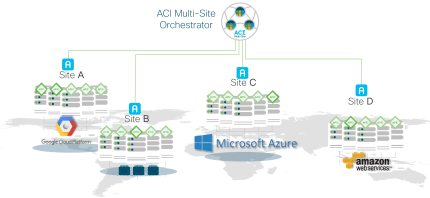

Future – ACI Extensions to Multi-Cloud

ACI Multi-Pod

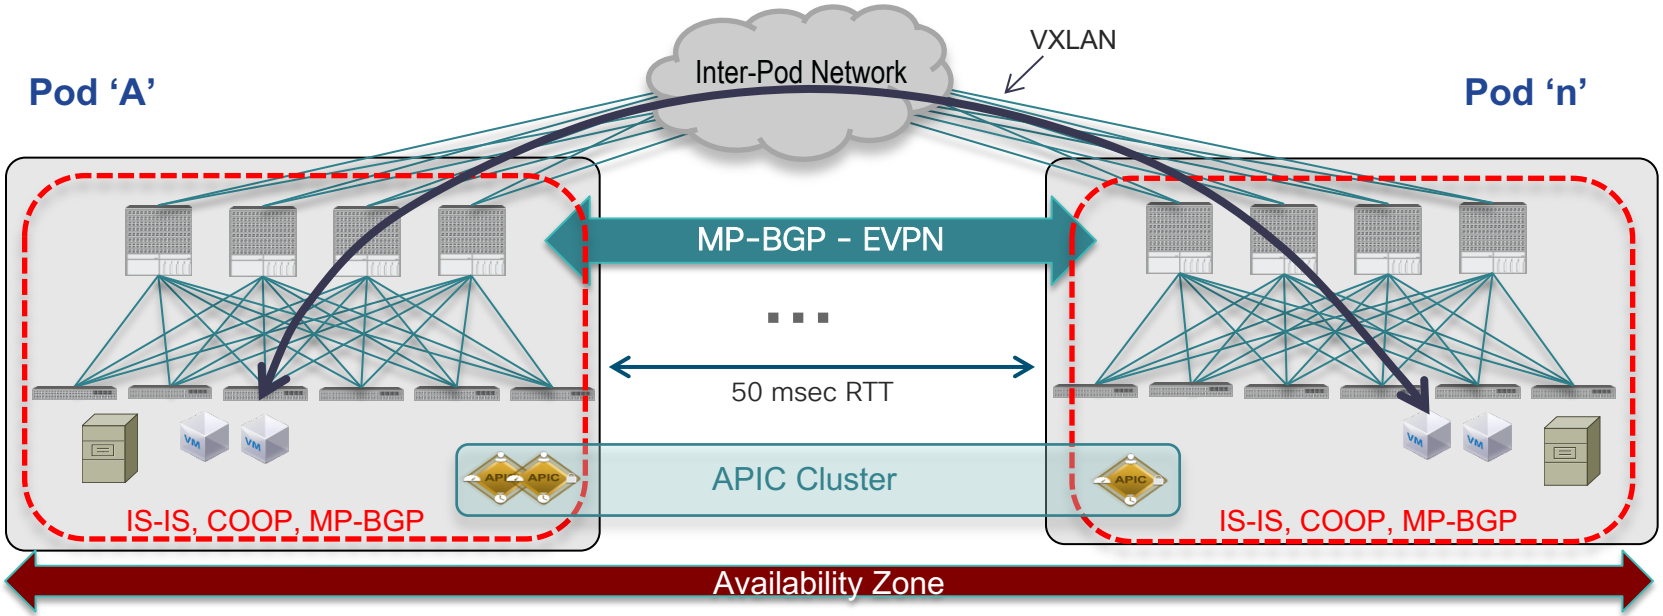

- Multiple ACI Pods connected by an IP Inter-Pod L3 network, each Pod consists of leaf and spine nodes
- Managed by a single APIC Cluster
- Single Management and Policy Domain
- Forwarding control plane (IS-IS, COOP) fault isolation
- Data Plane VXLAN encapsulation between Pods
- End-to-end policy enforcement

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Pod
Most Common Use Cases

- Need to scale up a single ACI fabric above 200 leaf nodes supported in a single Pod
- Handling 3-tiers physical cabling layout (for example traditional N7K/N5K/N2K deployments)

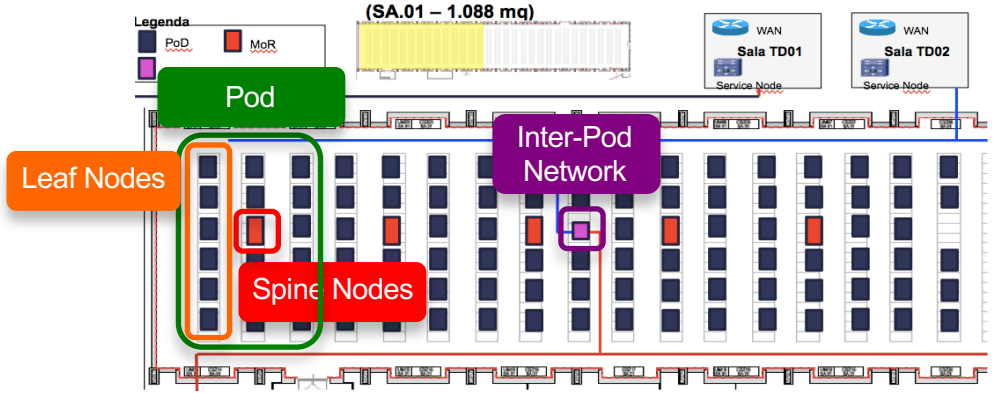

- True Active/Active DC deployments
  - Single VMM domain across DCs (ESXi Metro Cluster, vSphere HA/FT, DRS,…)
  - Deployment of clustered network services (FWs, SLBs)
  - Application clustering (L2 BUM extension across Pods)

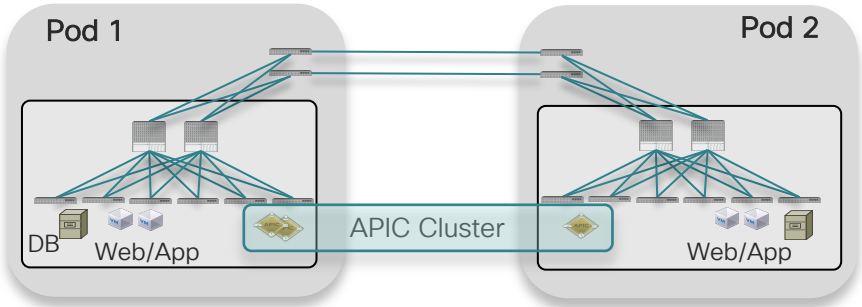

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Single Availability Zone with Maintenance & Configuration Zones

Scoping ‘Network Device’ Changes

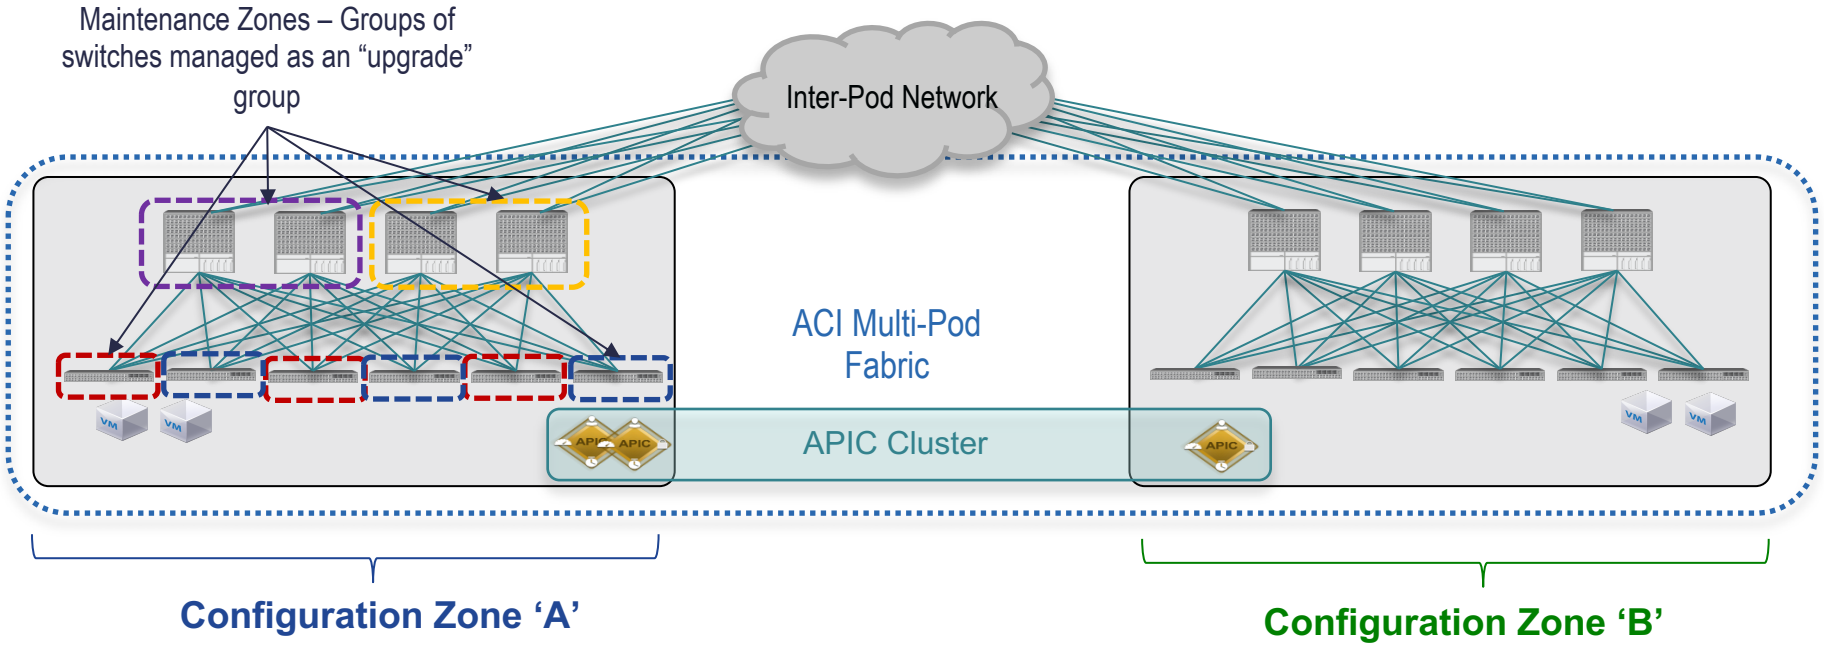

- Configuration Zones can span any required set of switches, simplest approach may be to map a configuration zone to an availability zone, applies to infrastructure configuration and policy only

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Pod

Inter-Pod Network (IPN) Requirements

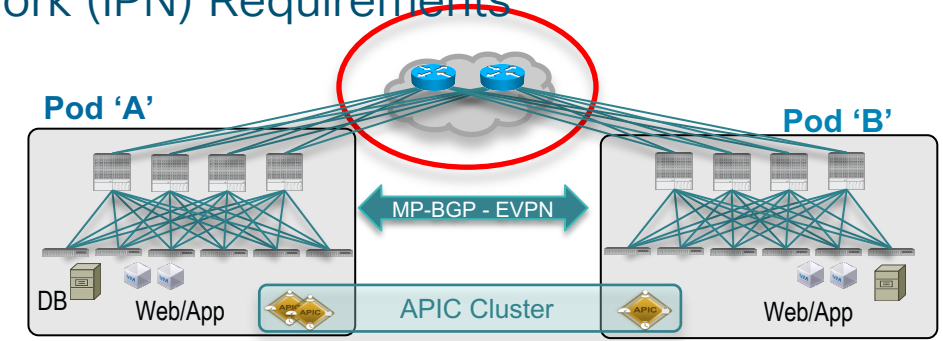

- Not managed by APIC, must be separately configured (day-0 configuration)
- IPN topology can be arbitrary, not mandatory to connect to all spine nodes
- Main requirements:
  - ✓ Multicast BiDir PIM → needed to handle Layer 2 BUM* traffic
  - ✓ OSPF to peer with the spine nodes and learn VTEP reachability
  - ✓ Increase MTU support to handle VXLAN encapsulated traffic
  - ✓ DHCP-Relay

* Broadcast, Unknown unicast, Multicast

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Pod

Supported Topologies

Intra-DC

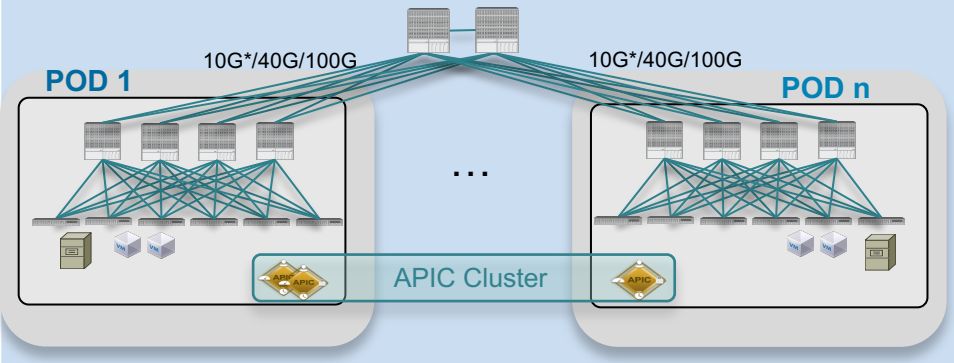

Two DC sites directly connected

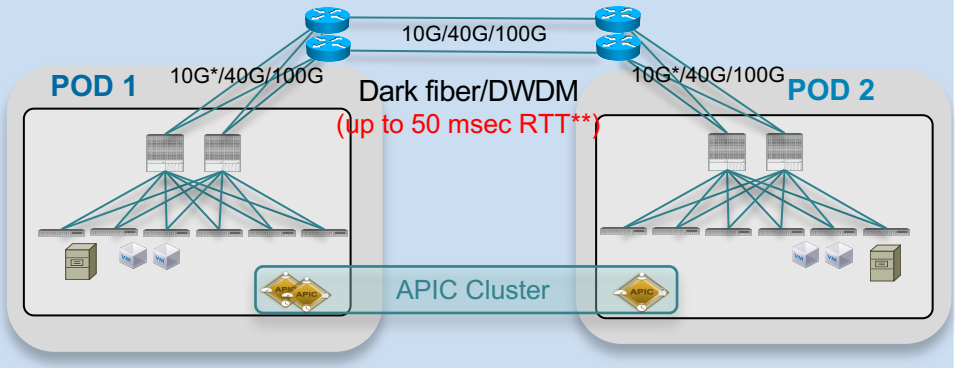

3 (or more) DC Sites directly connected

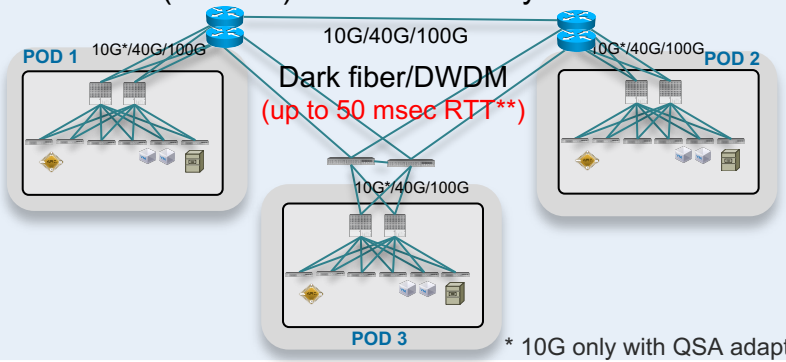

Multiple sites interconnected by a generic L3 network

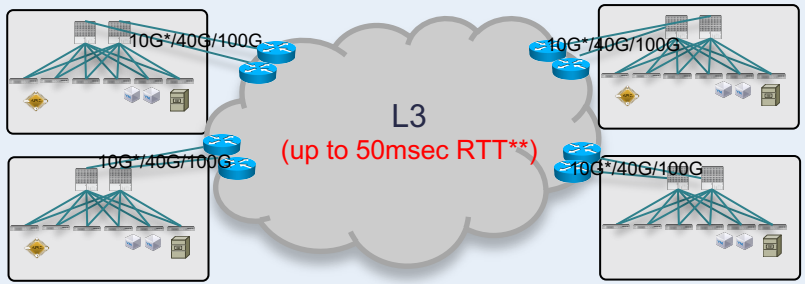

* 10G only with QSA adapters on EX/FX spines

** 50 msec support added in SW release 2.3(1)

Agenda

- ACI Foundations
- ACI Anywhere
  - ACI Multi-Pod
  - ACI Multi-Site
  - Physical Remote Leaf
  - Virtual Remote Leaf (vPod)
  - ACI with Public Cloud
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site
Overview

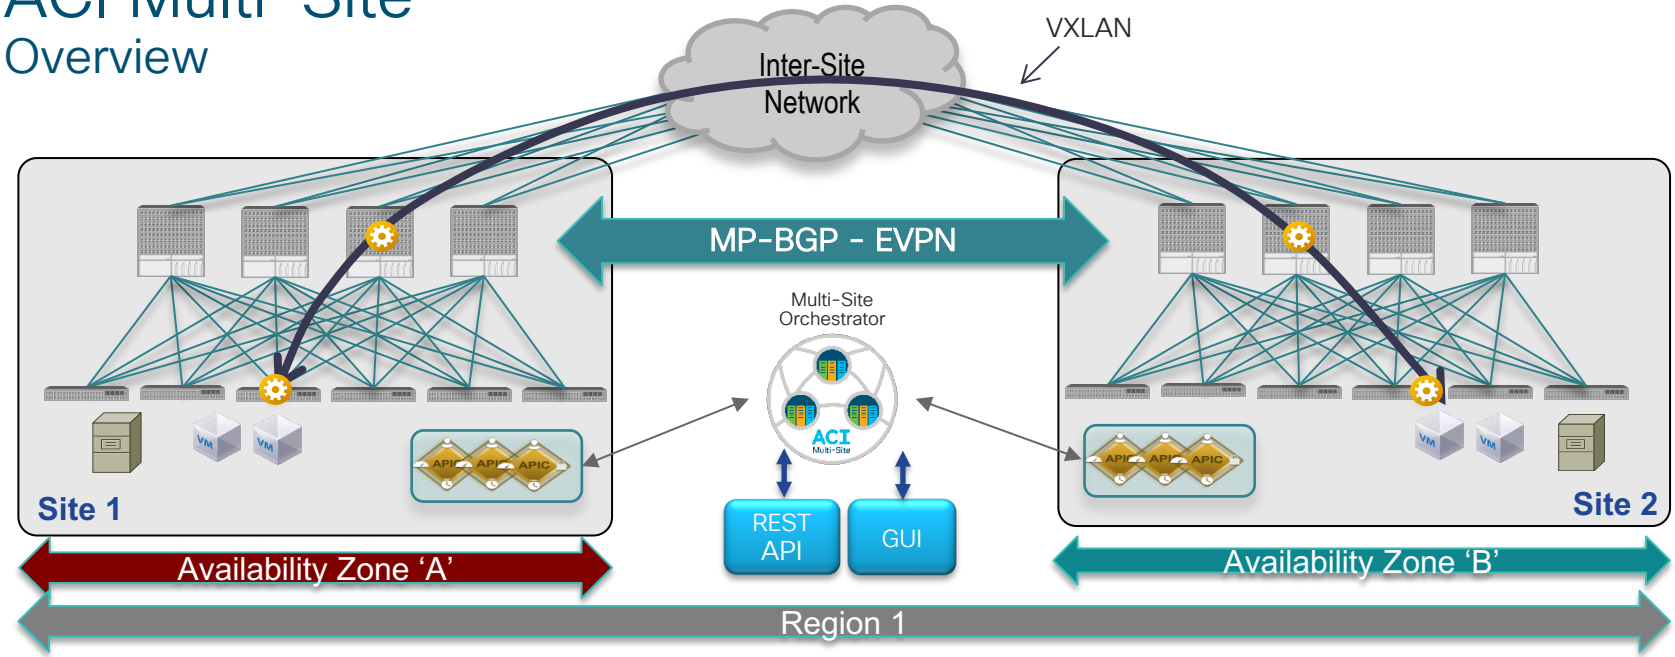

- Separate ACI Fabrics with independent APIC clusters
- No latency limitation between Fabrics
- ACI Multi-Site Orchestrator pushes cross-fabric configuration to multiple APIC clusters providing scoping of all configuration changes
- MP-BGP EVPN control plane between sites
- Data Plane VXLAN encapsulation across sites
- End-to-end policy definition and enforcement

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site
Most Common Use Cases

- Scale-up model to build a very large intra-DC network (above 400 leaf nodes)

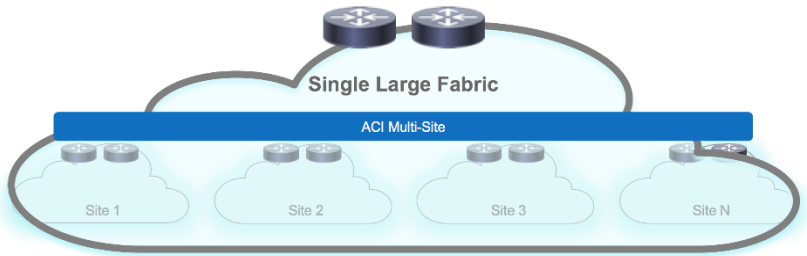

- Data Center Interconnect (DCI)
  - Extend connectivity and policy between 'loosely coupled' DC sites
  - Disaster Recovery and IP mobility use cases

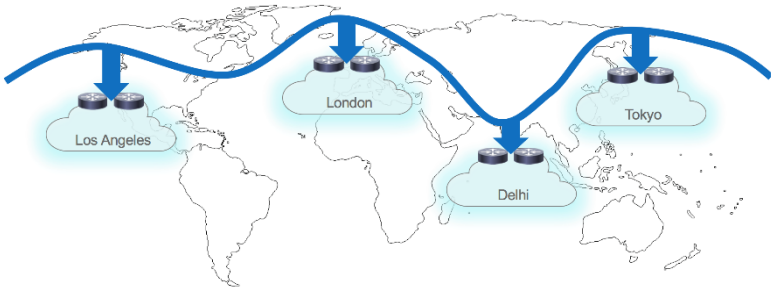

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site
Software and Hardware Requirements

- ACI Multi-Site introduced from release 3.0(1)
- Support all ACI leaf switches (1st Generation, -EX and -FX)
- Only -EX spine (or newer) to connect to the ISN
- New 9364C/9332C non modular spine (64/32 40G/100G ports) also supported
- 1st generation spines (including 9336PQ) not supported
- Can still leverage those for intra-site leaf to leaf communication

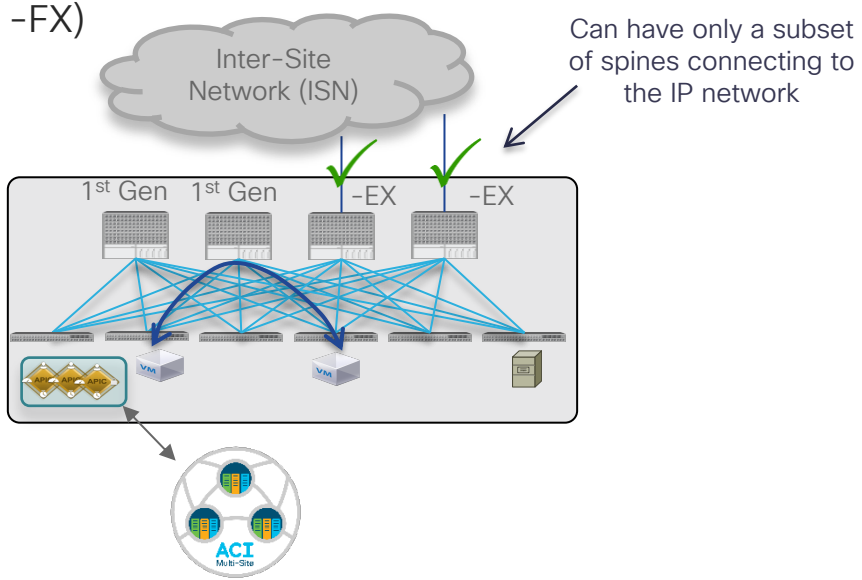

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site

Network and Identity Extended between Fabrics

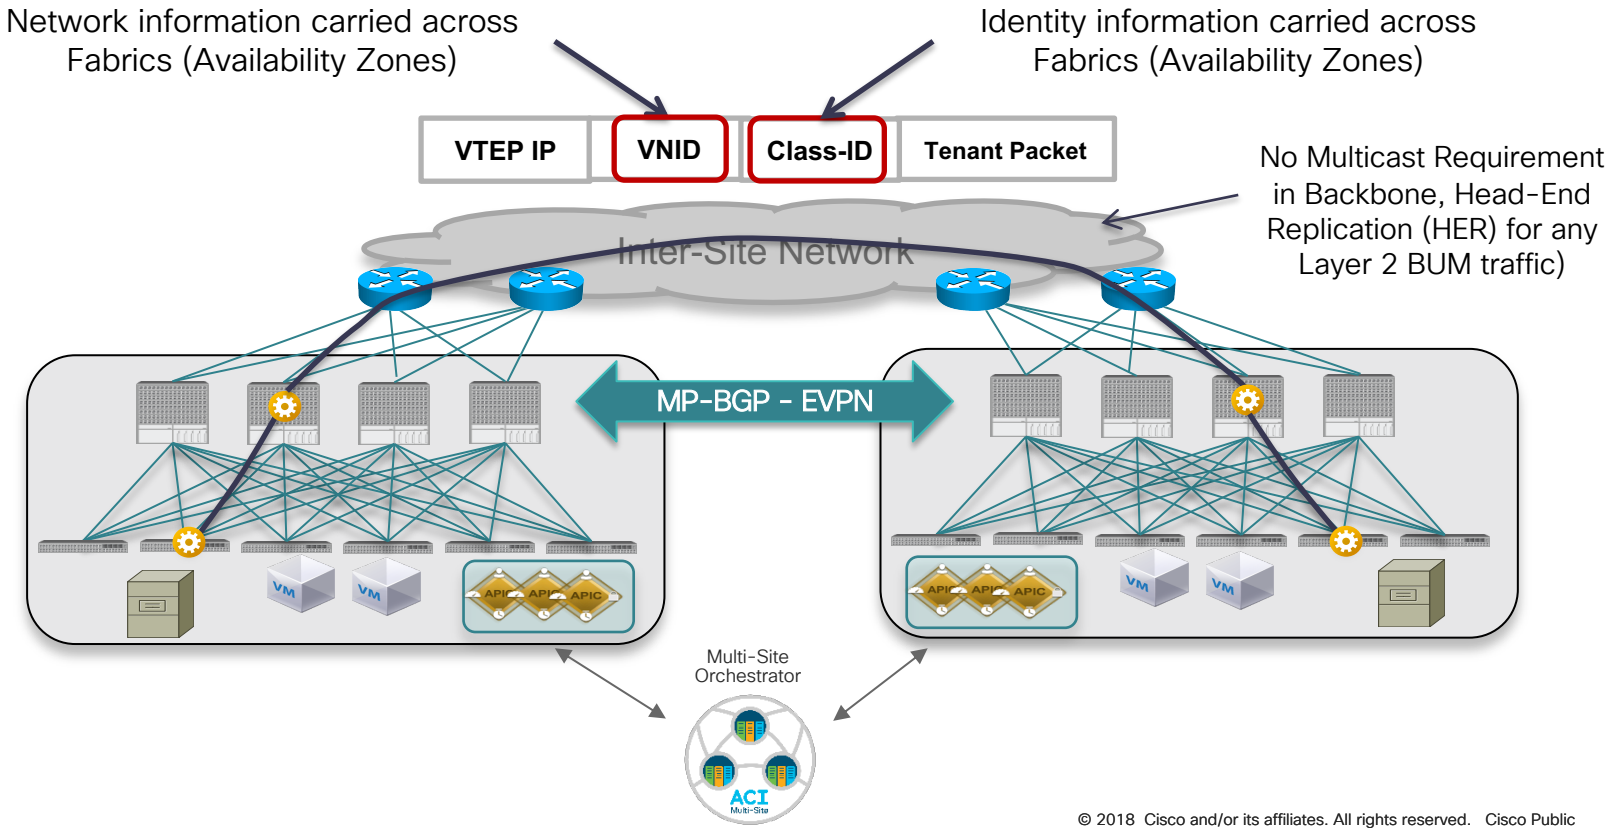

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site
Inter-Site Policies and 'Shadow' EPGs

- Inter-Site policies defined on the ACI Multi-Site Orchestrator are pushed to the respective APIC domains
  - End-to-end policy consistency
  - Creation of 'shadow' EPGs to locally represent the policies
- Policies enforced at the ingress leaf node at steady state

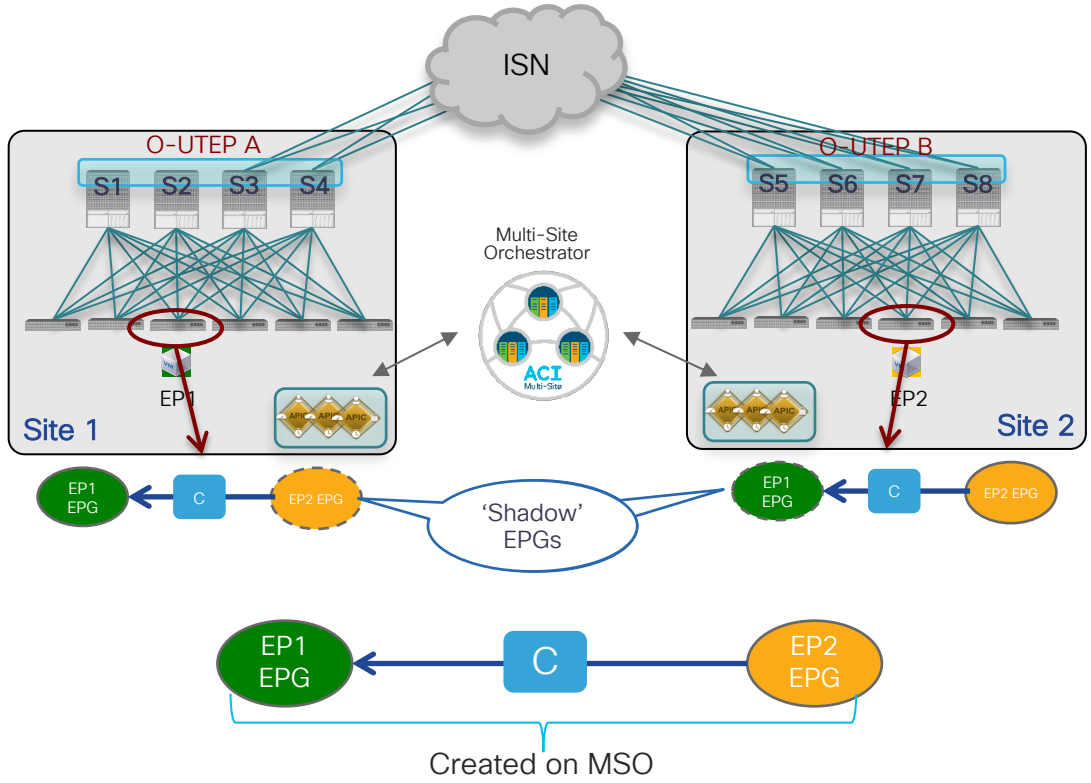

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site
Namespace Normalization

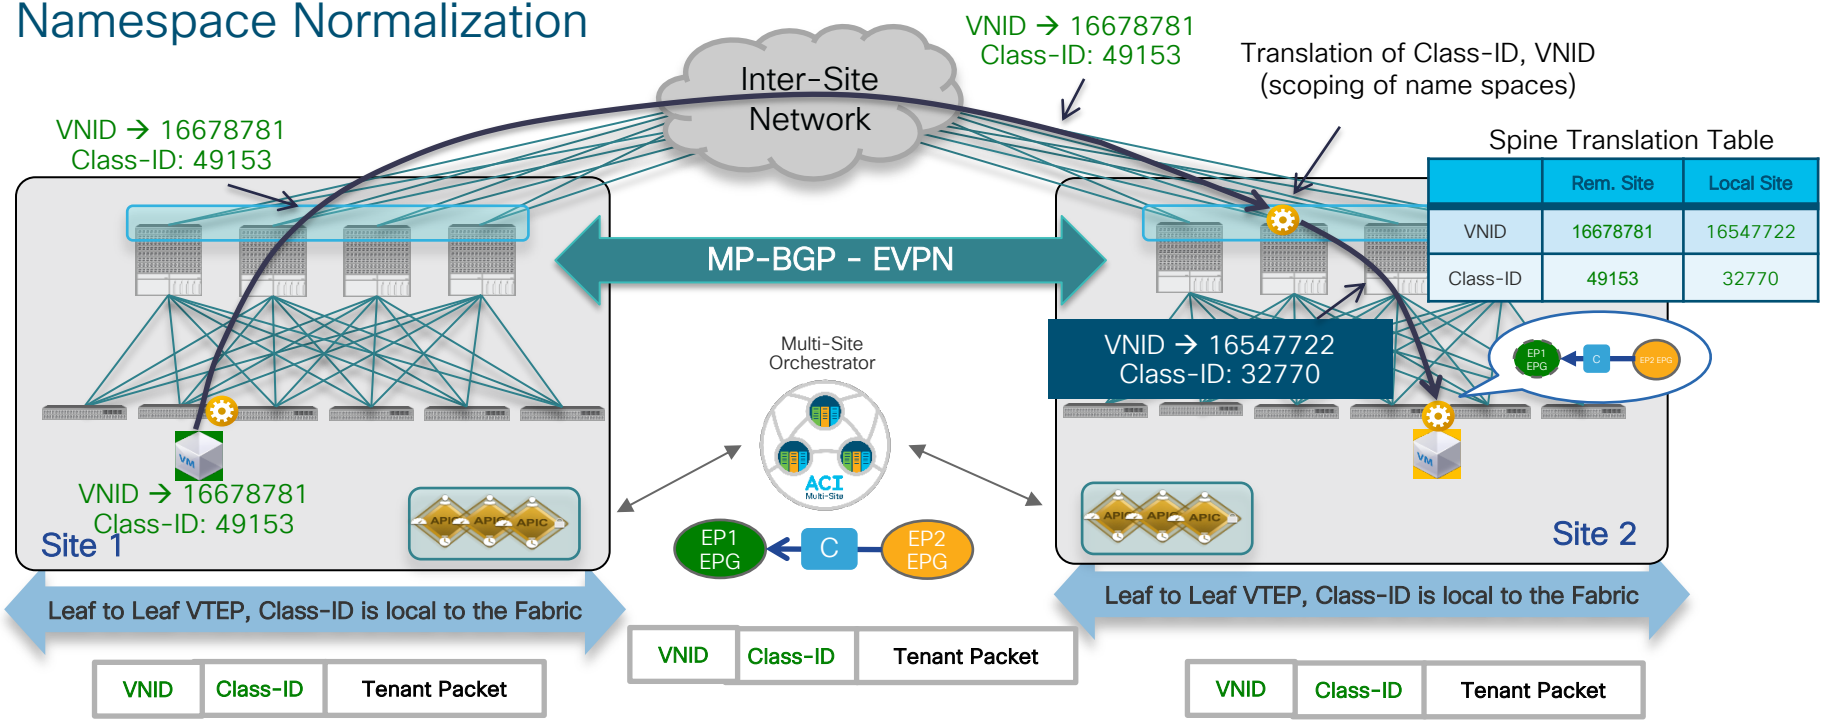

- Maintain separate name spaces with ID translation performed on the spine nodes
- Requires specific HW on the spine to support for this functionality
- Multi-Site Orchestrator instructs local APIC to program translation tables on spines

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Multi-Site Networking Options
Per Bridge Domain Behavior

Layer 3 only across sites

1

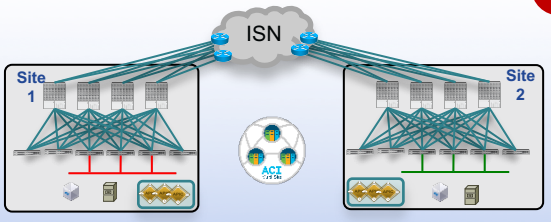

- Bridge Domains and subnets not extended across Sites
- Layer 3 Intra-VRF or Inter-VRF communication (shared services across VRFs/Tenants)

MSO GUI (BD)

IP Mobility without BUM flooding

2

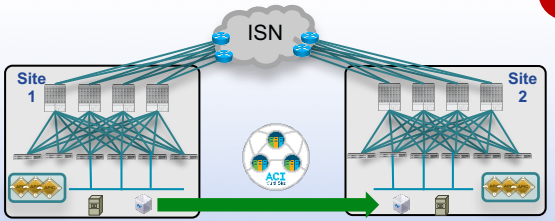

- Same IP subnet defined in separate Sites
- Support for IP Mobility ('cold' and 'live'* VM migration) and intra-subnet communication across sites
- **No Layer 2 BUM flooding across sites**

MSO GUI (BD)

Layer 2 adjacency across Sites

3

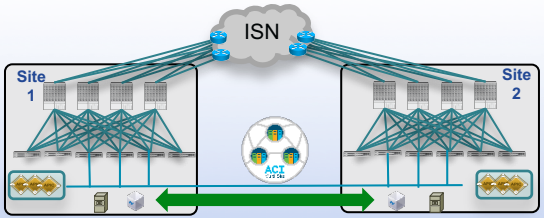

- Interconnecting separate sites for fault containment and scalability reasons
- Layer 2 domains stretched across Sites, support for application clustering
- **Layer 2 BUM flooding across sites**

MSO GUI (BD)

*'Live' migration officially supported from ACI release 3.2

Agenda

- ACI Foundations
- ACI Anywhere
  - ACI Multi-Pod
  - ACI Multi-Site
  - Physical Remote Leaf
  - Virtual Remote Leaf (vPod)
  - ACI with Public Cloud
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Business Value

Extending the ACI policy model outside the main datacenter to remote sites distributed over IP Backbone

Extending ACI fabric policy and L2/L3 connectivity to a small DR site without requiring the deployment of a full-blown ACI Fabric

Centralized Policy Management and Control Plane for remote locations

Small form factor solution at locations with space constraints

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Architecture Overview

A Remote Leaf pair gets associated with the Spines of one specific Pod in the Main DC

APIC and Spine Nodes (Proxy function) remain in the primary Pod

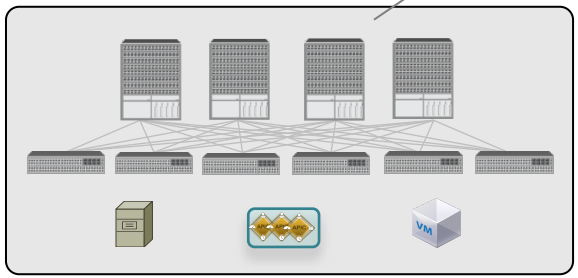

ACI Main DC

IPN Requirements in 4.0(1)

- 300 msec maximum RTT
- 100 Mbps minimum BW
- 1550B minimum MTU
- No PIM-Bidir required

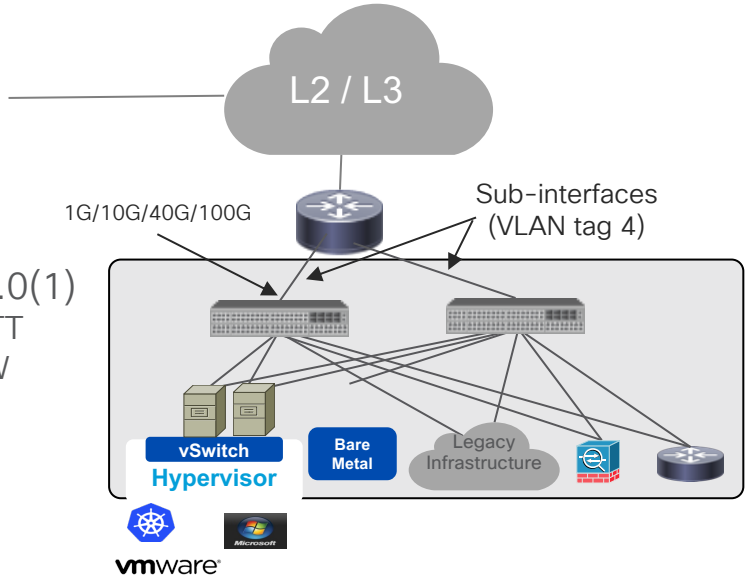

Remote Leaf Pair: a pair of Nexus 9300 nodes connected to a L3 Network via uplink ports and fully managed by a centralized APIC cluster

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Hardware and Software Support

ACI Main DC

Supported Spines

Fixed
- 9364C/9332C

Modular
- 9732C-EX
- 9736C-FX

Remote Location

Supported Leaf
- N93180YC-EX
- N93108TC-EX
- N93180LC-EX
- N93180YC-FX
- N9K-C93108TC-FX
- N9K-C9348GC-FXP
- N93240-YC-FX2
- N9336-FX2

All hardware from -EX onwards is supported

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Endpoint Connectivity Considerations

Recommended Design Option

RL Nodes not part of a vPC Domain

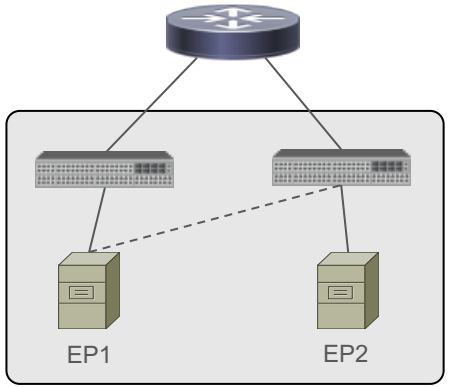

ACI 3.1(1) Release

- Dual attached host with single active uplinks (MAC pinning, Active/Standby teaming, etc.)
- Single attached hosts only

RL Nodes part of a vPC Domain

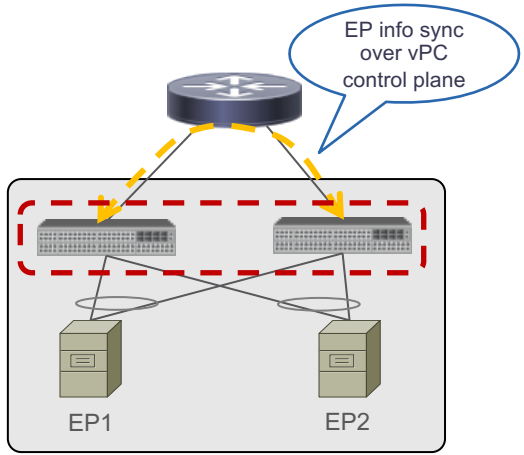

ACI 3.1(1) Release

- Dual attached host with Active/Active links (LACP)

RL Nodes part of a vPC Domain

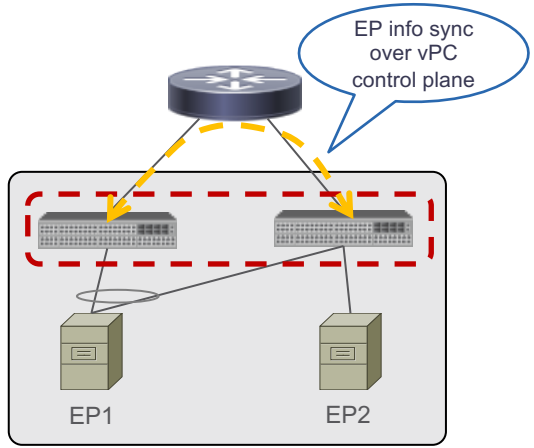

ACI 3.2(1) Release

- Dual attached host with Active/Active links (LACP)
- Single attached hosts (orphan ports)

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Automatic RL Discovery

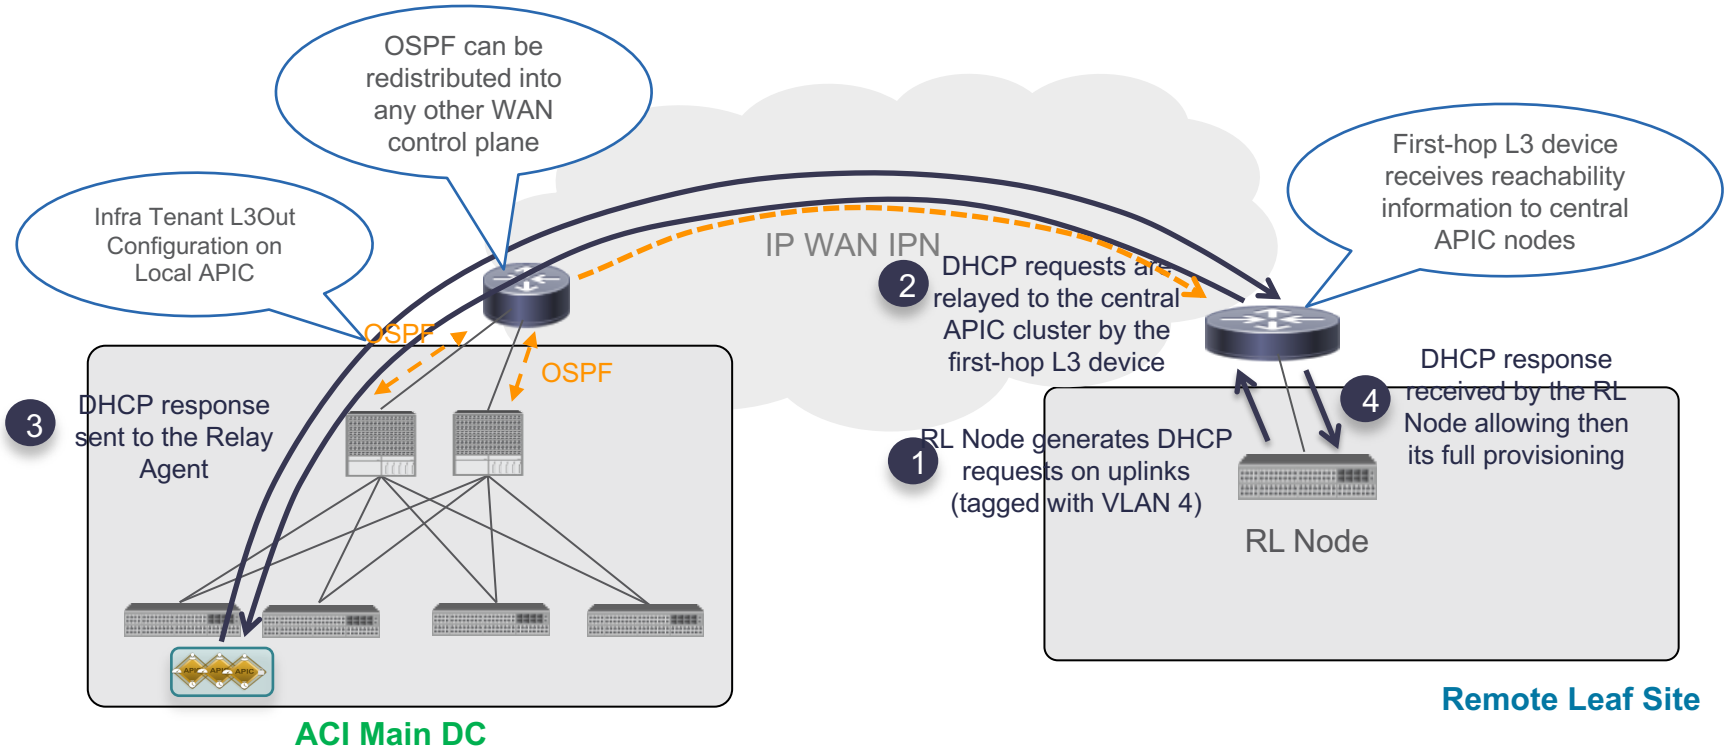

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Establishing End-to-End IP Connectivity

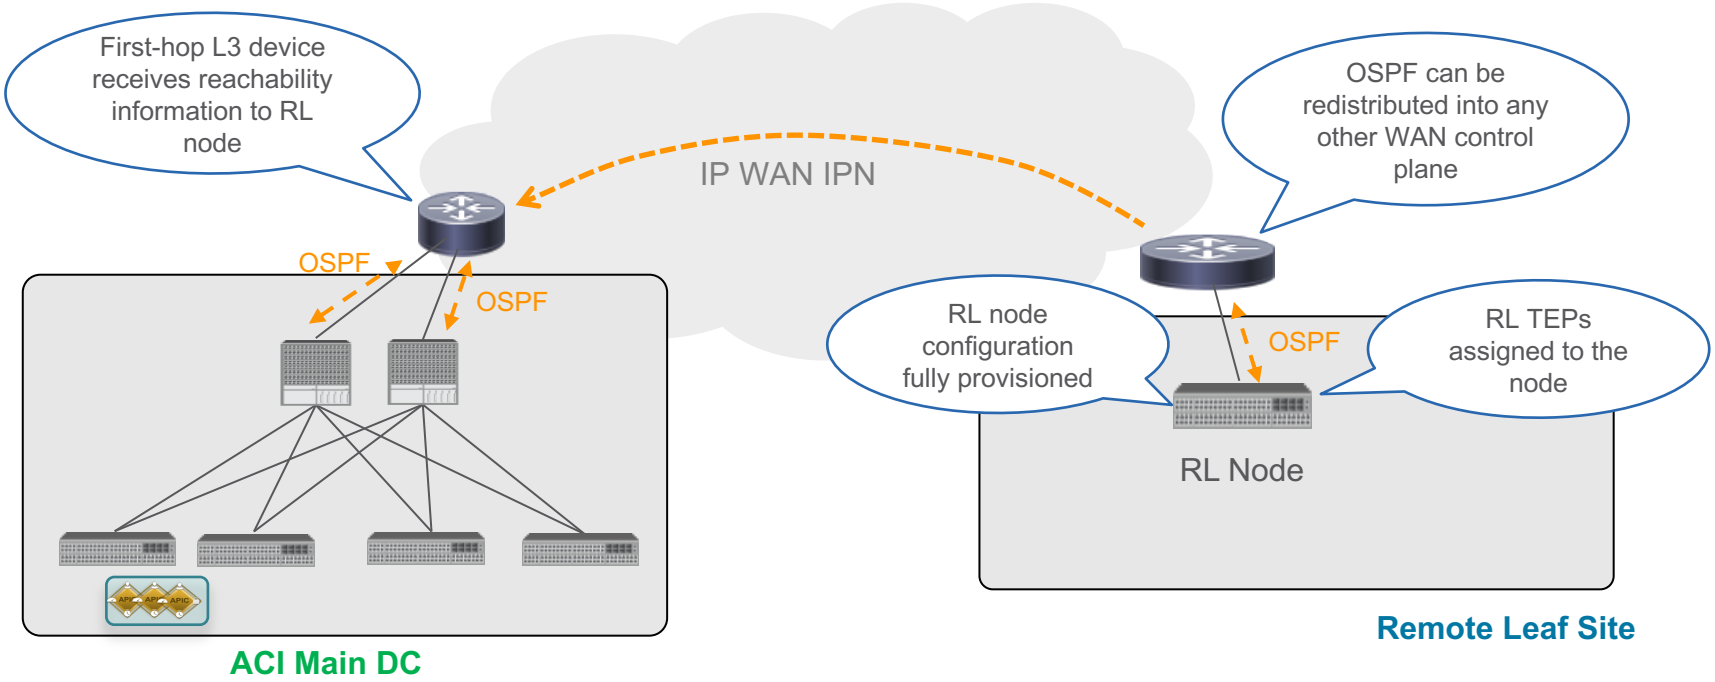

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
COOP for Announcing Remote Endpoint Information

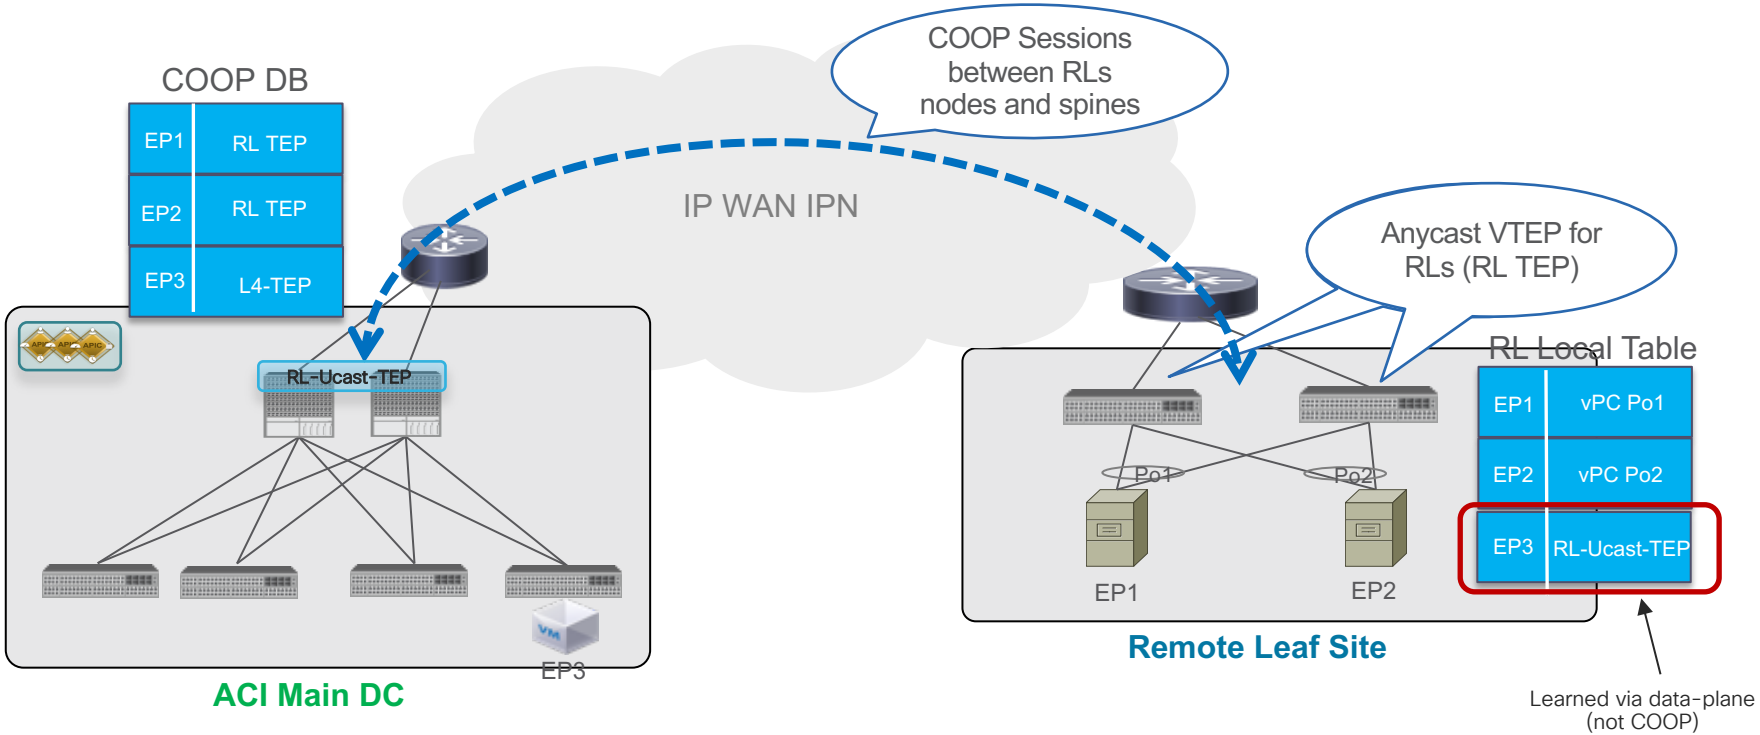

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Remote and Local Endpoints Communication

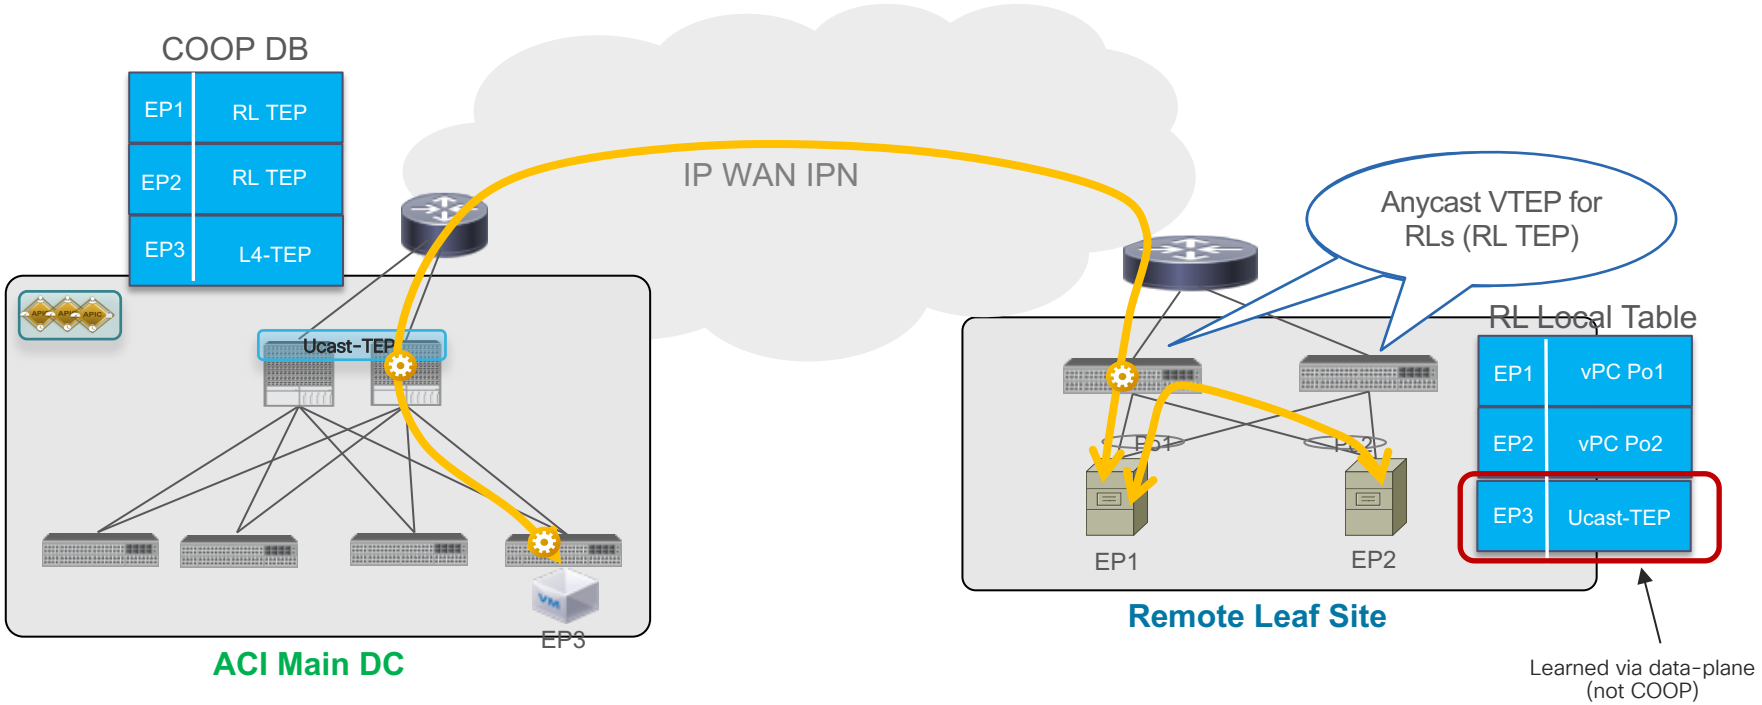

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Hair-Pin Communication between RL Pairs

Current behavior as ACI release 4.0(2)

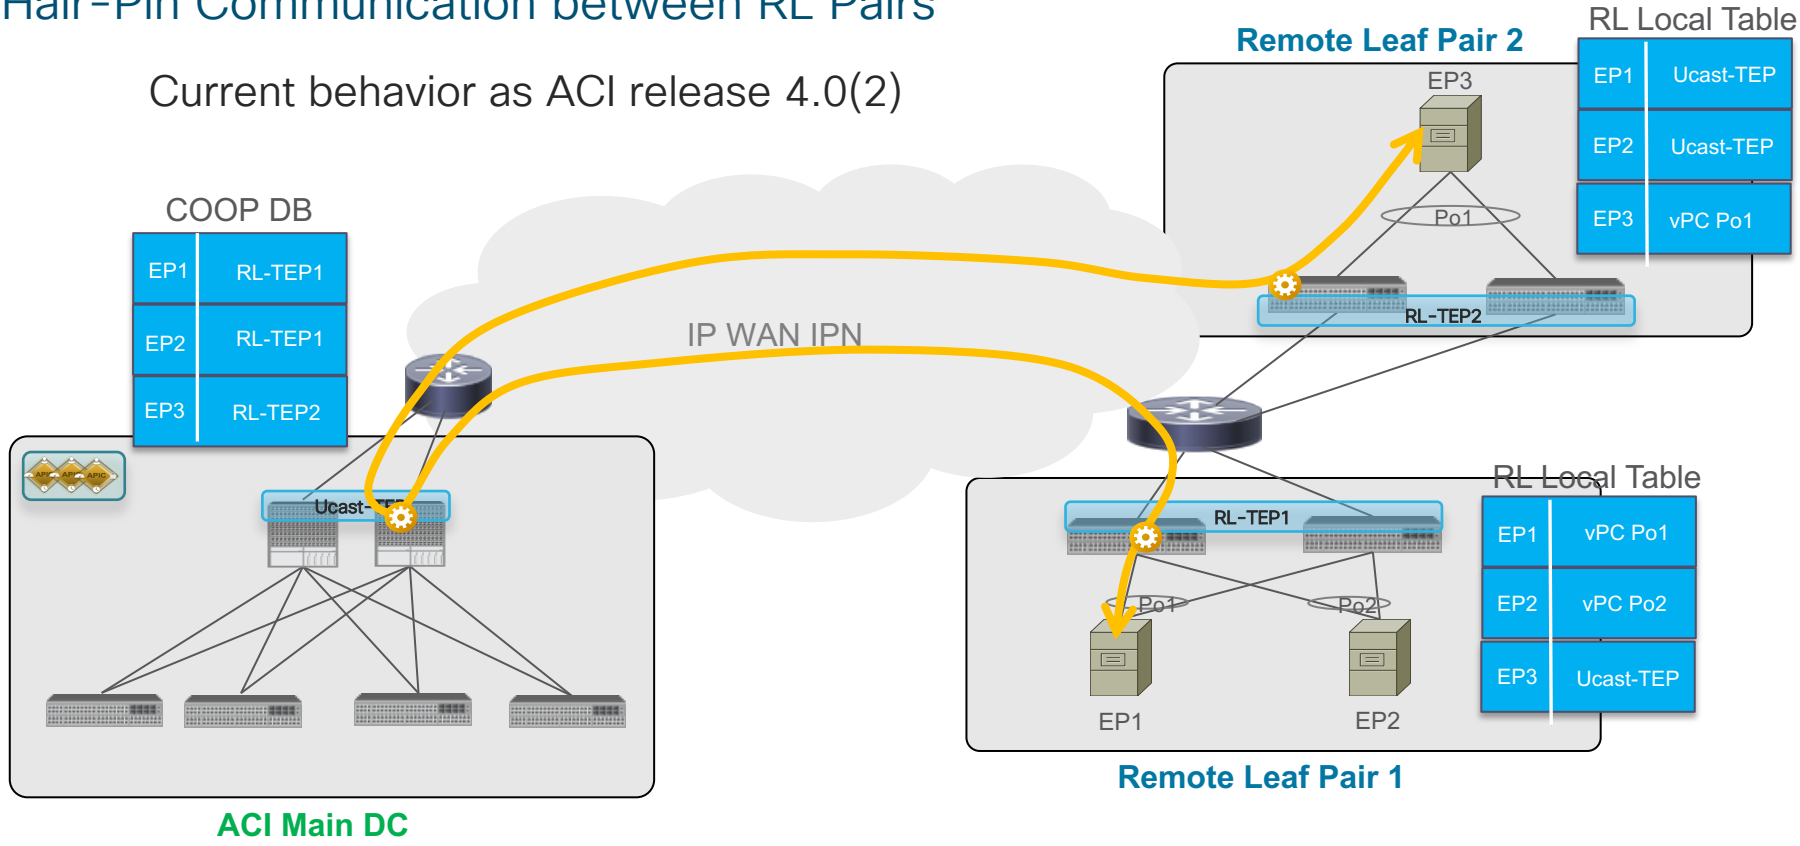

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Remote Physical Leaf
Direct Communication between RL Pairs

ACI 4.1(1) Release

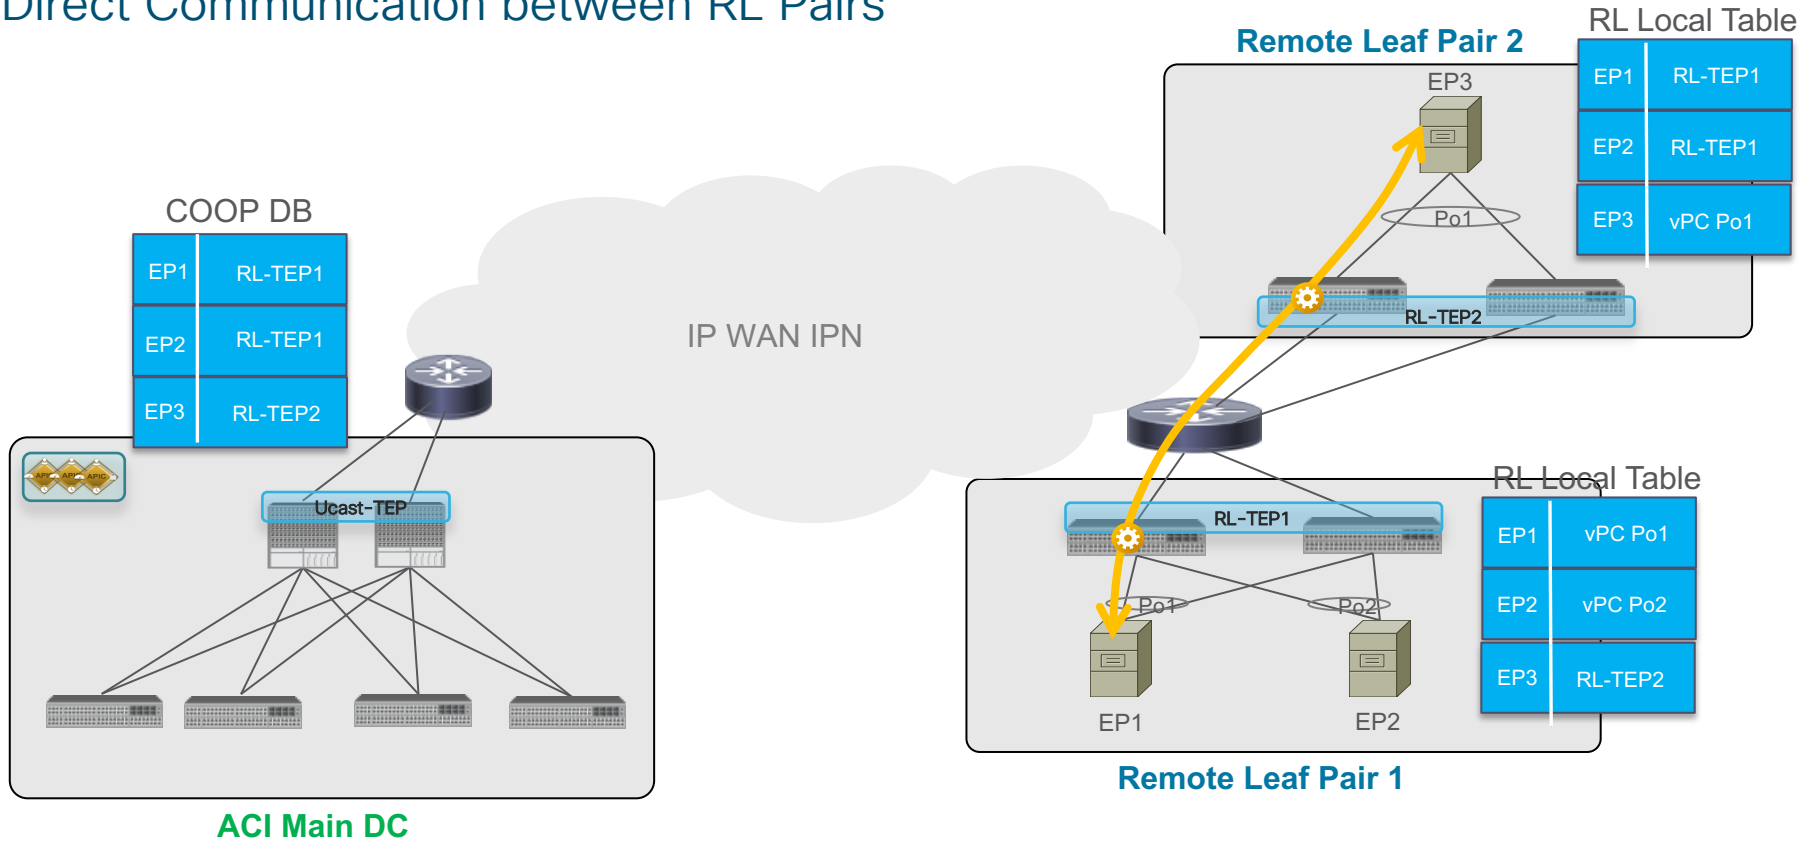

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- ACI Foundations
- ACI Anywhere
  - ACI Multi-Pod
  - ACI Multi-Site
  - Physical Remote Leaf
  - Virtual Remote Leaf (vPod)
  - ACI with Public Cloud
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Virtual Pod
Extend ACI to Virtual Remote Locations

ACI 4.0(1)
Release - LA

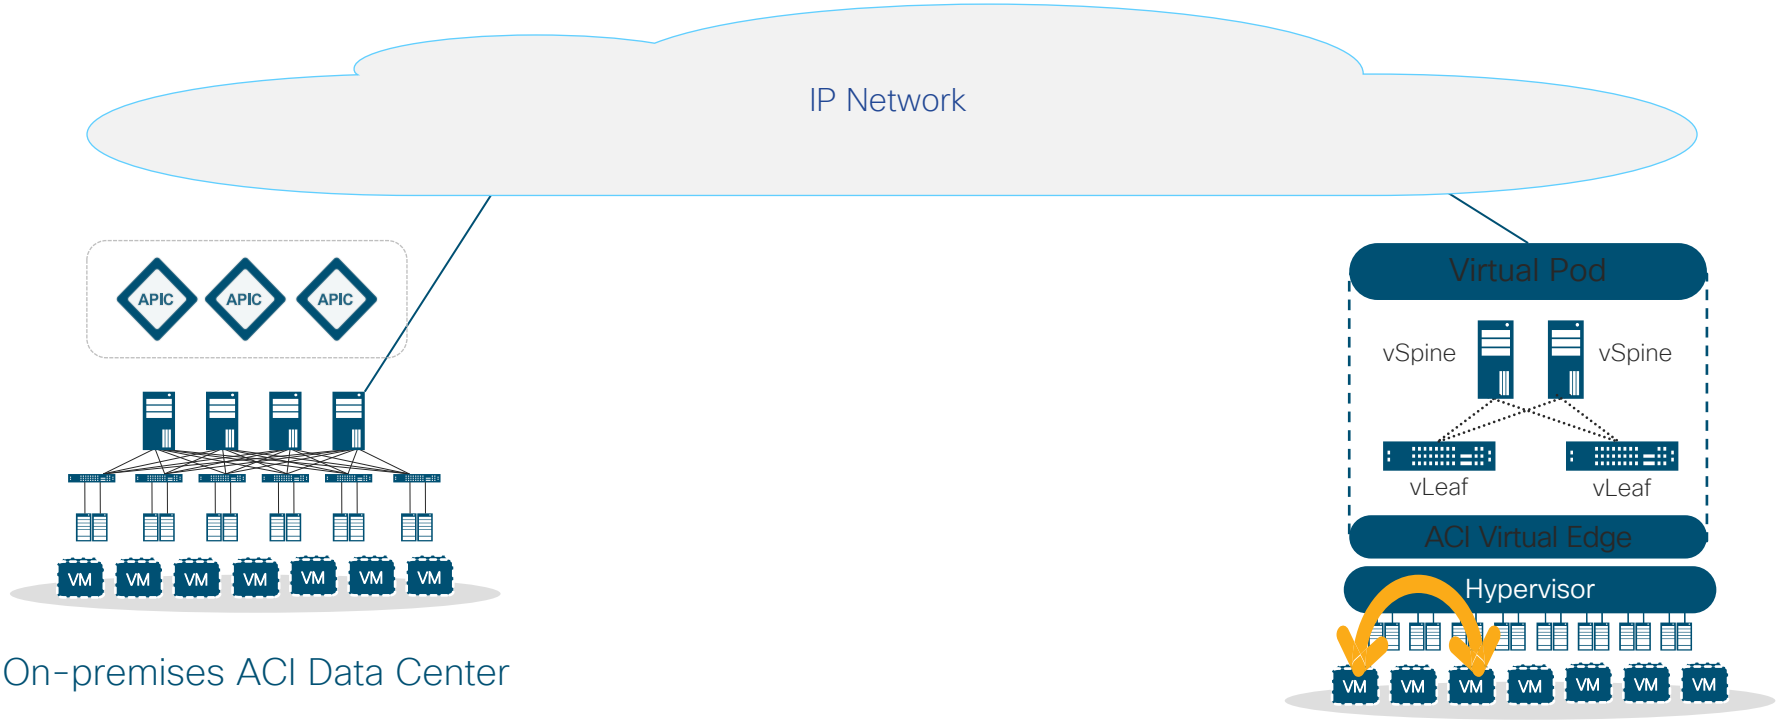

On-premises ACI Data Center

Bare Metal Clouds
(IBM, OVH, etc.)

Remote
Data Centers

Co-location Facilities
(Equinix, CoreSite etc.)

Brownfield
Deployments

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

ACI Virtual Pod (vPod)

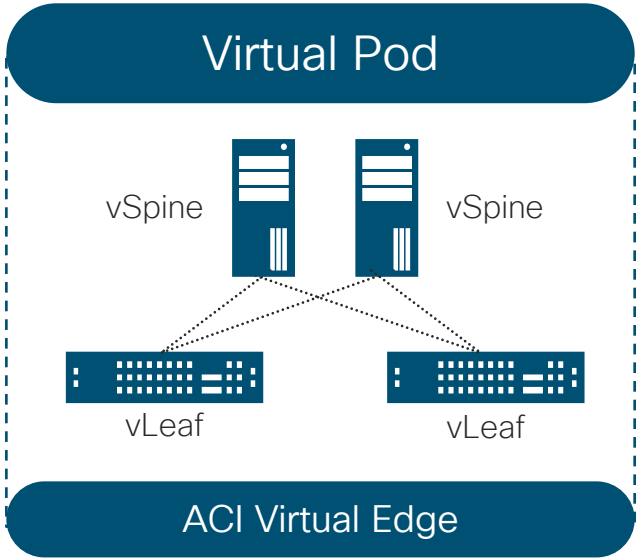

Management Cluster (vSpine + vLeaf)

- vSpine and vLeaf: Run ACI control plane function
- vLeaf: Distribute APIC policies to ACI Virtual Edge

ACI Virtual Edge (vPod Mode)

- Implements ACI data plane function and policy enforcement
- iVXLAN for communication within vPod and across Pods

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco ACI Virtual Edge
Decoupled From Hypervisor Kernel API Dependencies

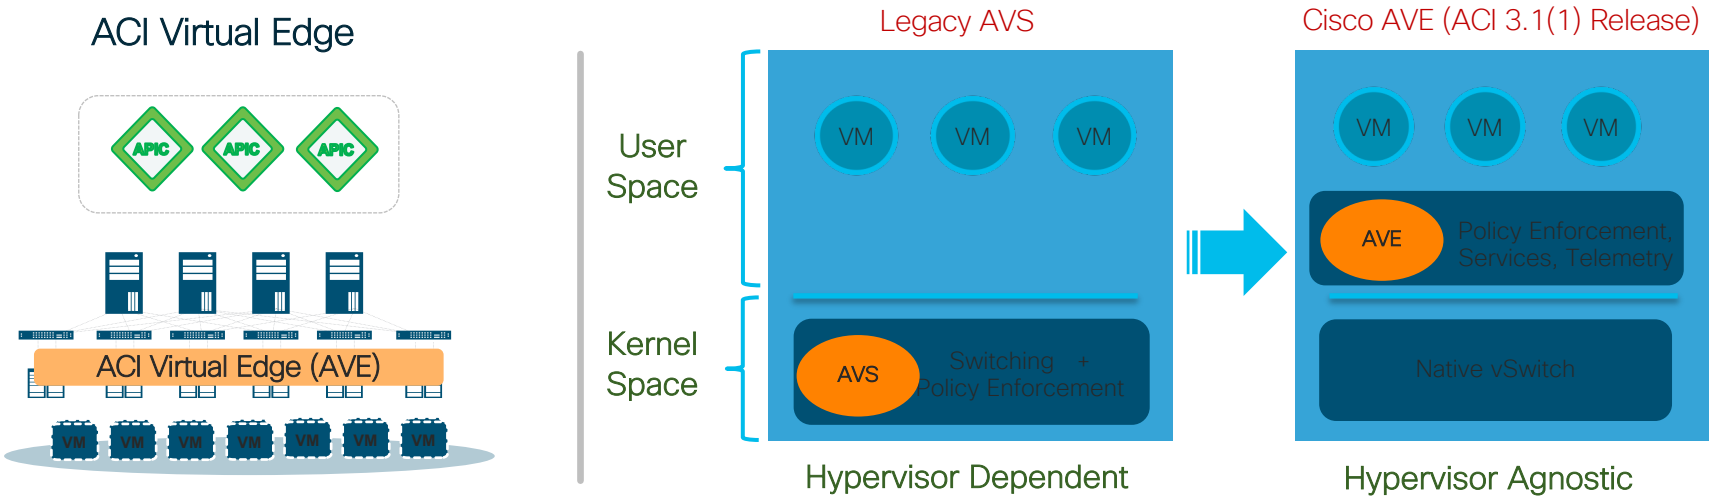

Maintain Existing Operational Models | Simple Transition/Migration AVS => AVE | Policy Consistency Across Multiple Hypervisors | AVS/AVE Feature Parity

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

vPod Components
vSpine, vLeaf, and Remote Virtual Leaf

vSpine + vTOR

- Run as container services inside VMs at the vPod location (collocated for availability)
- **vTOR:** Distribute APIC policies to AVE forwarders (DME/PE on vLeaf <-> Opflex on AVE)
- **vSpine**: Centralized endpoint and LPM database (COOP and BGP)
- ***Not*** in forwarding data path

Remote Virtual Leaf (AVE)

- Implements ***'all'*** ACI data path functions
- Use iVXLAN for communication within Remote site as well as between the vPod and other Pods
- Gen 2 of the AVE (ACI Virtual Edge)

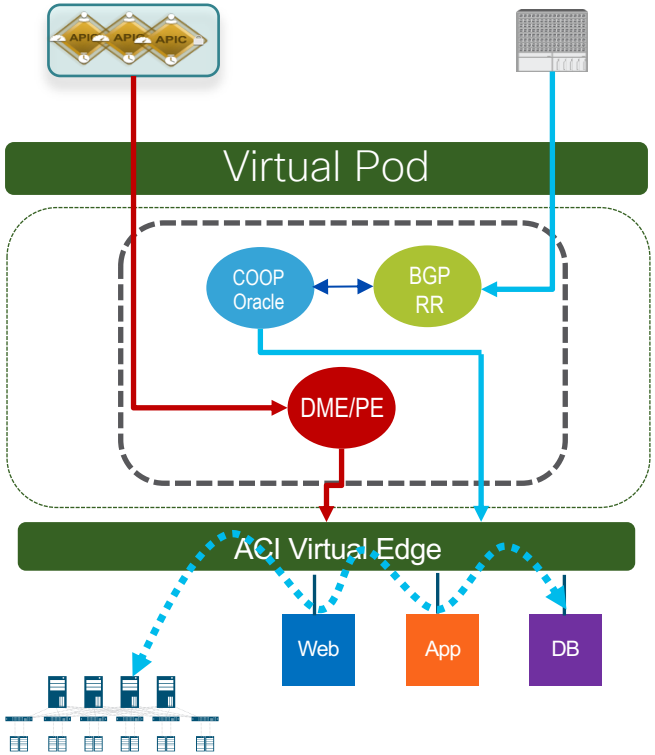

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Support for Multiple vPods

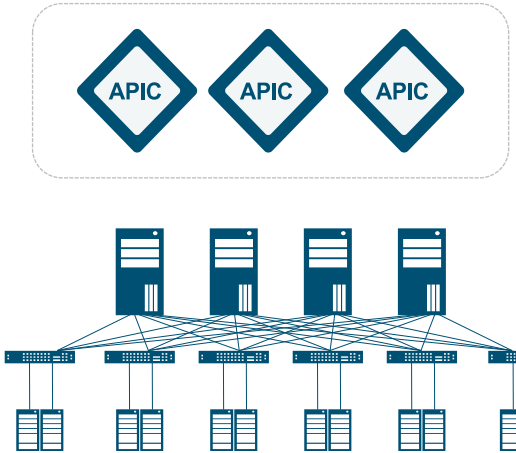

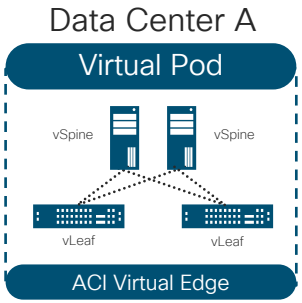

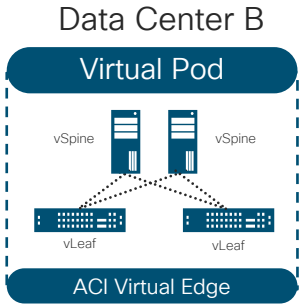

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

vMotion Across Pods

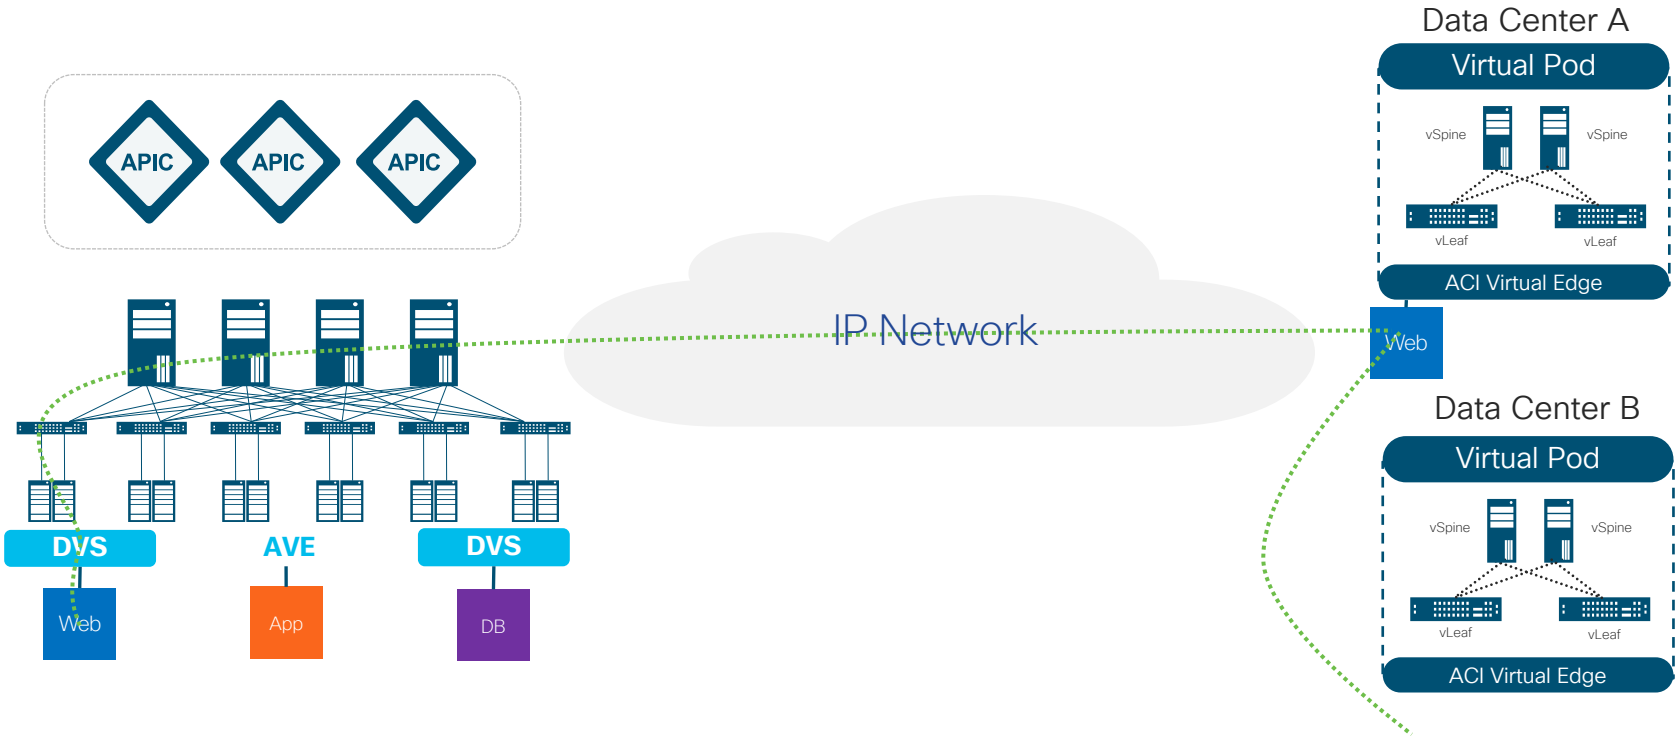

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

vPod Control Plane

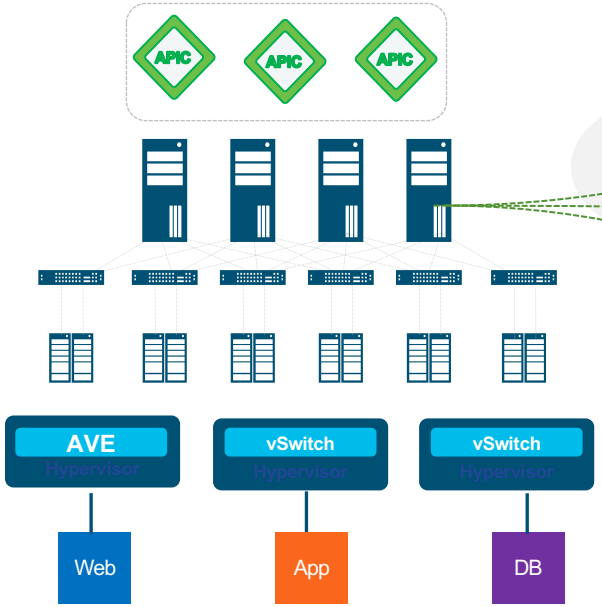

IP Network

Logical Connection To Spine
(BGP-EVPN)

vPod A
Virtual Pod
vSpine
vSpine
vToR
vToR
ACI Virtual Edge

vPod B
Virtual Pod
vSpine
vSpine
vToR
vToR
ACI Virtual Edge

vPod C
Virtual Pod
vSpine
vSpine
vToR
vToR
ACI Virtual Edge

vPod Data Plane

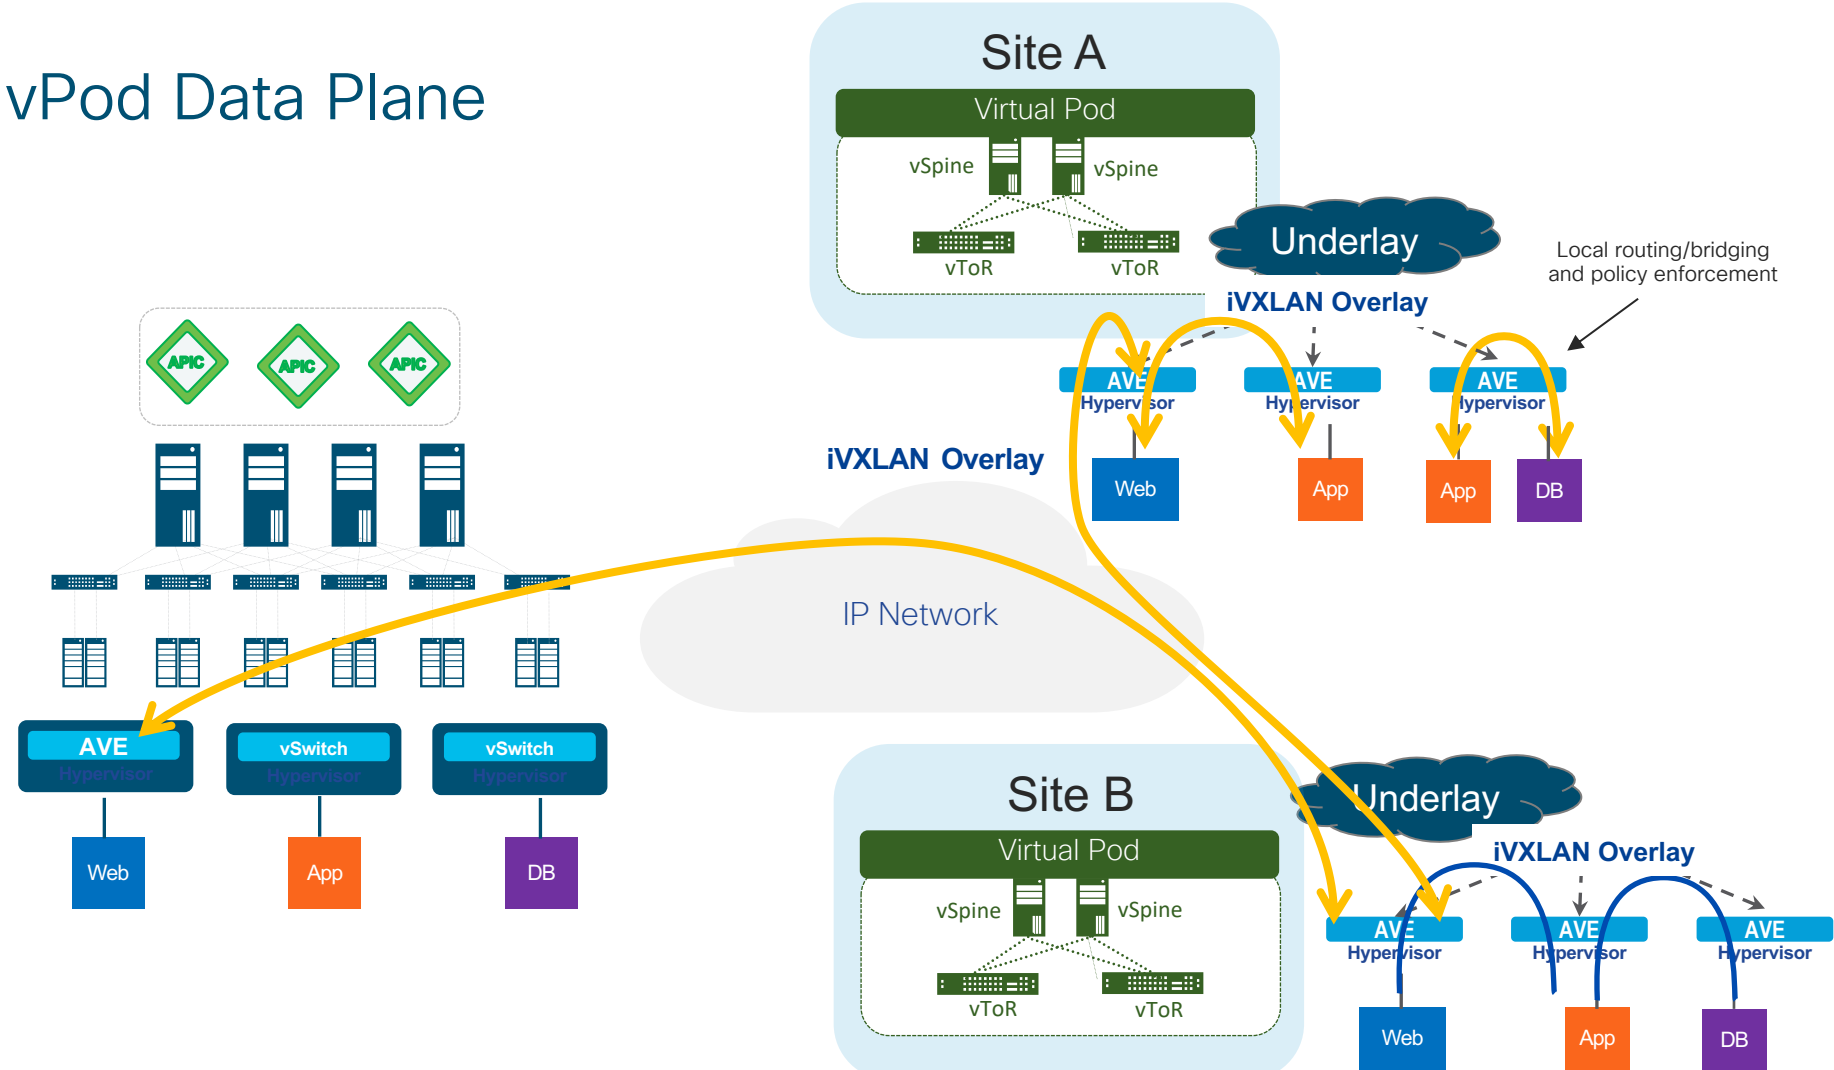

Phase 1 vPOD connectivity to “outside”

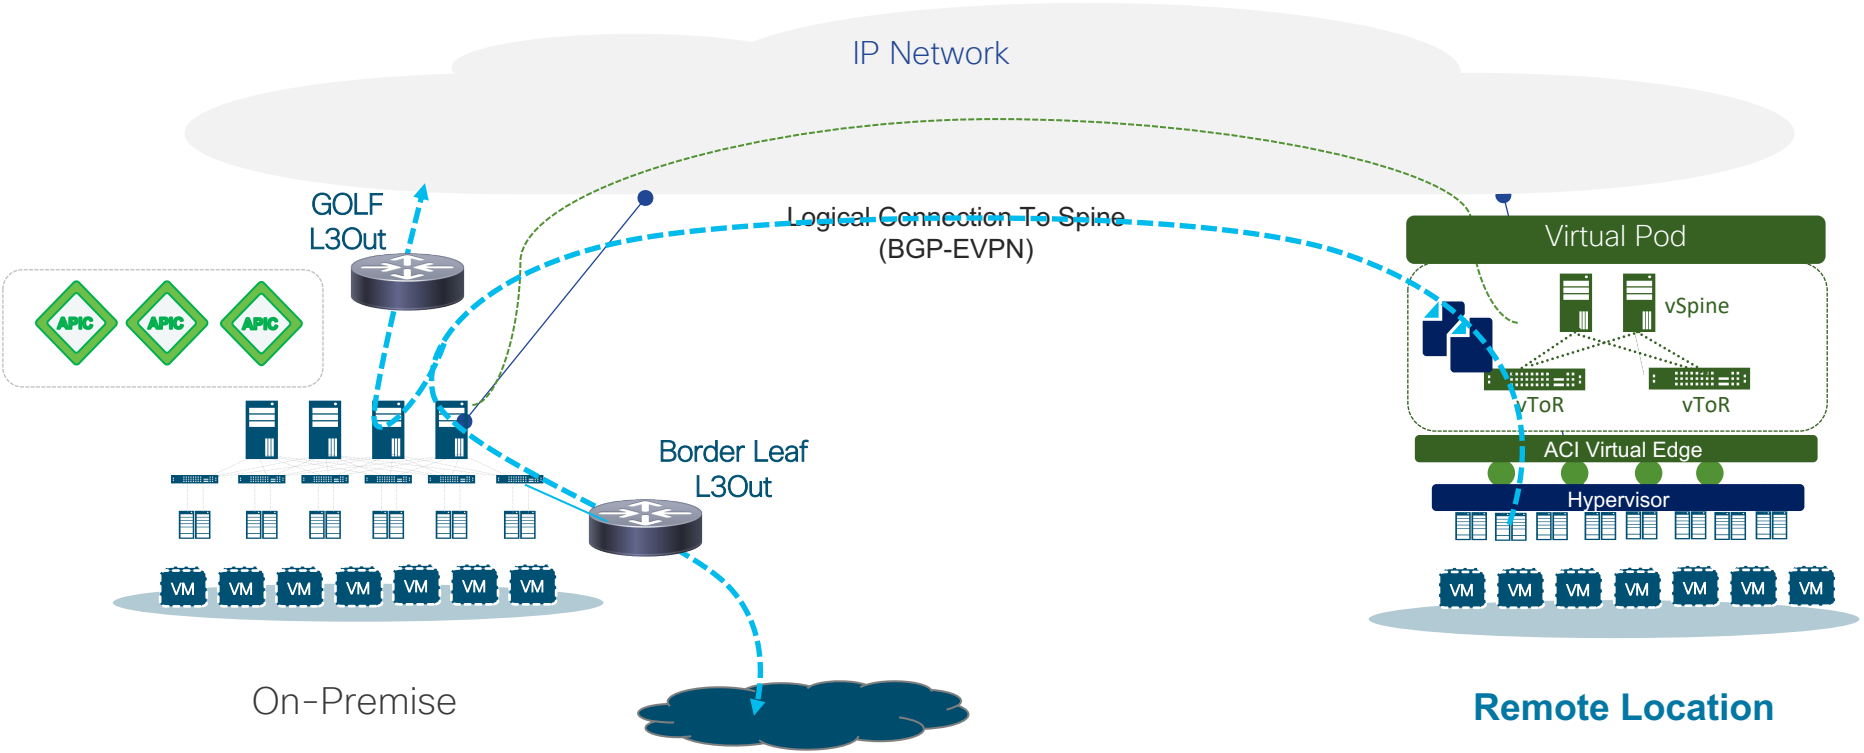

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- ACI Foundations
- ACI Anywhere
  - ACI Multi-Pod
  - ACI Multi-Site
  - Physical Remote Leaf
  - Virtual Remote Leaf (vPod)
  - ACI with Public Cloud
- Q&A

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Fundamentals

- **Regions** – Think of it as multiple data center with more than one physical location. Pod or site could be used for ACI
- **Availability Zones (AZ)** – Set of buildings, Internet uplinks and power. Think of it as a data center but may contains more than one physical location. Path or node attachment could be used in ACI
- **Virtual Private Cloud (VPC)** – Set of subnets with one ore more CIDR blocks running in a single region across multiple data centers (AZ). Similar to VRF
- **Subnet** – Range of IP addresses. Each subnet must reside within one AZ and can't span zones. Minimum subnet size is /28. BD Subnet

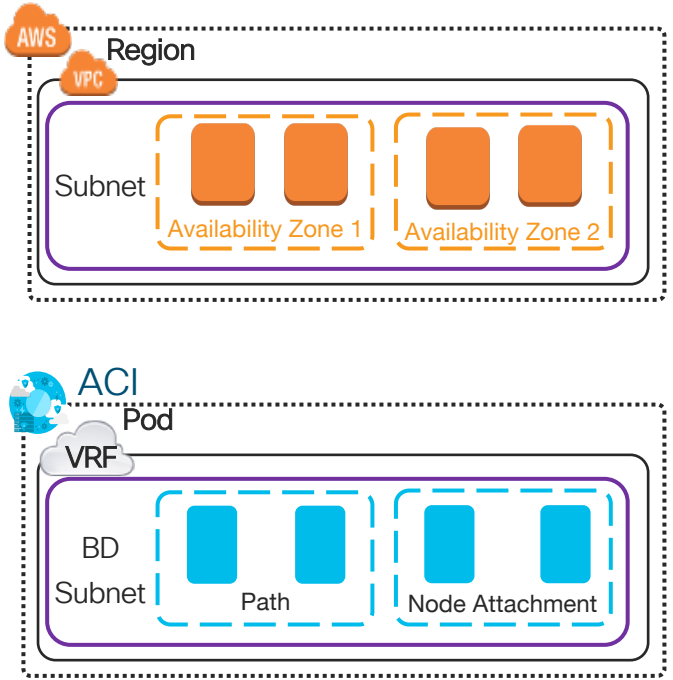

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

Fundamentals (Cont.)

- **Security Group** – Act as a firewall for associated EC2 instance (VM), controlling both inbound and outbound traffic at network interface (EP) level. Equivalent to EPG with white-list
- **Security Group Rule** – Rules applied to inbound traffic (ingress) or outbound traffic (egress). Combination of contracts and filters in ACI
- **Network ACL** – Used to deny / permit select traffic at a subnet level. Network ACLs are stateless. In ACI, it is similar to taboo and grey-list contracts
- **Route Table** – Can be associated with multiple subnets. Acts like a source-based policy-based routing (PBR) rule.

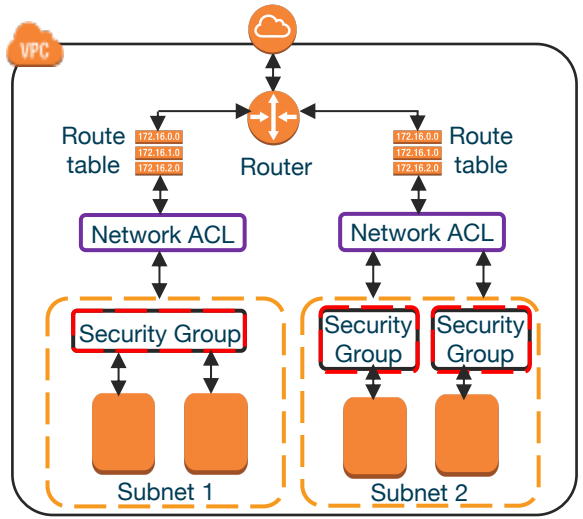

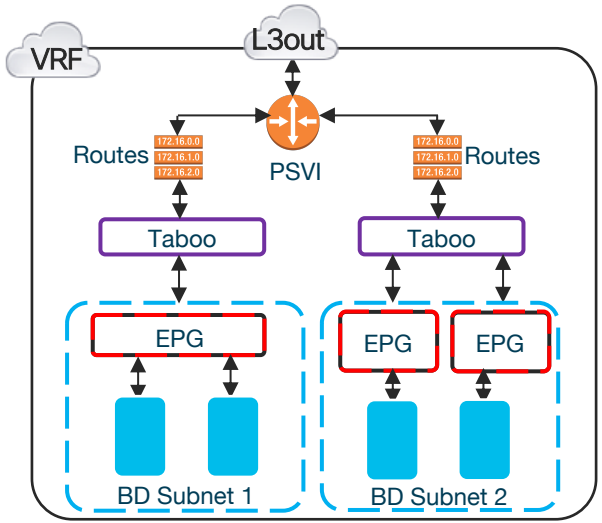

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

For your info & reference

Policy Mapping – AWS

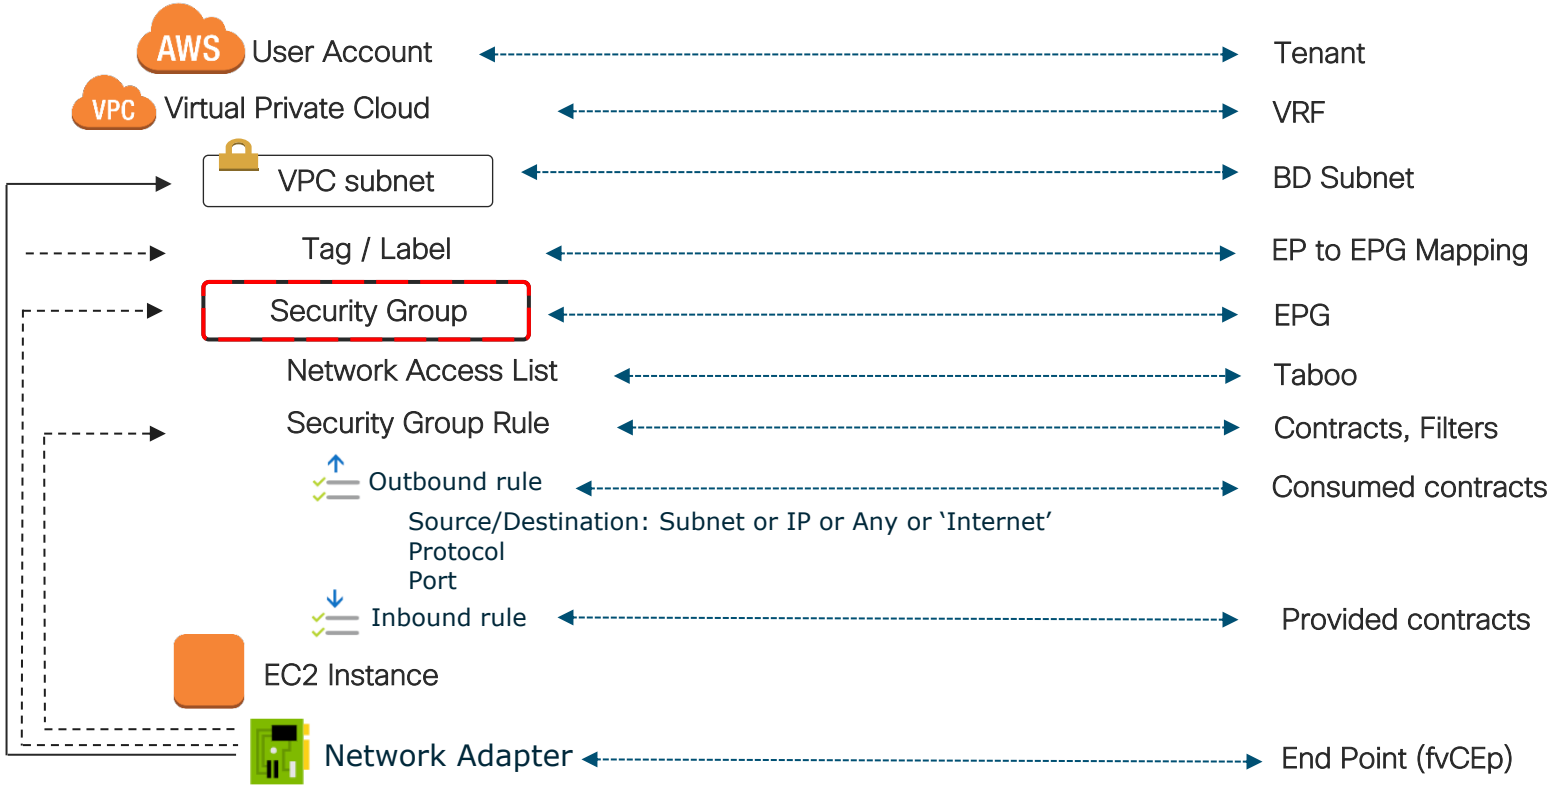

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cloud Infra

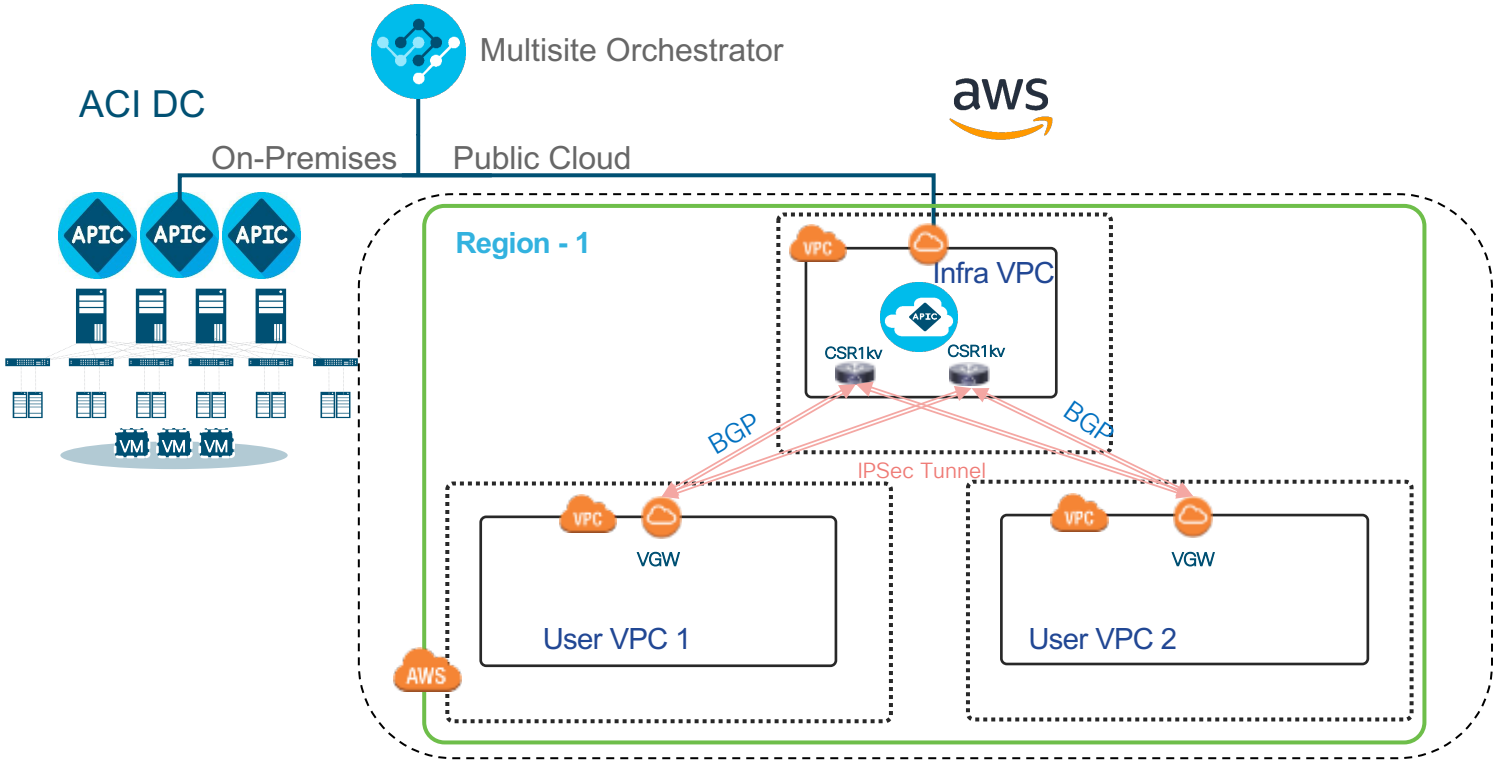

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

Infra VPC

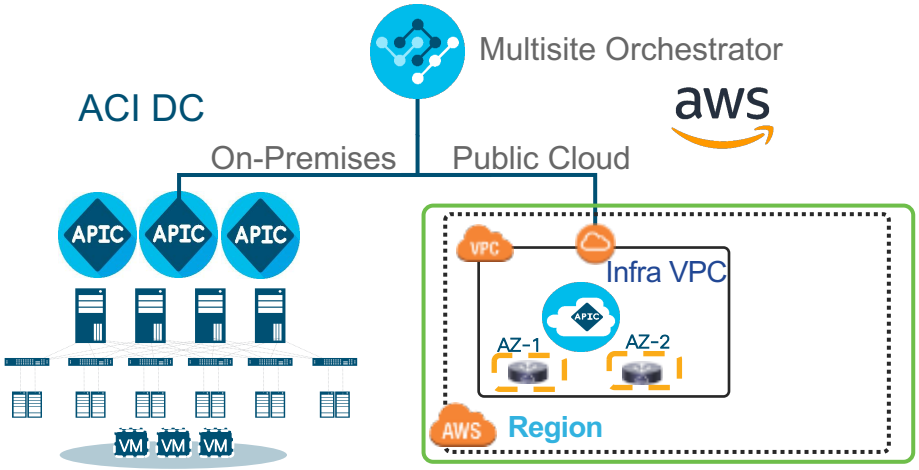

- Infra VPC will be used as transit VPC to connect between on-premises ACI Fabric and AWS region or connect one AWS region with another AWS region
- You would need an AWS account which will act similar to Fabric admin
- Need to have proper IP subnet planned ahead of the deployment

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

User VPC

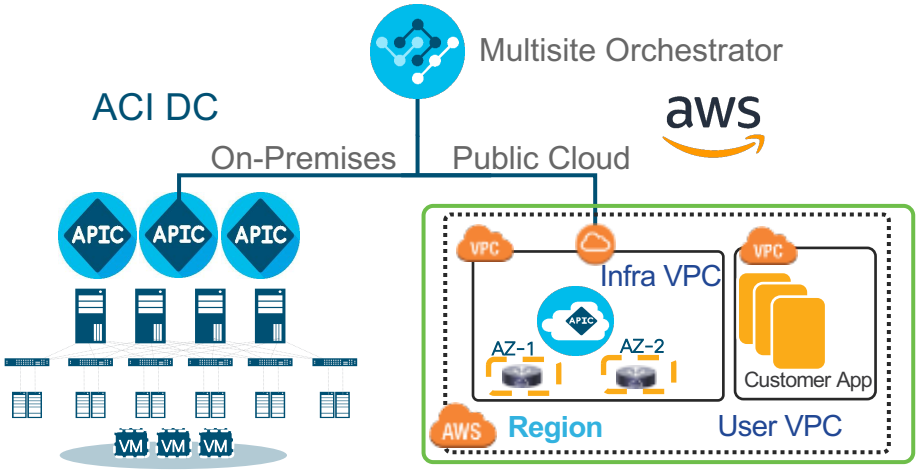

- User VPC will be created by Cloud APIC where all application policy will be enforced
- Need an AWS account which will act as Tenant admin before creating user VPC
- IP subnets need to be unique within a User VPC
- User VPC communicates with another User VPC through the Infra VPC

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Public

Thank you