



You make **possible**



Application Centric Infrastructure (ACI)

The Intent Driven Data Center

Max Ardica – Principal Engineer, Cisco

Mike Herbert – Principal Engineer, Cisco

Takuya Kishida – Technical Marketing Engineer, Cisco

Carlos Pereira – Distinguished Systems Engineer, Cisco

Dang Ngo – Technical Solutions Architect, Cisco

TECACI-2009

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



Who Are We??



Max Ardica
Principal Engineer -
IBNG



Takuya Kishida
TME - IBNG



Dang Ngo
Technical Solutions
Architect



Mike Herbert
Principal Engineer -
IBNG



Carlos Pereira
Distinguished
Systems Engineer

Agenda for the Day

(Yes your brain will hurt by the end)

- 8:30 – 9:00: Introduction to the Intent Based Data Center (ACI)
- 9:00 – 10:30: ACI Fabric Design Best Practices Part I
- 10:45 – 11:15: ACI Fabric Design Best Practices Part II
- 11:15 – 12:45: ACI Anywhere Design and Principles Part I
- 14:30 – 15:00: ACI Anywhere Design and Principles Part II
- 15:00 – 16:00: ACI and Cisco's Multi-Domain Architecture
- 16:00 – 16:30: How to Operate ACI Fabric(s) Part I
- 16:45 – 18:45: How to Operate ACI Fabric(s) Part II

A key question for the Day

When do I get caffeine?

- 8:30 - 10:30: Technical Seminar
- 10:30 - 10:45: Coffee Break
- 10:45 - 12:45: Technical Seminar
- 12:45 - 14:30: Lunch
- 14:30 - 16:30: Technical Seminar
- 16:30 - 16:45: Coffee Break
- 16:45 - 18:45: Technical Seminar



Introduction to the Intent/Application Based Data Center

Why do we build Networks?



Experience versus Risk

Expect Great
Experience



End-Users
(consumers)



Business & Operations
must balance Risk and
Experience



Applications
(providers)


Experience versus Risk

Between users and applications, there is a network

Expect Great
Experience



End-Users



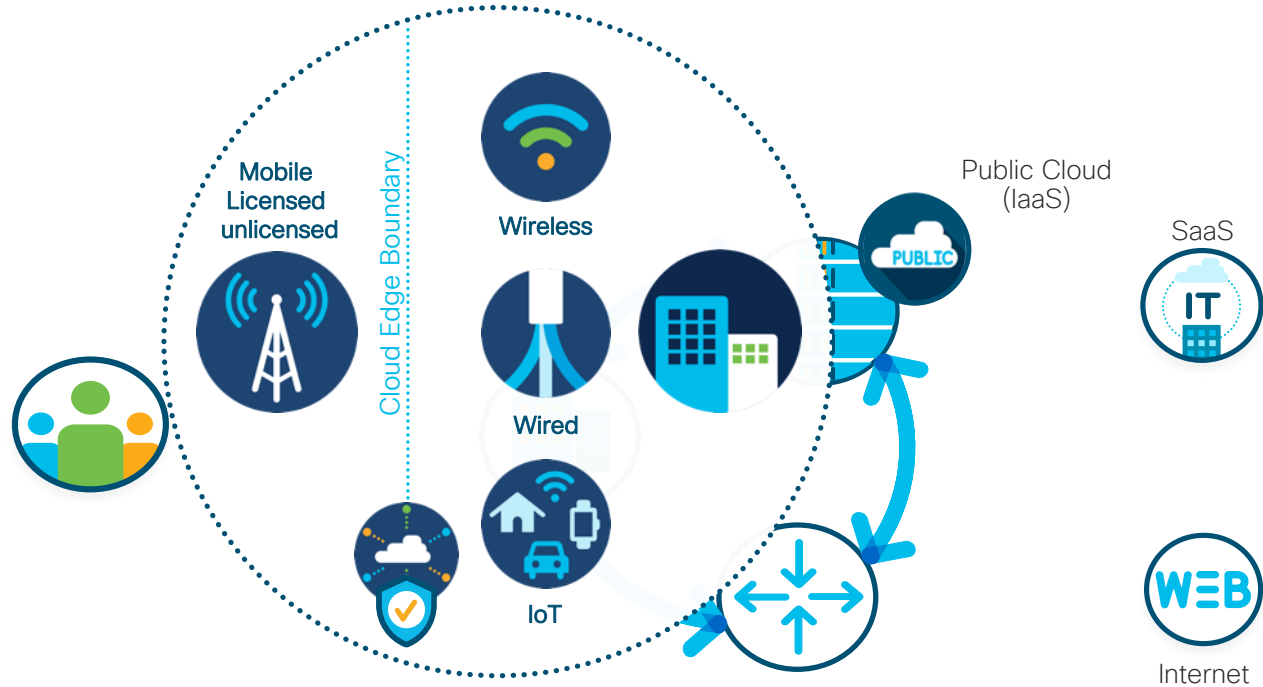
Business & Operations
must balance Risk and
Experience



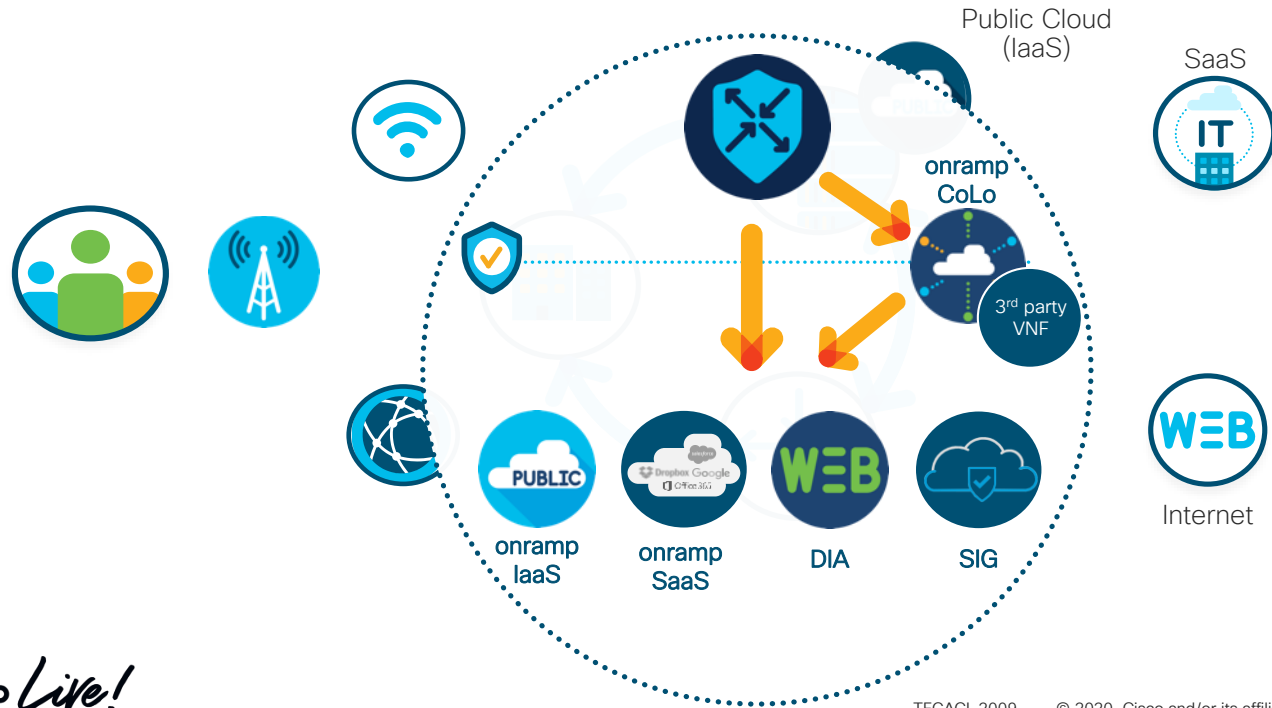
Applications

It's a multi-access world

With trust boundary's and a cloud edge function

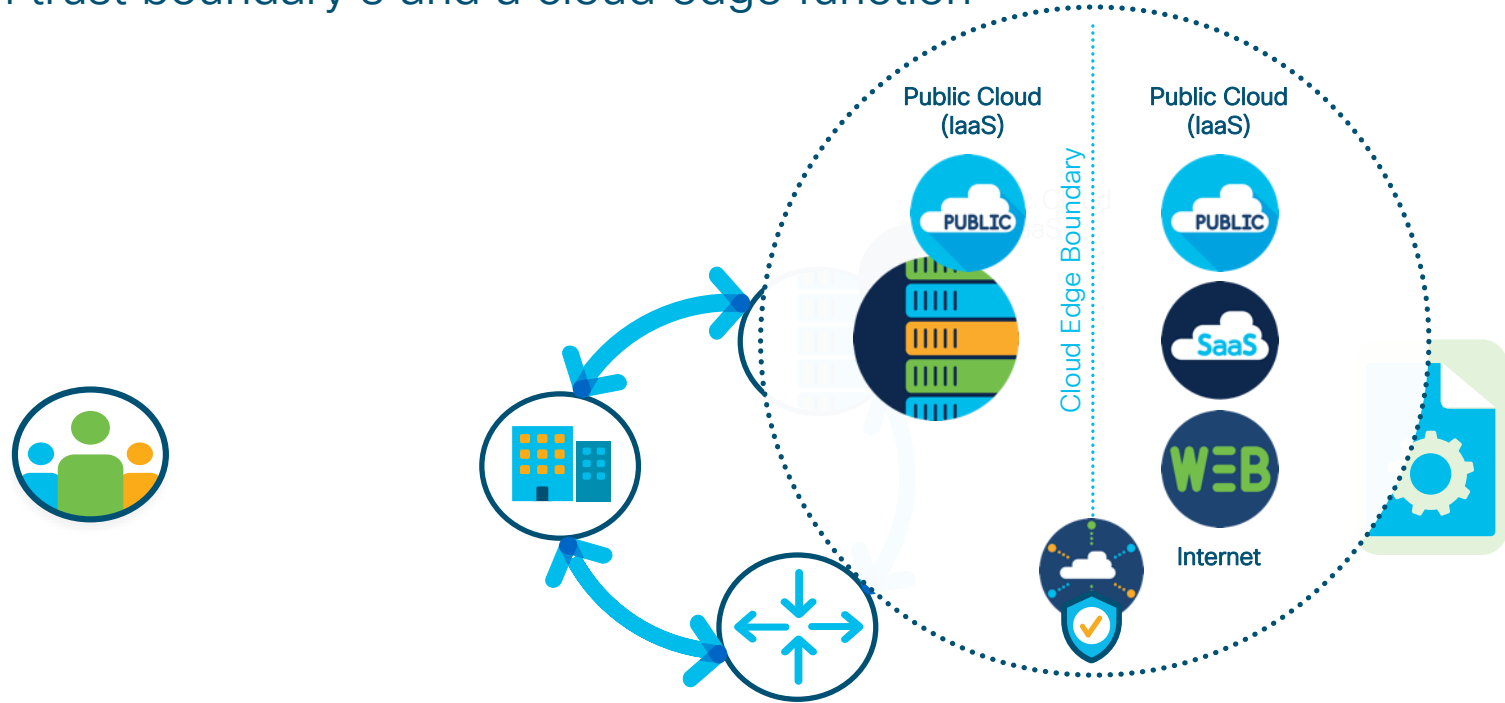


With Hybrid Cloud Connectivity World

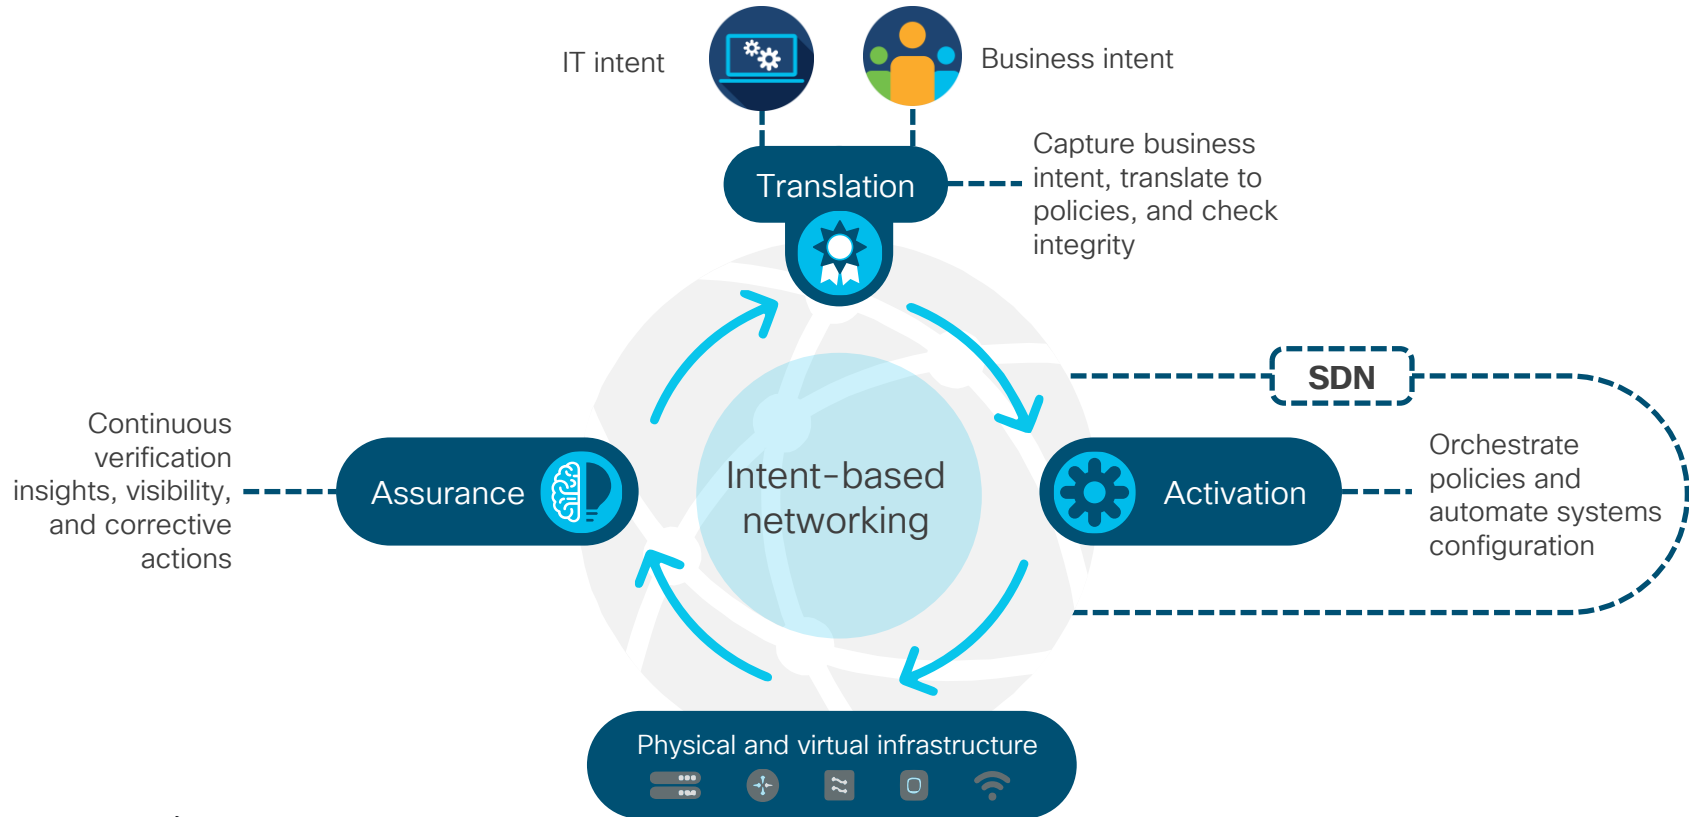


It's a multi-cloud world

With trust boundary's and a cloud edge function

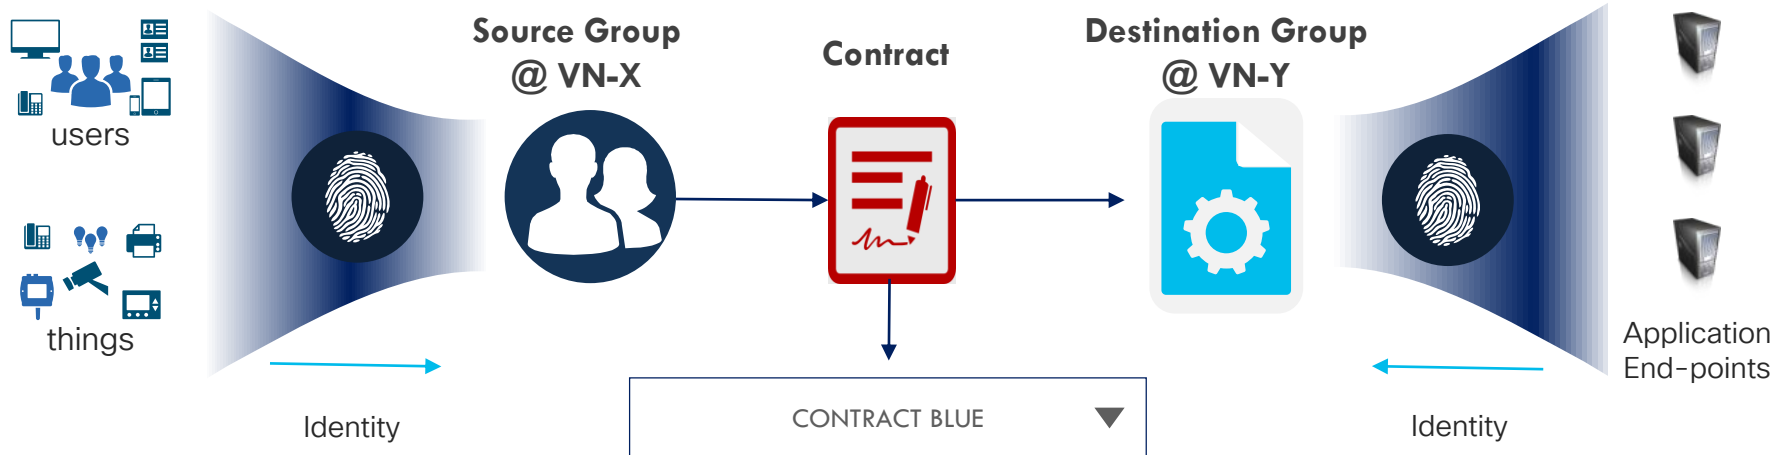


Introducing Intent Based Networking



This is the Intent of the Infrastructure

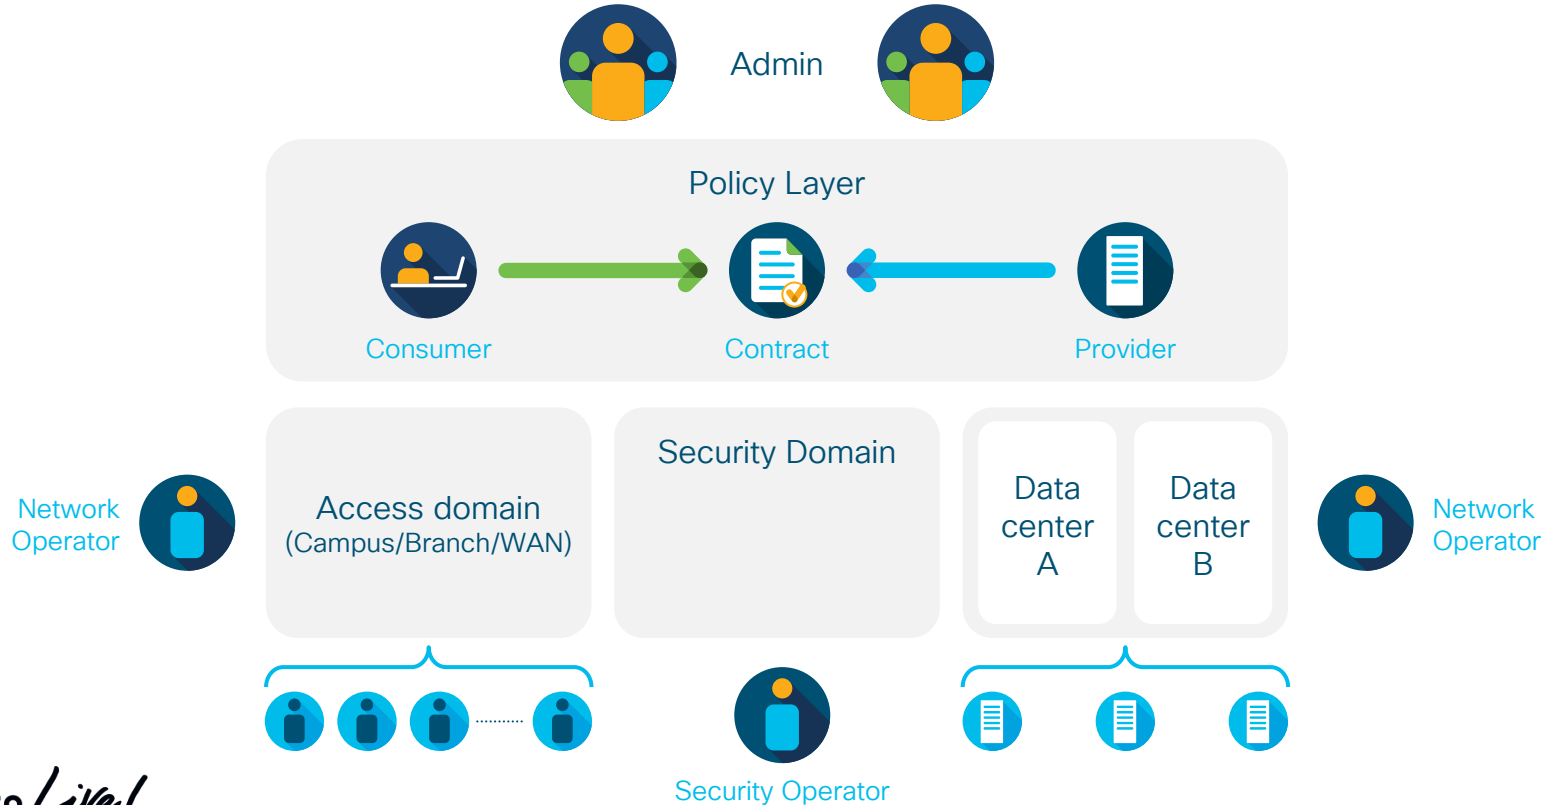
User to Application Policies



Classifier	Action
Port Number	Permit
IP Address	Deny
Application Type	Copy
	Gold Service

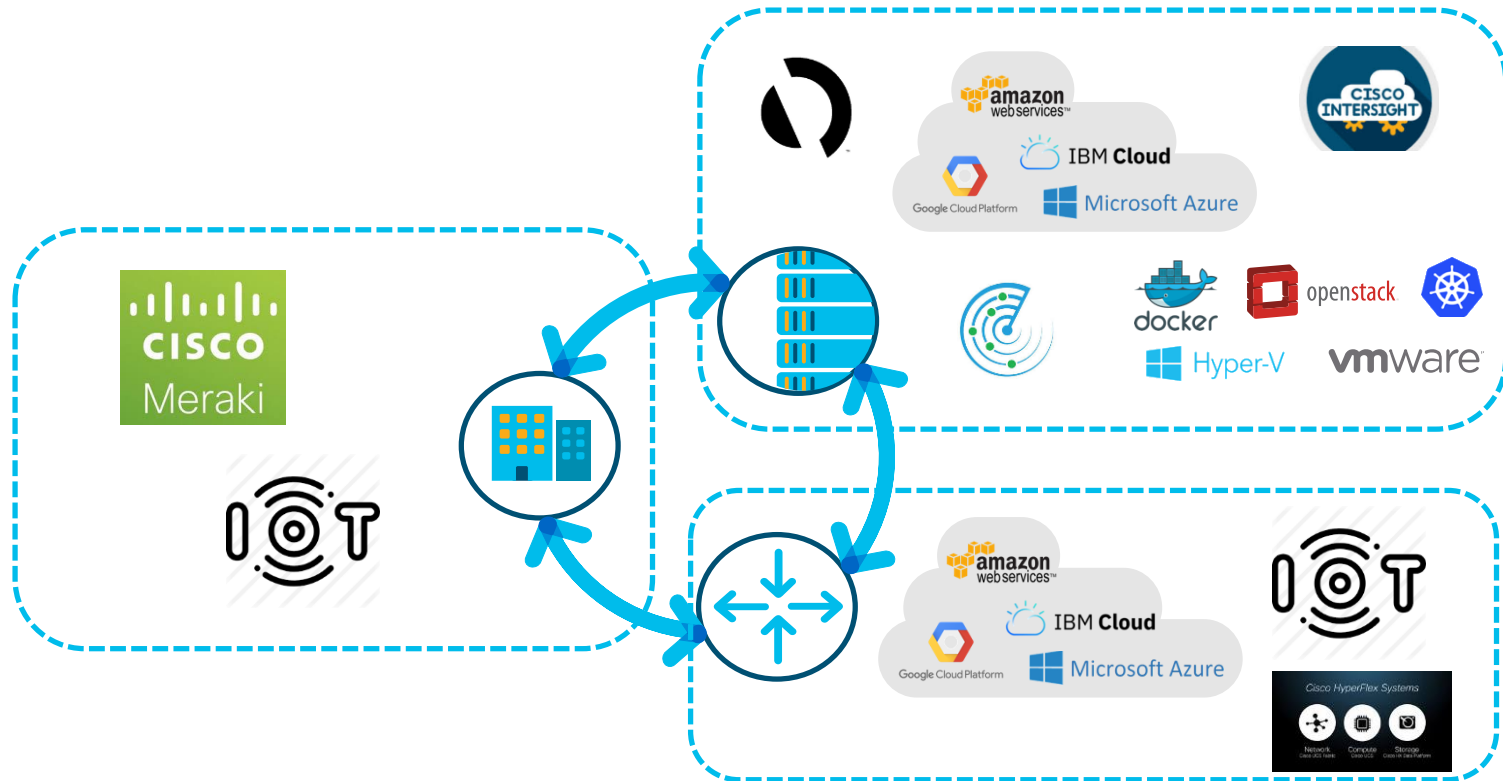
X-Domain Intent Based Networking

Unified policy language across domains



X-Domain Intent - Architectural View

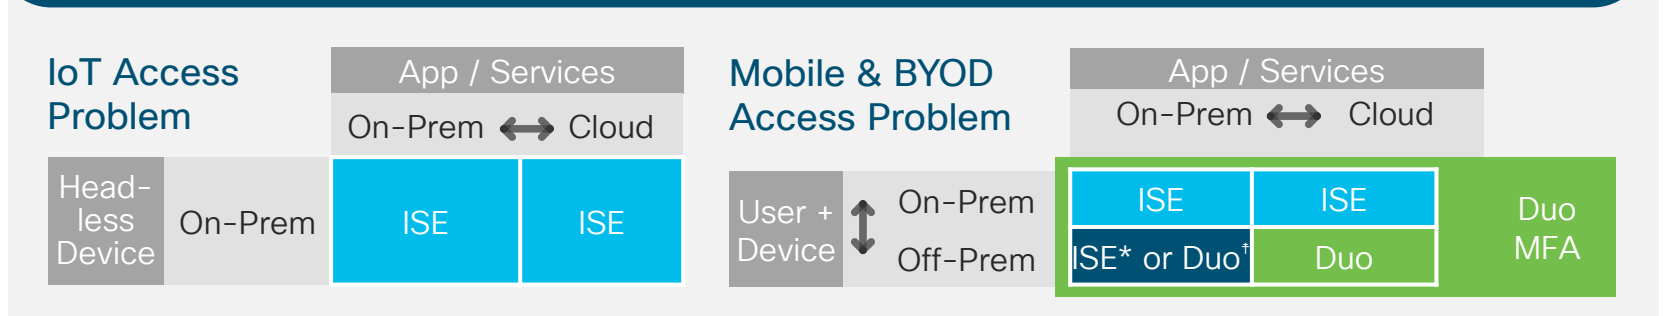
Horizontal and Vertical 'Multi-Domain'



X-Domain IBN is part of a Zero Trust Policy Framework



Unified Access across Hybrid IT Enterprises

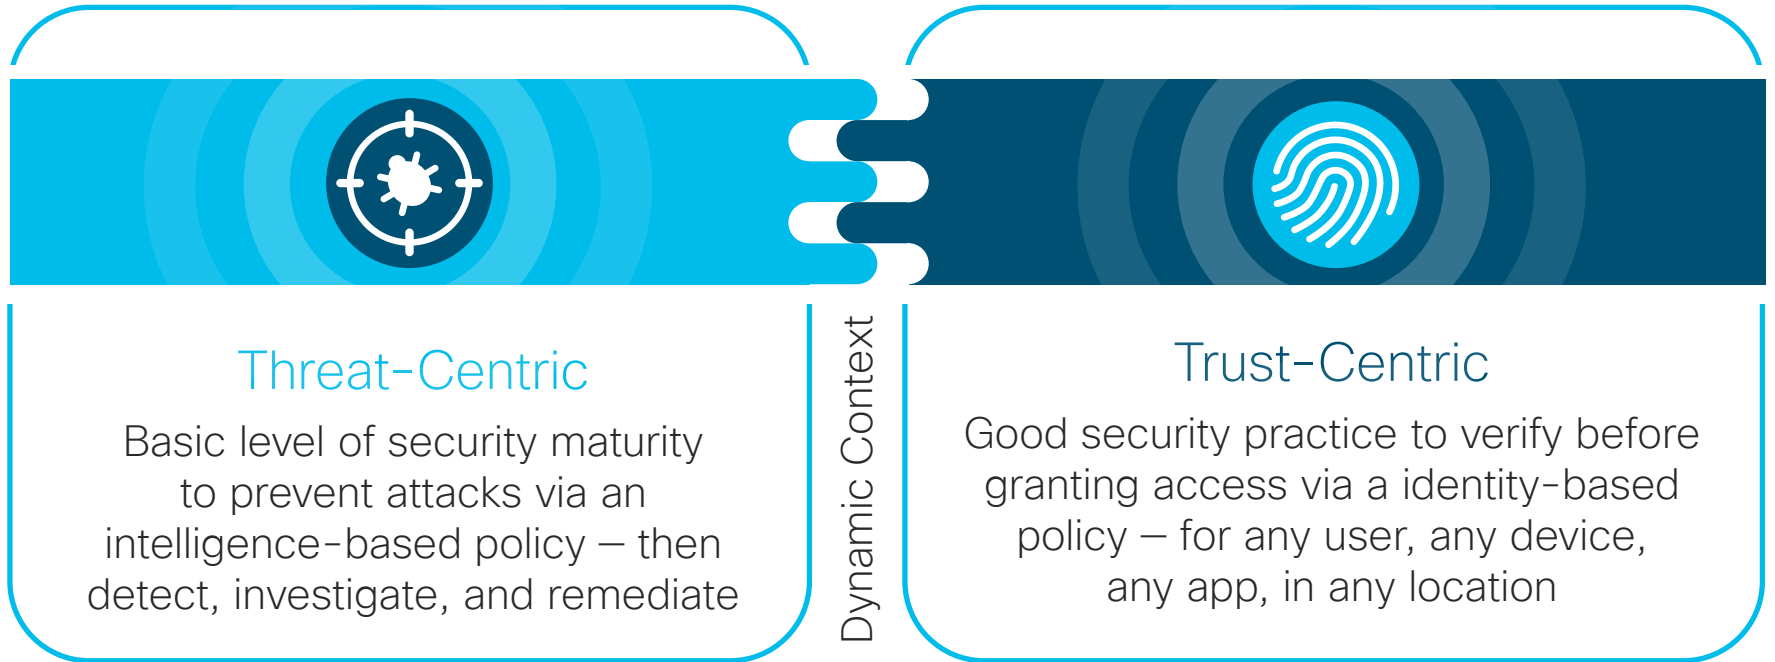


* Integrated with AnyConnect

† Duo Network Gateway (i.e. reverse proxy)

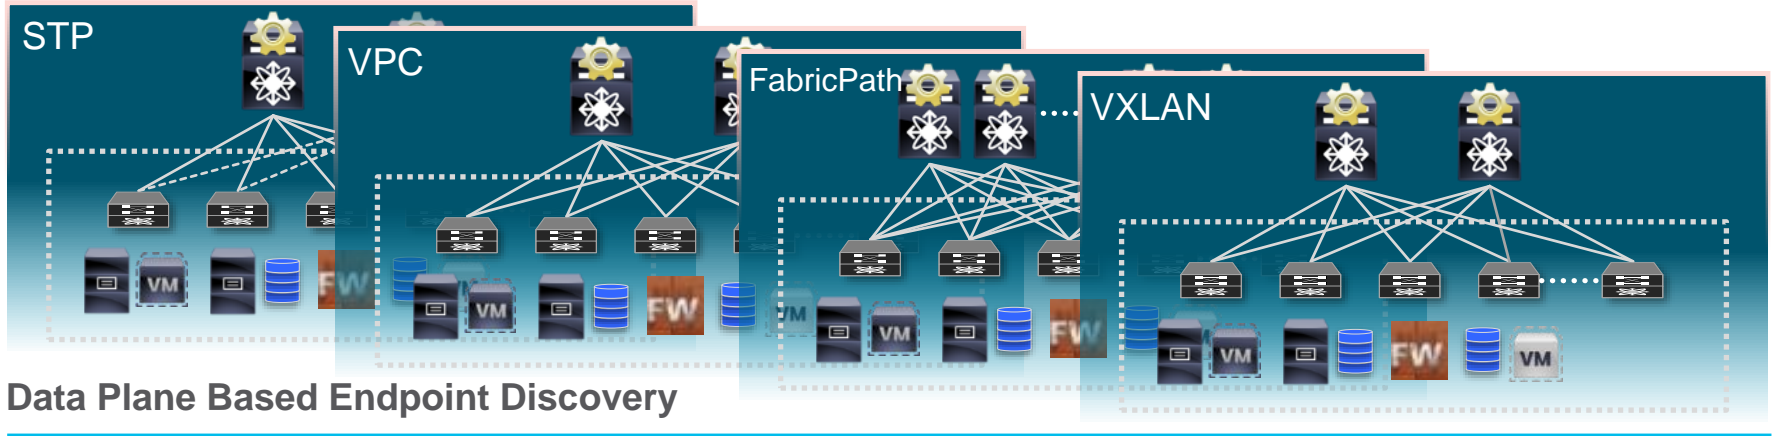
IBN X-Domain is a component of security

Complementary security elements





What is ACI?

Next Gen Forwarding & Networking




Control Plane Based Endpoint Location Tracking



ACI Anywhere


Extending the Reach of the New Network

ACI Single Pod Fabric



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

ACI Multi-Site




ACI 3.1/4.0 - Remote Leaf and vPod extends an Availability Zone (Fabric) to remote locations

ACI Multi-Cloud




ACI 1.0 - Leaf/Spine Single Pod Fabric

ACI Multi-Pod Fabric



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

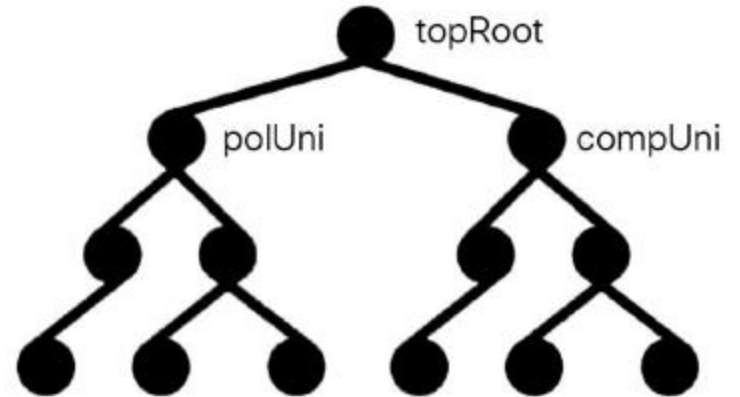
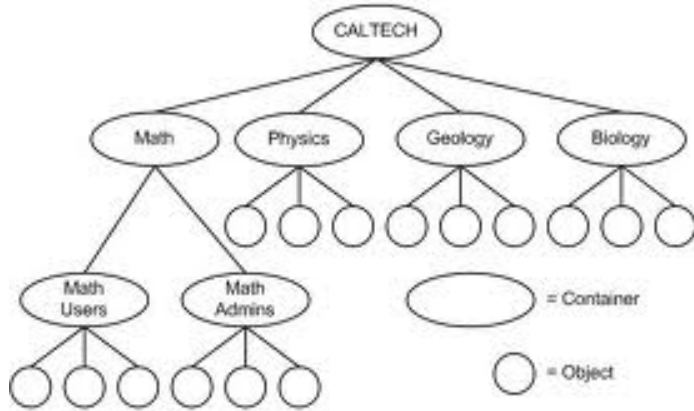
ACI Remote Leaf



ACI 4.1 - ACI Extensions to Multi-Cloud

Directory Enabled Networking

Data Base Defined Networking





Common Operational Properties - AD, LDAP, ...


System Management, Change Management, System Integrity, Correlation

APIC as an Operations Platform

More than just the Controller DataBase



Framework for Cisco + 3rd Party Apps Running Across all APIC cluster nodes





The DME Data Base runs on APIC servers (physical or virtual) and is replicated across APIC nodes (3 copies)

Dedicated “Service Engine” Nodes can be added to an APIC Cluster


Backup Servers for the APIC Cluster

APIC as an Operations Platform

More than just the Controller DataBase



Framework for Cisco + 3rd Party Apps Running Across all APIC cluster nodes



The DME Data Base runs on APIC servers (physical or virtual) and is replicated across APIC nodes (3 copies)

Dedicated “Service Engine” Nodes can be added to an APIC Cluster

Backup Servers for the APIC Cluster

Identity Based Networking

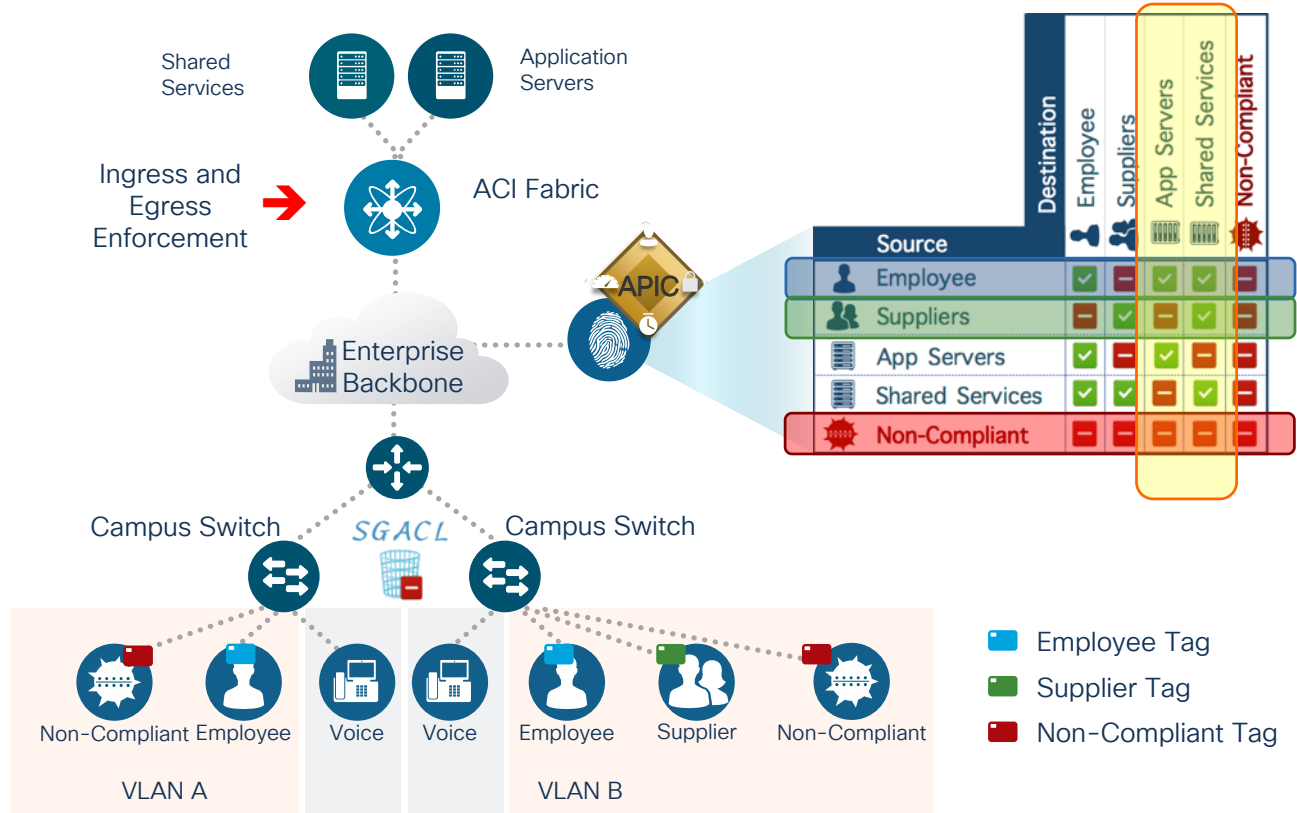


VTEP IP	Group Policy	VNID	Tenant Packet
---------	--------------	------	---------------

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

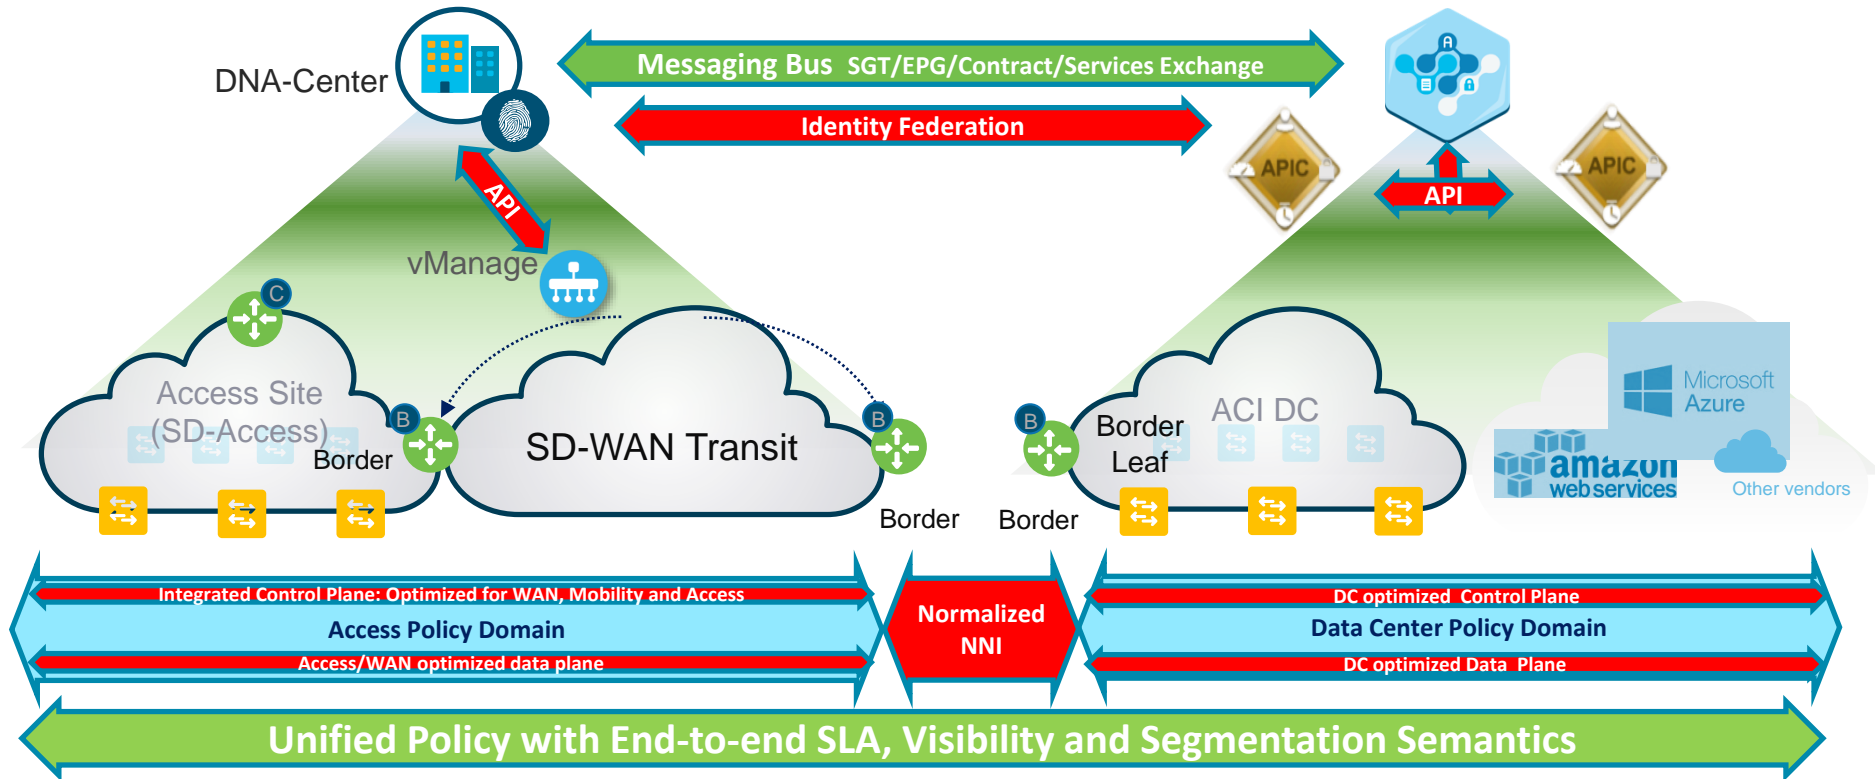
End Point Group Tags (EPG's, SGT's) are encoded in a VXLAN header

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



X-Domain Intent Based Networking

Enhanced Data Plane and Messaging BUS Services Plane



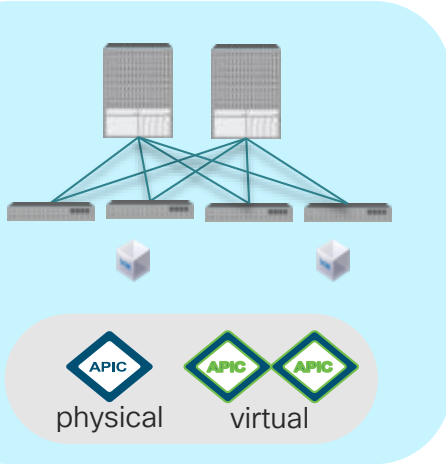
ACI Fabric Design

Agenda


- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

Single Fabric/Pod Topology Options


ACI mini



Standard (2-Tier)



Multi-Tier (3-Tier)



- From release 4.0
- For small deployment
- 1 APIC + 2 virtual APICs
- Up to 2 spine + 4 leaf switches
- Replacing virtual APICs with physical APICs
 - Expand to a full ACI fabric

- From the first release (1.0)
- Most popular and standard ACI topology

- From release 4.1
- For DCs with cabling restrictions

Not sure what to pick? Go with 2-Tier!


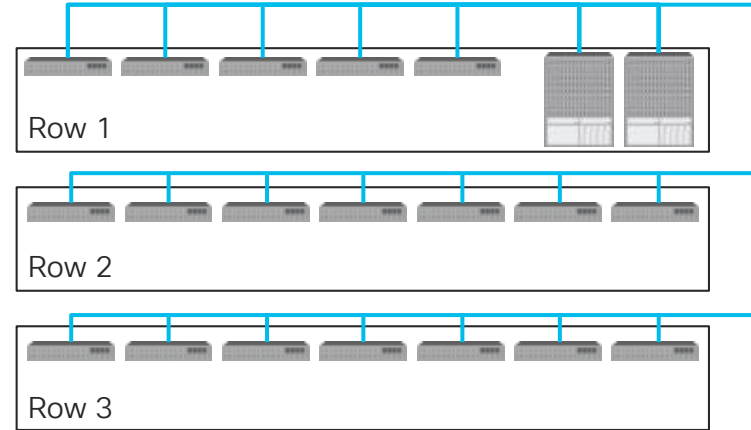
CISCO *Live!*

More about ACI Mini:


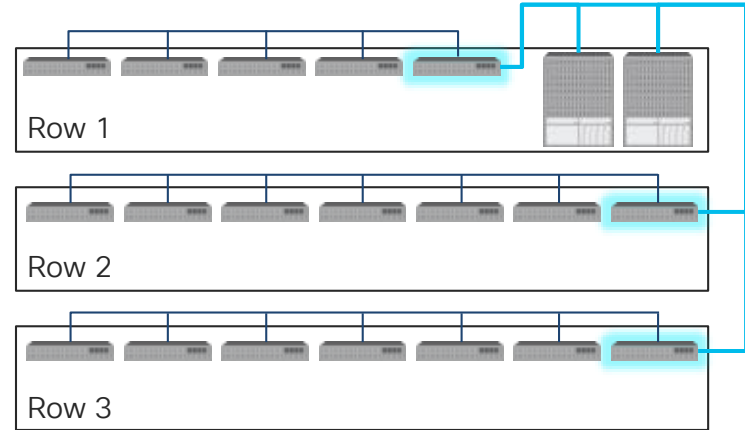
<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/Cisco-Mini-ACI-Fabric-and-Virtual-APICs.html>

Standard (2-Tier) vs. Multi-Tier (3-Tier)

Standard (2-Tier)



Multi-Tier (3-Tier)




Agenda

- Single Fabric/Pod Topology Options
- **ACI Design Considerations**
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

Network Centric & Application Centric Basic

Network Centric


- 1 VLAN = 1 EPG = 1 BD



- Similar to traditional network
 - VLAN as a broadcast domain
 - Easy to connect to legacy networks
- Typically simple or no contracts (no ACLs)

Application Centric

- Multiple EPGs per BD



- Multiple security domains (EPGs) in one broadcast domain
- Flexible network and security design

Details are later



Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - [ACI Fabric Bring Up](#)
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

ACI Fabric Bring Up

- CCO Document:
- [Setting Up an ACI Fabric: Initial Setup Configuration Example](#)
- https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/white_papers/Cisco-ACI-Initial-Deployment-Cookbook.html
- Breakouts:
- How to setup an ACI fabric from scratch - BRKACI-2004

Best Practice Global settings

- **Enforce Subnet Check Globally**
 - Prevent endpoint learning outside of BD subnets
- **Enforce EPG VLAN Validation**
 - Prevent overlapping VLAN pool in one EPG
- **Enable Domain Validation**
 - Enforce domain association to an EPG
 - This cannot be disabled.
You may have configurations that were working even if they were incorrect. Before enabling it make sure you verify the domain assignment to the EPGs and the associated AEPs.
- IP Aging Enabled
- Rogue EP Protection (more details later)
- MCP per VLAN enabled

The screenshot displays the Cisco ACI System Settings interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'Virtual Networking', and 'L4-L7 Serv'. Below this, a secondary navigation bar shows 'QuickStart', 'Dashboard', 'Controllers', 'System Settings', and 'Smart Licens'. The main content area is titled 'System Settings' and features a left-hand menu with various configuration options: Quota, APIC Connectivity Preferences, System Alias and Banners, System Response Time, Global AES Passphrase Encryption..., BD Enforced Exception List, Fabric Security, Control Plane MTU, Endpoint Controls, Fabric-Wide Settings, and Port Tracking. The 'Fabric-Wide Settings Policy' configuration panel is open on the right, showing a 'Properties' section with several checkboxes. A red box highlights three specific settings: 'Enforce Subnet Check' (checked), 'Enforce EPG VLAN Validation' (checked), and 'Enforce Domain Validation' (checked). Other visible settings include 'Disable Remote EP Learning' (unchecked), 'Enable Remote Leaf Direct Traffic Forwarding' (unchecked), 'Opflex Client Authentication' (checked), and 'Reallocate Gipo' (unchecked).

Best Practice Global settings

- Enforce Subnet Check Globally
 - Prevent endpoint learning outside of BD subnets
- Enforce EPG VLAN Validation
 - Prevent overlapping VLAN pool in one EPG
- Enable Domain Validation
 - Enforce domain association to an EPG
- **IP Aging Enabled**
 - Endpoint IPs tied to a MAC can age out individually
- Rogue EP Protection (more details later)
- MCP per VLAN enabled

The screenshot displays the Cisco ACI System Settings interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'Virtual Networking', and 'L4-L7 Services'. The main navigation menu on the left lists various settings, with 'Endpoint Controls' highlighted at the bottom. The main content area shows the 'Endpoint Controls' configuration page, with 'Ip Aging' highlighted in a red box. The 'Administrative State' is set to 'Enabled'.

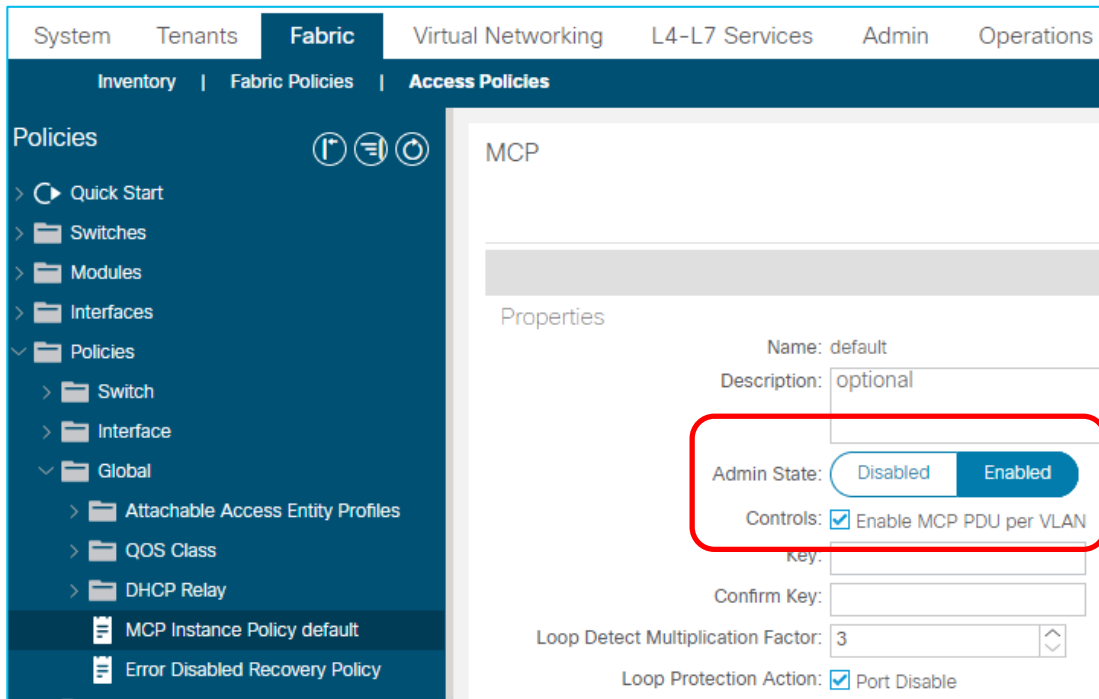
Best Practice Global settings

- Enforce Subnet Check Globally
 - Prevent endpoint learning outside of BD subnets
- Enforce EPG VLAN Validation
 - Prevent overlapping VLAN pool in one EPG
- Enable Domain Validation
 - Enforce domain association to an EPG
- IP Aging Enabled
 - Endpoint IPs tied to a MAC can age out individually
- **Rogue EP Protection (more details later)**
 - Mitigate impacts of endpoint flap issues
- MCP per VLAN enabled

The screenshot displays the Cisco ACI System Settings interface. The left-hand navigation pane shows the 'System Set' menu with 'Endpoint Controls' highlighted and circled in red. The main content area is titled 'Endpoint Controls' and features a sub-tab 'Rogue EP Control' also circled in red. Below this, the 'Properties' section is visible, showing the 'Administrative State' set to 'Disabled' (with 'Enabled' as an alternative), and three configuration fields: 'Rogue EP Detection Interval' set to 60, 'Rogue EP Detection Multiplication Factor' set to 4, and 'Hold Interval (sec)' set to 1800.

Best Practice Global settings

- Enforce Subnet Check Globally
 - Prevent endpoint learning outside of BD subnets
- Enforce EPG VLAN Validation
 - Prevent overlapping VLAN pool in one EPG
- Enable Domain Validation
 - Enforce domain association to an EPG
- IP Aging Enabled
 - Endpoint IPs tied to a MAC can age out individually
- Rogue EP Protection (more details later)
 - Mitigate impacts of endpoint flap issues
- MCP per VLAN enabled
 - Granular Mis Cabling Protection



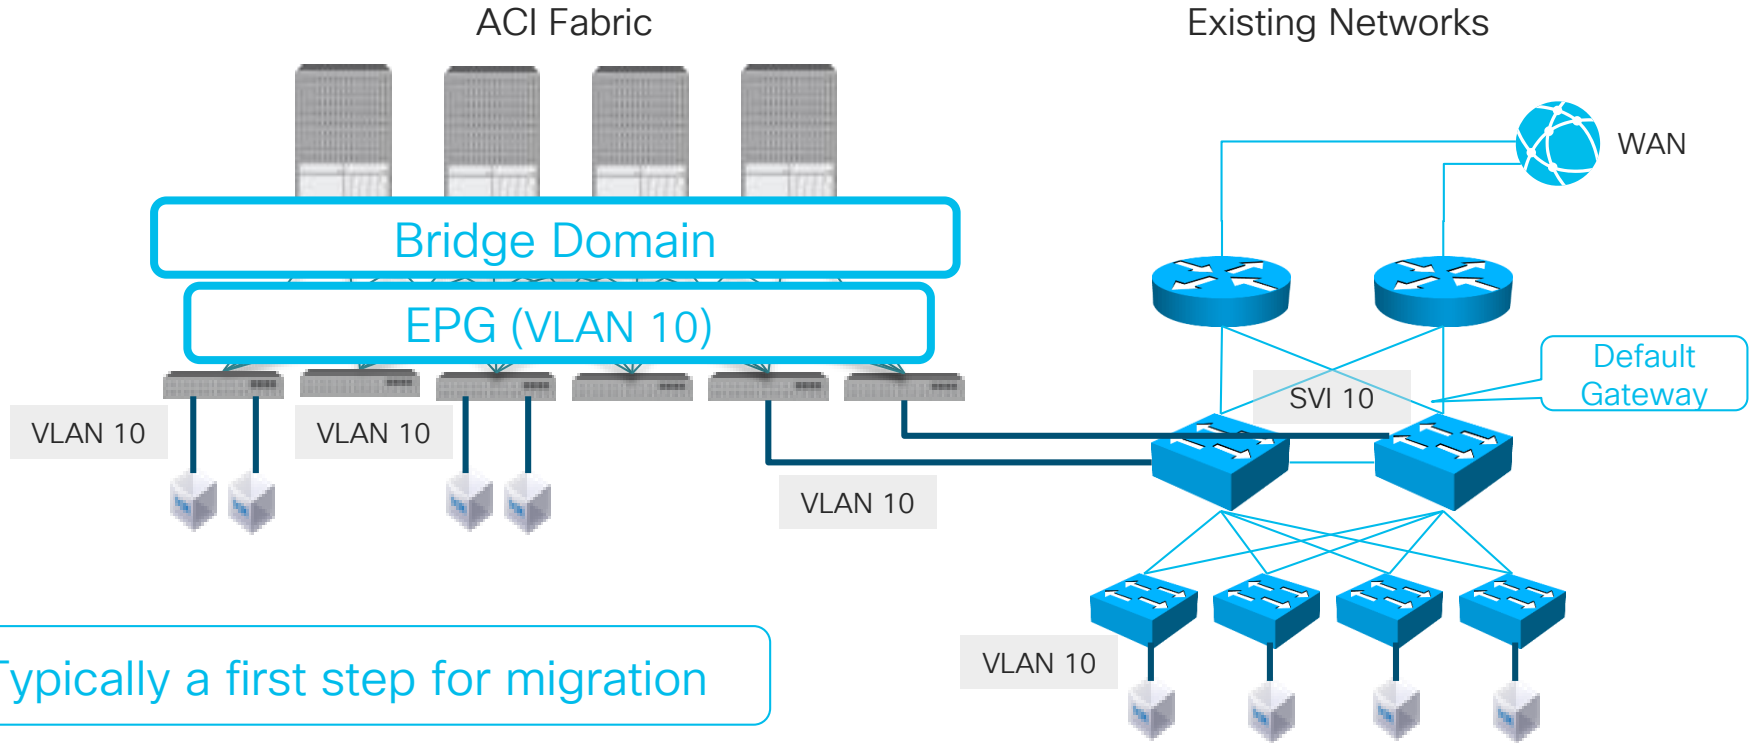
The screenshot shows the Cisco DNA Center interface for configuring Fabric Policies. The 'Fabric' tab is selected, and the 'Access Policies' section is active. The 'Policies' sidebar on the left shows a tree view with 'Global' expanded to 'MCP Instance Policy default'. The main content area displays the 'MCP' configuration page. The 'Admin State' is set to 'Enabled', which is highlighted with a red box. The 'Controls' section has 'Enable MCP PDU per VLAN' checked, also highlighted with a red box. Other visible settings include 'Name: default', 'Description: optional', 'Loop Detect Multiplication Factor: 3', and 'Loop Protection Action: Port Disable'.

Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - [L2 Connectivity to Existing Networks](#)
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization


Layer 2 Connectivity to Existing Networks

with Network Centric Design (1 BD = 1 VLAN)



Typically a first step for migration

Layer 2 BD mode recommendations

Rule of thumb 
Flood everything in L2 BD

The default gateway is
on external device



Disable Unicast Routing (= L2 BD)

Unicast Routing:

best practice when connecting with
external networks

Make ACI behave in the same way as
the external networks
=> Flood everything

L2 Unknown Unicast:	<input checked="" type="radio"/> Flood	<input type="radio"/> Hardware Proxy	
L3 Unknown Multicast Flooding:	<input checked="" type="radio"/> Flood	<input type="radio"/> Optimized Flood	
IPv6 L3 Unknown Multicast:	<input checked="" type="radio"/> Flood	<input type="radio"/> Optimized Flood	
Multi Destination Flooding:	<input checked="" type="radio"/> Flood in BD	<input type="radio"/> Drop	<input type="radio"/> Flood in Encapsulation
PIM:	<input type="checkbox"/>		
IGMP Policy:	select an option <input type="button" value="v"/>		
ARP Flooding:	<input checked="" type="checkbox"/>		

Optimize L2 Unknown Unicast (L2UU) in ACI BD

Flood mode

L2 Unknown Unicast: Flood Hardware Proxy

- Traffic with unknown destination MAC is flooded
 - Same as traditional switch

== Note ==

If there are any silent hosts or sensitive applications, use Flood just to be on the safe side.

Hardware Proxy mode

L2 Unknown Unicast: Flood Hardware Proxy

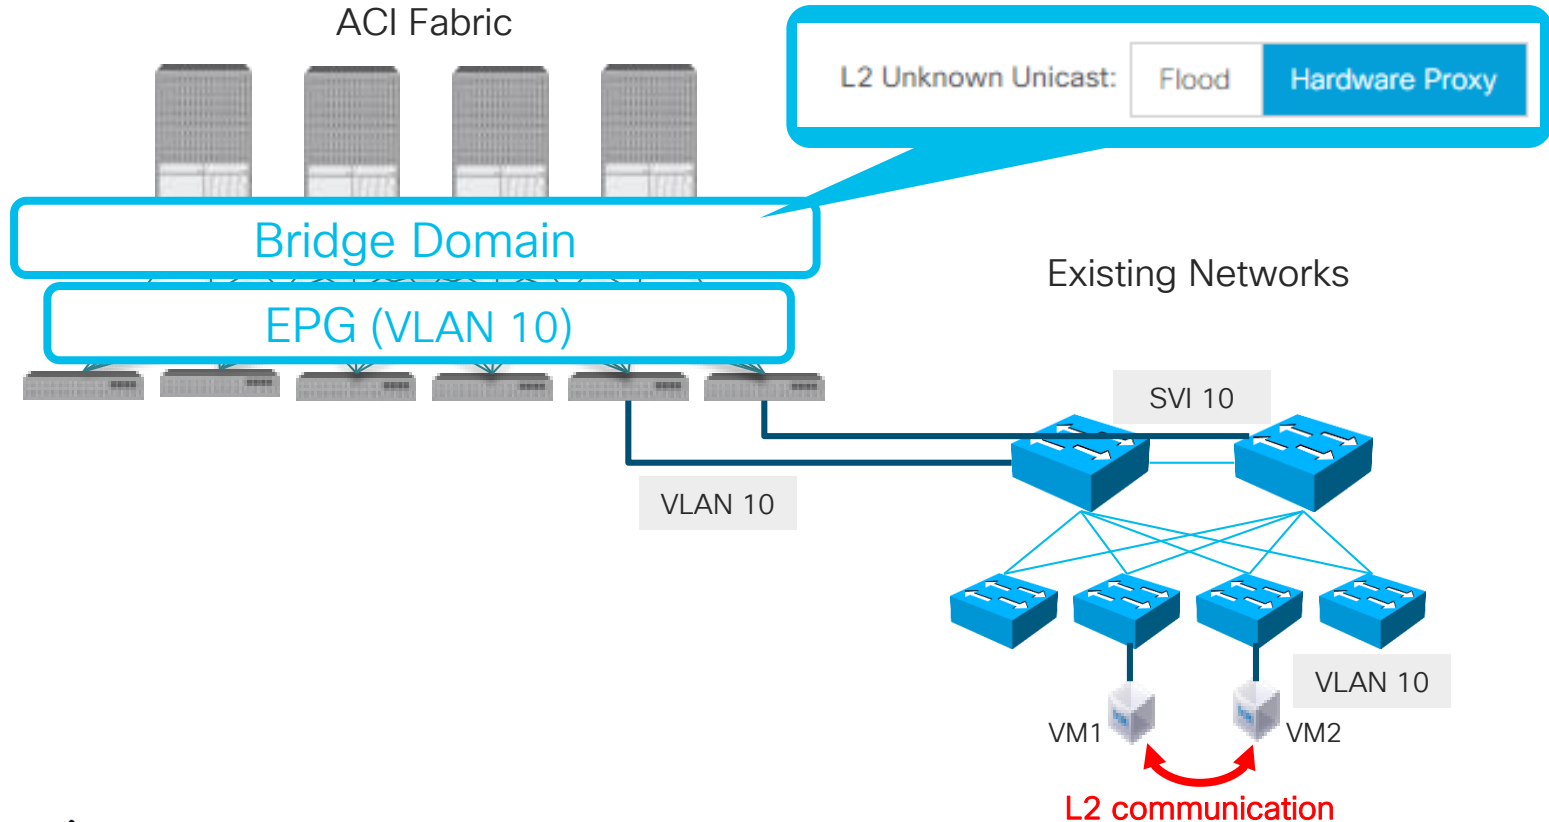
- Traffic with unknown destination MAC is sent to Spine Proxy.
 - Saves bandwidth
 - No unnecessary remote endpoint learning everywhere

== Caution ==

If the spine also does not know the MAC address, the packet is **dropped**. No silent host detection.

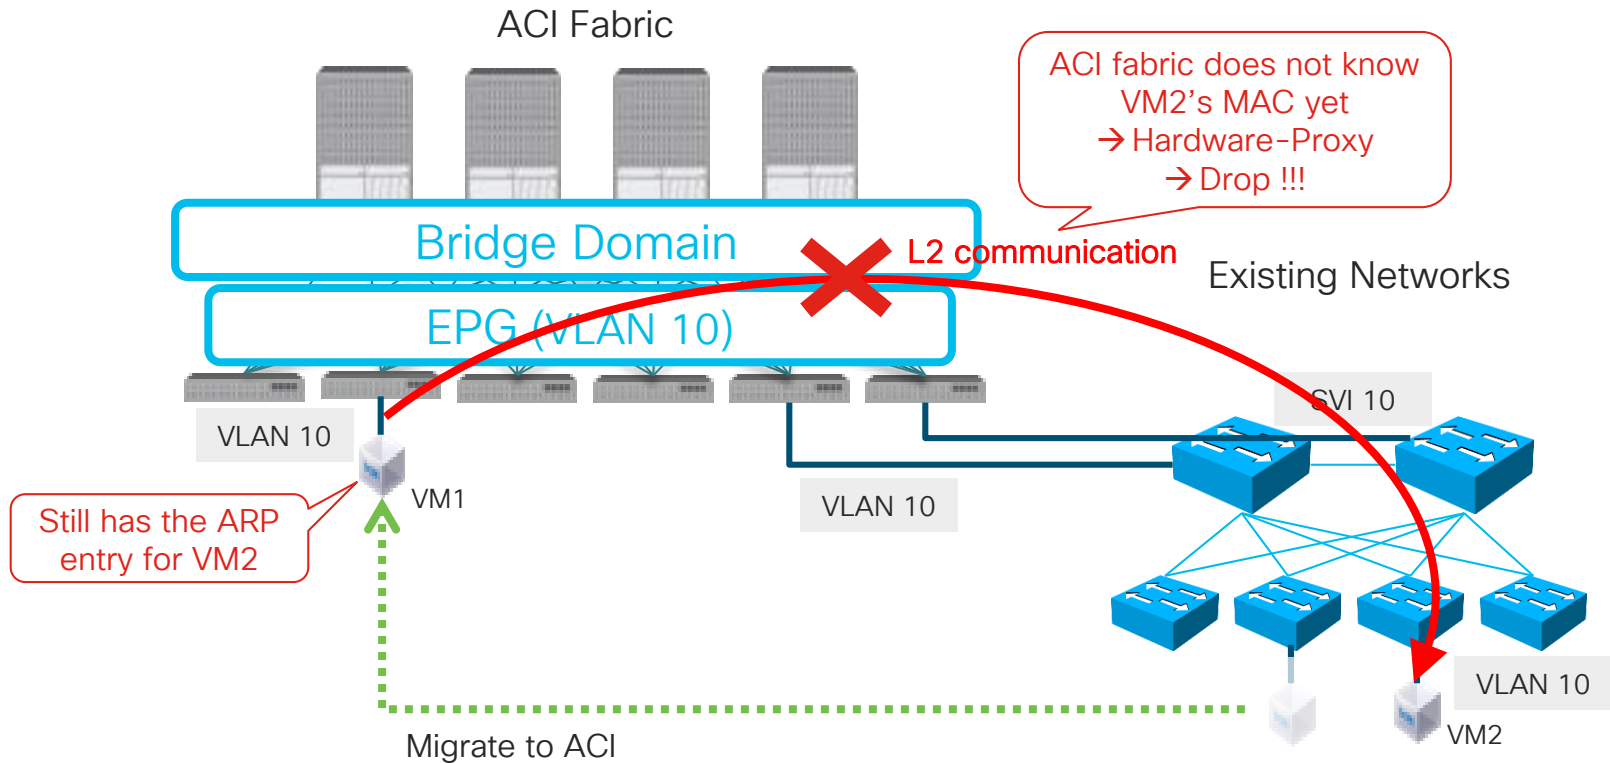
L2UU Hardware-Proxy drop example

Migration Scenario



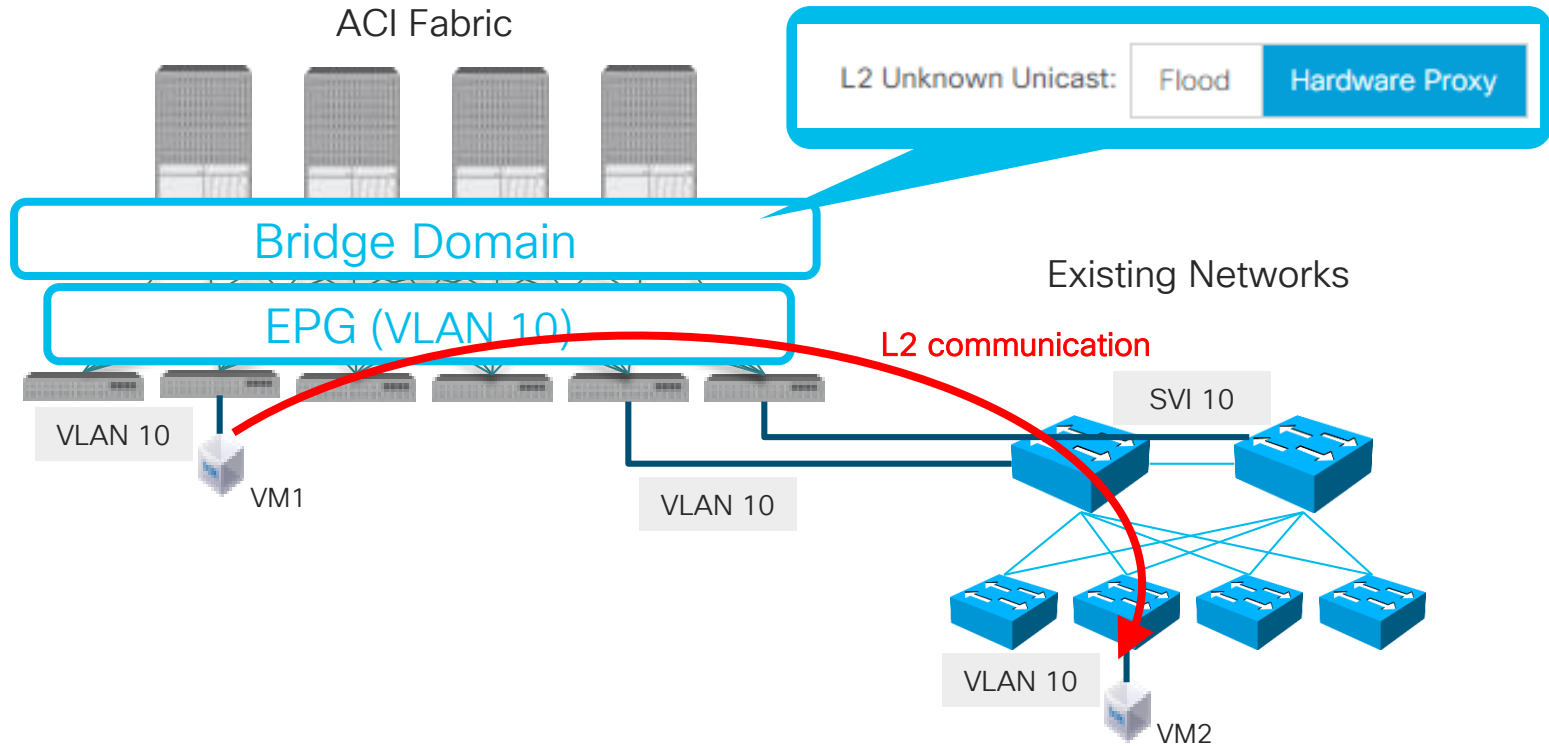
L2UU Hardware-Proxy drop example

Migration Scenario



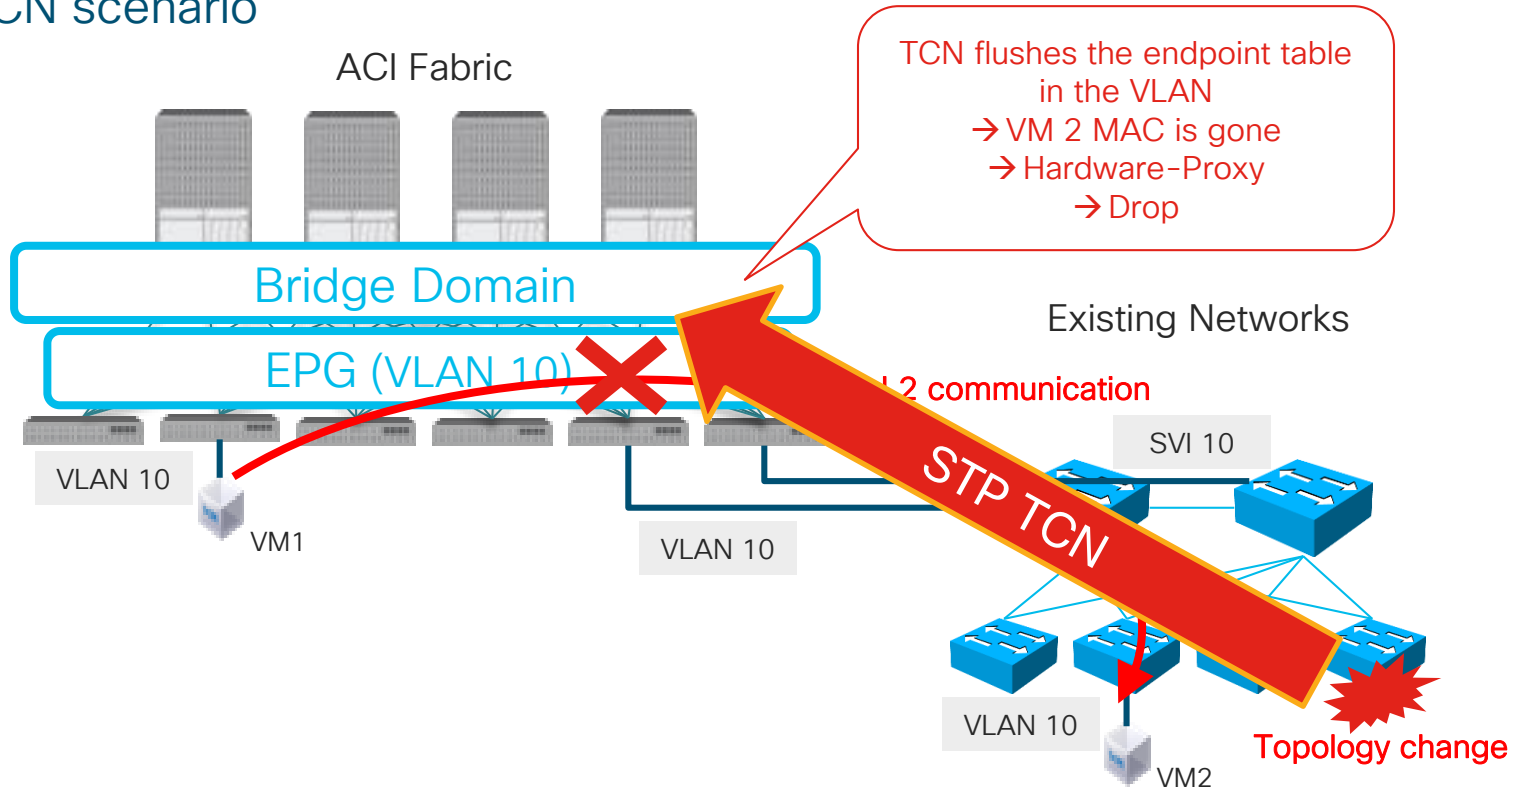
L2UU Hardware-Proxy drop example 2

STP TCN scenario




L2UU Hardware-Proxy drop example 2

STP TCN scenario



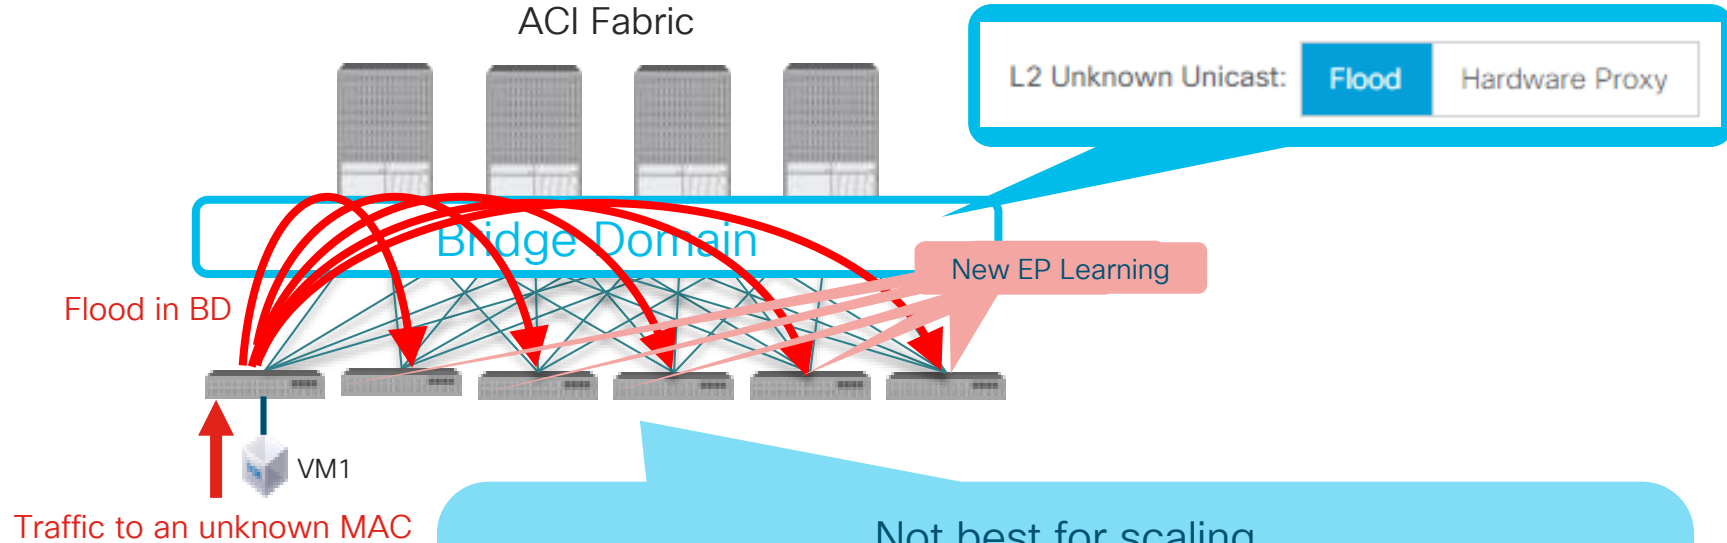
Rule of thumb (again)

Rule of thumb 

Flood everything

- in L2 BD
- when connecting external networks

Down side of L2UU Flood




Not best for scaling

1. Non-optimal bandwidth usage
2. All leaf switches learn remote endpoint
 - Easier to fill up the endpoint table
 - No more new remote endpoints mean another flood

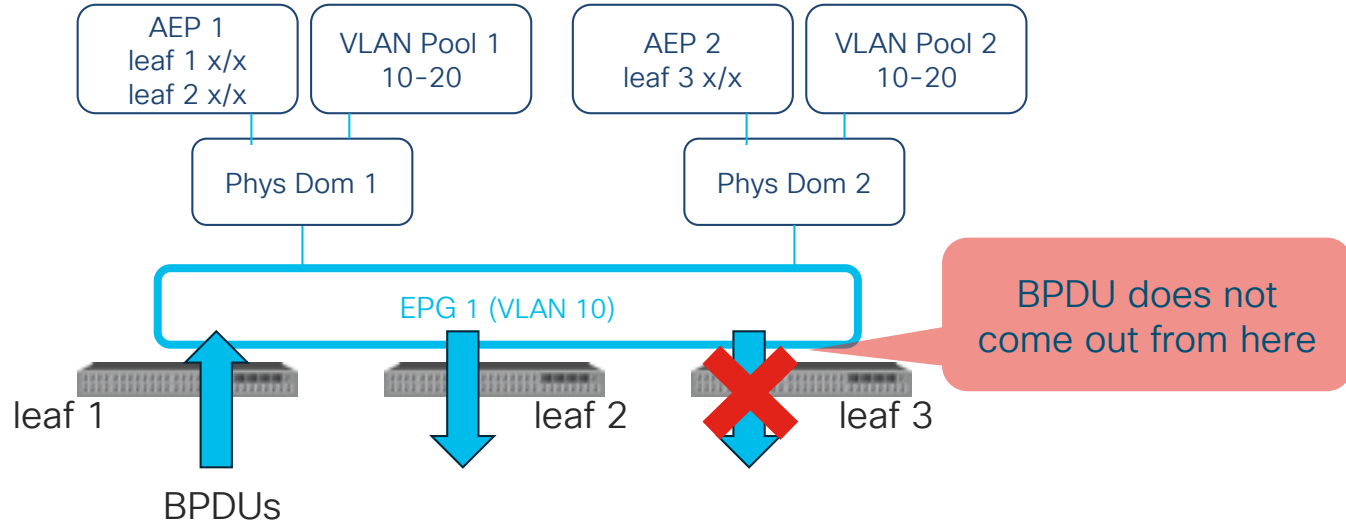
Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - [Loop prevention in ACI](#)
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

ACI forwarding BPDUs allows the external switches which run STP to prevent loops



STP flooding domain and ACI VLAN




Each encap VLAN is assigned a VxLAN ID based on a VLAN Pool.
This VxLAN ID is used to flood STP BPDUs.

- Different VLAN Pool = different STP flooding domain

STP flooding domain and ACI VLAN (cont.)

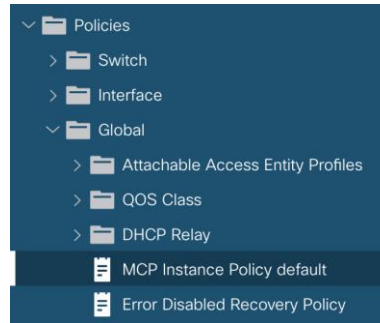
Recommended to create a domain and VLAN pool per pod

- Isolate STP domain per pod
- Prevent unnecessary BPDU, TCN propagations to other pods



Miscabling Protocol (MCP)

- After ports go up, MCP waits before sending MCP PDUs
- This is so that Spanning-Tree IF present can converge
- *If during that time there is a temporary loop*, Rogue EP will quarantine the IP/MAC of the hosts on the BD that are experiencing the loop.
- This protects the fabric from the effects of the temporary loop



Name: default

Description: optional

Admin State: Disabled Enabled

Controls: Enable MCP PDU per VLAN

Key:

Confirm Key:

Loop Detect Multiplication Factor:


Loop Protection Action: Port Disable

Initial Delay (sec):

Transmission Frequency (sec): (msec):

If a loop is occurring endpoints would be continuously flapping

- If a Loop occur like in this picture the host MAC keeps flapping between leafs like leaf1 and leaf4
- ACI has multiple ways to detect loops with this type of scenario
- ACI can detect loops from EP moves and it can disable BD learning when too many EP moves occur, or it can disable the ports where the move is happening or it can quarantine the specific EP that are flapping.



EP move dampening, EP Loop Protection, Rogue EP Detection

	EP move dampening	EP Loop Protection	Rogue EP Detection
Scope	per BD	Global	Global
Detection	aggregate number of all moves per BD in a second	number of moves of an individual endpoint between two ports	number of moves of an individual endpoint in the specified interval
Detects MAC and/or IP move	Detects MAC moves, IP moves	Detects MAC moves	Detects MAC moves, IP moves
Possible actions	BD learn disable per leaf	port disable or BD learn disable per leaf	Programs static entry to disable learning for the specific entry

Hardening the ACI fabric to reduce the chance and/or impact of Layer 2 Loops

MCP globally enabled

 MCP Instance Policy default

Admin State: Disabled Enabled

Rogue EP Control


Properties

Administrative State: Disabled Enabled

Rogue EP Detection Interval: 30

Rogue EP Detection Multiplication Factor: 9

Hold Interval (sec): 600

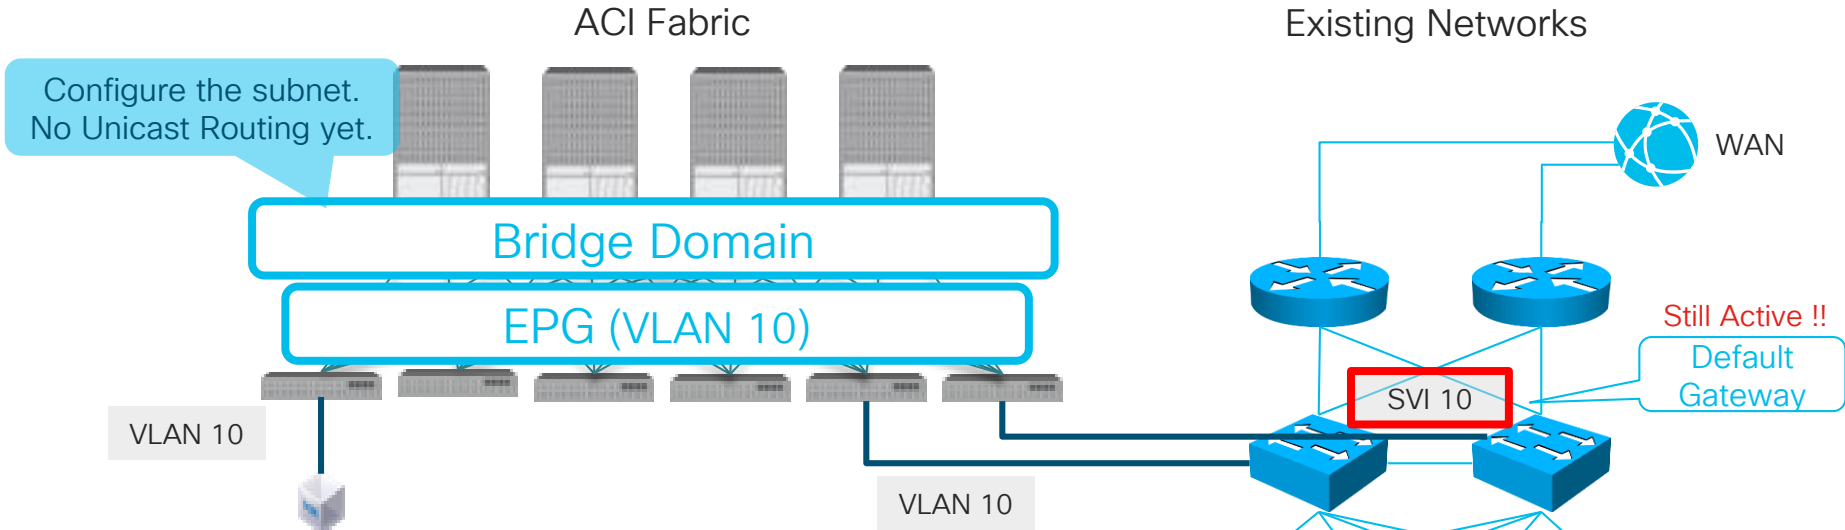


Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - [Moving the default gateway to ACI](#)
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

Migrate the default gateway to ACI

with Network Centric Design (1 BD = 1 VLAN)



Unicast Routing:

Operational Value for Unicast Routing: false

Custom MAC Address: 00:22:BD:F8:19:FF

Virtual MAC Address: Not Configured

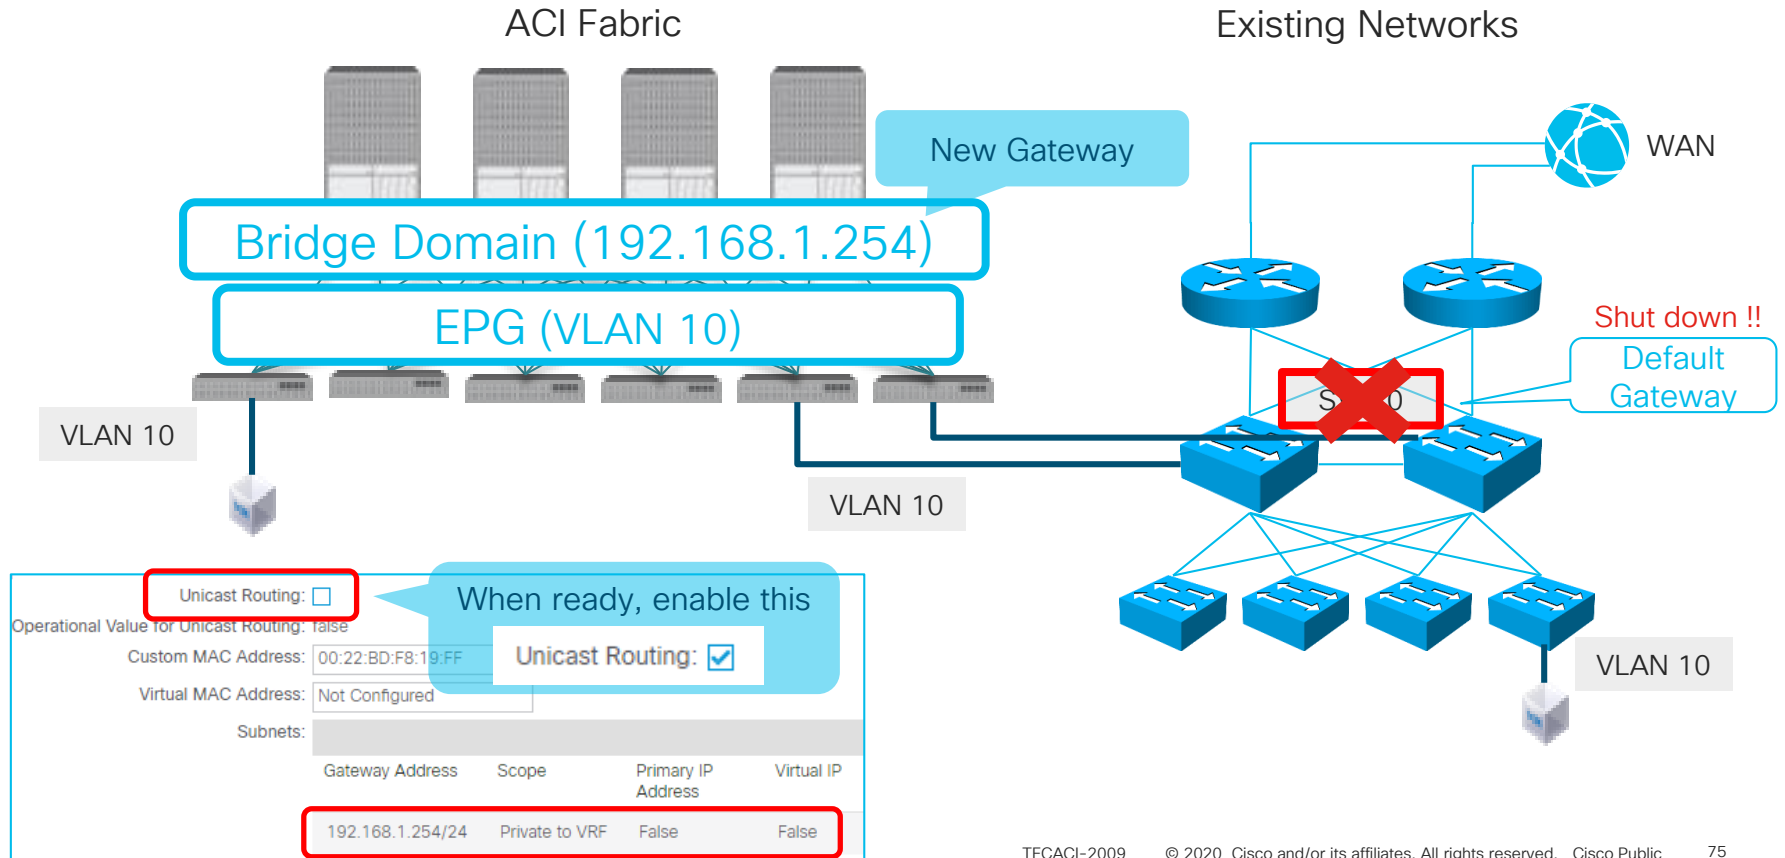
Subnets:

Gateway Address	Scope	Primary IP Address	Virtual IP
192.168.1.254/24	Private to VRF	False	False

Change the MAC to the current gateway MAC for seamless switchover

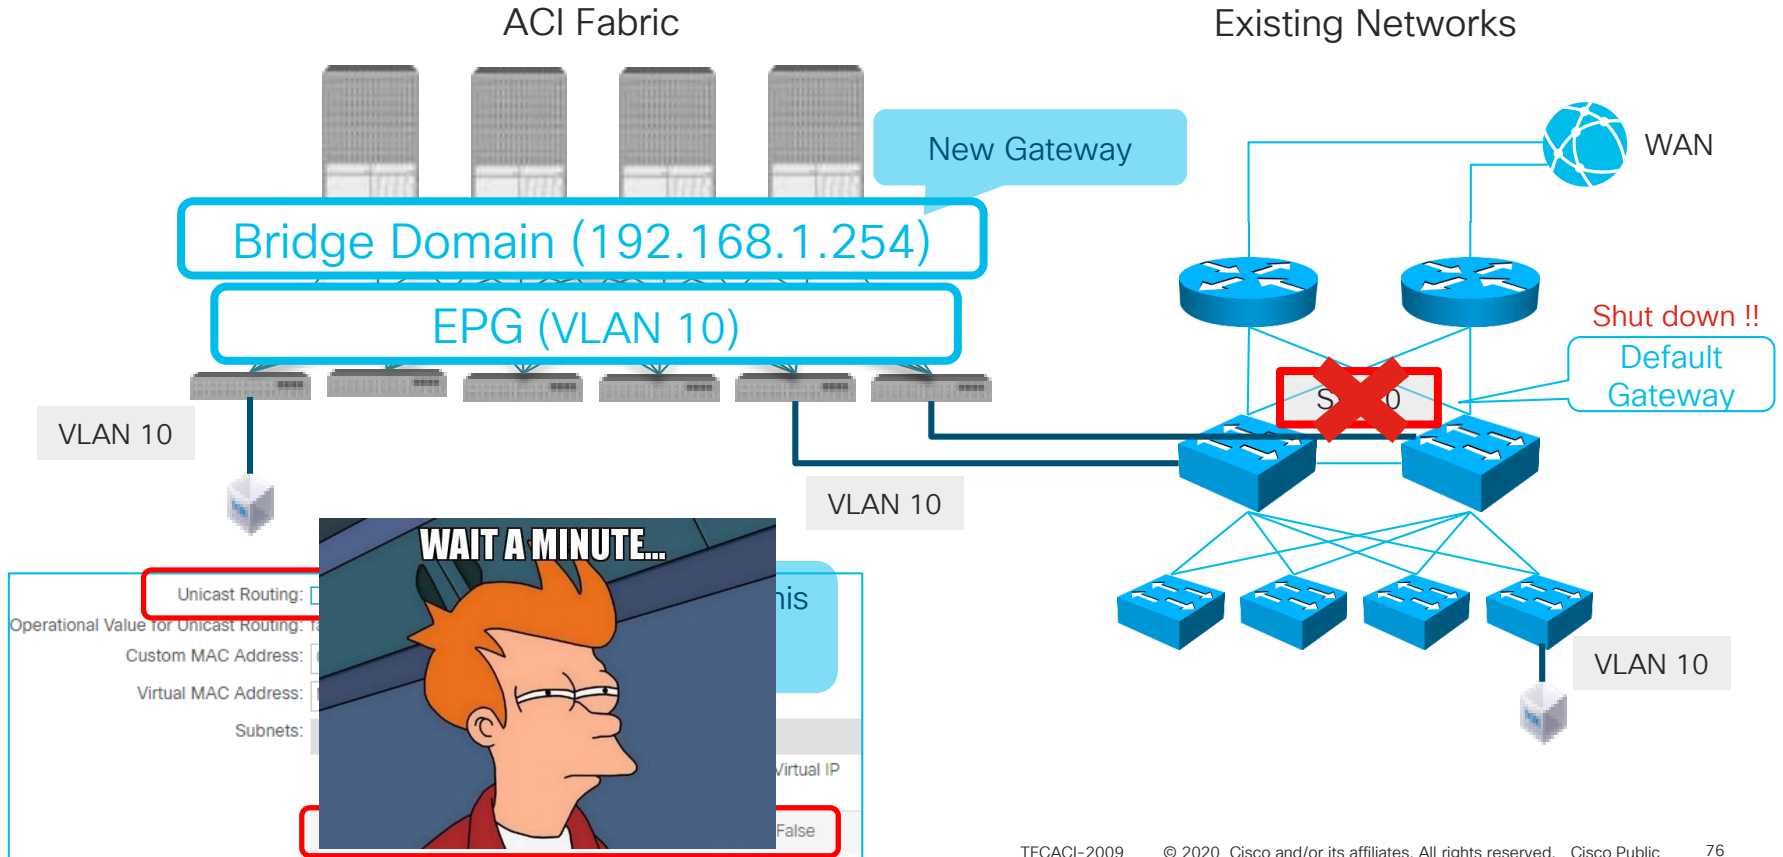
Migrate the default gateway to ACI

with Network Centric Design (1 BD = 1 VLAN)



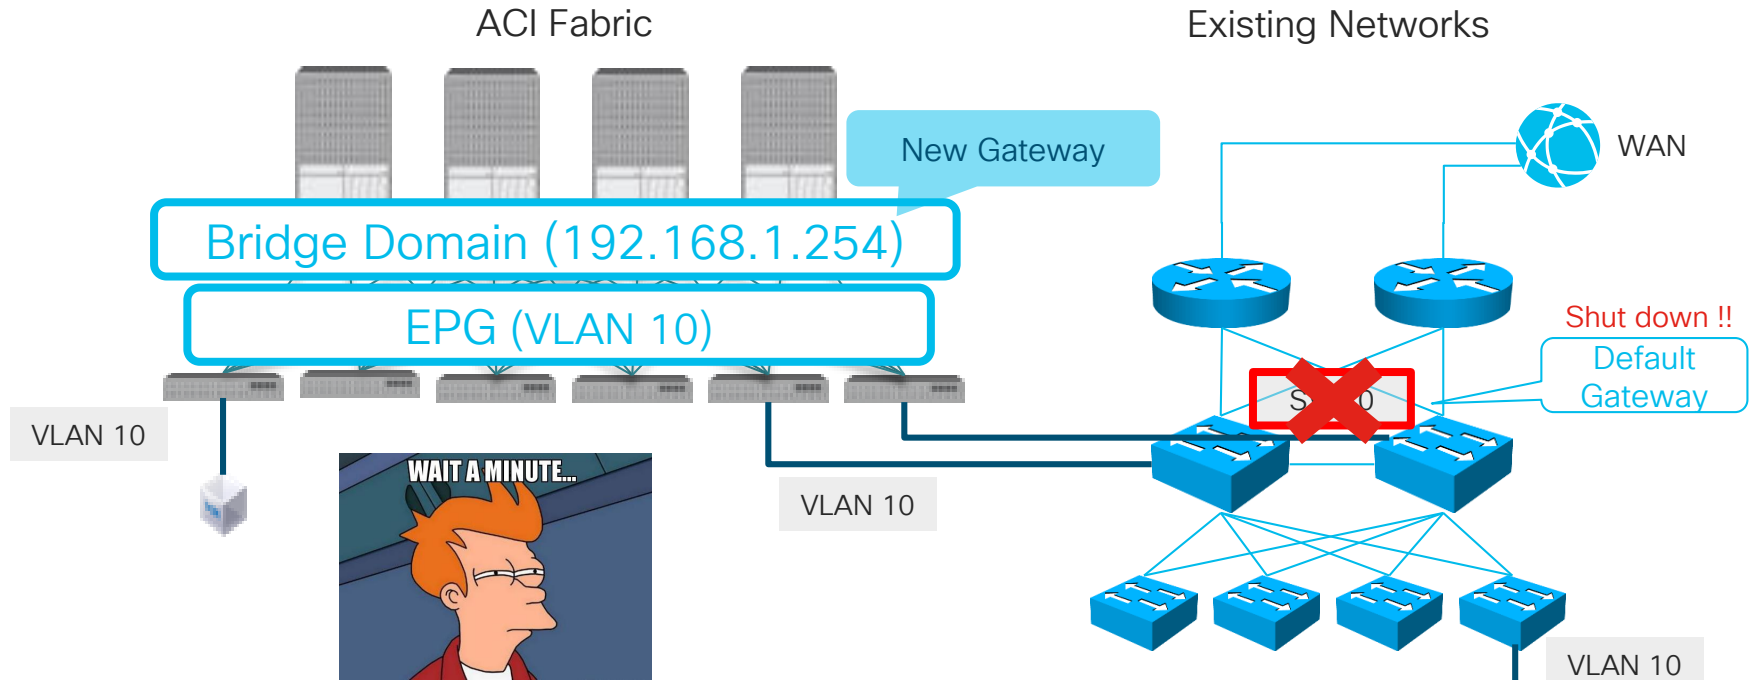
Migrate the default gateway to ACI

with Network Centric Design (1 BD = 1 VLAN)



Migrate the default gateway to ACI

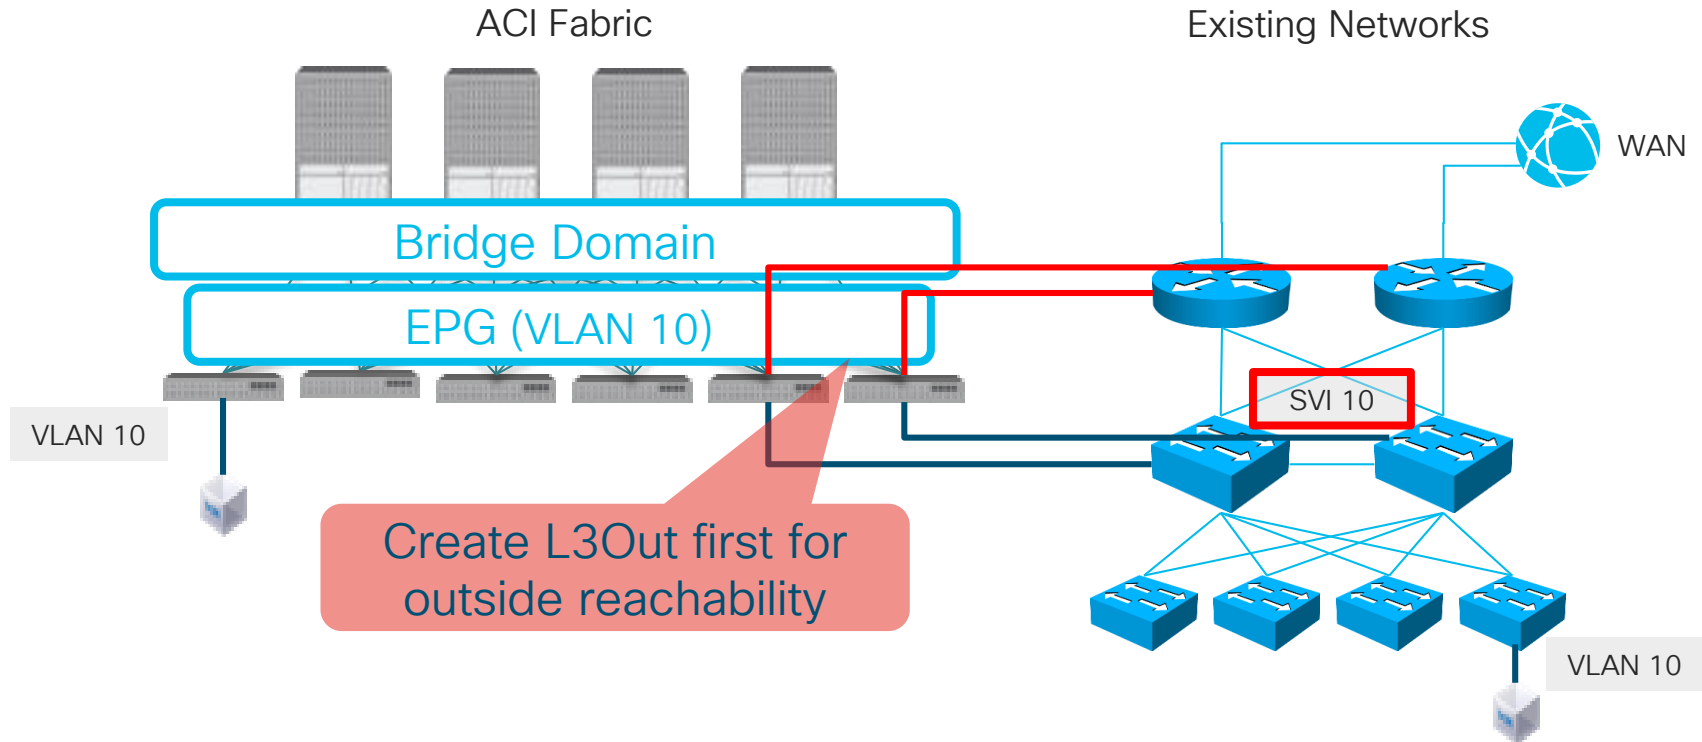
with Network Centric Design (1 BD = 1 VLAN)



how am I supposed to go out to WAN now?

Migrate the default gateway to ACI

with Network Centric Design (1 BD = 1 VLAN)



Enabling Unicast Routing enables Endpoint IP learning as well


Keep in mind:

“ACI does data-plane learning for IP as well”

Bridging Traffic → MAC learning

Routing Traffic → MAC & IP learning

ARP Request → MAC & IP learning



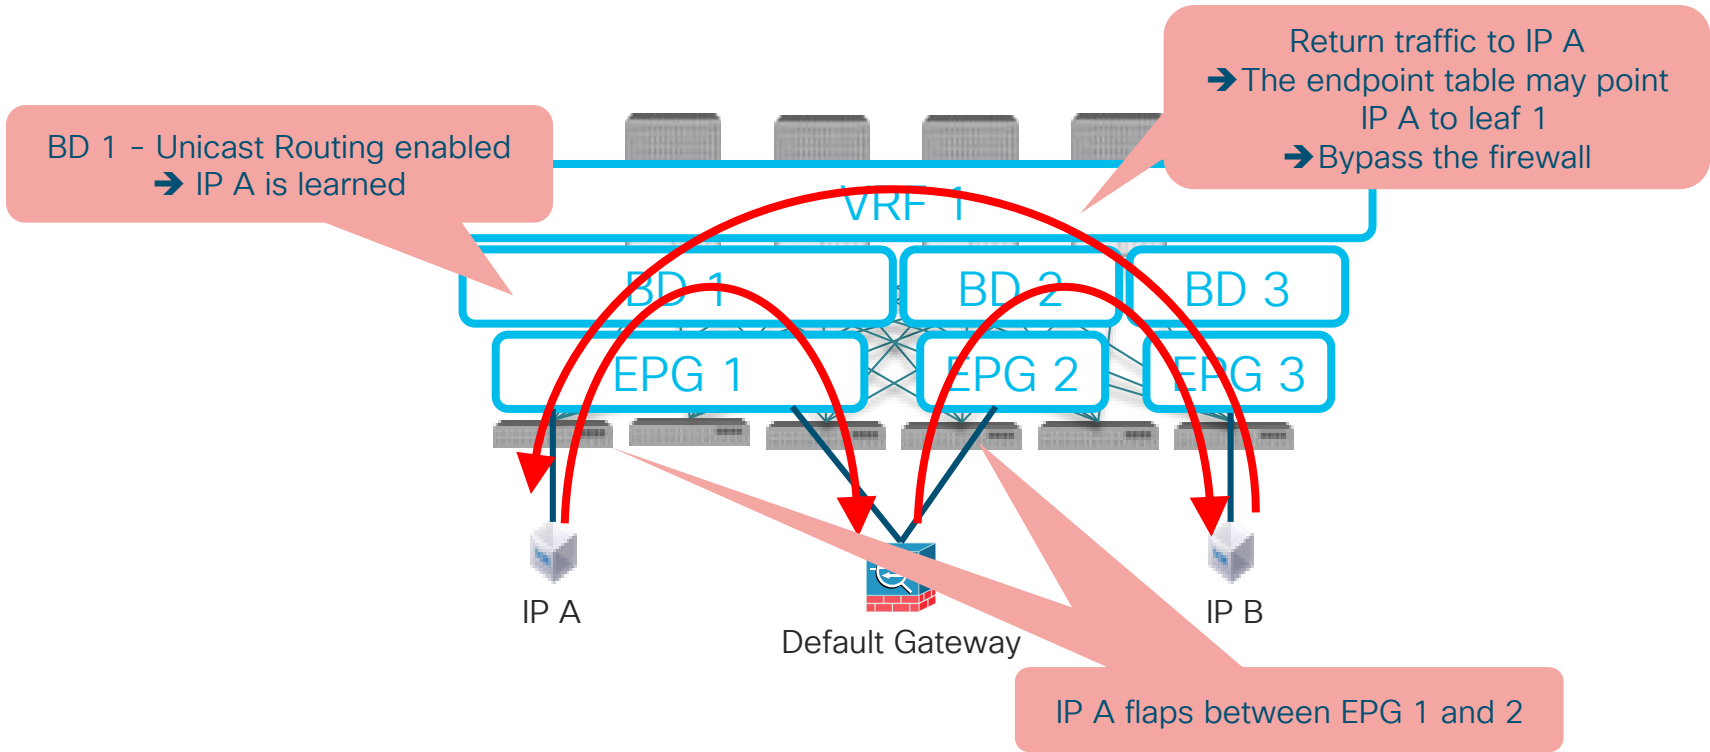
Does your NIC teaming work well with ACI?

If you have servers with Active/Active teaming
=> Tuning may be required => more on this later

Make sure you understand how endpoint learning works:

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>

Why not Unicast Routing on L2 BD?



Caution - Just one example of many other corner cases

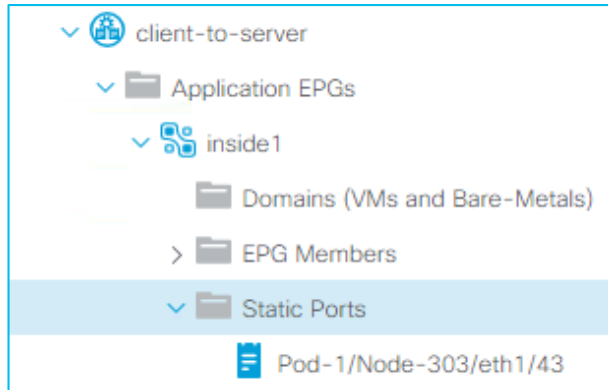
Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - [Connecting servers \(Physical, VMM Integration, UCSM Integration\)](#)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

Two Server Connectivity options

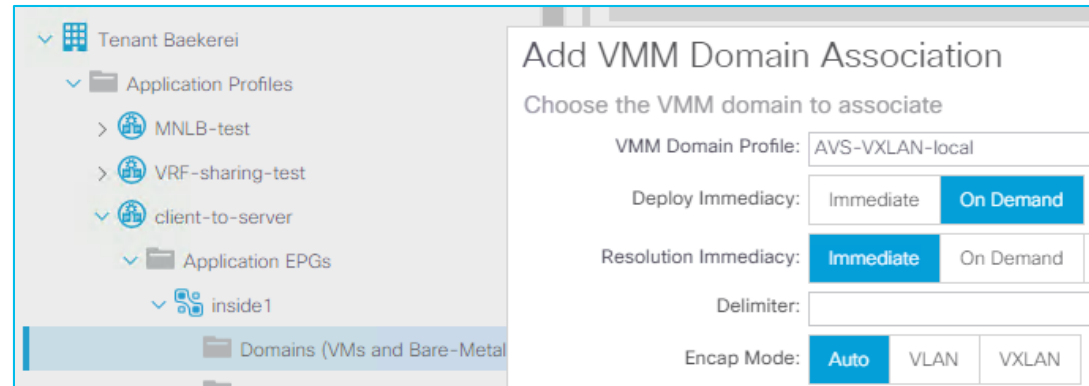
Manual (Static Path Binding)

- Associate a physical domain to an EPG
- Manually assign a VLAN (EPG) on a leaf interface



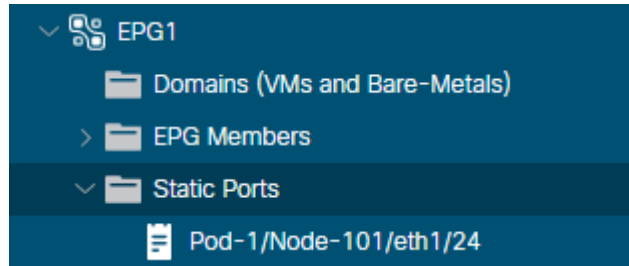
VMM Integration (Dynamic Path Binding)

- Associate a VMM domain to an EPG.
- ACI-VMM dynamically figures out a leaf interface and VLAN via CDP/LLDP/opflex



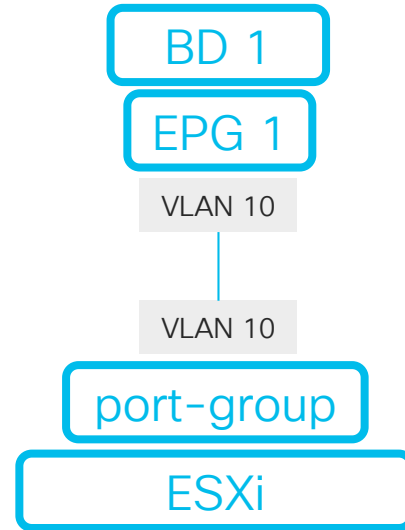
Physical Domain (physdom)

Can be used for both physical and virtual servers.



Configure Manually

Network Admin



Server Admin



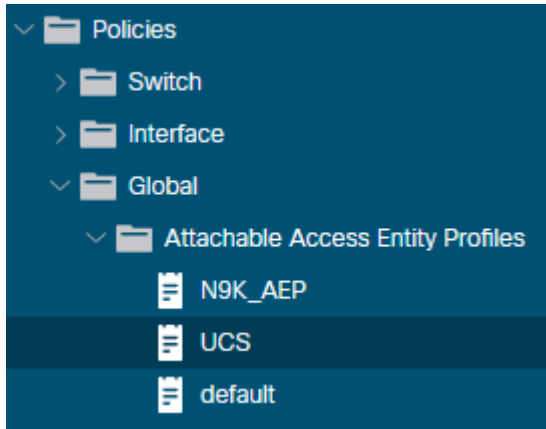
This can be tedious and error prone...

Bulk Static Path Binding via AEP

Assign VLAN 10 as EPG 1 on all interfaces in AEP UCS

- Design AEP (group of interfaces) carefully

AEP



Tenant




Application profile



EPG

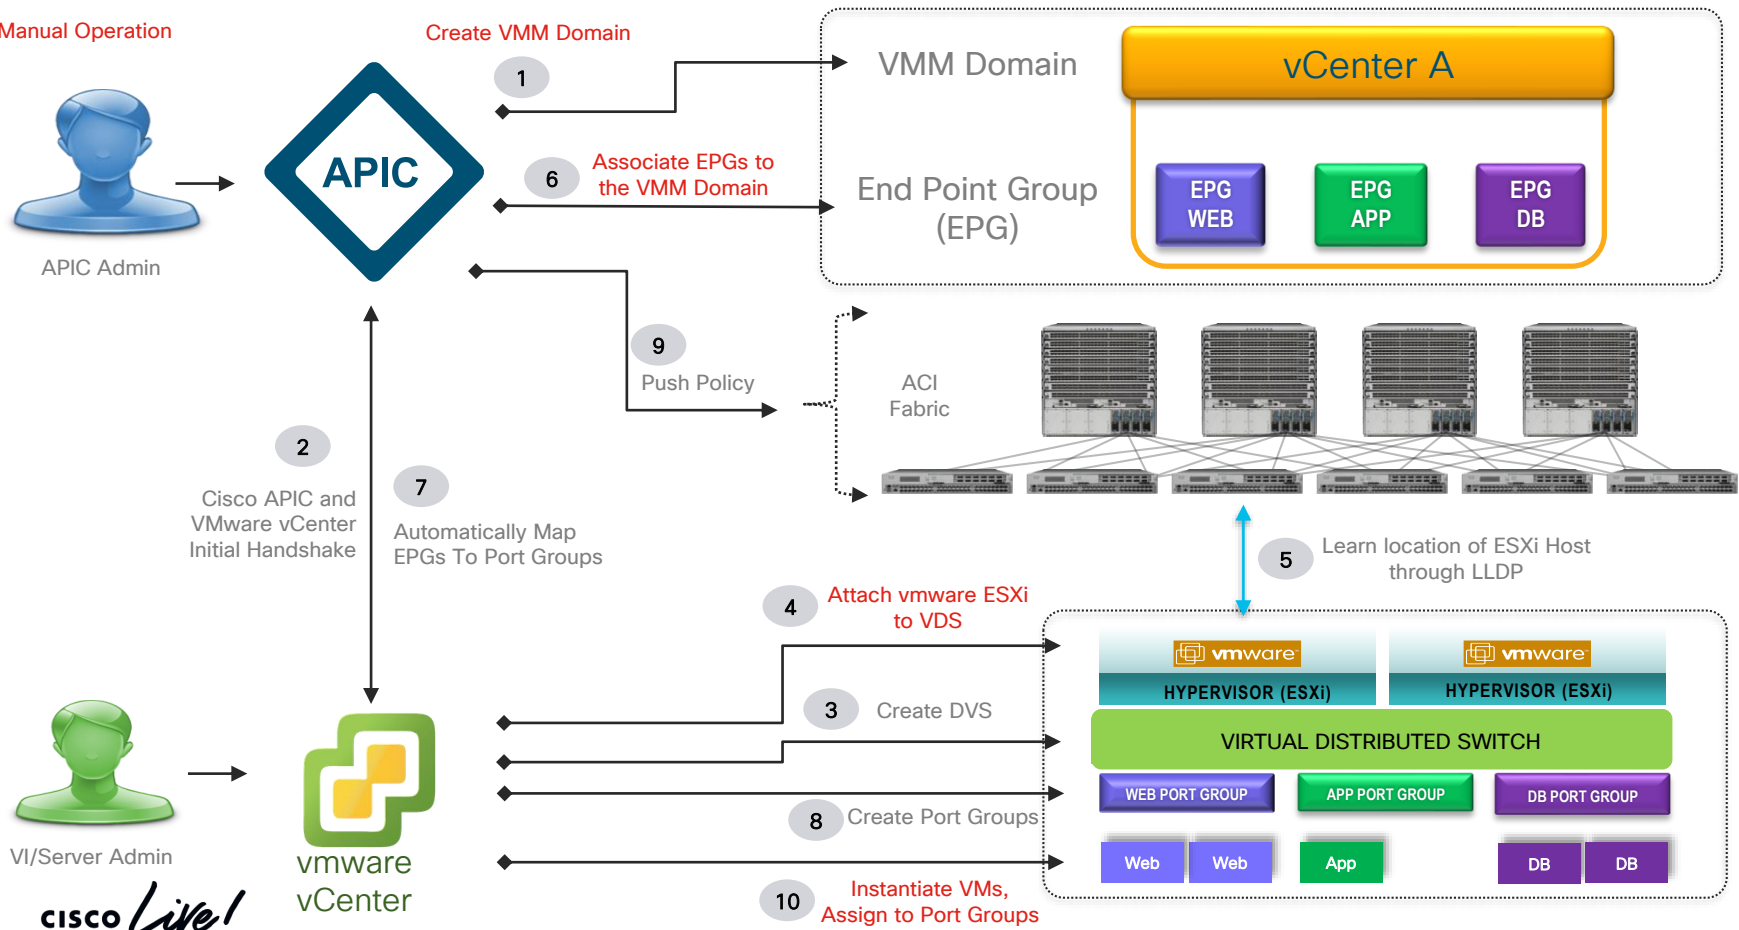


VLAN



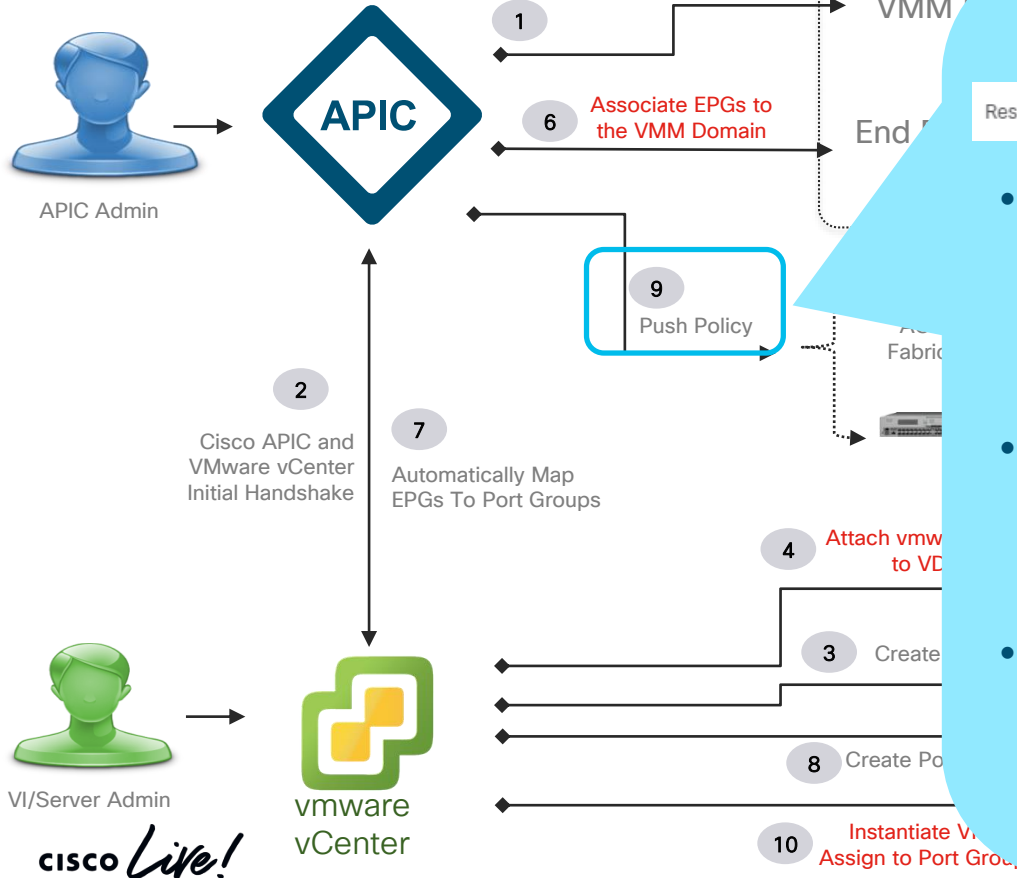
VMM Domain (vmware vCenter example)

Red - Manual Operation



VMM Domain (vmware vCenter example)

Red - Manual Operation




Resolution Immediacy

Resolution Immediacy: Immediate On Demand Pre-provision

- Immediate
When the presence of ESXi on VMM integrated DVS is confirmed via LLDP/CDP
- On Demand
When a VM is attached to the port-group via vCenter
- Pre-provision (recommended)
When the VMM domain is associated to the EPG
Similar to Bulk Static Path Binding via AEP

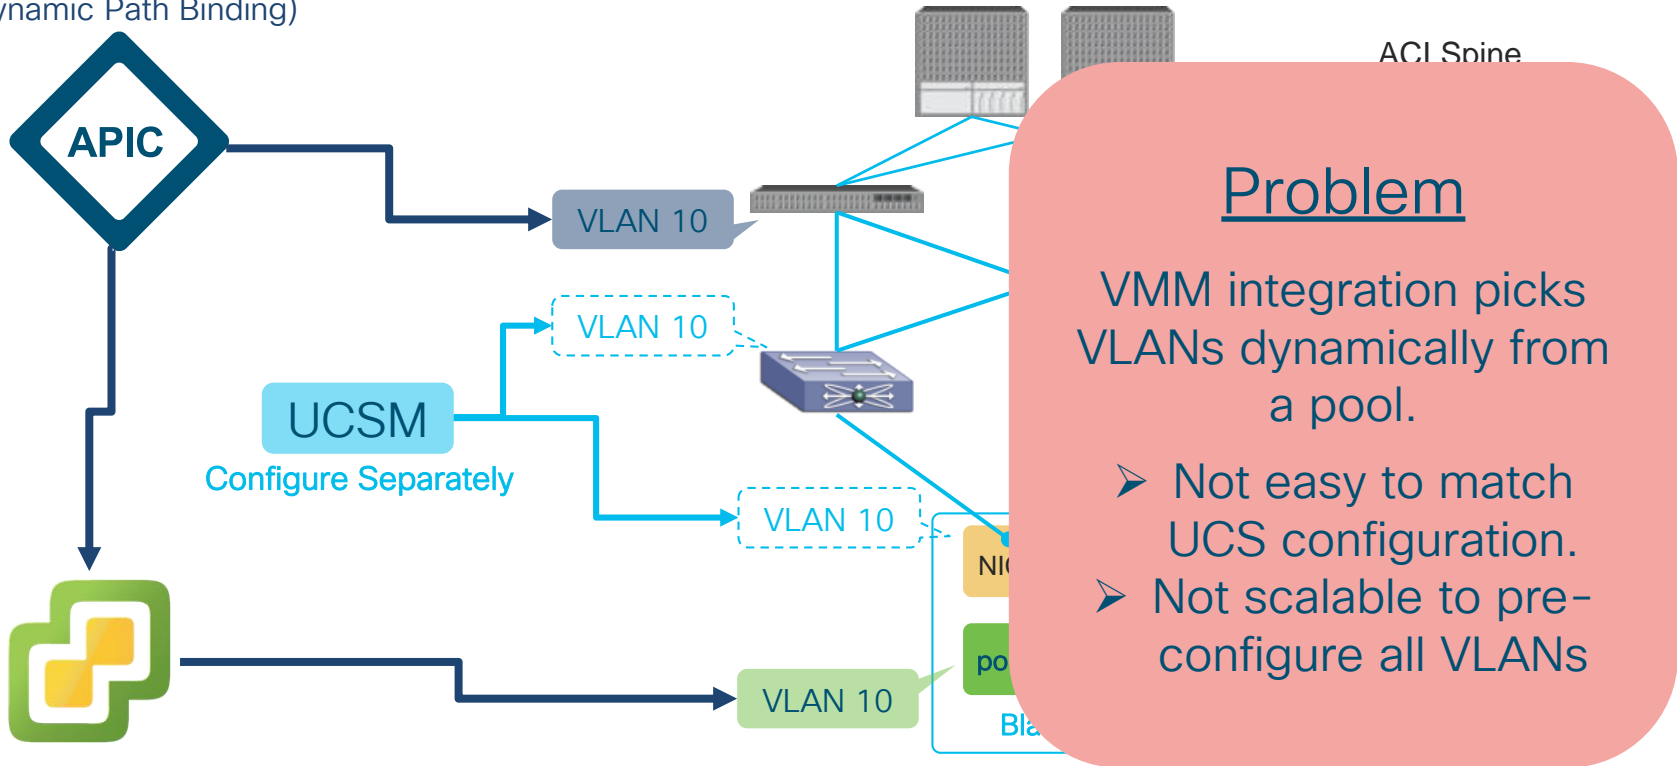
ACI and UCS Fabric Interconnect (UCSM)

VMM Integration
(Dynamic Path Binding)



ACI and UCS Fabric Interconnect (UCSM)

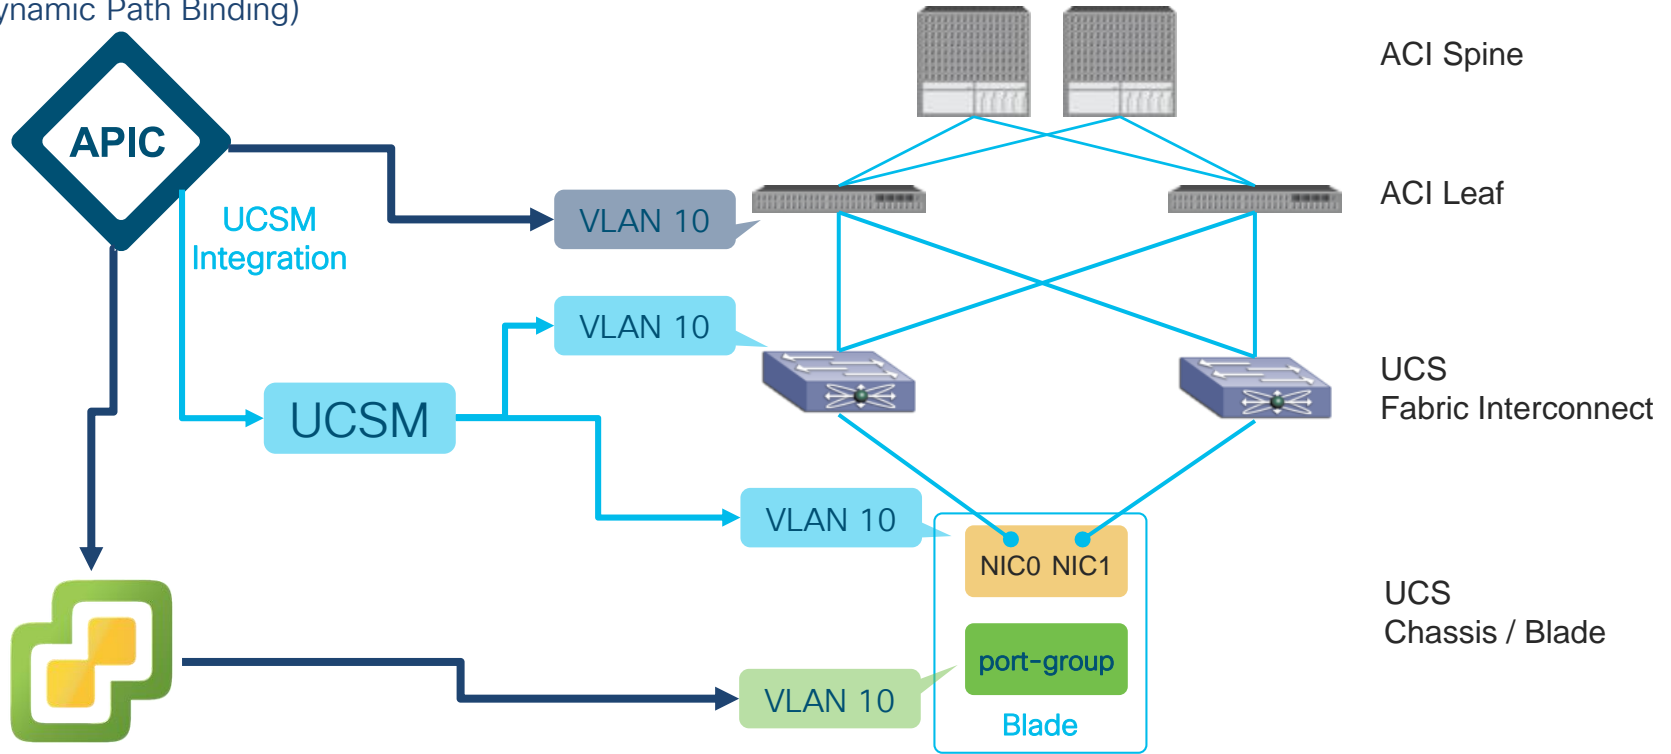
VMM Integration
(Dynamic Path Binding)



ACI and UCSM Integration

From ACI 4.1

VMM Integration
(Dynamic Path Binding)



UCS FI has a range of reserved VLANs

Quoting the UCSM Network Management Guide:

“ ”

You cannot create VLANs with IDs from 4030 to 4047 or from 4093 to 4095. These ranges of VLAN IDs are reserved.

“ ”

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Network-Mgmt/3-2/b_UCSM_Network_Mgmt_Guide_3_2/b_UCSM_Network_Mgmt_Guide_3_2_chapter_0101.html

ACI (UCSM Integration) does not verify if the VMM domain uses those reserved VLANs.

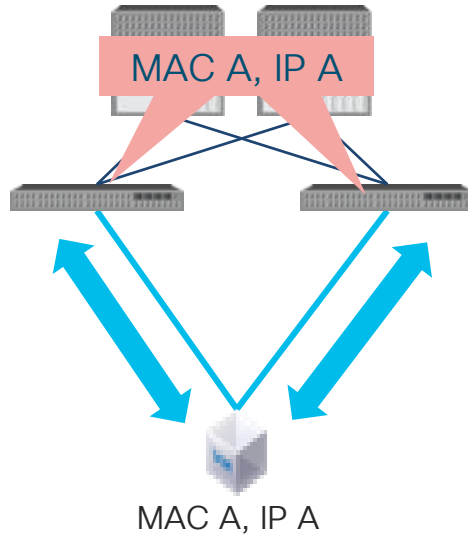
➔ Ensure not to use such VLANs in your pool.

Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - [Teaming Options](#)
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

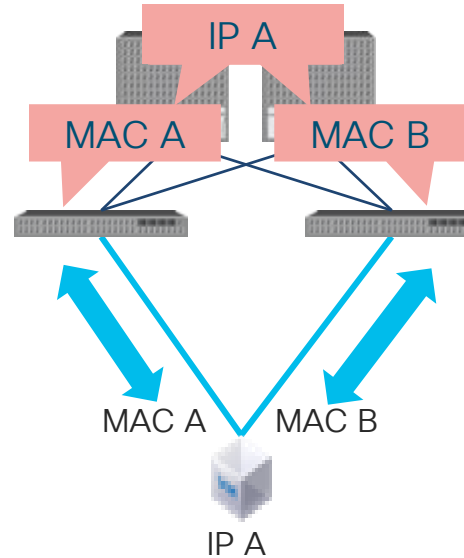
NIC Teaming Issues with ACI

MAC/IP flap
between nodes/ports



Same problem as any
other switch

IP flap
between MAC




ACI specific due to IP
data-plane learning

General NIC Teaming Types

All three work well with ACI

Microsoft - Switch Independent Hyper-V
VMware - Route Based on Originating Port ID


Active / Standby



One uplink is used for all VMs
Another uplink is only for failover


cisco *Live!*

Active / Active (Link Aggregation (LAG))



All VMs use all uplinks as
logically one uplink

Active / Active (MAC Pinning)

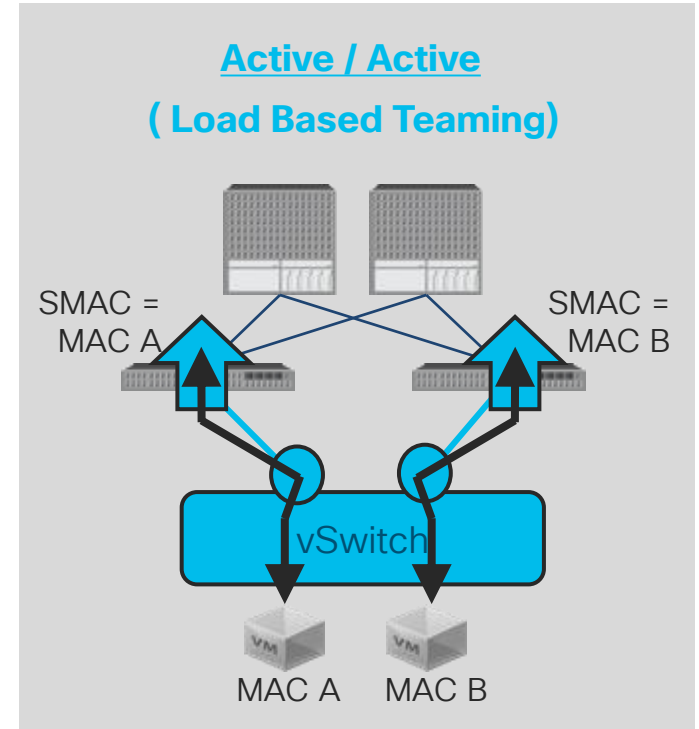


Each VM is assigned (pinned)
to a dedicated uplink for all
traffic

NIC Teaming Types to avoid

VMware – Load Based Teaming (LBT)

- **How it works:**
 - Similar to MAC Pinning (Route Based on Originating Port ID)
 - Re-pin every 30 sec based on the link utilization
- **Reasons why not to use it:**
 - MAC/IP flap between nodes/ports
 - 30 sec interval could be acceptable (not preferred just like any other switches)
- **Solution:**
 - LAG such as LACP or static port-channel



NIC Teaming Types to avoid

Microsoft – Switch Independent with Address Hash or Dynamic Load

- **How it works:**

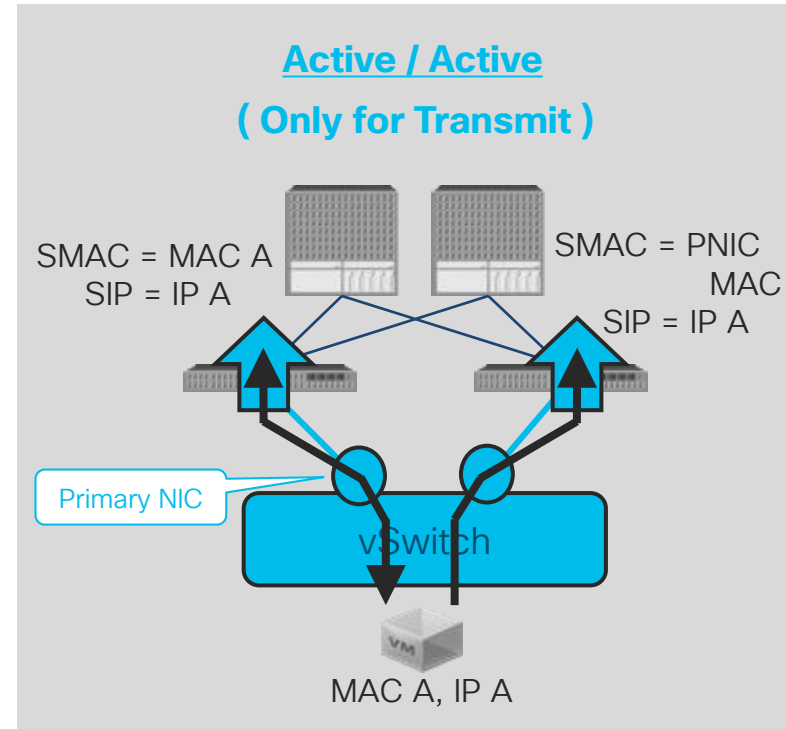
- Outgoing traffic on the primary NIC
sMAC is original MAC (could be primary NIC's MAC)
- Outgoing traffic on non-primary NIC
sMAC is each NIC's MAC
- Return traffic is only on the primary NIC
Controlling it via ARP (IP A -> MAC A)

- **Reasons why not to use it:**

- IP flaps between MACs
- No load balancing for return traffic

- **Solution:**

- LAG such as LACP or static port-channel
- Switch Independent with Hyper-V
- Disable VRF IP data-plane learning on ACI

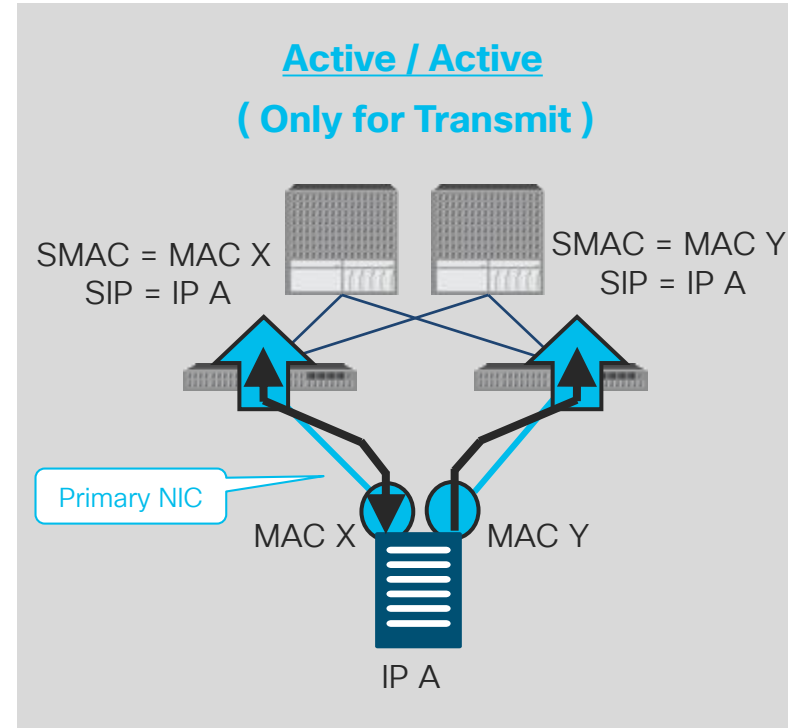


NIC Teaming Types to avoid

HP – Transmit Load Balancing Teaming

- **How it works:**
 - Outgoing traffic
sMAC is each NIC's MAC
 - Return traffic is only on the primary NIC
Controlling it via ARP (IP A -> MAC X)
- **Reasons why not to use it:**
 - IP flaps between MACs
 - No load balancing for return traffic
- **Solution:**
 - LAG such as LACP or static port-channel
 - Active / Standby (HP – Network Fault Tolerance, NFT)
 - Disable VRF IP data-plane learning on ACI

<https://support.hpe.com/hpsc/doc/public/display?docId=c01984706>



VRF IP Data-Plane Learning knob on ACI

From ACI 4.0

- **Enabling VRF IP Data-Plane Learning (default)**

- ✓ Quick detection of endpoint move without using CPU resource
- ✓ Traffic optimization with IP remote endpoint (less flood, etc.)
- × IP flaps between MACs with some NIC teaming options

- **Disabling VRF IP Data-Plane Learning**

- × Relying on CPU and ARP for IP learning (the same as traditional switch)
- × Non optimal forwarding with no IP remote endpoint learning
- ✓ No IP flaps between MACs with some NIC teaming options

The screenshot shows the ACI GUI configuration for a VRF-Tenant. The navigation path is: Application Profiles > Networking > Bridge Domains > VRFs > VRF-Tenant. The configuration options are:

- Create Route Target Profile:
- DNS labels: (enter names separated by comma)
- Route Tag Policy:
- IP Data-plane Learning:** Disabled **Enabled**
- Enable GOLF-OPFLEX MODE:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html#_Toc18440064

Which Teaming Options to Use

	Server uses all pNICs for client-to-server traffic	Server uses all the pNICs for server-to-client traffic	One VM uses all pNICs for client-to-server traffic	One VM uses all pNICs for server-to-client traffic	ACI configuration required
Port-channeling/LACP	Yes	Yes	Yes	Yes	vPC
Active/Standby Teaming	No	No	No	No	Nothing special
MAC pinning	Yes	Yes	No	No	Nothing special
Hyper-V Wwitch Independent with Hyper-V port	Yes	Yes	No	No	Nothing special
Hyper-V Switch Independent with Dynamic Load	Yes	Yes	No	Yes	dataplane learning tuning
A/A TLB teaming	No	Yes	No	Yes	dataplane learning tuning
Vmware LBT Teaming	Yes	Yes	No	No	Don't configure too aggressive timers with Rogue EP detection

Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - [Allowing Traffic through ACI](#)
 - Configuring L3 Connectivity to the Outside
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

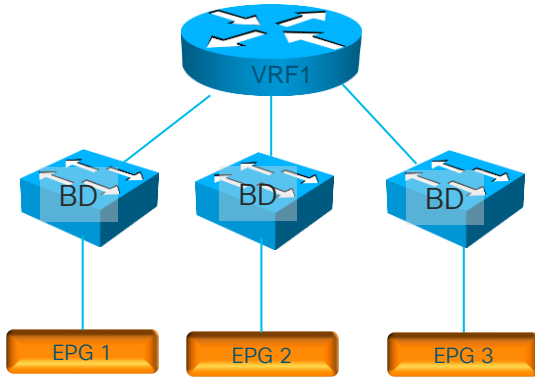
How to allow all traffic

Many start migration to ACI without enforcing security in ACI first

VRF unenforced

Policy Control Enforcement Preference:

Enforced Unenforced



No contract at all.

➤ EPG means nothing but VLAN(s)

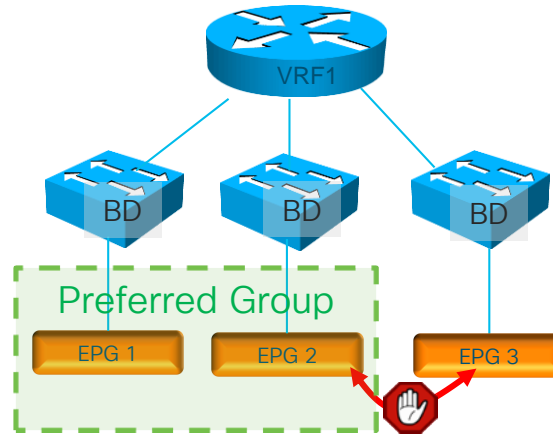
cisco *Live!*



Preferred Group

Preferred Group Member:

Exclude Include




Implicit allow for all EPGs in the preferred Group
(recommended)

vzAny with permit

▼ VRF-1

Multicast

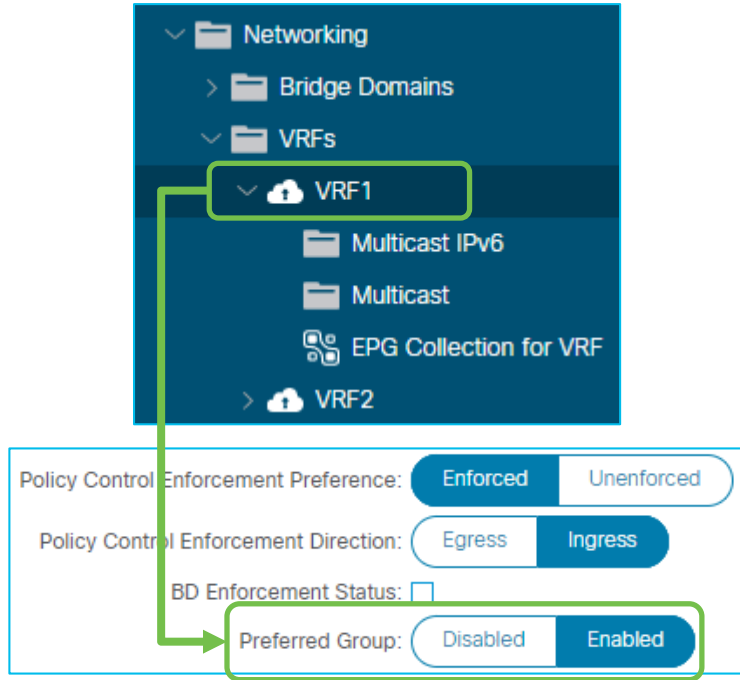
EPG Collection for VRF



Manually creating permit any any

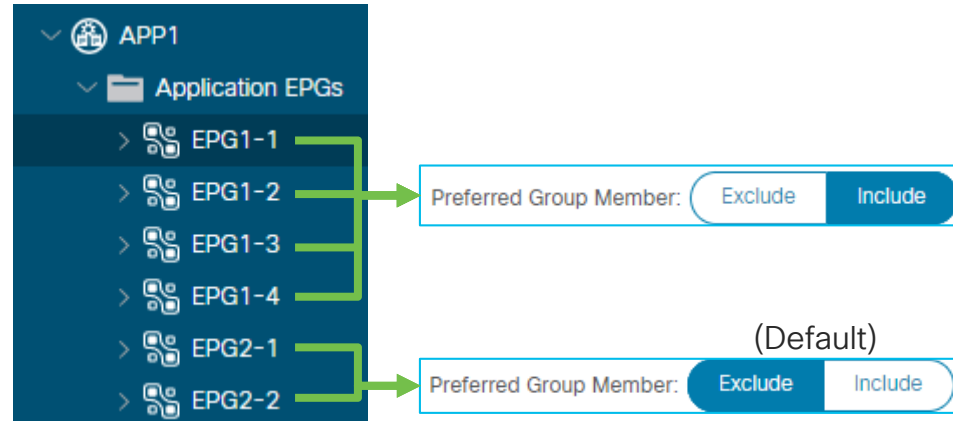
Preferred Group Configuration

1. Enable Preferred Group in the VRF



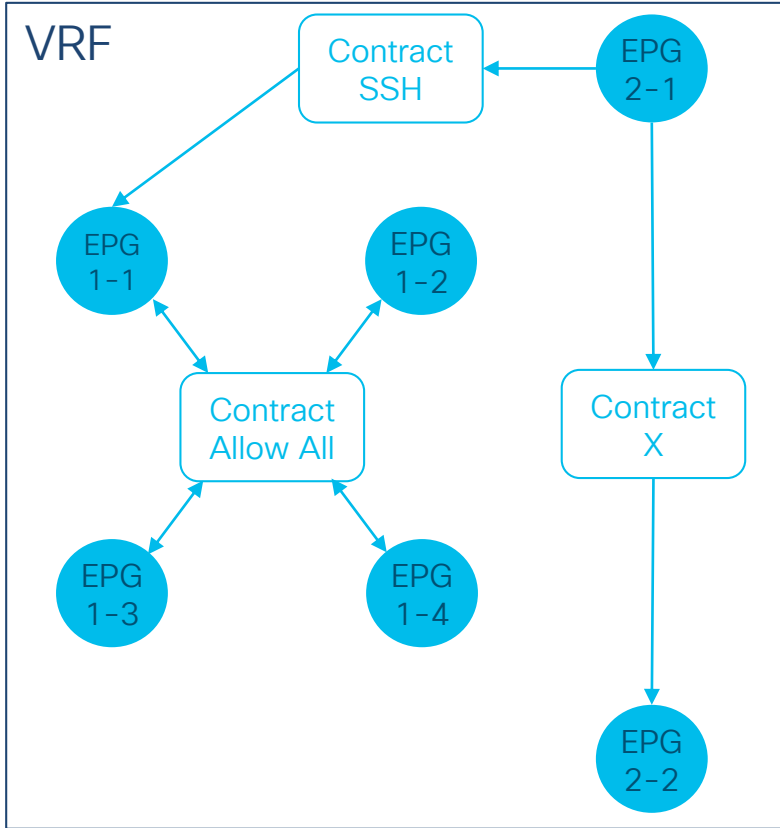
The screenshot shows the Cisco GUI configuration for VRF1. The 'Networking' tree is expanded to 'VRFs' > 'VRF1'. A green box highlights 'VRF1'. Below, the 'Policy Control Enforcement Preference' is set to 'Enforced', 'Policy Control Enforcement Direction' is 'Ingress', and 'BD Enforcement Status' is unchecked. A green box highlights 'Preferred Group' which is set to 'Enabled'.

2. Include each EPG in the Preferred Group

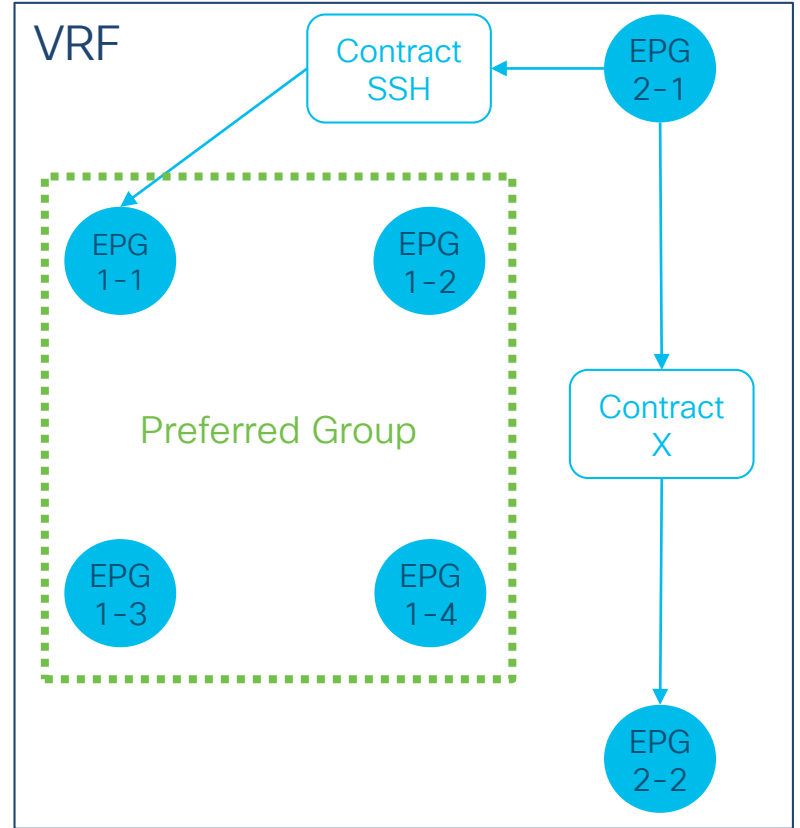


The screenshot shows the Cisco GUI configuration for Application EPGs under APP1. The 'Application EPGs' list includes EPG1-1, EPG1-2, EPG1-3, EPG1-4, EPG2-1, and EPG2-2. Green arrows point from each EPG to a 'Preferred Group Member' control. For EPG1-1 through EPG1-4, the 'Include' button is selected. For EPG2-1 and EPG2-2, the 'Include' button is selected, and the text '(Default)' is shown above the control.

Without Preferred Group

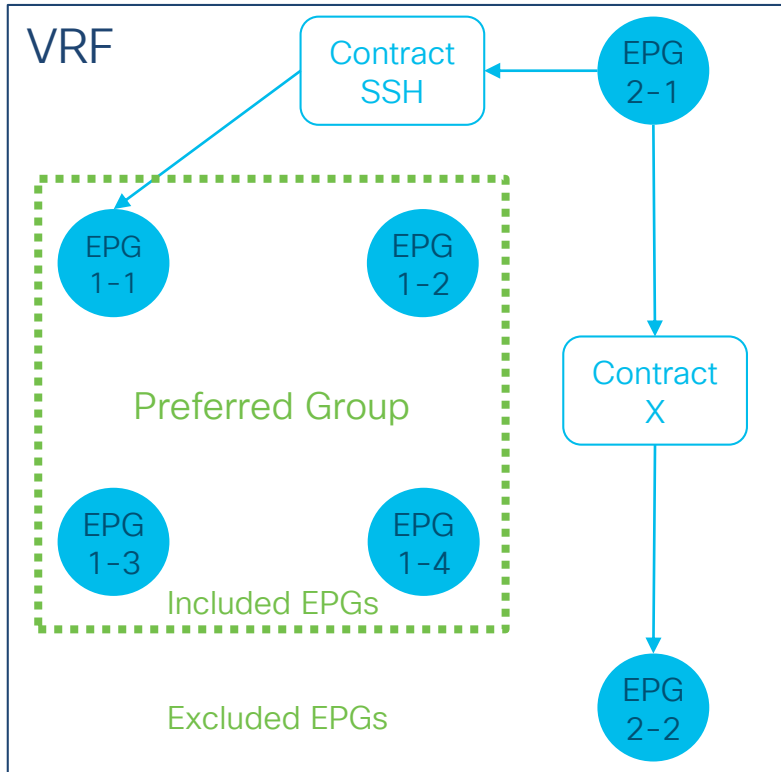


With Preferred Group



Preferred Group

How does the zoning-rule look like?



1. Your contracts (priority 7)

Excluded EPG ↔ Excluded EPG

Excluded EPG ↔ Included EPG

Source	Destination	Filter	Action
EPG-2-1	EPG-1-1	SSH	permit
EPG-1-1	EPG-2-1	SSH-r	permit
EPG-2-1	EPG-2-2	X	permit
EPG-2-2	EPG-2-1	X-r	permit

2. Implicit deny for all Excluded EPGs (priority 18)

Excluded EPGs → Any

Any → Excluded EPGs

Source	Destination	Filter	Action
EPG-2-1	any	implicit	deny
any	EPG-2-1	implicit	deny
EPG-2-2	any	implicit	deny
any	EPG-2-2	implicit	deny

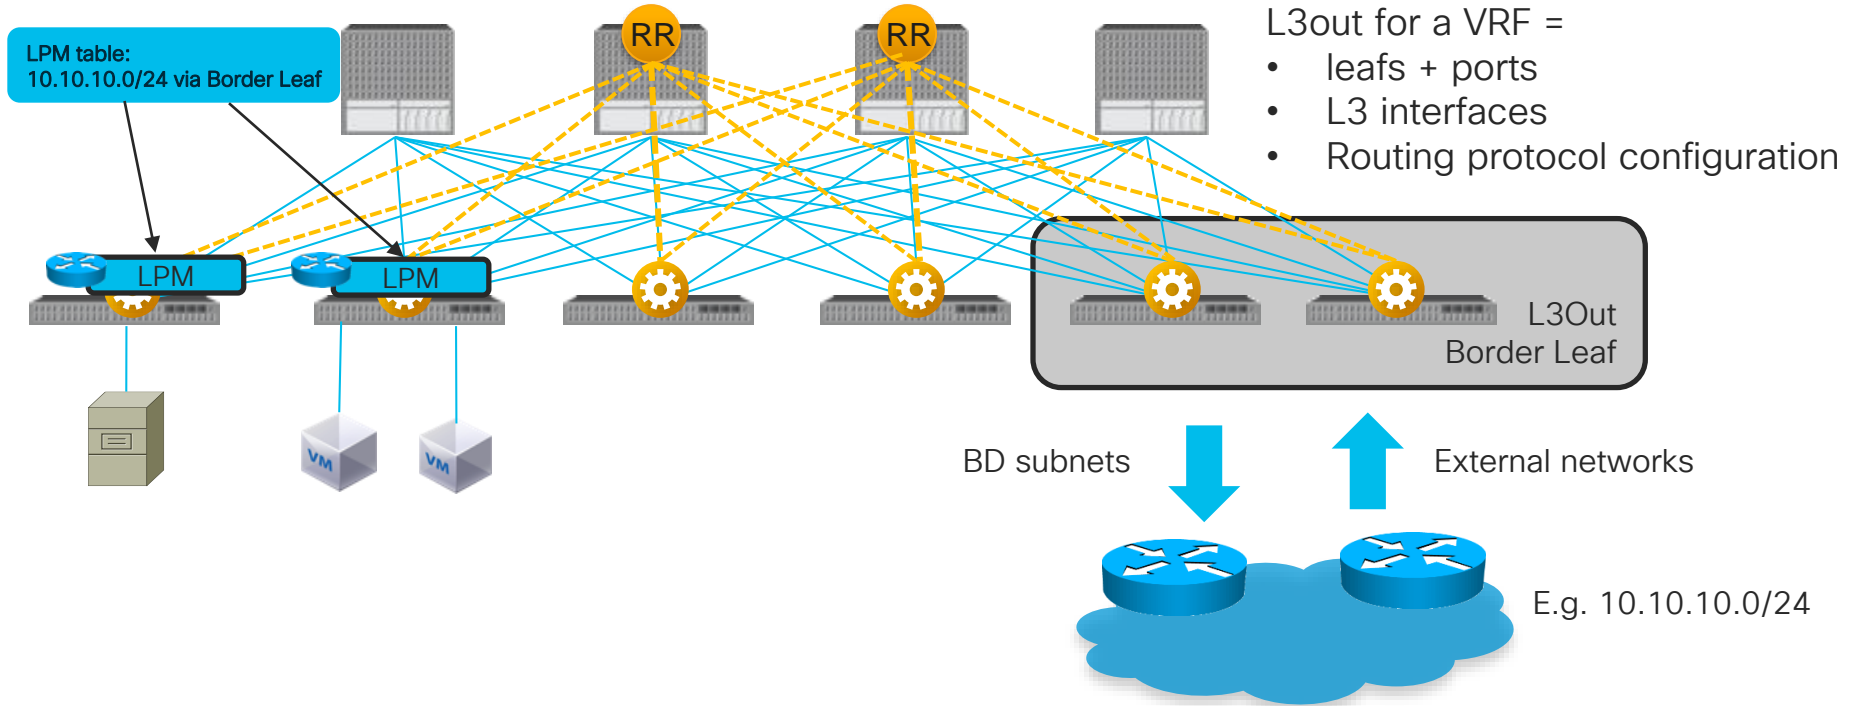
3. Implicit permit all Included EPGs (priority 20)

Source	Destination	Filter	Action
any	any	implicit	permit

Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - [Configuring L3 Connectivity to the Outside](#)
- Application Centric Design
 - Network Centric & Application Centric
 - Contract Priority and Optimization

If ACI is the default gateway, L3Out is required for outside reachability



L3Out Interface Types

Supported L3Out Interface Types:

- Routed Interface
- Routed Sub-Interface
- SVI

Rule of thumb

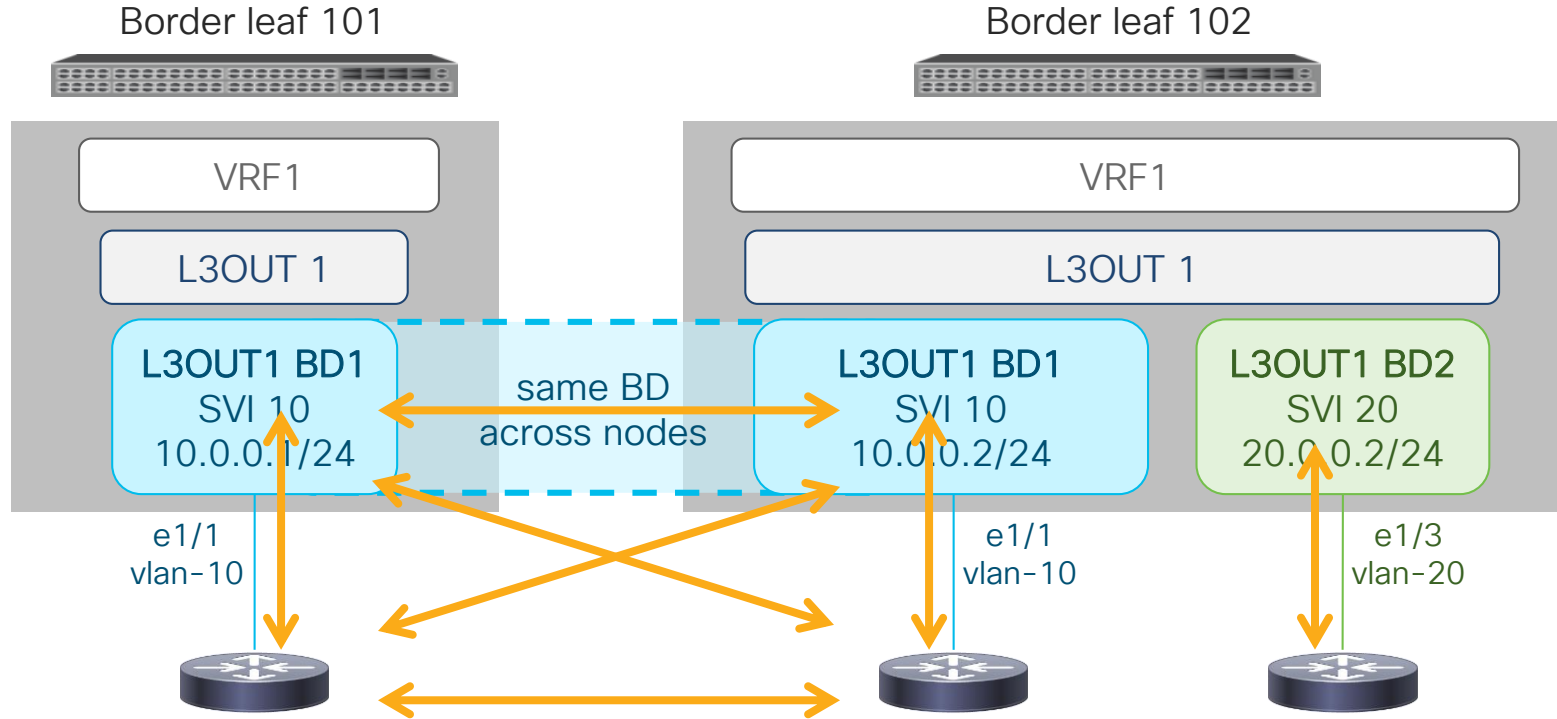


Use Routed (Sub-)Interface
if possible
SVI is typically required only
for FW, LB, etc.

Why? fast routing, next-hop convergence/detection is 101 for routing, right?

L3Out with SVI

↔ Routing Protocol Neighbors

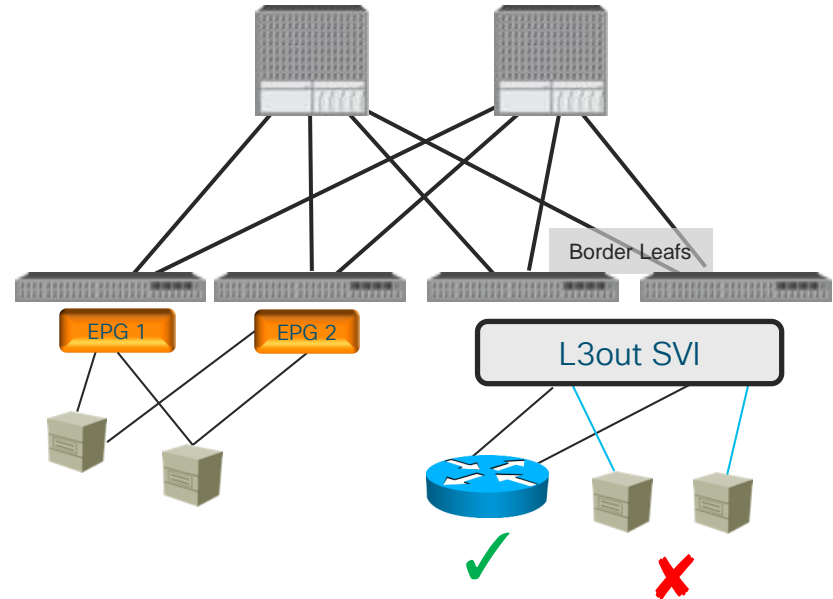


L3Out exception:

The same access VLAN means the same broadcast domain



The L3Out is for routing devices (routers, FW, etc.)

- The L3out is meant to attach Routing devices
- It is not meant to attach servers directly on the SVI of a L3Out
- Servers should be attached to EPGs and Bridge Domains
- There are multiple reasons for this:
 - The L2 domain created by a L3out with SVIs is not equivalent to a regular Bridge Domain
 - The L3ext classification is designed for hosts multiple hops away (more on this later)



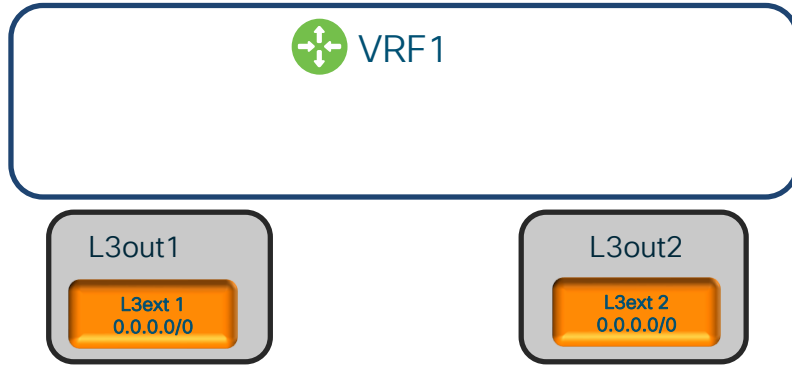
Avoid using 0.0.0.0/0 as a L3external if possible

- Classifying traffic from the L3Out with a L3Out subnet 0.0.0.0/0 is possible but it can lead to misconfigurations
- If two L3exts in the same VRF have a 0.0.0.0/0
- If EPG1 is allowed to talk to L3ext1 => EPG1 can also talk to L3ext2

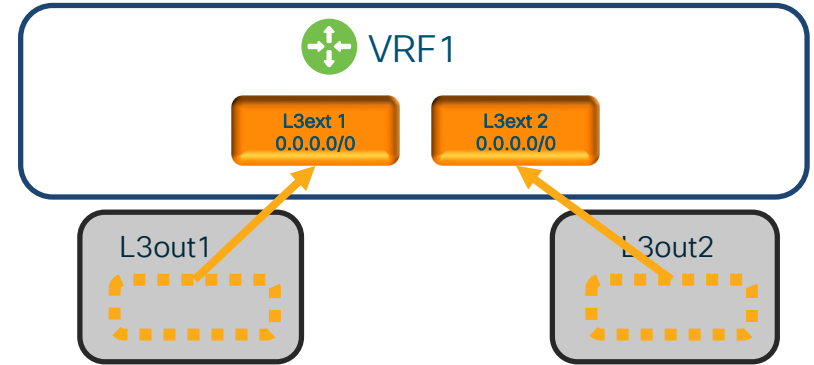



L3Out subnet design

How you configure



How you should imagine



Rule of thumb  L3Out subnet is per VRF

L3Out Interface subnet as L3Out Subnets?

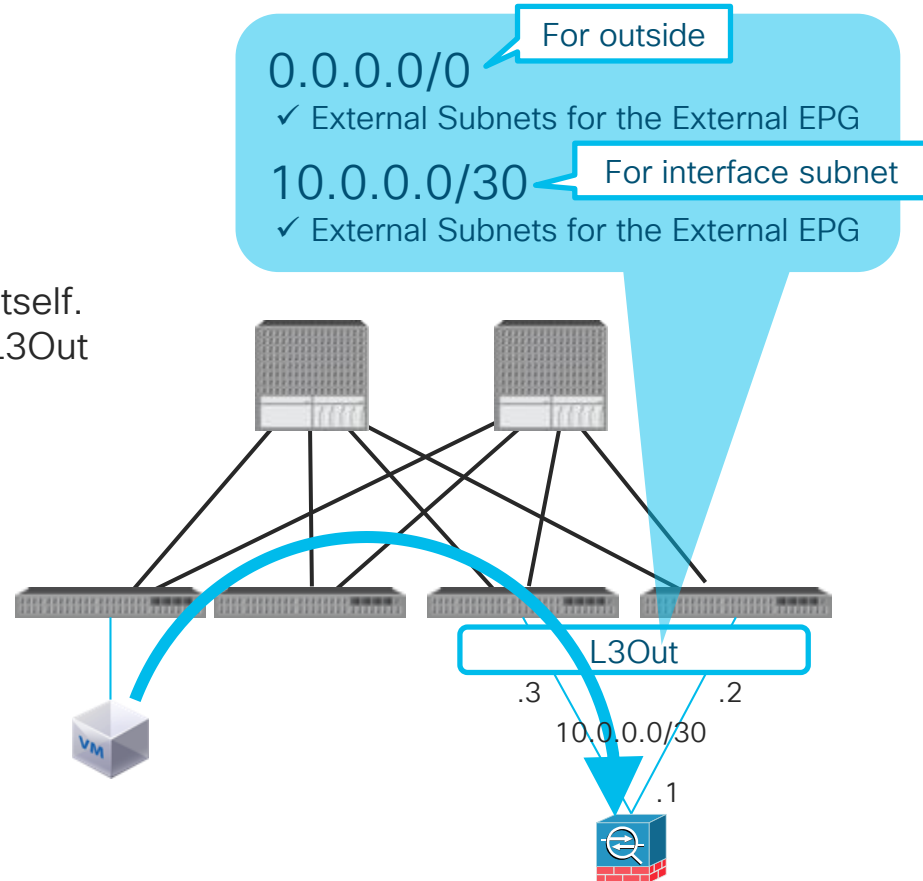
L3Out for L4-L7 devices (FW, LB)

Typically traffic is destined to the IP of the FW or LB itself.

- L3Out subnet should include such subnet (L3Out interface subnet)

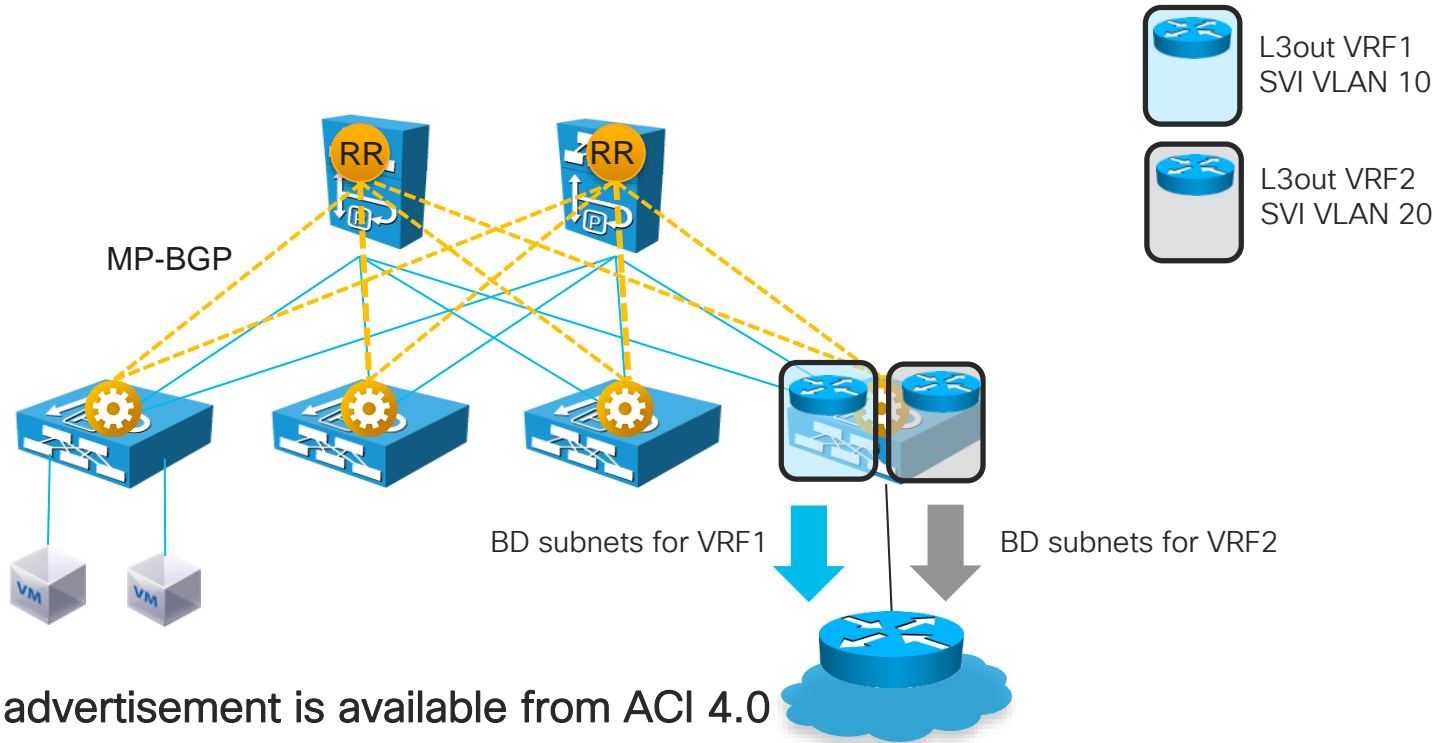
Why?

- By default, L3Out Interface Subnet uses a special pcTag (class ID) 1. This may lead an unintended result. (See CSCuz12913 for details if needed)

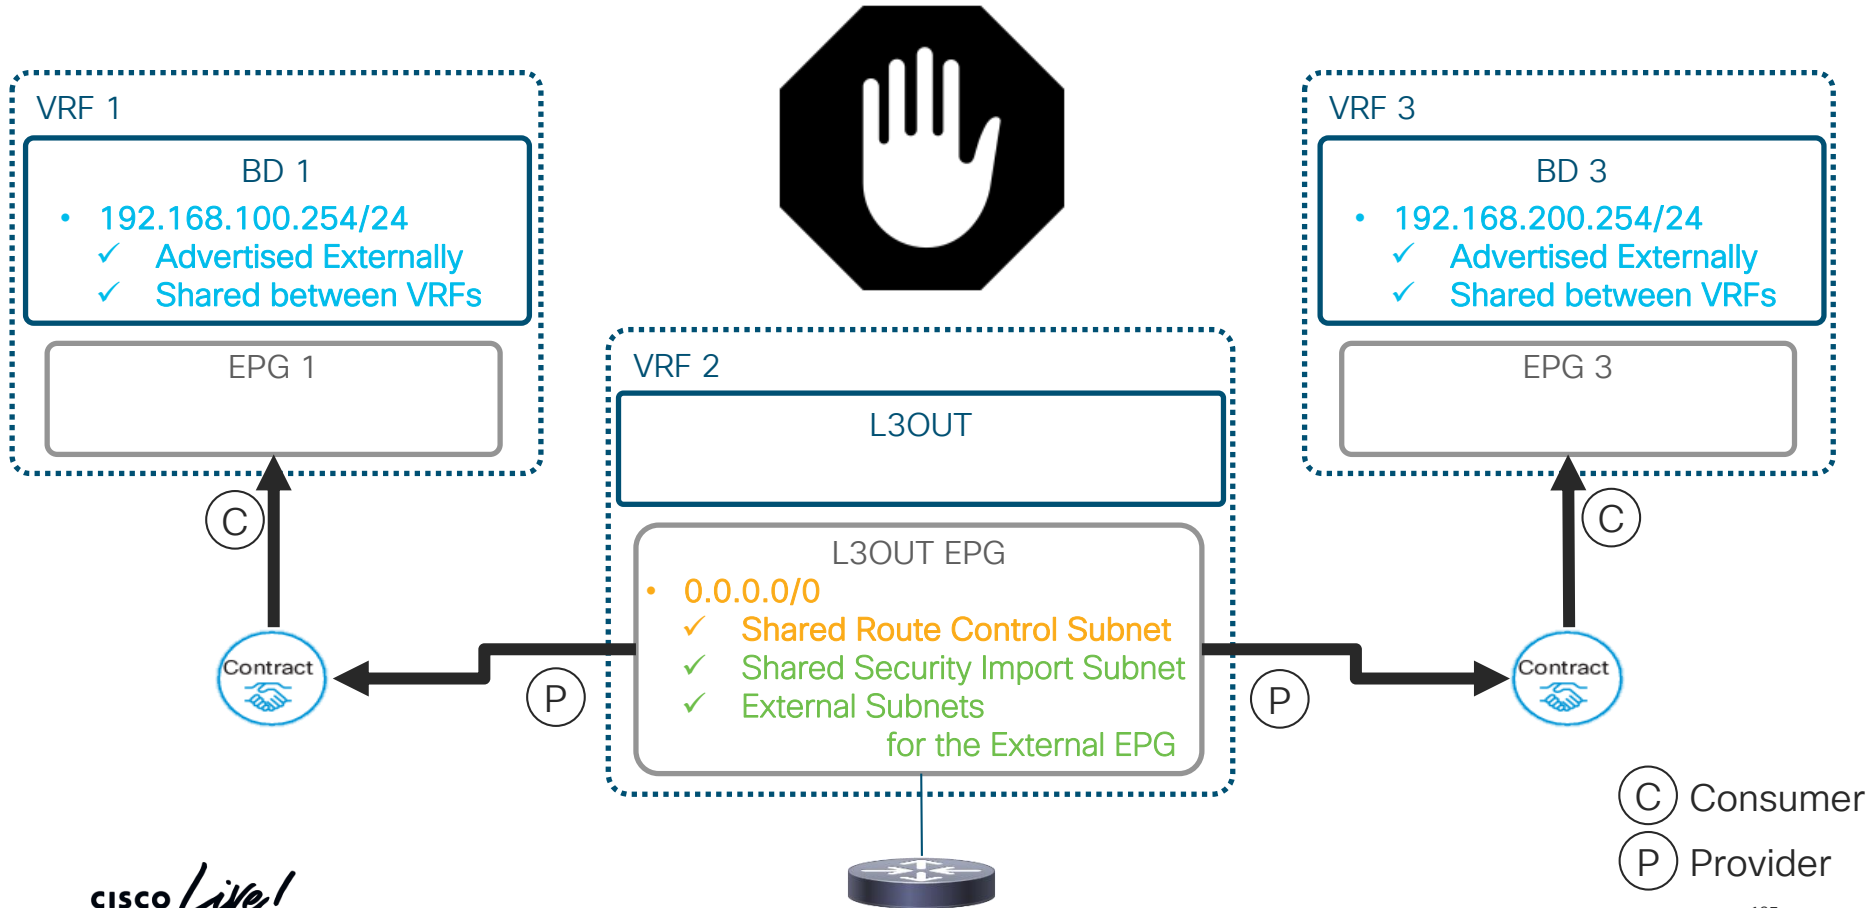


Multi Tenancy Design and L3Out:

One L3Out per VRF = VRF-Lite

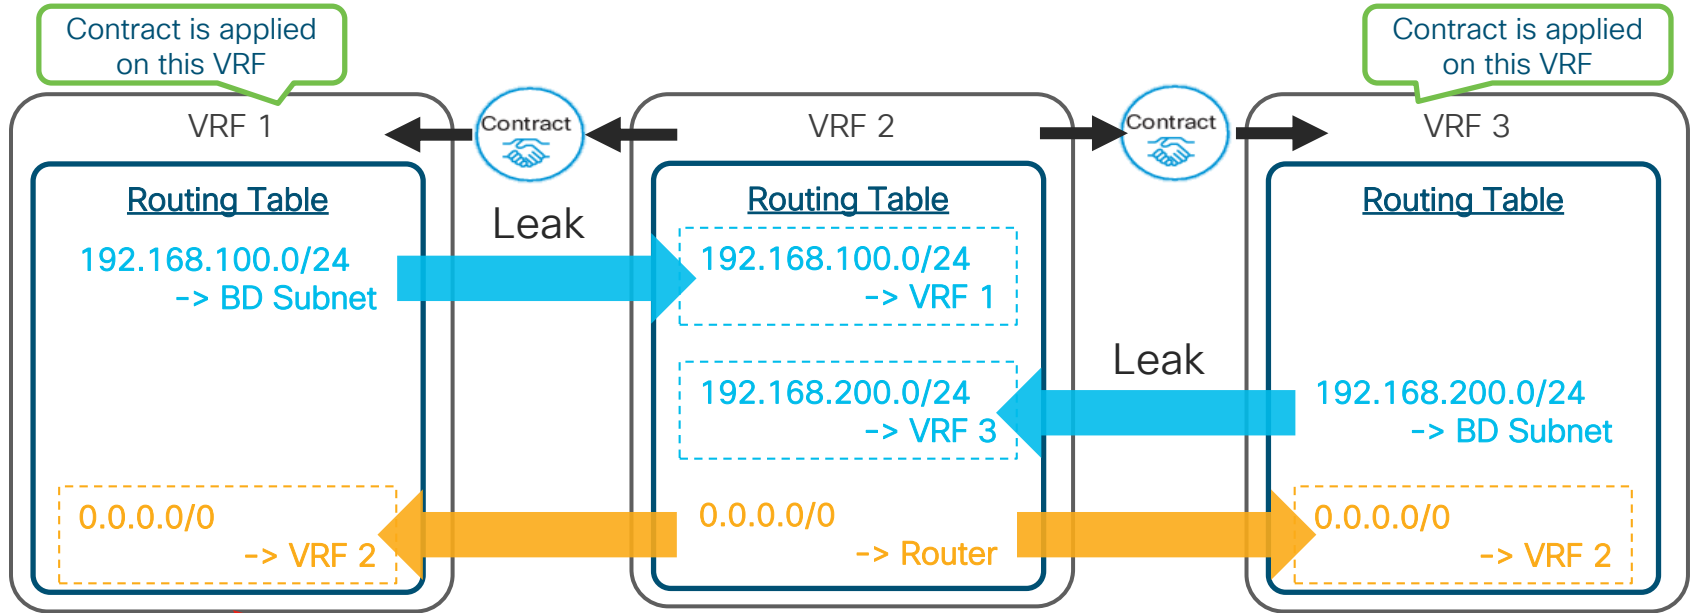


Multi Tenancy with Shared L3Out



Multi Tenancy with Shared L3Out

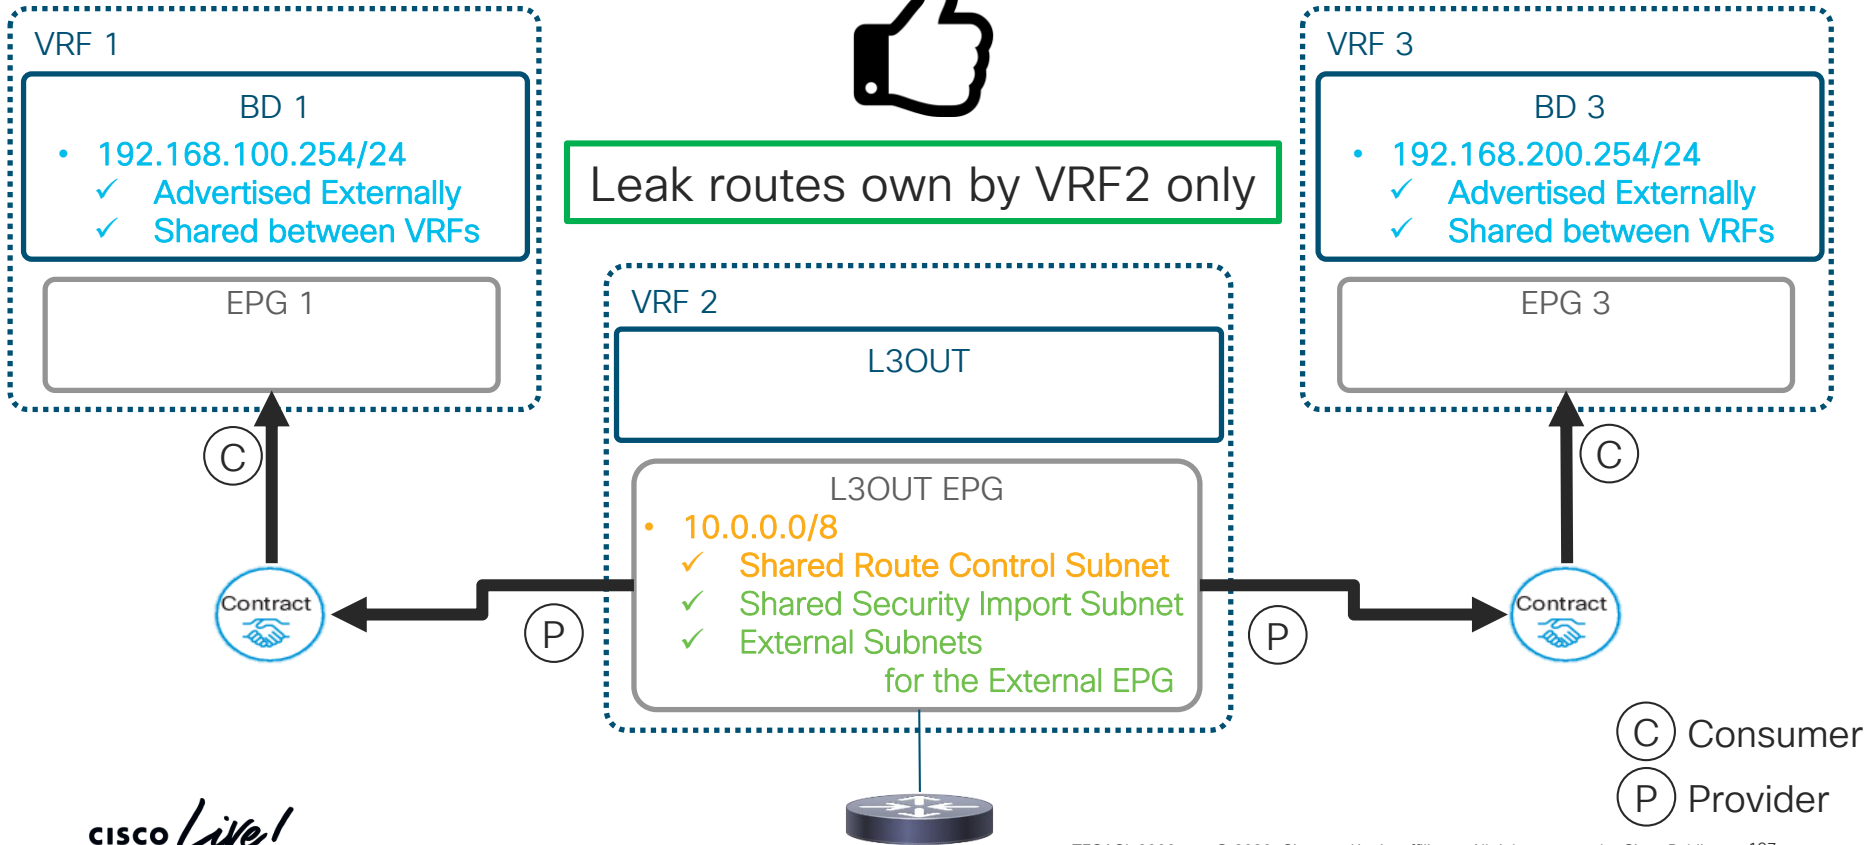
And its pitfall



VRF 1 and VRF 3 can talk to each other

Multi Tenancy with Shared L3Out

Avoid the pitfall




Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- **Application Centric Design**
 - **Network Centric & Application Centric**
 - Contract Priority and Optimization

Network Centric & Application Centric

Network Centric


- 1 VLAN = 1 EPG = 1 BD



- Similar to traditional network
 - VLAN as a broadcast domain
 - Easy to connect to legacy networks
- Typically simple or no contracts (no ACLs)

Application Centric

- Multiple EPGs per BD



- Multiple security domains (EPGs) in one broadcast domain
- Flexible network and security design


Now is the time for details 🤔

Network Centric & Application Centric



Could I re-use the same VLAN with Application Centric Design?

✔ Supported

✘ Not Supported



Need granular security groups



Not Supported


The same VLAN can be used for only one EPG in the same leaf

Network Centric & Application Centric Design



Could I re-use the same VLAN with Application Centric Design? (cont.)

✔ Supported


✘ Not Supported



Need granular security groups



The same VLAN can be re-used for any EPGs on **different** leaf. EPG (security group) deployment is physically limited.



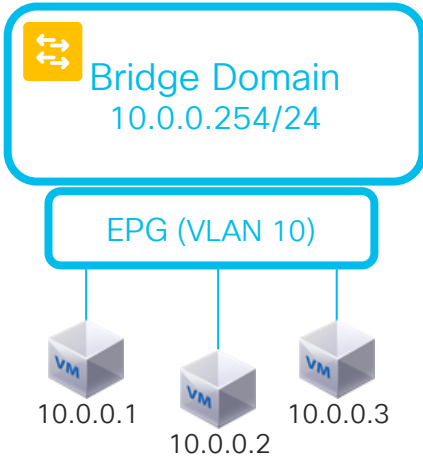
Local VLAN scope allows the same VLAN to be re-used on the same leaf. Lower scalability Complex Domain, VLAN Pool design

VLAN Scope: Global scope Port Local scope

Network Centric & Application Centric Design


Network Centric

A security group
in 1 subnet



More granular security?


Multiple security groups
in 1 subnet



What if multiple subnets
need to share the same
security rules?

Application Centric

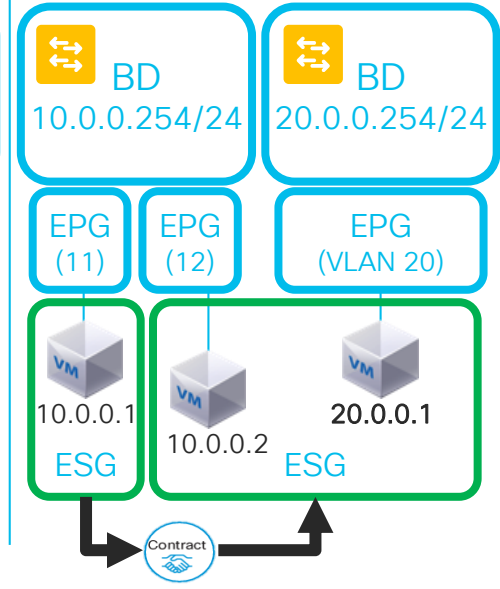
Security groups across
subnets



Sharing a broadcast
domain leads another
security concern

== Future ==
Endpoint Security Group
(ESG)

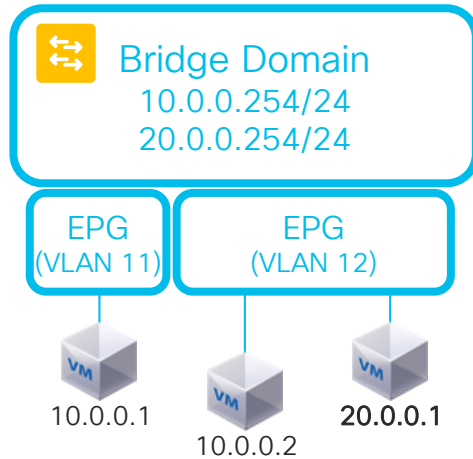
Security groups across
broadcast domains



Network Centric & Application Centric Design

Considerations for multiple subnets in one BD


Security groups across subnets



- Spine-Proxy to minimize the impact of flood




Be aware of the risk for L2 communication drop.



Silent Host Detection (ARP glean) can handle unknown targets. Unicast Routing and BD subnet are required.

OR

- Flood in Encap to minimize the impact of flood



From 3.1 and on 2nd (or newer) leaf, all traffic is flooded only within each encap VLAN (not EPG).

Agenda

- Single Fabric/Pod Topology Options
- ACI Design Considerations
 - ACI Fabric Bring Up
 - L2 Connectivity to Existing Networks
 - Loop prevention in ACI
 - Moving the default gateway to ACI
 - Connecting servers (Physical, VMM Integration, UCSM Integration)
 - Teaming Options
 - Allowing Traffic through ACI
 - Configuring L3 Connectivity to the Outside
- **Application Centric Design**
 - Network Centric & Application Centric
 - **Contract Priority and Optimization**

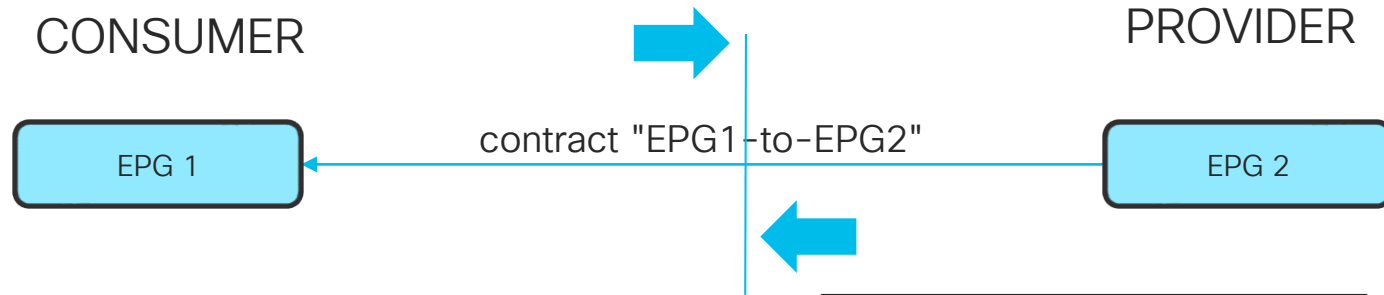
Communication between EPGs is configured via "contracts"

EPG 1 consumes "HTTP"

EPG 2 provides "HTTP"



You can then add subjects to the contract and define the filters for each direction



subject1: HTTP permit
subject2: Various other filters
subject3: etc..

Filter Chain For Consumer to Provider

Service Graph:

QoS Priority:

Target DSCP:

Filters

Name

Filter Chain For Provider to Consumer

Service Graph:

QoS Priority:

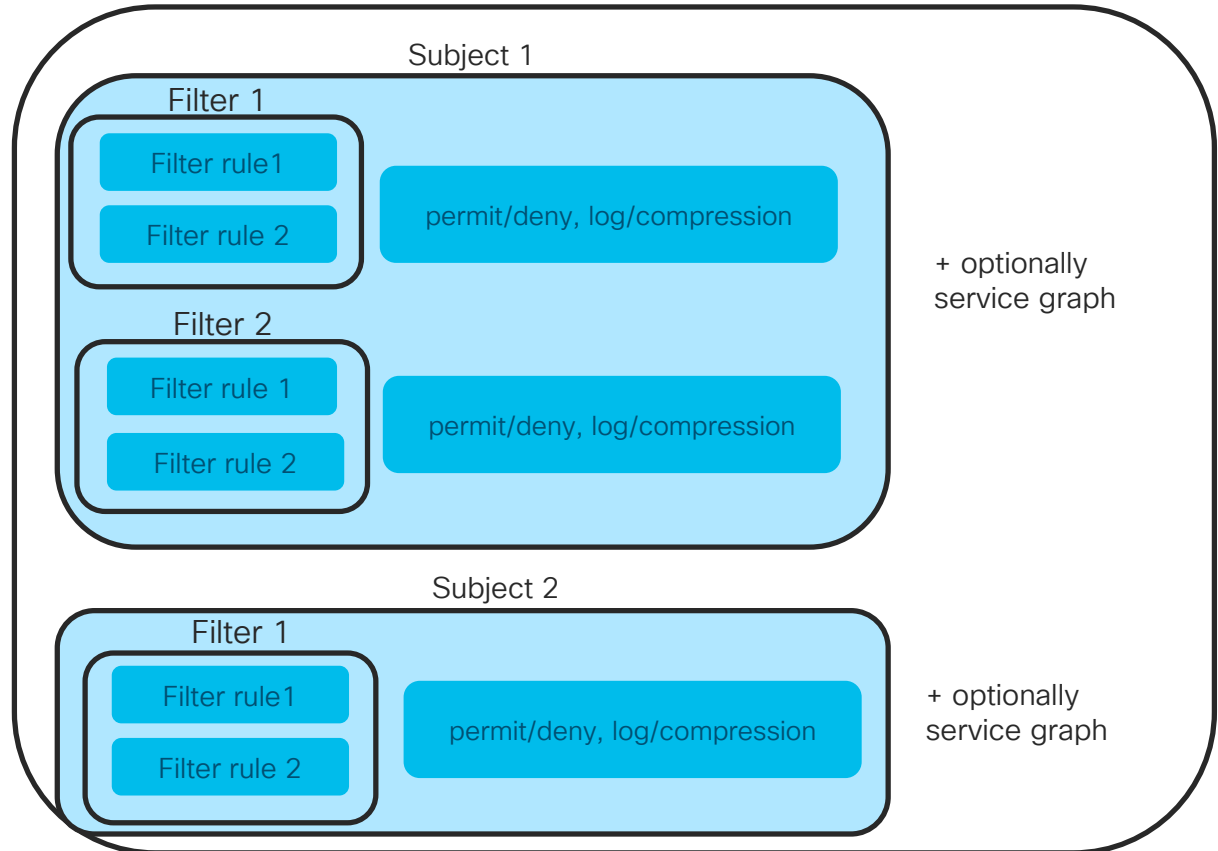
Target DSCP:

Filters

Name	Directives
------	------------

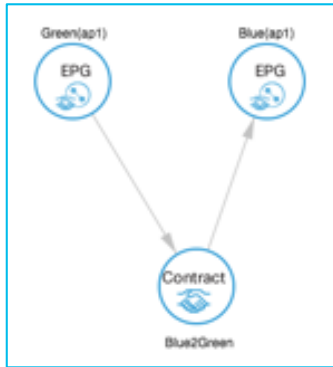
A Contract defines a set of filter rules in each "direction"

Contract



- Contract subjects contain a list of filters
- Each Filter consists of multiple filter rules
- Filter rules match protocols, src and dst ports
- Filters are attached to subjects as permit or deny rules
- Filters are attached to subjects with a directive (log, none)

The contract filters are programmed in the Policy Cam on the Leaf



VRF - vrf1

Policy **Operational** Stats Health Fault

Associated EPGs Associated External Rou

↻ ↓

Name	Description	State	Issues	QoS	Encap	PC Tag
ap1/Blue		applied		Unspecified		16387
ap1/Green		applied		Unspecified		32771



```
leaf1# show zoning-rule scope 2162697 | egrep -E "Scope|32771|16387"  
Rule ID SrcEPG DstEPG FilterID operSt Scope Action Priority  
4616 16387 32771 5 enabled 2162697 permit src_dst_any(8)  
4617 32771 16387 5 enabled 2162697 permit src_dst_any(8)
```

Contract Priority 101

When a packet matches multiple contracts:

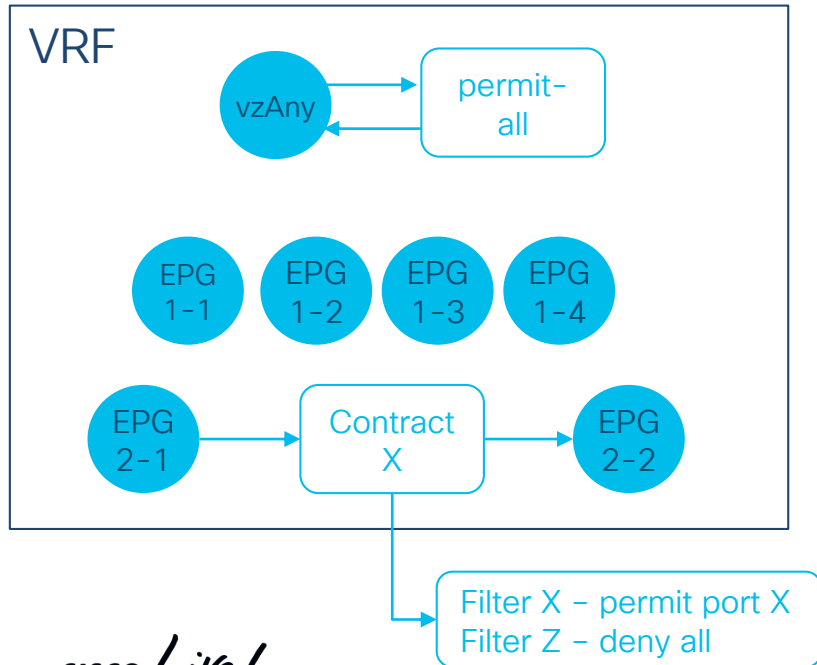
1. Specific EPG > vzAny
2. Specific match parameter wins
3. Deny > Redirect > Permit

Note - “Redirect” is used for Service Graph PBR

Contract Priority 101

Requirements:

1. Allow EPG2-1 to EPG2-2 on port X
2. All other communications are allowed



Src	Dst	Filter	Action
EPG2-1	EPG2-2	X	permit
EPG2-1	EPG2-2	all	deny
any	any	all	permit
any	any	implicit	deny


↑ Higher Priority

Redirect all traffic to a FW

With Service Graph PBR and vzAny

1. Create a PBR Service Graph that points to the FW
2. Create a permit all contract
3. Apply the PBR Service Graph to the permit all contract
4. vzAny provides and consume the permit all contract

Note - The servers' default gateway needs to be the BD



becomes



The contract is provided and consumed under vzAny

Tenant Baekerei

- Quick Start
- Tenant Baekerei
 - Application Profiles
 - Networking
 - Bridge Domains
 - VRFs
 - VRF-1
 - anyEPG-to-anyEPG (highlighted)
 - client-to-server
- External Bridged Networks
- External Routed Networks
 - Route Maps/Profiles
 - Set Rules for Route Maps
 - Match Rules for Route Maps
- 123

vzAny

Policy C

General Subj

Properties

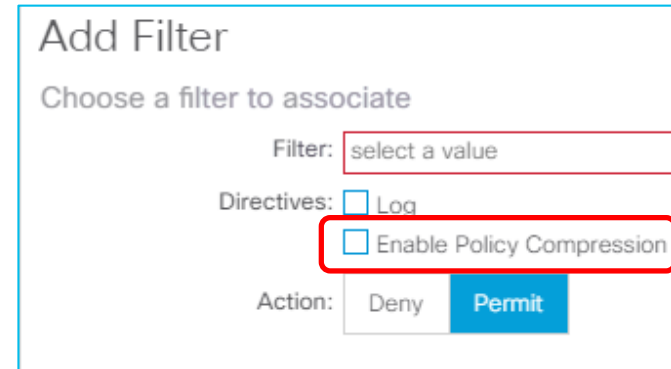
Name	Tenant	Type	QoS Class	Match Type
anyEPG-to-anyEPG	Baekerei	Contract	Unspecified	AtleastOne

Consumed Contracts:

Name	Tenant	Type	QoS Class
anyEPG-to-anyEPG	Baekerei	Contract	Unspecified

Contract Optimization

- Contracts (policy TCAM) Usage becomes a bottleneck easier in Application Centric Design
- There are many optimizations (depending on the leaf hw):
 - Change the ASIC profile to "policy TCAM" intensive profiles
 - Use L4 port range operations. It uses one entry only in TCAM.
 - Enable Policy Compression (at the cost of granularity of statistics)
 - Bidirectional rules are combined into one (from 3.2).
 - Rules used by the same EPG are combined into one. This is called indirection (from 4.0).



Add Filter

Choose a filter to associate

Filter:

Directives: Log Enable Policy Compression

Action:

Design Whitepapers

- <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737592.pdf>
- https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/multicast/b_Using_Layer3_Multicast/b_Using_Layer3_Multicast_chapter_00.html

For More Reading

- Design Guide: <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737909.pdf>
- About endpoint learning and BD settings: <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html>
- About the L3Out:
- <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/guide-c07-743150.html>
- About Migration:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/migration_guides/migrating_existing_networks_to_aci.html
 - Getting Started Guide: <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/getting-started/Cisco-APIC-Getting-Started-Guide-401.html>
 - Step by Step ACI deployment: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/white_papers/Cisco-ACI-Initial-Deployment-Cookbook.html
 - Virtualization Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/virtualization/Cisco-ACI-Virtualization-Guide-411/Cisco-ACI-Virtualization-Guide-411_chapter_010.html
 - Verified Scalability Guide for ACI 4.1: <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/verified-scalability/Cisco-ACI-Verified-Scalability-Guide-411.html>

ACI Anywhere, Extending the ACI Fabric

Agenda

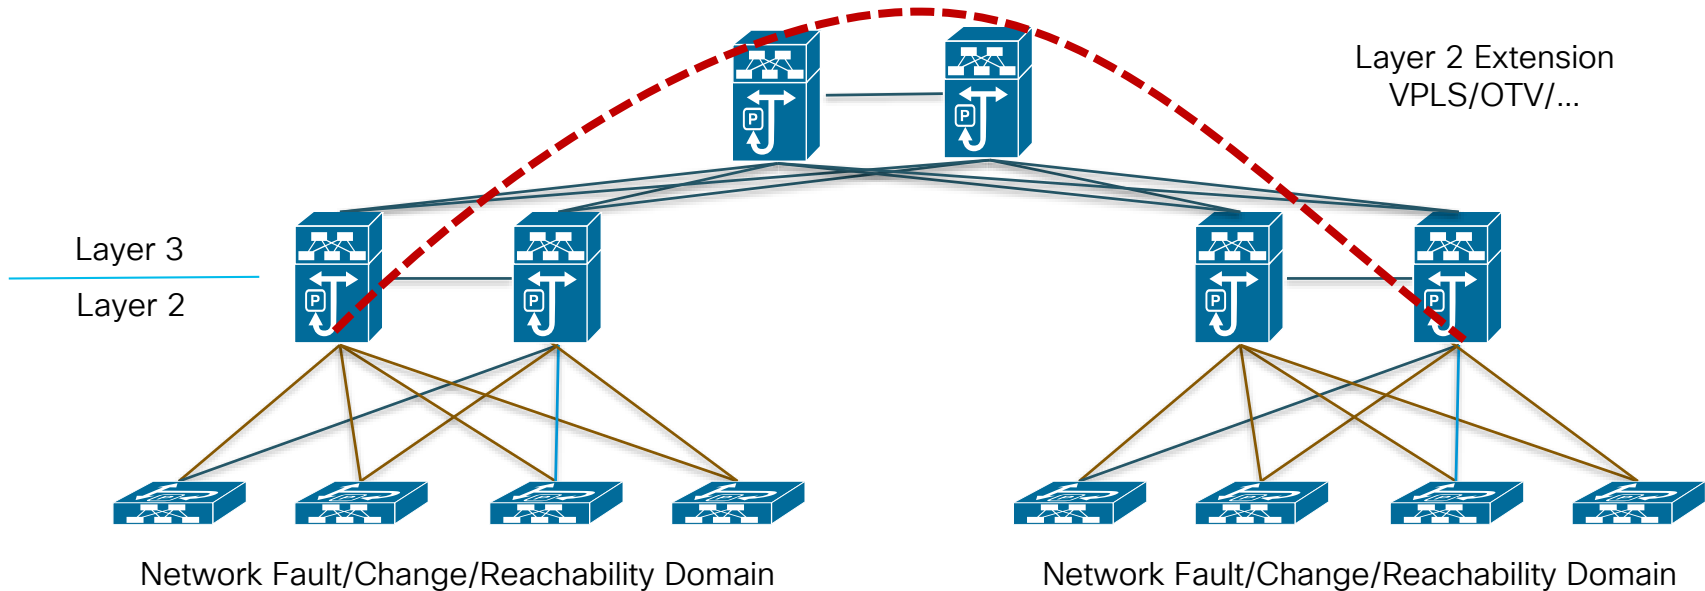
- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - Mapping use cases to the proper solutions
 - Active/Active DC → Multi-Pod
 - Disaster Recovery → Multi-Site
 - Migration/Coexistence with Legacy DC Networks and 'Disaggregated DCs' Model → Physical Remote Leaf
 - Baremetal Cloud Integration → Virtual Pod (vPod)
 - Extending ACI to the Cloud
 - Connecting the users to the Multi-Cloud DC
 - ACI and SDA Integration
 - ACI and SDWAN Integration

Overall Design Principles

Systems Availability

Best of Breed prior to 2014

- Distinct Network Domains for availability
 - Extension of Layer 2 works but complicates change and fault isolation
 - Sizing of each domain is a balance between need, risk and cost



Data Center Interconnect Solutions

Yesterday

A Tale from the Past

Ethernet

Over dark fiber or protected D-WDM

- *vPC double-sided (caution with DWDM SLA)*
 - Dual site interconnection
- *OTV or VXLAN EVPN*
 - Dual/Multiple sites interconnection

MPLS

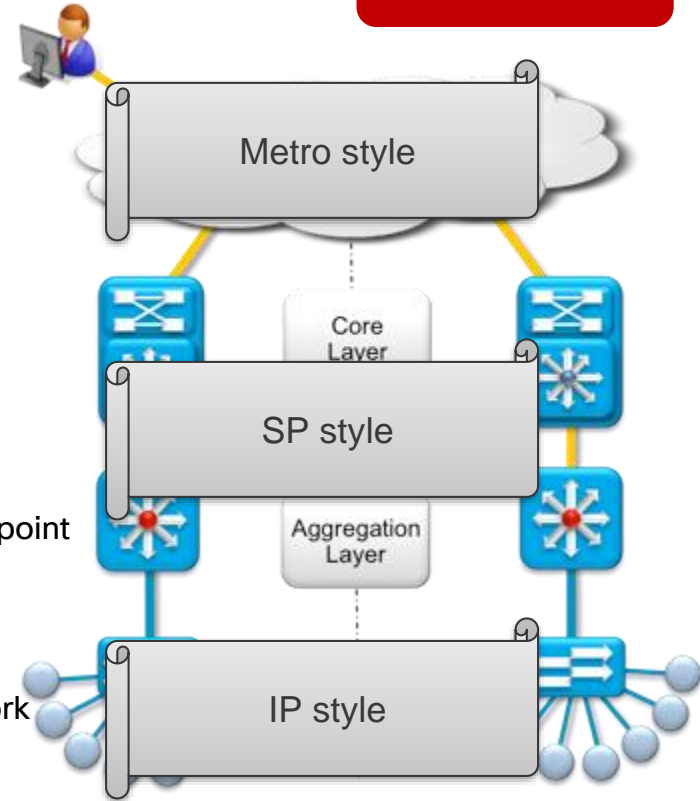
MPLS Transport

- *MPLS-EVPN*
 - SP, Point to Multipoint
- *PBB-EVPN*
 - Large scale & Multi-tenants, Point to Multipoint

IP

IP Transport

- *OTV*
 - Interconnect Traditional-based DC Network
- *VXLAN EVPN*
 - interconnect VXLAN-based Fabric
 - Layer 2 Ext. only and/or Layer 3 Ext. (multitenancy)
- *LISP*
 - For Subnet extension and Path Optimization



Does ACI Change Anything?

YES

Reachability is now Decoupled from Fault and
Change

Reachability is Decoupled from Topology as well


Data Center Interconnect Solutions

Today


ACI Simplifies the Deployment of DCI

- Common Control/Data Plane options used across different architectures
- Consistent security policies end-to-end


ACI Multi-Pod Fabric



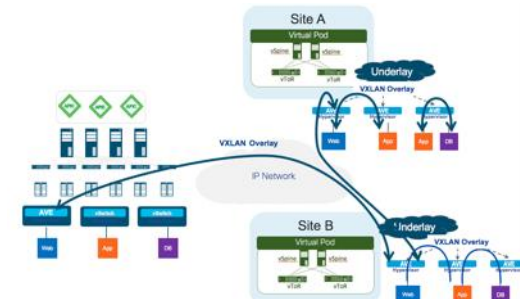
ACI Multi-Site



ACI Physical Remote Leaf



ACI Virtual Remote Leaf (vPod)



cisco *Live!*

Multi-Pod or Multi-Site?

That is the question...



And the answer is...

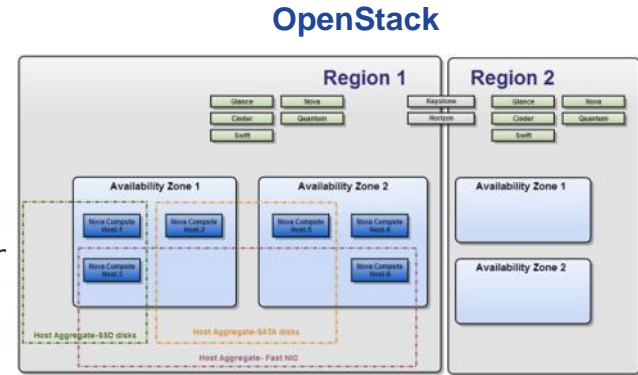
BOTH!



Framework for Multi-Cloud High Availability Design


Regions and Availability Zones

- Regions - Each Region has its own full OpenStack deployment, including its own API endpoints, networks and compute resources
- Availability Zones - Inside a Region, compute nodes can be logically grouped into Availability Zones, when launching a new VM instance, we can specify AZ or even a specific node in a AZ to run the VM instance




- Regions - Separate large geographical areas, each composed of multiple, isolated locations known as Availability Zones
- Availability Zones - Distinct locations within a region that are engineered to be isolated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same region

Amazon Web Services




Typical Requirement

Creation of Two Independent Fabrics/AZs



Fabric 'A' (AZ 1)

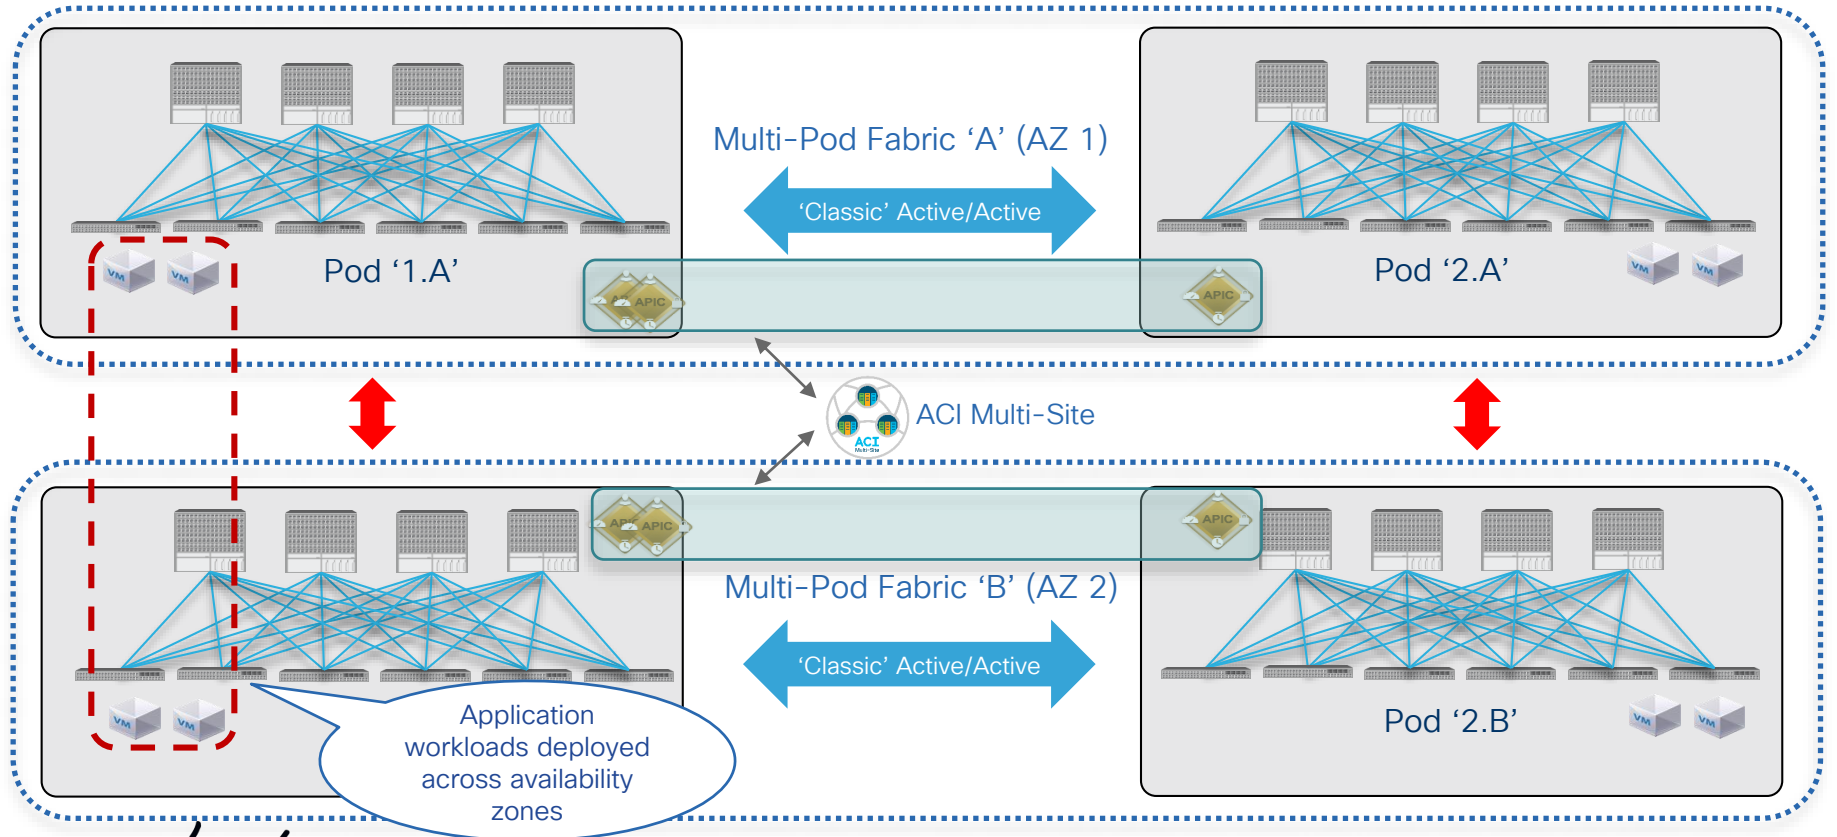


Fabric 'B' (AZ 2)

Application workloads deployed across availability zones

Typical Requirement

Creation of Two Independent Fabrics/AZs



Agenda

- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - Mapping use cases to the proper solutions
 - Active/Active DC → Multi-Pod
 - Disaster Recovery → Multi-Site
 - Migration/Coexistence with Legacy DC Networks and 'Disaggregated DCs' Model → Physical Remote Leaf
 - Baremetal Cloud Integration → Virtual Pod (vPod)
 - Extending ACI to the Cloud
 - Connecting the users to the Multi-Cloud DC
 - ACI and SDA Integration
 - ACI and SDWAN Integration

Active/Active DC Deployment

Metro Virtual Data Centre

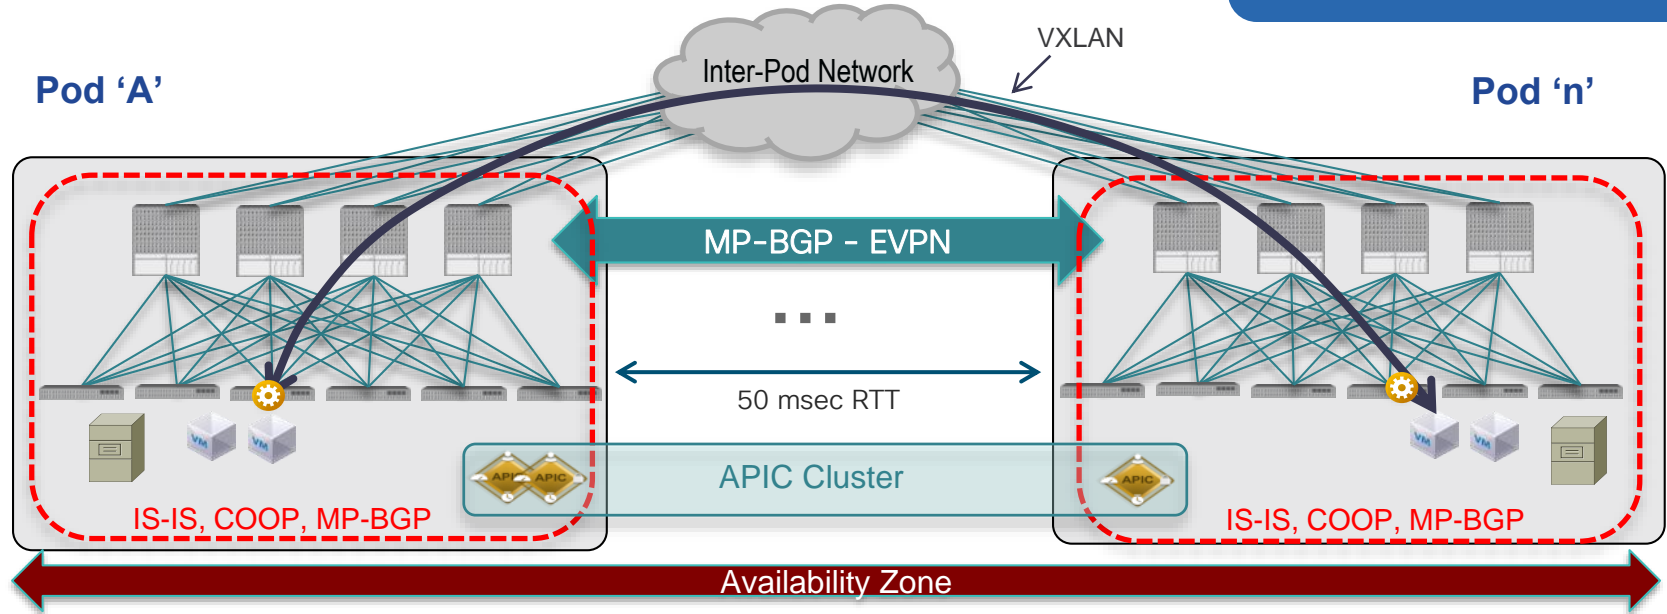
- High-availability application and data solution architecture which leverages a **dual data centre physical infrastructure (tightly coupled DC for short distances)**
- Management and interaction of applications in a paired data centre environment
- Disaster Avoidance and Prevention by pro-actively migrates seamlessly Virtual Machines with **no interruption (vMotion for example)**
- **Active-active** capability and workload rotation to accelerate incident response time and increase confidence
 - Deployment of an ESXi Metro Cluster with vSphere HA, Fault Tolerance (FT), DRS
 - Service Nodes (FW, SLB) clustered across DCs (Active/Standby, Active/Active)



ACI Multi-Pod

The Ideal Architecture for Active/Active DC Deployments

For More Information on
ACI Multi-Pod:
[BRKACI-2003](#)

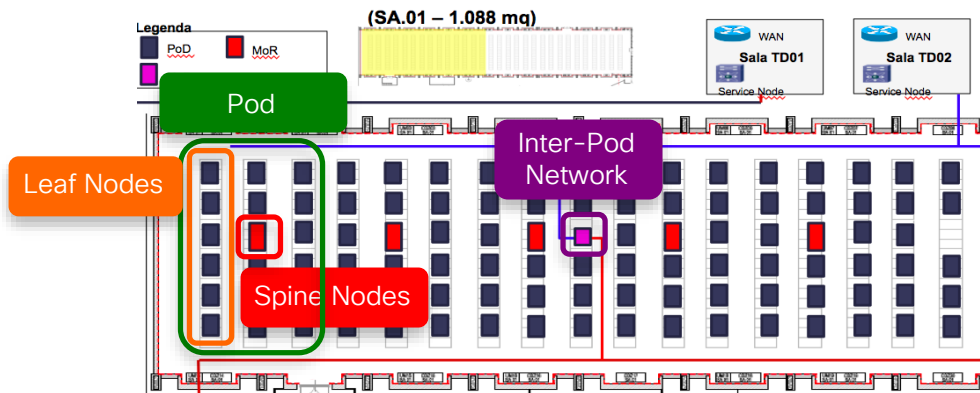


- Multiple ACI Pods connected by an IP Inter-Pod L3 network, each Pod consists of leaf and spine nodes
- Managed by a single APIC Cluster
- Single Management and Policy Domain
- Forwarding control plane (IS-IS, COOP) fault isolation
- Data Plane VXLAN encapsulation between Pods
- End-to-end policy enforcement

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 (stretched ESXi Metro Cluster, vSphere HA/FT, DRS initiated workload mobility,...)


Deployment of Active/Standby or Active/Active clustered network services (FWs, SLBs) across DCs

Application clustering (L2 BUM extension across Pods)



ACI Multi-Pod APIC Cluster Deployment

APIC – Distributed Multi-Active Data Base



- All Services in ACI run against their own portions of a Database
- Services and Database Processes are active on all nodes (not active/standby)
- The Data Base is distributed as active + 2 backup instances (shards) for every attribute

APIC Cluster Deployment with Multi-Pod

Deployment Recommendations

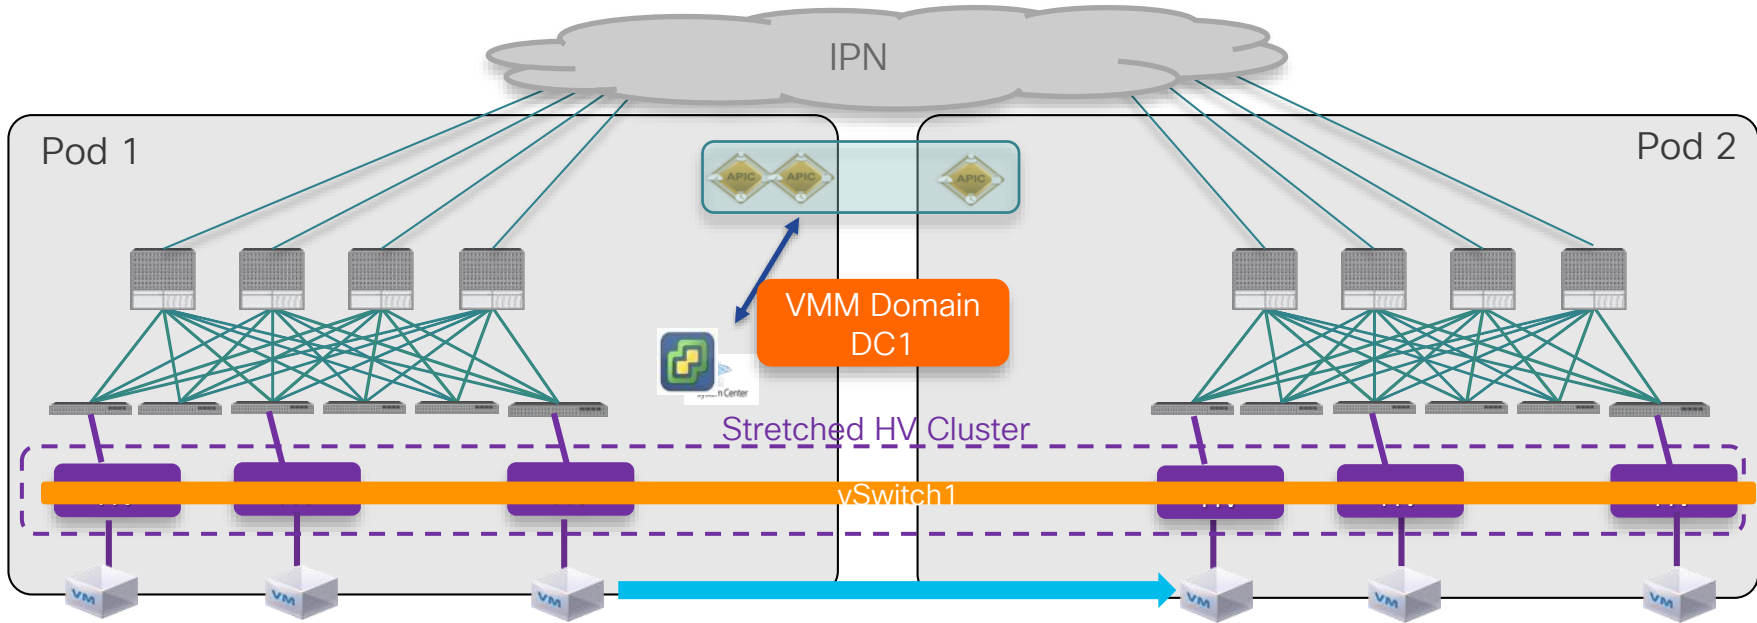
- **Main recommendation:** deploy a 3 nodes APIC cluster when less than 80 leaf nodes are deployed across Pods
- From ACI release 4.1(1) can deploy 4 nodes to support up to 200 leaf nodes across Pods
- When 5 (or 7) nodes are really needed for scalability reasons, follow the rule of thumb of never placing more than two APIC nodes in the same Pod (when possible):

	Pod1	Pod2	Pod3	Pod4	Pod5	Pod6
2 Pods*						
3 Pods						
4 Pods						
5 Pods						
6+ Pods						

CISCO Live! *'ID Recovery' procedure possible for recovering of lost information

Multi-Pod and Virtual Machine Manager (VMM) Integration

ACI Multi-Pod and VMM Integration

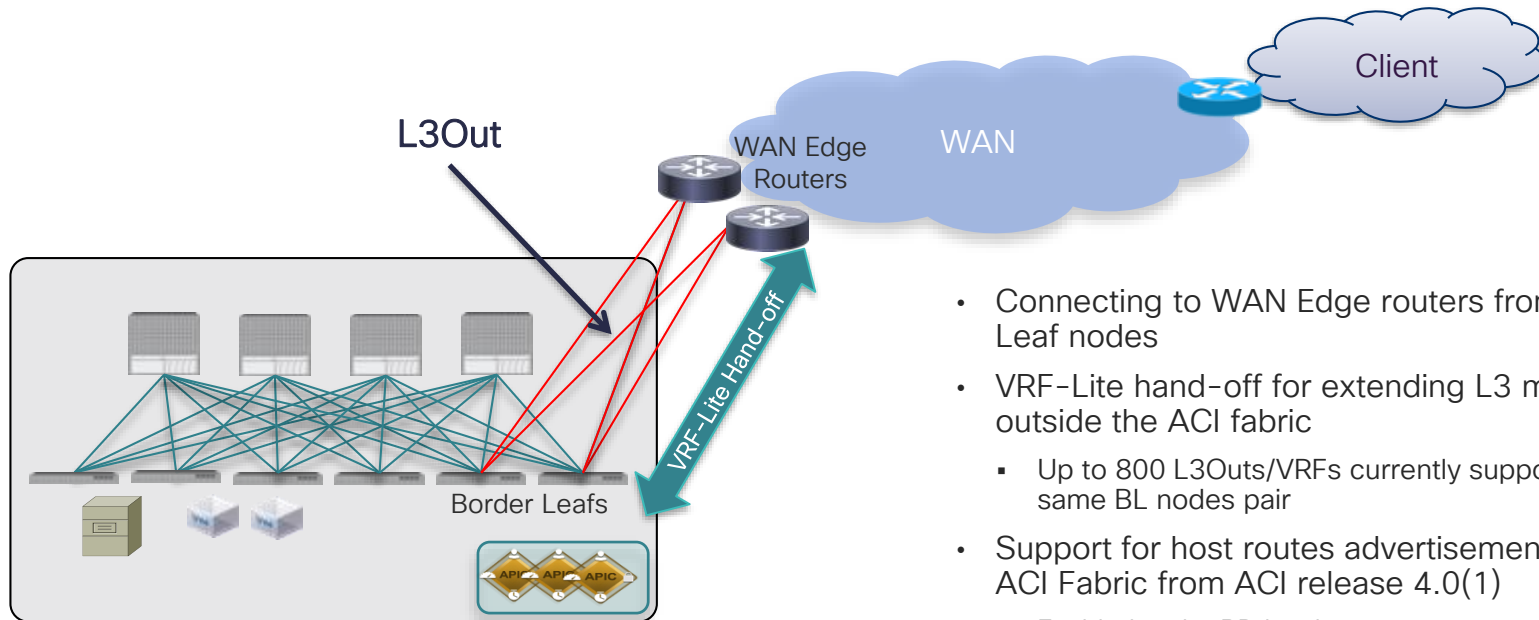


- Cluster of Hypervisors stretched across Pods
 - Single VMM domain created across Pods
 - Logical switch extended across the hypervisors part of the same stretched cluster
- Support for all intra-cluster functions (vSphere HA/FT, DRS, etc.)

Multi-Pod Connectivity to the External L3 Domain

Connecting to the External Layer 3 Domain

'Traditional' L3Outs on the BL Nodes (Recommended Option)

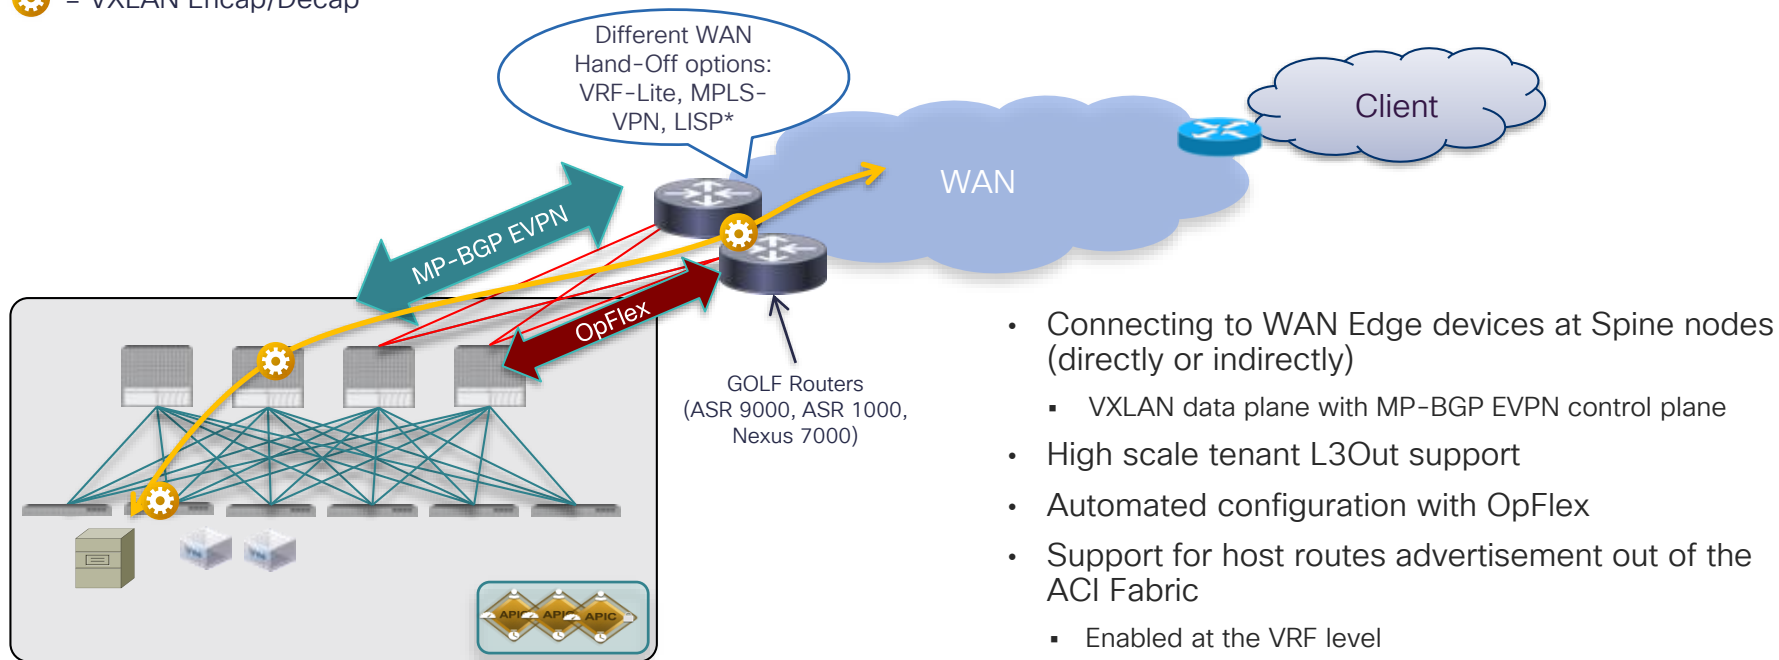


- Connecting to WAN Edge routers from Border Leaf nodes
- VRF-Lite hand-off for extending L3 multi-tenancy outside the ACI fabric
 - Up to 800 L3Outs/VRFs currently supported on the same BL nodes pair
- Support for host routes advertisement out of the ACI Fabric from ACI release 4.0(1)
 - Enabled at the BD level
- Support for L3 Multicast and Shared L3Out

Connecting to the External Layer 3 Domain

'GOLF' L3Outs (VRF High Scale Use Cases)

⚙️ = VXLAN Encap/Decap

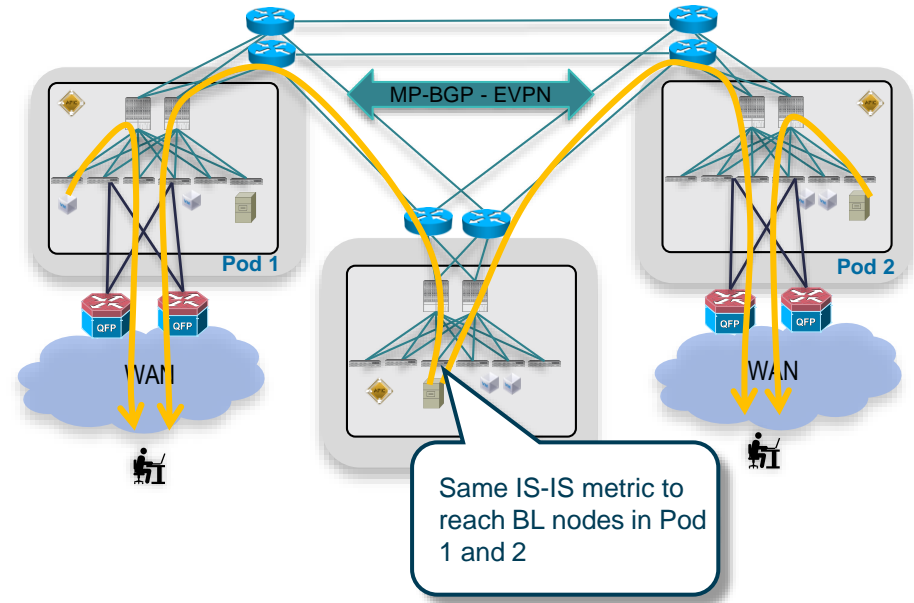


- Connecting to WAN Edge devices at Spine nodes (directly or indirectly)
 - VXLAN data plane with MP-BGP EVPN control plane
- High scale tenant L3Out support
- Automated configuration with OpFlex
- Support for host routes advertisement out of the ACI Fabric
 - Enabled at the VRF level
- No support for L3 Multicast or Shared L3Out (every tenant VRF requires its own L3Out)

Connecting Multi-Pod to Layer 3 Domain

Sharing L3Out Across Pods

- A Pod **does not need** to have a dedicated WAN connection (i.e. can offer transit services to other Pods)
- Multiple WAN connections can be deployed across Pods
- Outbound traffic: by default VTEPs always select WAN connection in the local Pod based on preferred metric
- Leaf nodes in Pods without local L3Outs will load-balance traffic between L3Outs in remote Pods

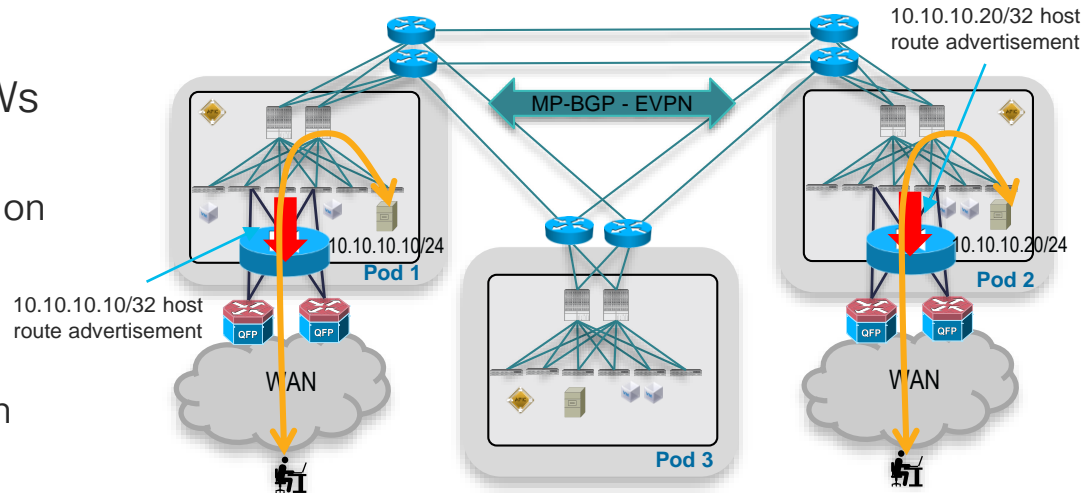


Connecting Multi-Pod to Layer 3 Domain

Use of Host-Route Advertisement on BL L3Outs

ACI 4.0(1)
Release

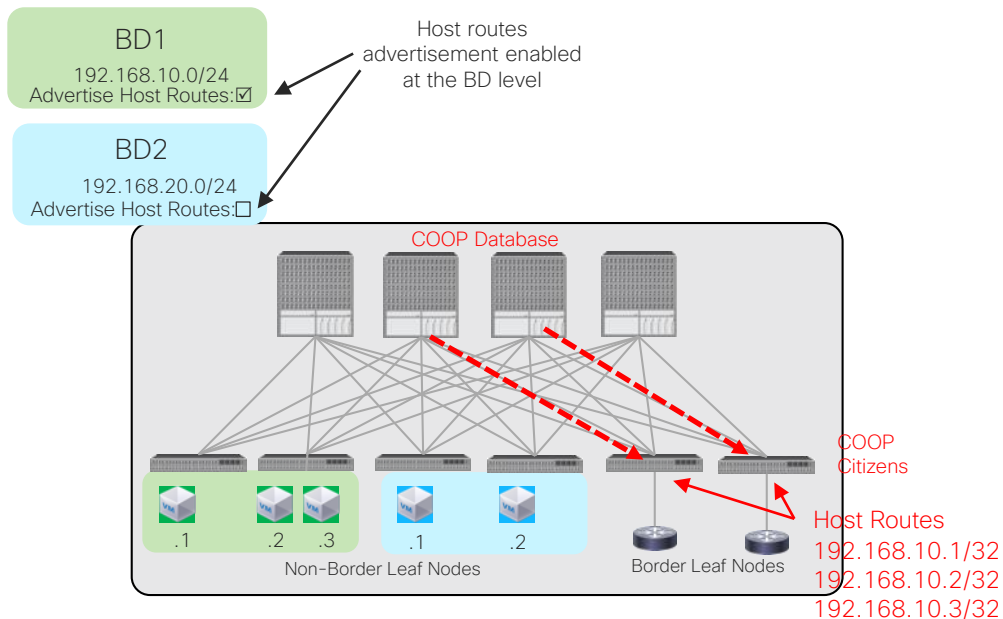
- Host routes advertisement is a best option to ensure all the deployed FWs are actively utilized
 - Support for host route advertisement on BL nodes available from ACI release 4.0(1)
 - Enabled at the BD level
 - Requires an L3Out connection in each Pod
 - Allows to keep symmetric inbound and outbound traffic paths



Host Routes Advertisement

ACI 4.0(1)
Release

Downloading the Host Routes to the Border Leaf Nodes




- Endpoint information is stored on the COOP database of the spines in the fabric
- When a BD is enabled with Host Routing the border leaf nodes download the host routes from the spines
- Enabling Host Routing on the Bridge Domain does not automatically advertise the host routes out of an L3Out, it must be explicitly configured (see a following slide)

Multi-Pod Network Services Integration Models


Multi-Pod and Network Services

Integration Models


Typical options for an Active/Active DC use case



- Active and Standby pair deployed across Pods
- No issues with asymmetric flows



- Active/Active FW cluster nodes stretched across Sites (single logical FW)
- Requires the ability of discovering the same MAC/IP info in separate sites at the same time
- Supported from ACI release 3.2(4d) with the use of Service-Graph with PBR



- Independent Active/Standby pairs deployed in separate Pods
- Use of Symmetric PBR to avoid the creation of asymmetric paths crossing different active FW nodes

ACI Multi-Pod

Where to Go for More Information



- ✓ ACI Multi-Pod White Paper

<http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html?cachemode=refresh>

- ✓ ACI Multi-Pod Configuration Paper

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739714.html>

- ✓ ACI Multi-Pod and Service Node Integration White Paper

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739571.html>

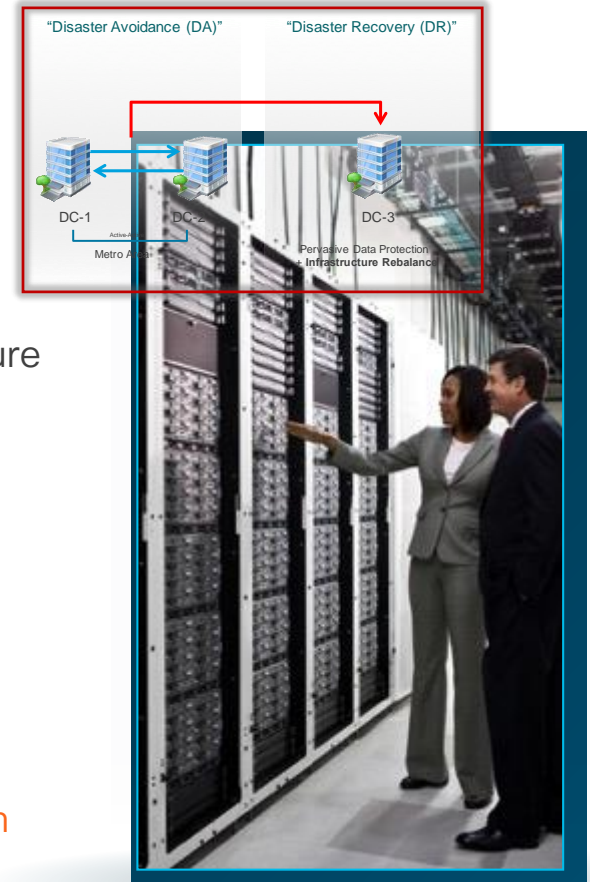
- ✓ BRKACI-2003

Agenda

- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - **Mapping use cases to the proper solutions**
 - Active/Active DC → Multi-Pod
 - **Disaster Recovery → Multi-Site**
 - Migration/Coexistence with Legacy DC Networks and 'Disaggregated DCs' Model → Physical Remote Leaf
 - Baremetal Cloud Integration → Virtual Pod (vPod)
 - Extending ACI to the Cloud
 - Connecting the users to the Multi-Cloud DC
 - ACI and SDA Integration
 - ACI and SDWAN Integration

Disaster Recovery Use Case

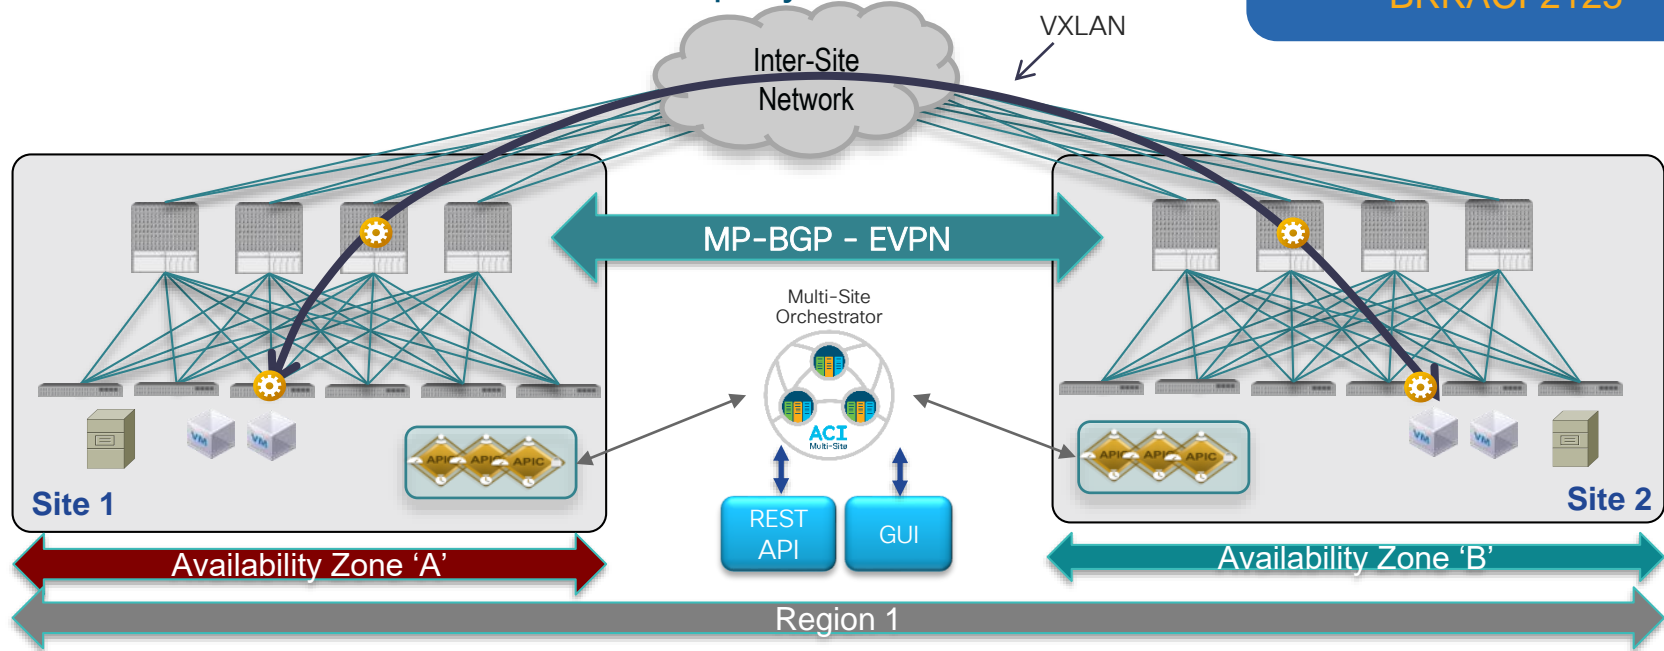
- Drive **cost** efficiencies through re-use of infrastructure and processes
- Integrate Disaster Recovery into day to day **operations**
- Make **capacity** growth sustainable through repurposed infrastructure and shared resources leveraging virtualization
- Provide a DR capability that is acceptable for any framework to all subsidiaries (**multi-tenancy**)
- Supports failover of passive services
 - VMware Site Recovery Manager (**SRM**)
 - Microsoft Cluster Services (**MSCS**)
 - IBM HA Clustering Multi-proc (**HACMP**) / (**PowerHA**)
 - Etc..
- Global Server Load Balancing, Route Health Injection or LISP (**Path Redirection**)



ACI Multi-Site

The Ideal Architecture for DR Deployments

For More Information on
ACI Multi-Site:
BRKACI-2125



- 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

cisco *Live!*

ACI Multi-Site

Hardware Requirements

- Support all ACI leaf switches (1st Generation, -EX and -FX)
 - Only -EX spine (or newer) to connect to the inter-site network
 - New 9364C non modular spine (64x40G/100G ports) supported for Multi-Site from ACI 3.1 release (shipping)
 - 1st generation spines (including 9336PQ) not supported
- Can still leverage those for intra-site leaf to leaf communication

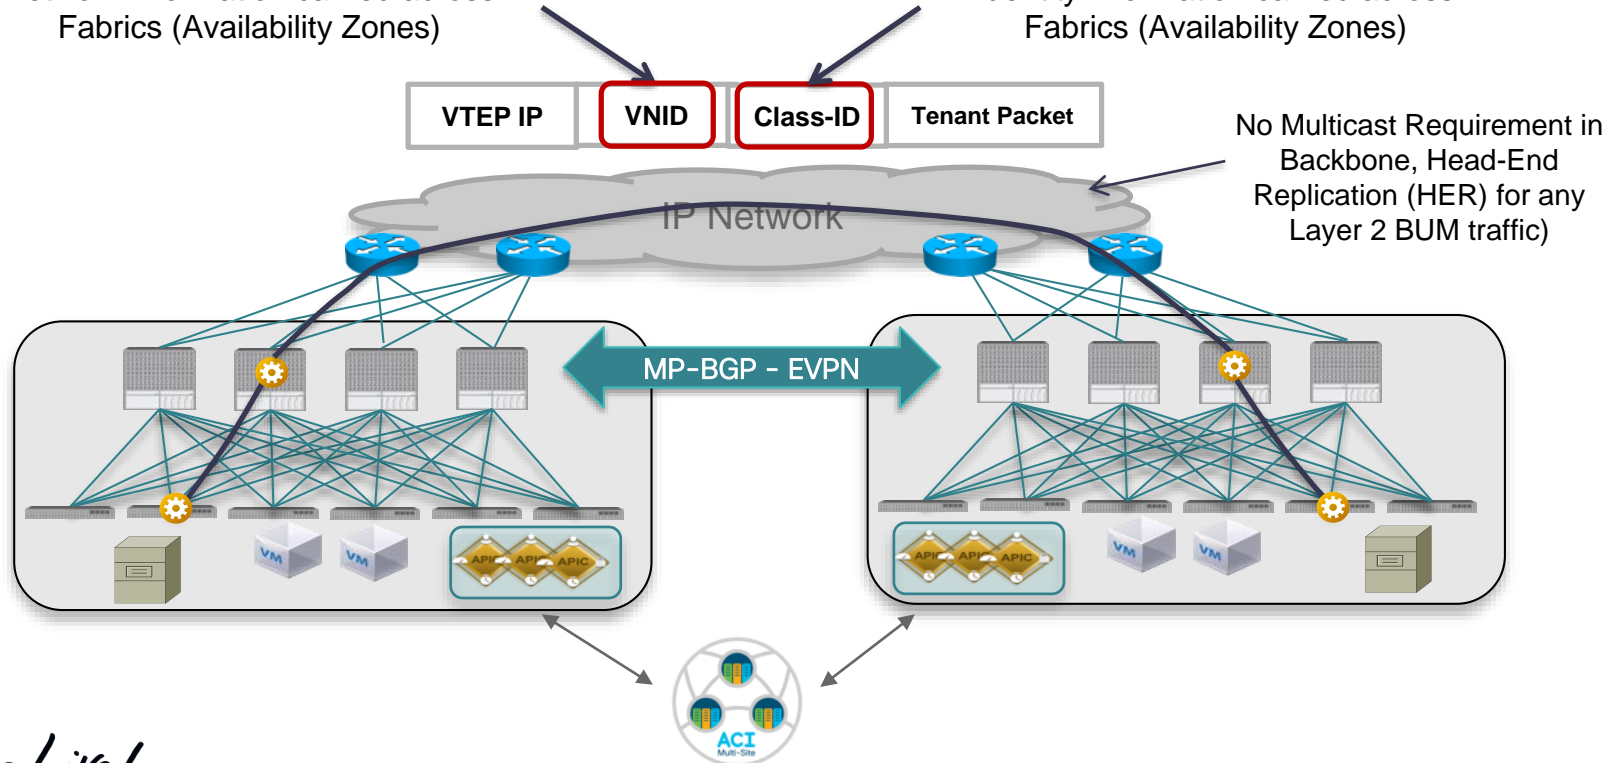


ACI Multi-Site

Network and Identity Extended between Fabrics

Network information carried across Fabrics (Availability Zones)

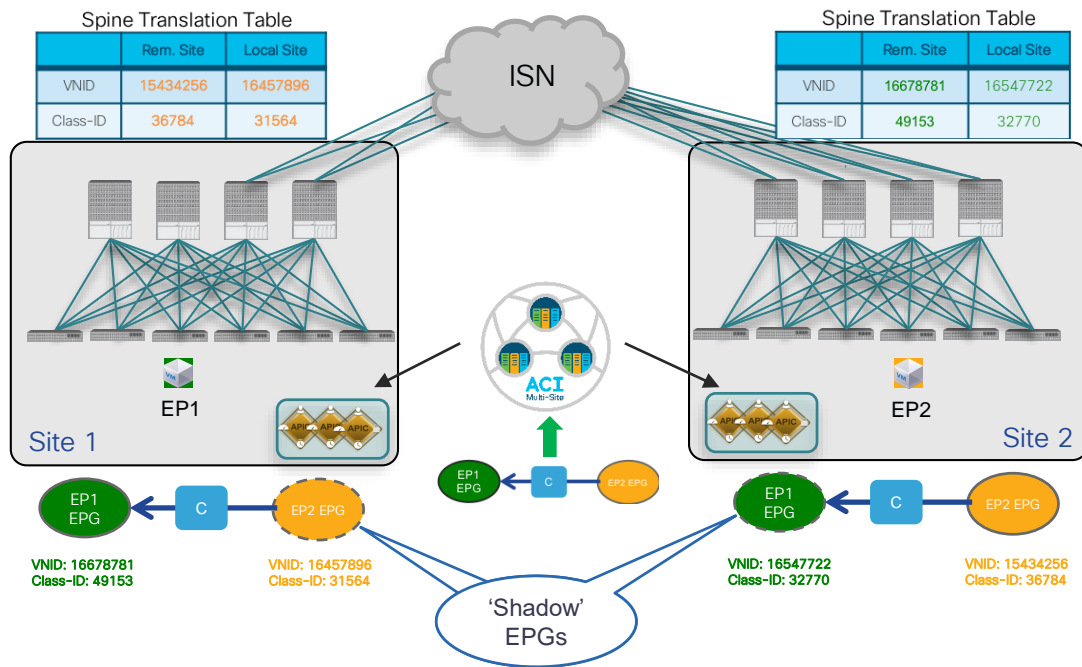
Identity information carried across Fabrics (Availability Zones)



ACI Multi-Site

Inter-Site Policies and Spines' Translation Tables

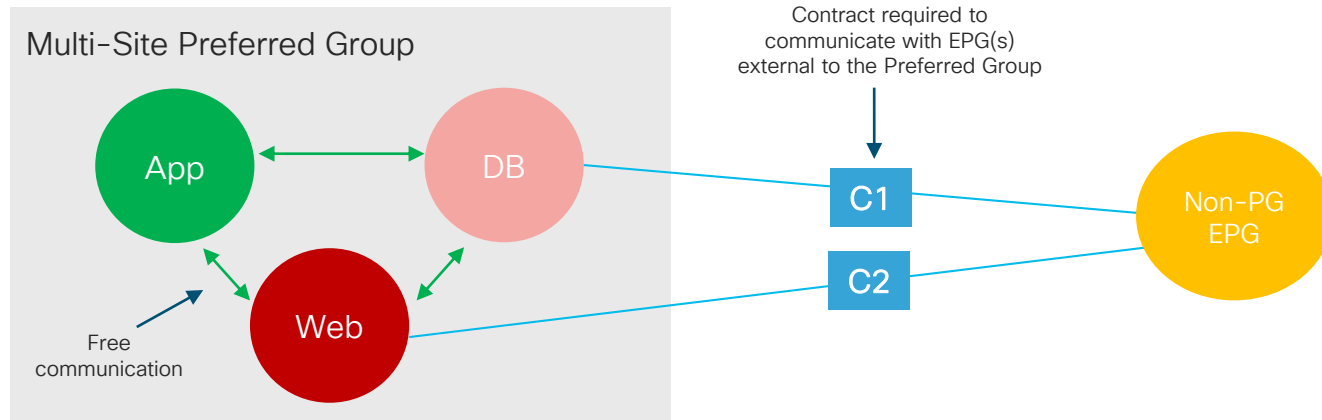
- Inter-Site policies defined on the ACI Multi-Site Orchestrator are pushed to the respective APIC domains
 - End-to-end policy consistency
 - Creation of 'Shadow' EPGs to locally represent the policies
- Inter-site communication requires the installation of translation table entries on the spines (namespace normalization)
- Up to ACI release 4.0(1) translation entries are populated only in two cases:
 - Stretched EPGs/BDs
 - Creation of a contract between not stretched EPGs



ACI Multi-Site

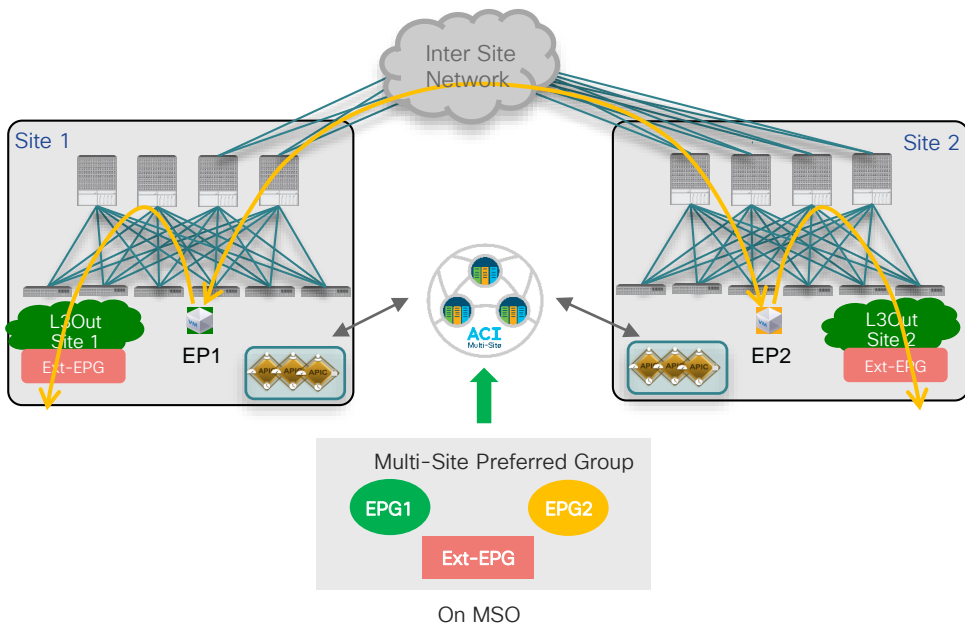
Removing Policy Enforcement: Preferred Groups

MSO 2.0(2)
Release



- "VRF unenforced" not supported with Multi-Site
- Multi-Site Preferred Group configuration from the Multi-Site Orchestrator is supported from MSO 2.0(2) release
 - Creates 'shadow' EPGs and translation table entries 'under the hood' to allow 'free' inter-site communication
 - 250 Preferred Groups supported as ACI release 4.1(1)
- Typically desired in legacy to ACI migration scenarios

Removing Policy Enforcement Preferred Groups for E-W and N-S Flows



- Adding internal EPGs and External EPGs (associated to L3Outs) to the Preferred Group allows to enable free east-west and north-south connectivity
- When adding the Ext-EPG to the Preferred Group:
 - Can't use 0.0.0.0/0 for classification, needs more specific prefixes
 - As workaround it is possible to use 0.0.0.0/1 and 128.0.0.0/1 to achieve the same result
 - Must ensure Ext-EPG is a stretched object

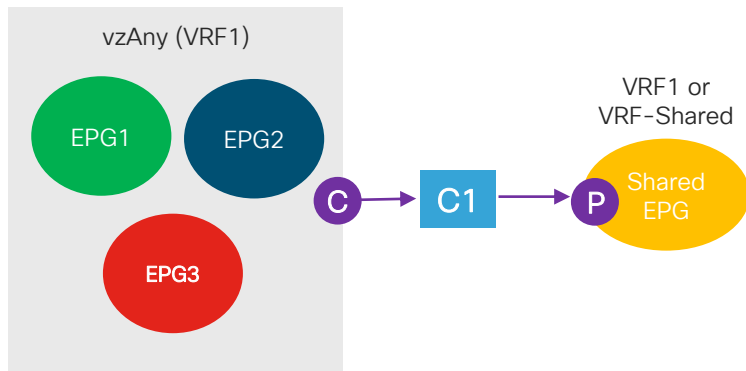
ACI Multi-Site

vzAny Support (MSO 2.2(4) Release)

MSO 2.2(4)
Release

What is vzAny? Logical object representing all the EPGs in a VRF

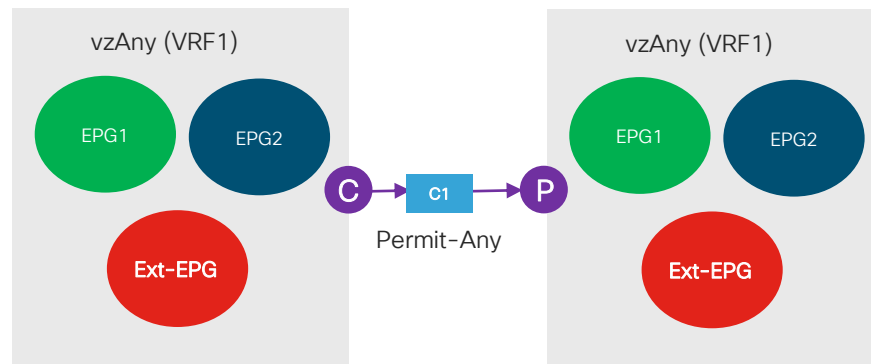
Use case 1: Many-to-One communication (Shared Services)



- Multiple EPGs part of a specific VRF1 consume the services provided by a shared EPG (part of VRF1 or of a VRF-shared)
- VRF-shared can be part of the same tenant or of a different tenant

CISCO *Live!*

Use case 2: Enable free communication inside a VRF

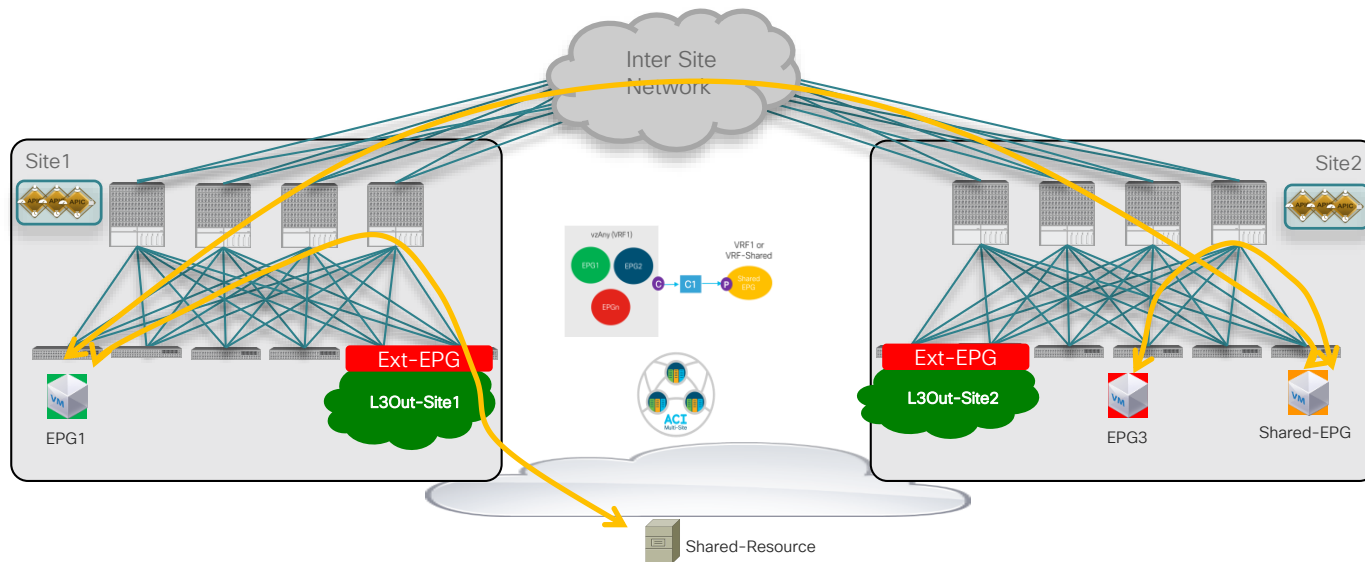


- vzAny provides and consumes a contract with an associated “Permit-any” filter
- Use ACI fabric only for network connectivity without policy enforcement
- Equivalent to “VRF unenforced”

ACI Multi-Site and vzAny

Many-to-One Communication (Shared Services)

MSO 2.2(4)
Release




- Proper translation entries are created on the spines of both fabrics to enable east-west communication
- Supported also for Shared Services behind an L3Out

ACI Multi-Site and vzAny

Enable Inter-Site Free Communication Inside a VRF

MSO 2.2(4)
Release




- Proper translation entries are created on the spines of both fabrics to enable east-west communication
- Supported also for connecting to the external Layer 3 domain

ACI Multi-Site Per Bridge Domain Behavior

Should be the behavior for the majority of BDs with Multi-Site


Layer 3 only across sites



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




IP Mobility without BUM flooding



- 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**



Layer 2 adjacency across Sites



- 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**



*'Live' migration officially supported from ACI release 3.2

ACI Multi-Site

Continuous Scale Improvements

	ACI Release 3.0	ACI Release 3.1	ACI Release 3.2	ACI Release 4.2
	MSO 1.0	MSO 1.1	MSO 1.2	MSO 2.1
Number Of Sites	5	8	10	12
Max Leafs (across sites)	250	800	1200	1600
Tenants	100	200	300	400
VRF	400	400	800	1000
BD	800	2,000	3,000	4,000
EPGs	800	2,000	3,000	4,000
Contracts	1,000	2,000	3,000	4,000
L3Out External EPGs	500	500	500	500
Isolated EPGs	N/A	400	400	400

For more information please refer to:

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/verified-scalability/Cisco-ACI-Verified-Scalability-Guide-422.pdf>

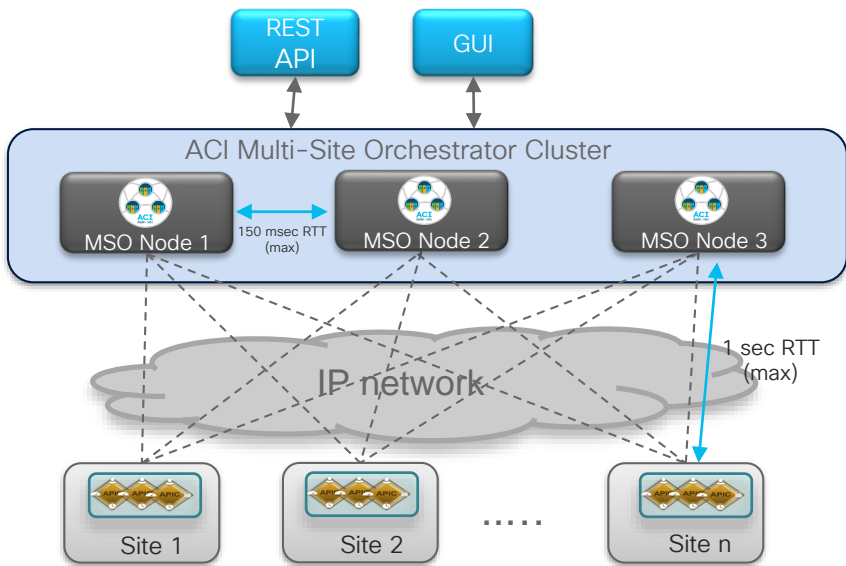


Multi-Site Multi-Site Orchestrator, Schemas and Templates



ACI Multi-Site


Multi-Site Orchestrator (MSO)



- Three MSO nodes are clustered and run concurrently (active/active)
 - Typical database redundancy considerations (minority/majority rules)
 - Up to 150 msec RTT latency supported between MSO nodes
- OOB Mgmt connectivity to the APIC clusters deployed in separate sites
 - Up to 1 sec RTT latency between MSO and APIC nodes
- Main functions offered by MSO:
 - Monitoring the health-state of the different ACI Sites
 - Provisioning of day-0 infrastructure configuration to establish inter-site EVPN control plane and VXLAN data plane
 - Defining and provisioning tenant policies
 - Day-2 operation functionalities

ACI Multi-Site Orchestrator

VM Based MSO Cluster



- Supported from the beginning (MSO release 1.0(1))
- Each Cisco ACI Multi-Site Orchestrator node is packaged in a VMware vSphere virtual appliance
- For high availability, you should deploy each Cisco ACI Multi-Site Orchestrator virtual machine on its own VMware ESXi host
- Requirements for MSO Release 1.2(x) and above:
 - VMware ESXi 6.0 or later
 - Minimum of eight virtual CPUs (vCPUs), 24 Gbps of memory, and 100 GB of disk space

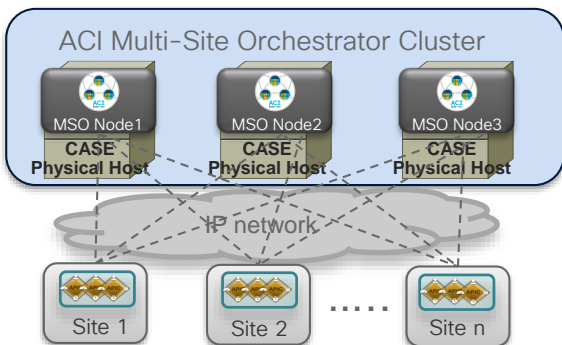
ACI Multi-Site Orchestrator

Cisco Application Service Engine (CASE) Based MSO Cluster


MSO 2.2(3)
Release

- CASE cluster available in different form factors (physical, on-prem VM, AWS instance)
- MSO is installed as an App on the CASE cluster
- Recommended MSO cluster deployment option going forward

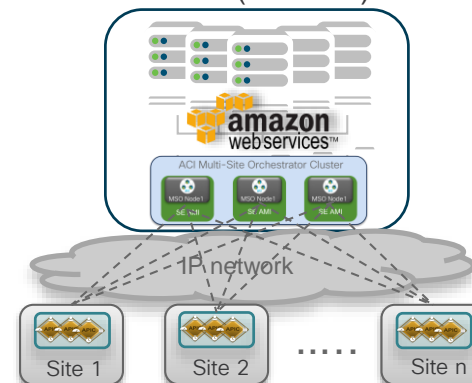
1 MSO App on CASE physical form factor



2 MSO App on CASE VM form factor (on premises)

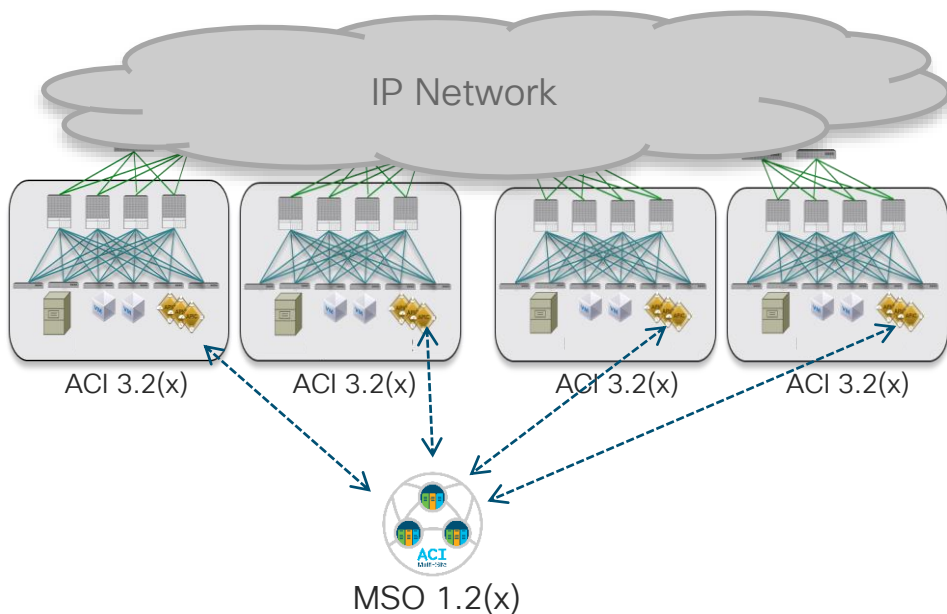


3 MSO App on CASE VM form factor (on AWS)



ACI Multi-Site

MSO and APIC Release Dependency (Pre-MSO 2.2(1) Release)




- Pre-MSO 2.2(1) release, MSO and ACI releases have to be aligned
 - For example MSO 1.2(x) is used with all sites running ACI release 3.2(x)
- Different ACI versions across sites are only supported during an ACI SW upgrade procedure
- The supported order of upgrade is:
 1. APIC firmware
 2. Switches firmware
 3. MSO

} For each Site
} As last task

ACI Multi-Site Interversion

Decoupling MSO and APIC Releases

MSO 2.2(1)
Release



- MSO “interversion” support is available with MSO release 2.2(1)
- Different ACI versions across sites can be supported at steady state
- MSO gets visibility into what functionalities are supported in each fabric (based on the specific ACI releases)
 - Preventing the deployment of unsupported functionalities

ACI Multi-Site Interversion

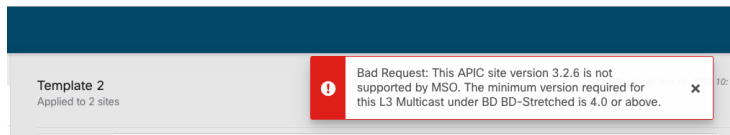
Decoupling MSO and APIC Releases

SW dependency for different ACI functionalities

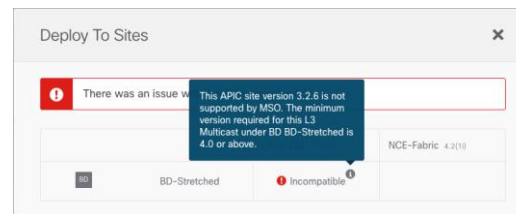
Feature	Minimum APIC Version
ACI Multi-Pod Support	Release 3.2(1)
Service Graphs (L4-L7 Services)	Release 3.2(1)
External EPGs	Release 3.2(1)
ACI Virtual Edge VMM Support	Release 3.2(1)
DHCP Support	Release 3.2(1)
Consistency Checker	Release 3.2(1)
CloudSec Encryption	Release 4.0(1)
Layer 3 Multicast	Release 4.0(1)
MD5 Authentication for OSPF	Release 4.0(1)
EPG Preferred Group	Release 4.0(2)
Host Based Routing	Release 4.1(1)
Intersite L3Out	Release 4.2(1)



Version check during the “Save” operation of a template

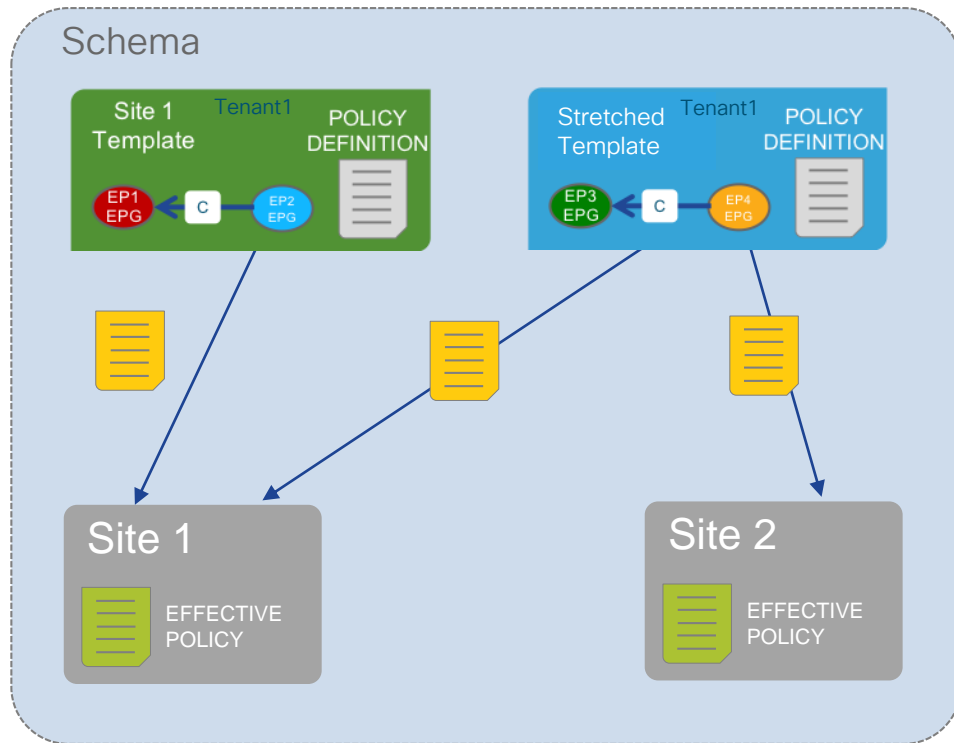


Version check at deployment time for a template



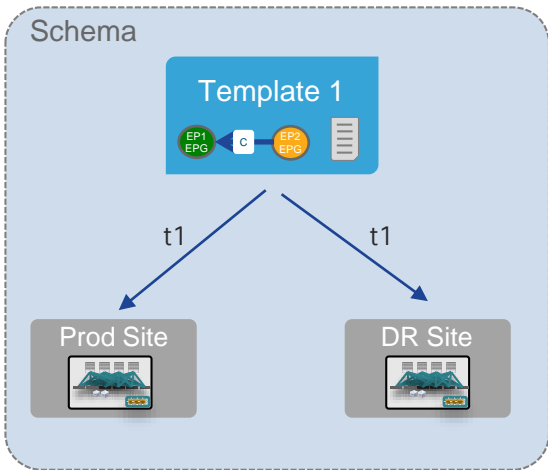
ACI Multi-Site MSO Schema and Templates

- Template = ACI policy definition (ANP, EPGs, BDs, VRFs, etc.)
- Schema = container of Templates sharing a common use-case
 - As an example, a schema can be dedicated to a Tenant
- The template is currently the atomic unit of change for policies
 - Such policies are concurrently pushed to one or more sites
- Scope of change: policies in different templates can be pushed to separate sites at different times

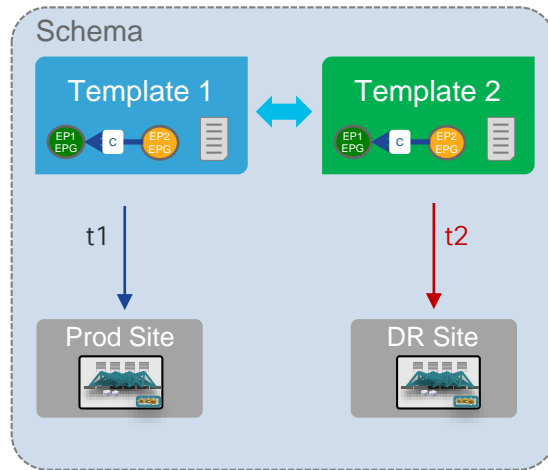


ACI Multi-Site

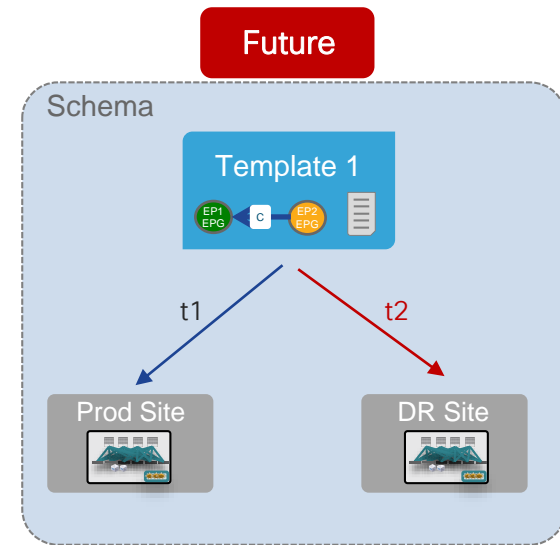
Schema and Templates Definition for the DR Use Case



- Single Template associated to Prod and DR Sites
- Any change applied to the template is pushed to both sites simultaneously
- Easiest way to keep consistent policies deployed across sites



- Separate Template associated to Prod and DR Sites (can use cloning)
- Changes made to a template can be applied only to the mapped site
- Requires sync between the two templates (manual or performed by an higher level Orchestration tool)

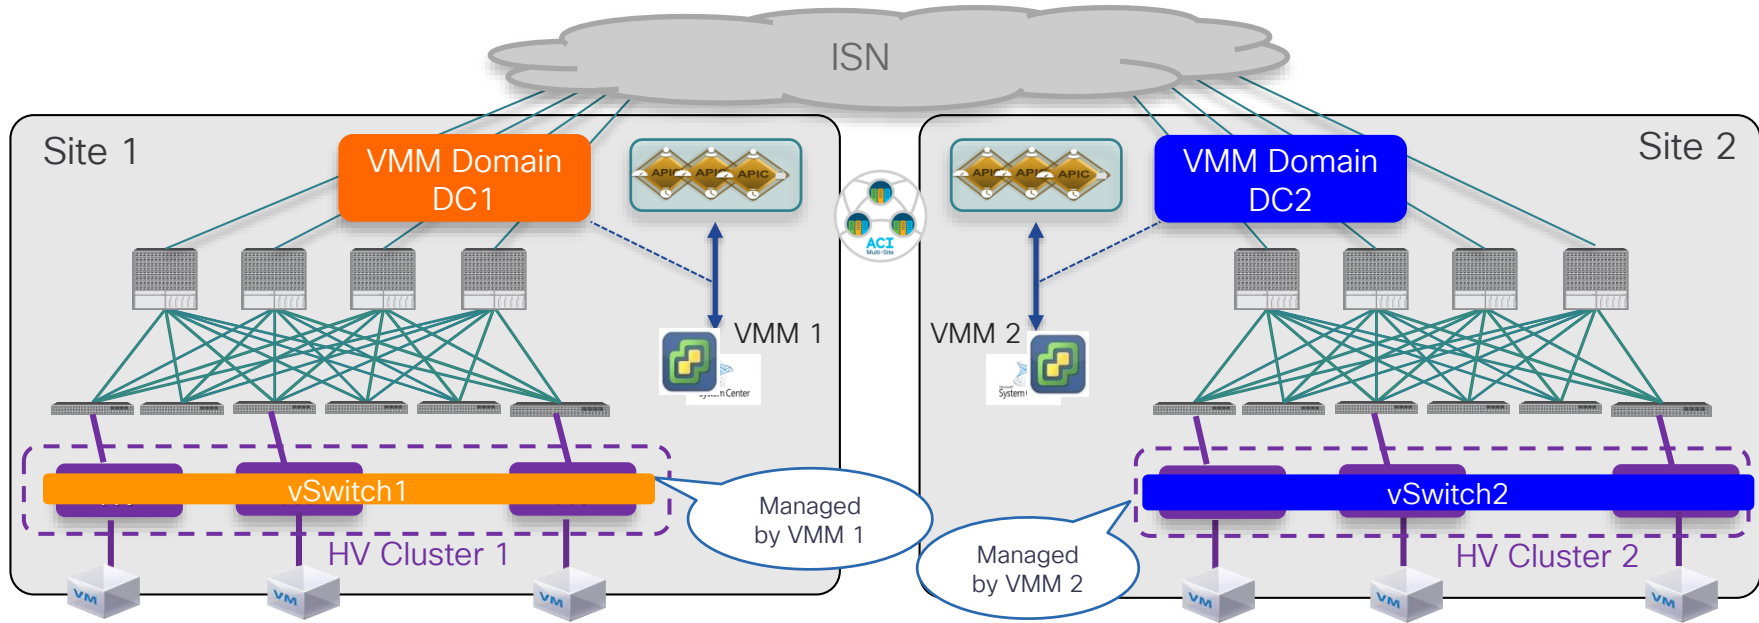


- Single Template associated to Prod and DR Sites
- Capability of independently apply changes to each site
- Brings together the advantages of the previous two options

Multi-Site and Virtual Machine Manager (VMM) Integration

ACI Multi-Site and VMM Integration

Option 1 – Separate VMM per Site



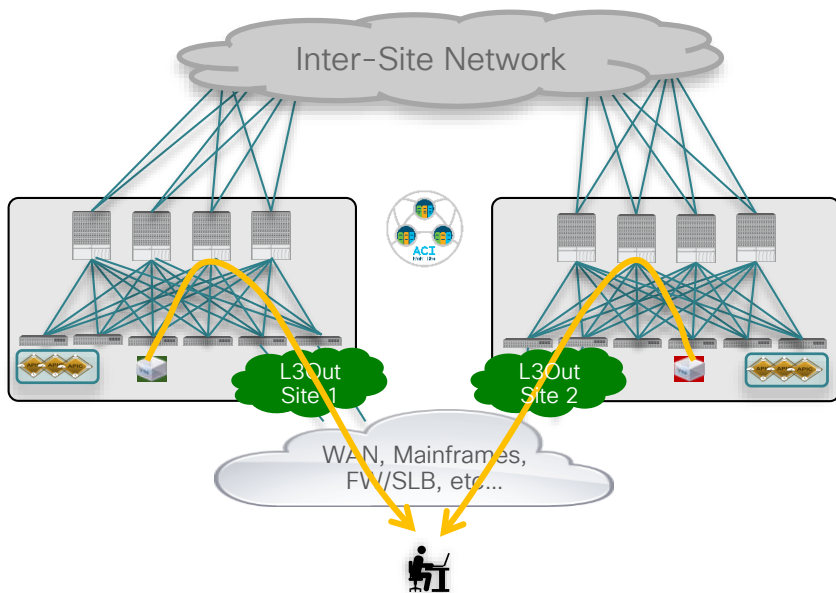
- Typical deployment model for an ACI Multi-Site
- Creation of separate VMM domains in each site, which are then exposed to the Multi-Site Orchestrator

Multi-Site Connectivity to the External L3 Domain


Problem Statement

Behavior before ACI Release 4.2(1)

Supported Design



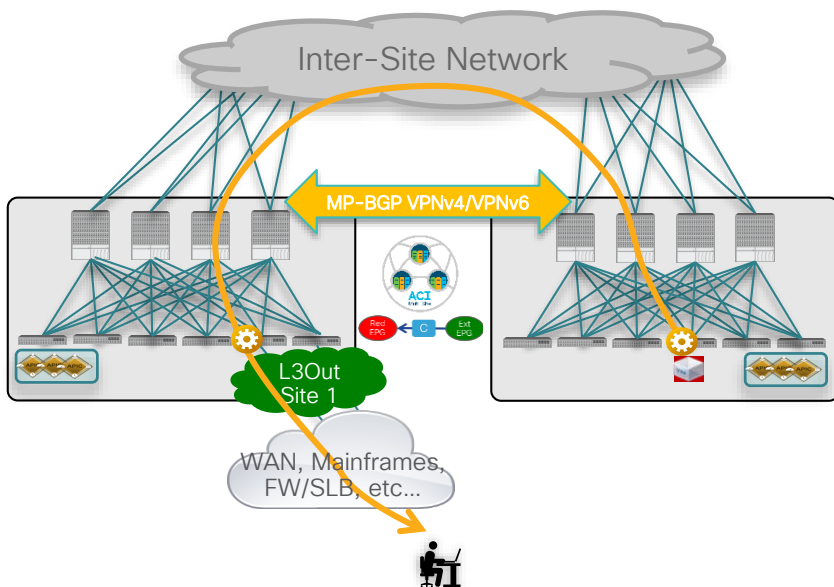
Not Supported Design



ACI Multi-Site and L3Out

Support of Intersite L3Out

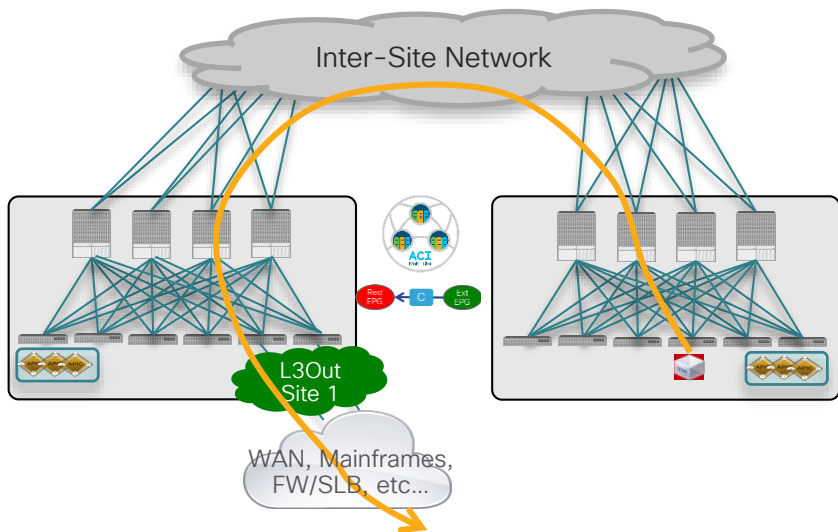
ACI 4.2(1)
Release




- Starting with ACI Release 4.2(1) it is possible for endpoints in a site to send traffic to resources (WAN, Mainframes, FWs/SLBs, etc.) accessible via a remote L3Out connection
- External prefixes are exchanged across sites via MP-BGP VPNv4/VPNv6 sessions between spines
- Traffic will be directly encapsulated to the TEP of the remote BL nodes
 - The BL nodes will get assigned an address part of an additional (configurable) prefix that must be routable across the ISN
 - This routable TEP pool can be configured on MSO or on APIC
- Same solution will also support transit routing across sites (L3Out to L3Out)

ACI Multi-Site and Intersite L3Out Supported Scenarios

ACI 4.2(1)
Release



- Endpoint to remote L3Out communication (intra-VRF)
- Endpoint to remote L3Out communication (inter-VRF)

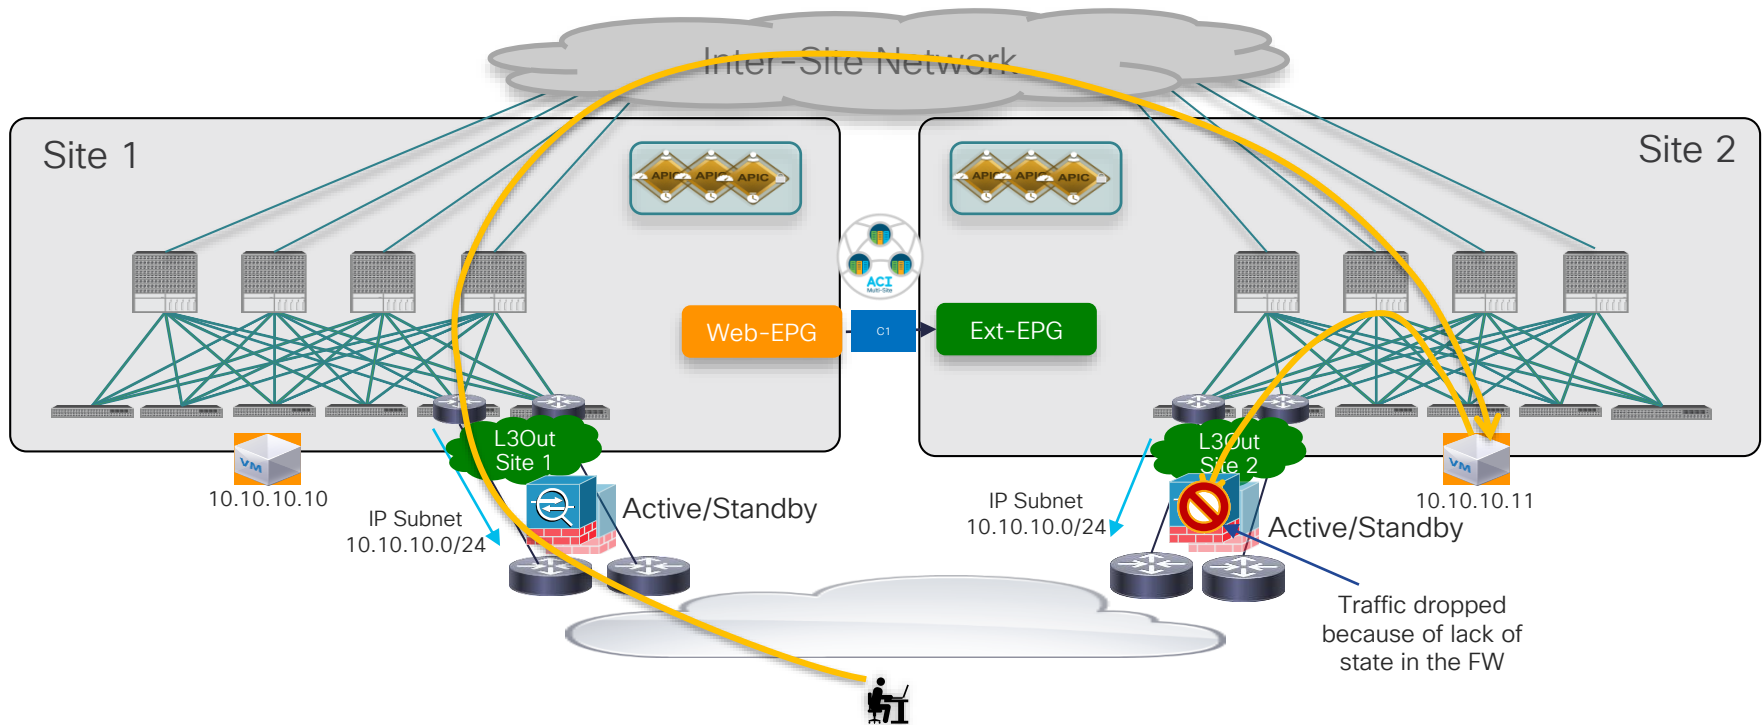


- Inter-site transit routing (intra-VRF)
- Inter-site transit routing (inter-VRF)

Solving Asymmetric Routing Issues with the External Network

Multi-Site and L3Out

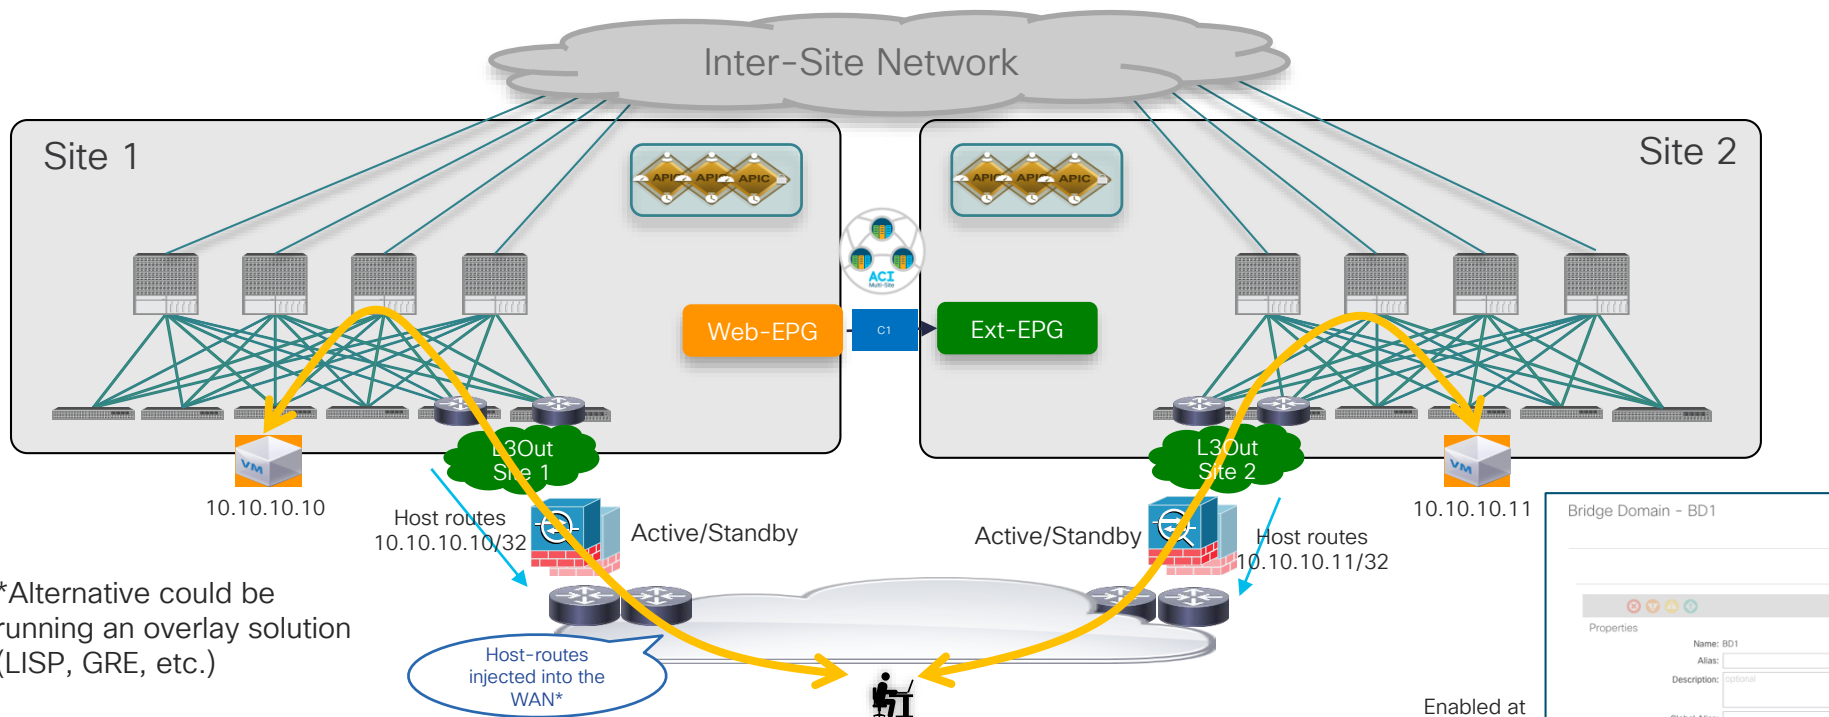
Endpoints Normally Use Local L3Outs for Outbound Traffic



Solving Asymmetric Routing Issues

Use of Host-Routes Advertisement

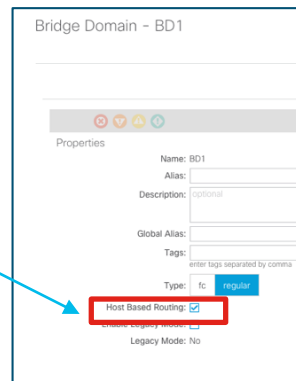
ACI 4.0(1)
Release



*Alternative could be running an overlay solution (LISP, GRE, etc.)

- Ingress optimization requires host-routes advertisement on the L3Out
 - Native support on ACI Border Leaf nodes available from ACI release 4.0(1)
 - Supported also on GOLF L3Outs (enabled at the VRF level)

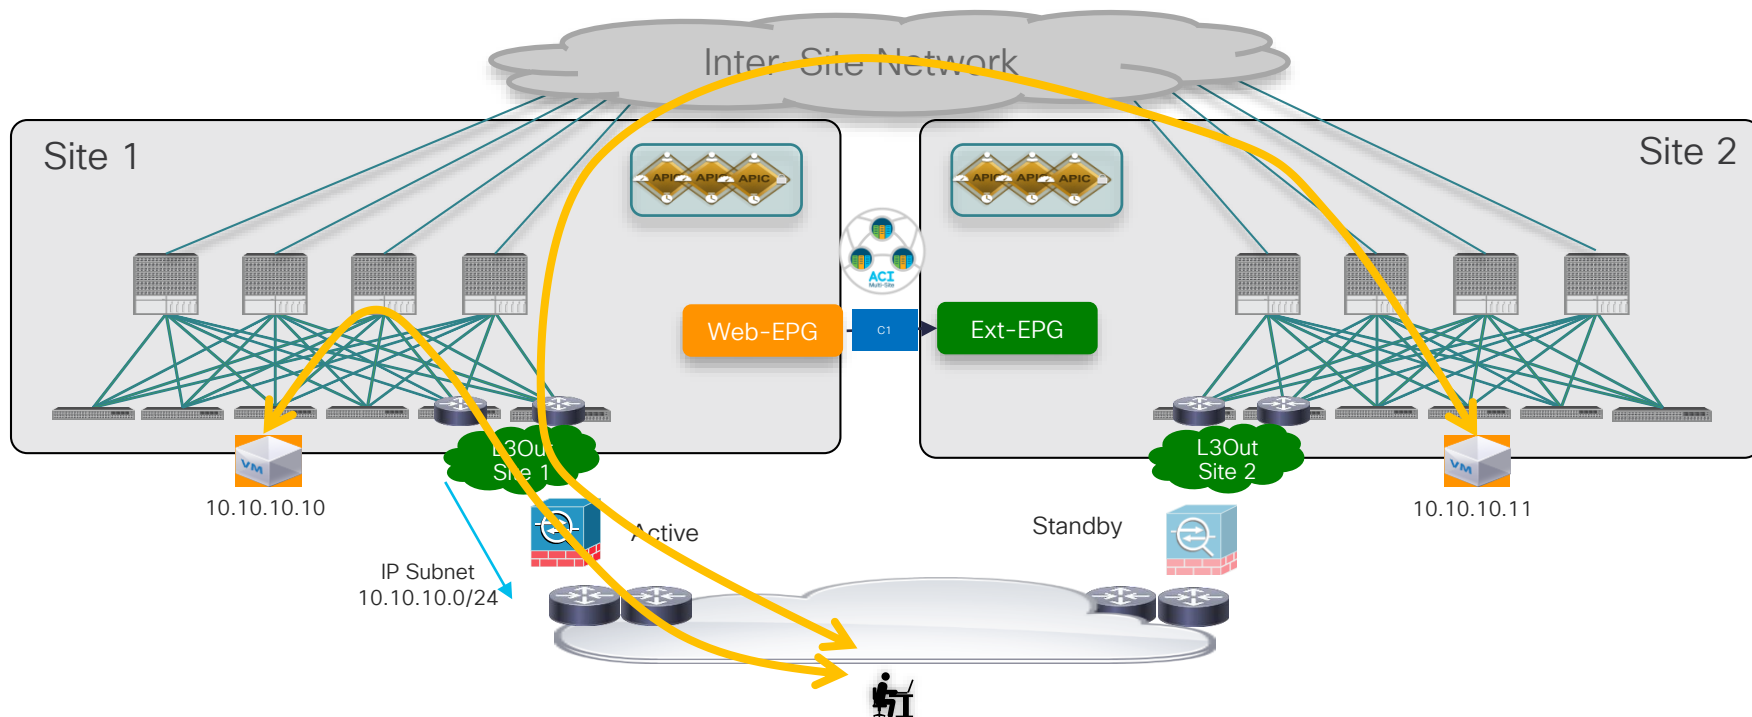
Enabled at the BD level



Solving Asymmetric Routing Issues

Use of Active/Standby FW Pair Deployed across Sites

ACI 4.2(1)
Release




- Inbound and outbound flows are forced through the site with the active perimeter FW node
 - Common scenario in a Multi-Pod deployment, less desirable with Multi-Site
- Requires Intersite L3Out support (ACI release 4.2(1))

Multi-Site Network Services Integration


Multi-Site and Network Services

Integration Models


Deployment options fully supported with ACI Multi-Pod



- Active and Standby pair deployed across Pods
- Currently supported only if the FW is in L2 mode or in L3 mode but acting as default gateway for the endpoints
- From ACI 4.2(1) will be also supported as perimeter FW



- Active/Active FW cluster nodes stretched across Sites (single logical FW)
- Requires the ability of discovering the same MAC/IP info in separate sites at the same time
- **Not supported**



- **Recommended deployment model for ACI Multi-Site**
- Option 1: supported from 3.0 for N-S if the FW is connected in L3 mode to the fabric → mandates the deployment of traffic ingress optimization
- Option 2: supported from 3.2 release with the use of Service Graph with Policy Based Redirection (PBR)


Multi-Site Integration with ACI Multi-Pod

ACI Multi-Pod and Multi-Site


Main Use Cases

ACI 3.2(1)
Release

- Adding a Multi-Pod Fabric as a 'Site' on the Multi-Site Orchestrator (MSO)

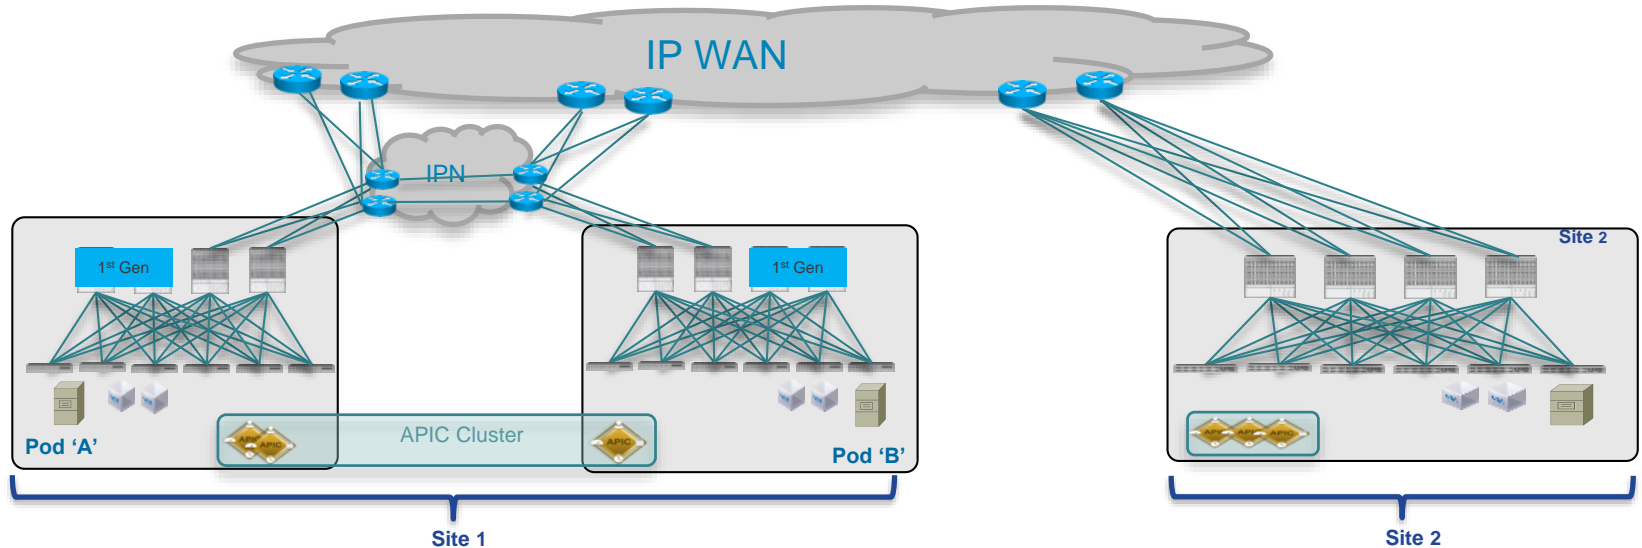


- Converting a single Pod Fabric (already added to MSO) to a Multi-Pod fabric



ACI Multi-Pod and Multi-Site

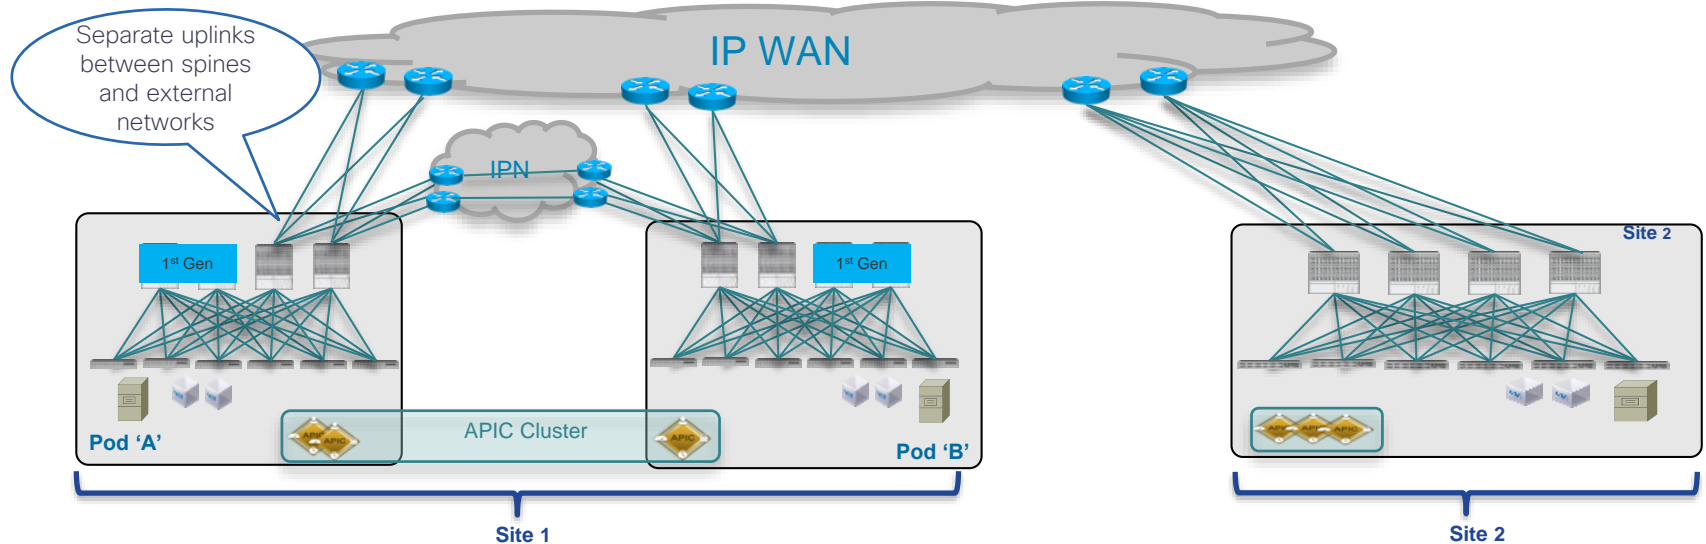
Connectivity between Pods and Sites



- Only 2nd generation spines must be connected to the external network
 - Need to add 2nd gen spines in each Pod (at least two per Pod) and migrate connections to the IPN from 1st gen spines to 2nd gen spines
- Single 'infra' L3Out and set of uplinks to carry both Multi-Pod and Multi-Site East-West traffic

Connectivity between Pods and Sites

Not Supported Topology



- Only 2nd generation spines must be connected to the external network
 - Need to add 2nd gen spines in each Pod (at least two per Pod) and migrate connections to the IPN from 1st gen spines to 2nd gen spines
- Single 'infra' L3Out and set of uplinks to carry both Multi-Pod and Multi-Site East-West traffic

ACI Multi-Site

Where to Go for More Information



- ✓ ACI Multi-Site White Paper

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html>

- ✓ Deploying ACI Multi-Site from Scratch

<https://www.youtube.com/watch?v=HJJ8lznodN0>

- ✓ BRKACI-2125

Agenda

- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - Mapping use cases to the proper solutions
 - Active/Active DC → Multi-Pod
 - Disaster Recovery → Multi-Site
 - Migration/Coexistence with Legacy DC Networks and ‘Disaggregated DCs’ Model → Physical Remote Leaf
 - Baremetal Cloud Integration → Virtual Pod (vPod)
 - Extending ACI to the Cloud
 - Connecting the users to the Multi-Cloud DC
 - ACI and SDA Integration
 - ACI and SDWAN Integration

ACI Remote Physical Leaf

Business Value and Use Cases

For More Information on
ACI Remote Leaf:
[BRKACI-2387](#)



Extending the ACI policy model outside the main datacenter to remote sites distributed over IP Backbone (Telco DCs, CoLo locations, etc.)



Extending ACI fabric policy and L2/L3 connectivity to a small DC site without requiring the deployment of a full-blown ACI Fabric or for migration/coexistence with legacy DC sites



Centralized Policy Management and Control Plane for remote locations

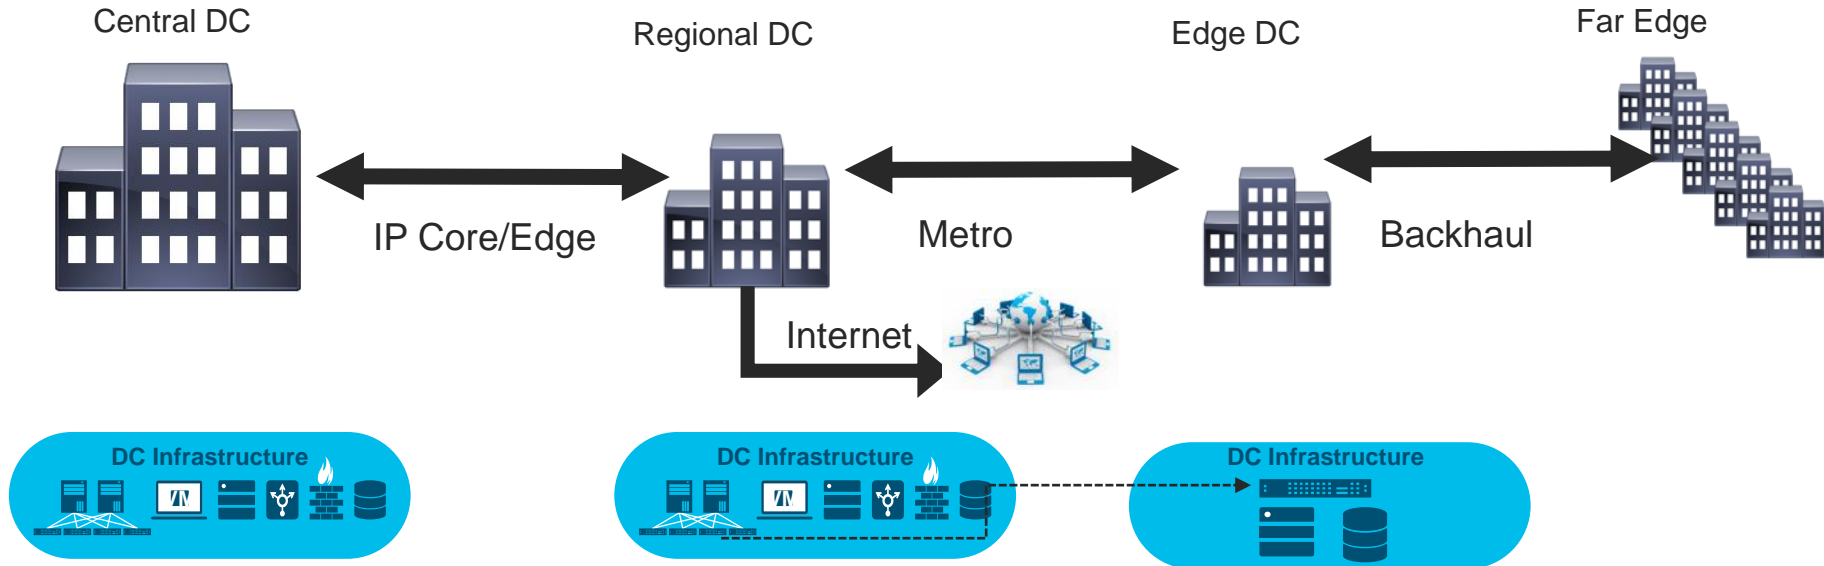


Small form factor solution at locations with space constraints

ACI Remote Physical Leaf


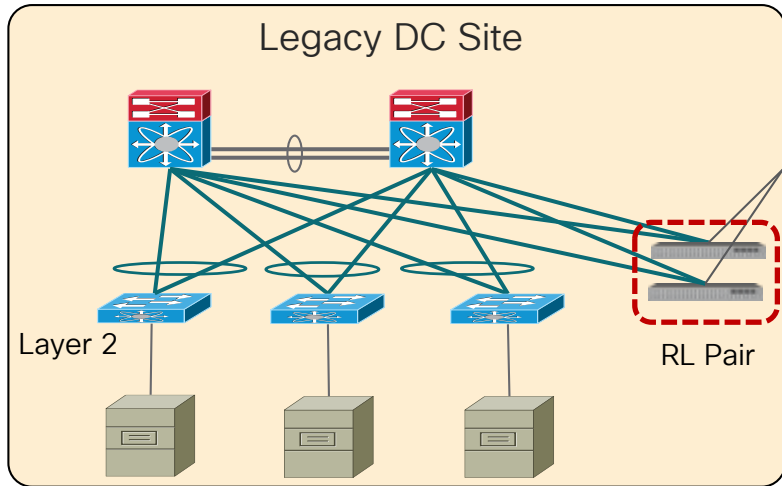
‘Disaggregated DC Model’ for 5G Deployment

- Increased SP DC Footprint with edge transformation, vRAN
- Application/Services are distributed (CUPS, CU-DU split, etc.)



ACI Remote Physical Leaf

Migration/Coexistence Use Case




- Connecting a greenfield ACI fabric to a legacy DC location
 - Single point of management
 - Coexistence
 - Application migration

ACI Remote Physical Leaf

Conceptual Architecture

ACI 3.1(1)
Release

APIC and Spine Nodes (Proxy function) remain at primary Pod(s)

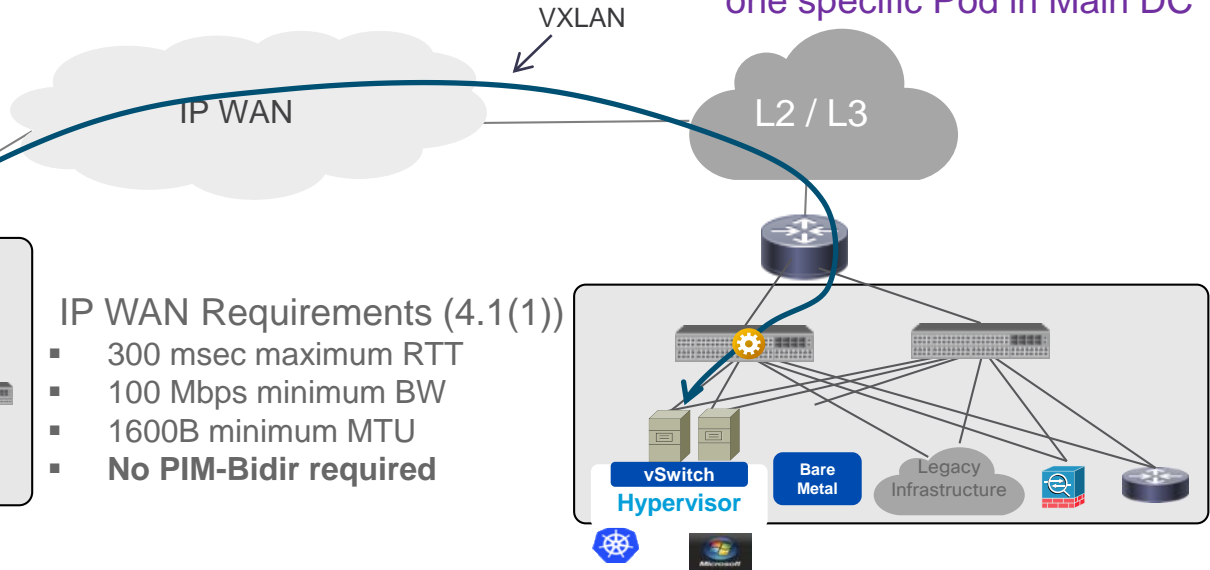


ACI Main DC

IP WAN Requirements (4.1(1))

- 300 msec maximum RTT
- 100 Mbps minimum BW
- 1600B minimum MTU
- No PIM-Bidir required

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



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

ACI Remote Physical Leaf

Hardware/Software Support

ACI Main DC

Supported Spines

Fixed

- 9364C/9332C

Modular

- 9732C-EX
- 9736C-FX
- 9736Q-FX

Remote Location

Supported Leaf

- N93180YC-EX
- N93108TC-EX
- N93180LC-EX
- N93180YC-FX
- N93180TC-FX
- N9348GC-FXP
- N9358GY-FXP
- N93240YC-FX2
- N9336C-FX2
- N93360YC-FX2
- N93216TC-FX2

All hardware from -EX onwards is supported



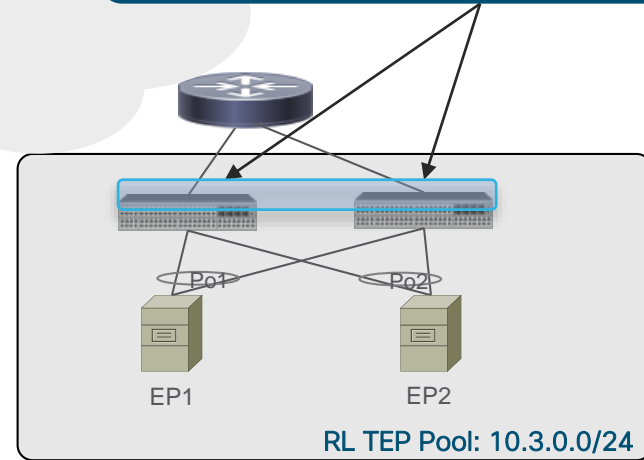
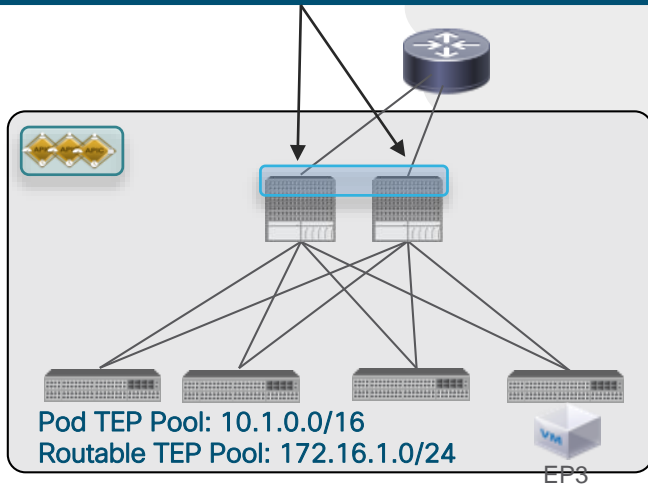
ACI Remote Physical Leaf

TEP Address Pools

Anycast IP on Spines of ACI Main DC
RL-Ucast-TEP
RL-Mcast-TEP

IP WAN IPN

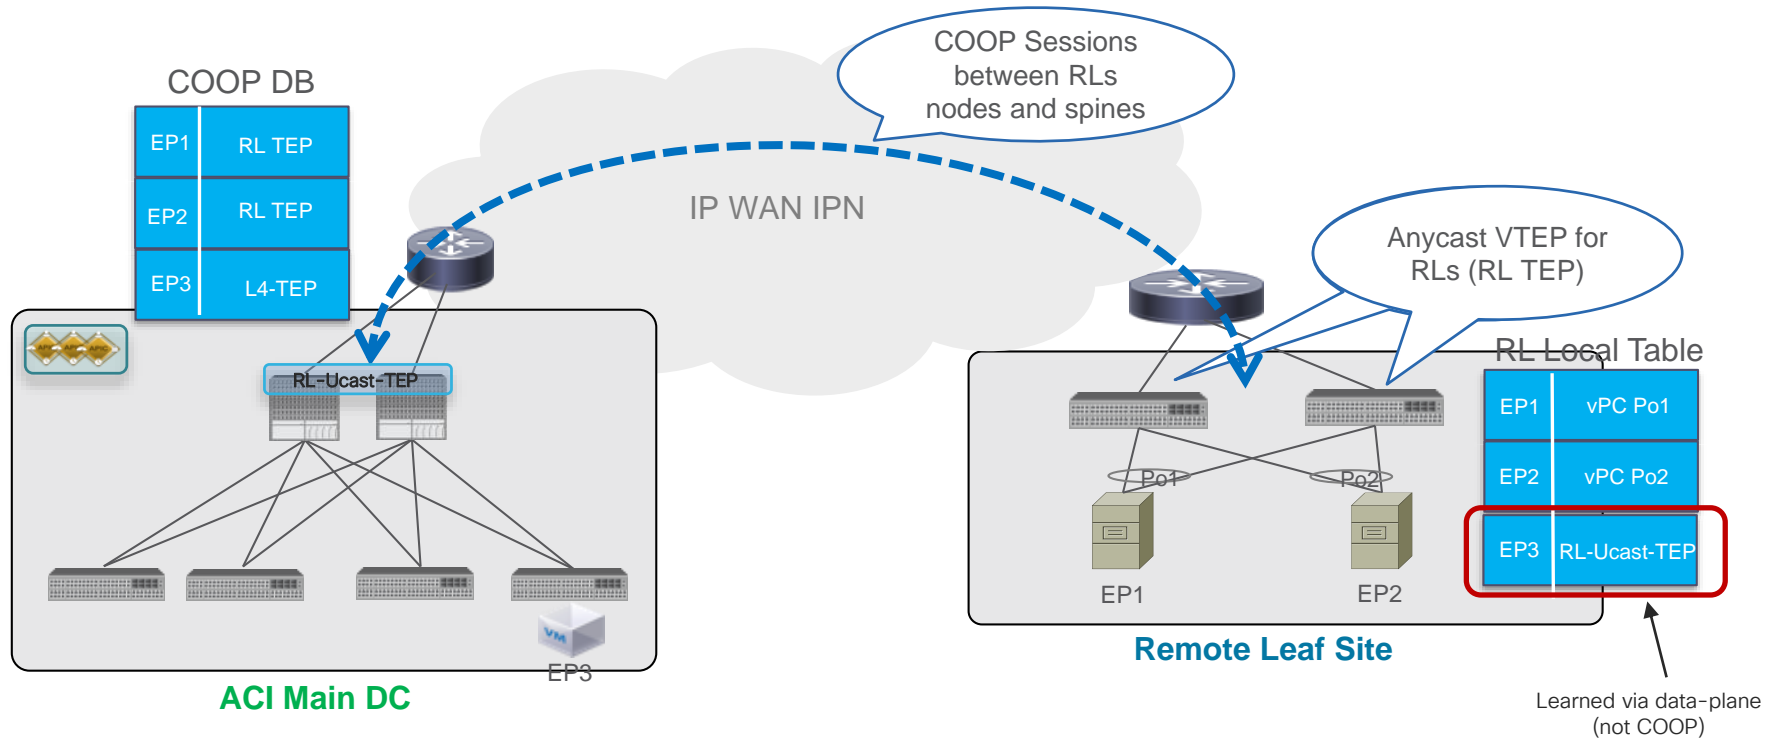
TEP Addresses are allocated from a specific RL
TEP Pool configured on APIC



RL-Ucast-TEP and RL-Mcast-TEP address are allocated from the routable TEP pool associated to the Pod

ACI Remote Physical Leaf

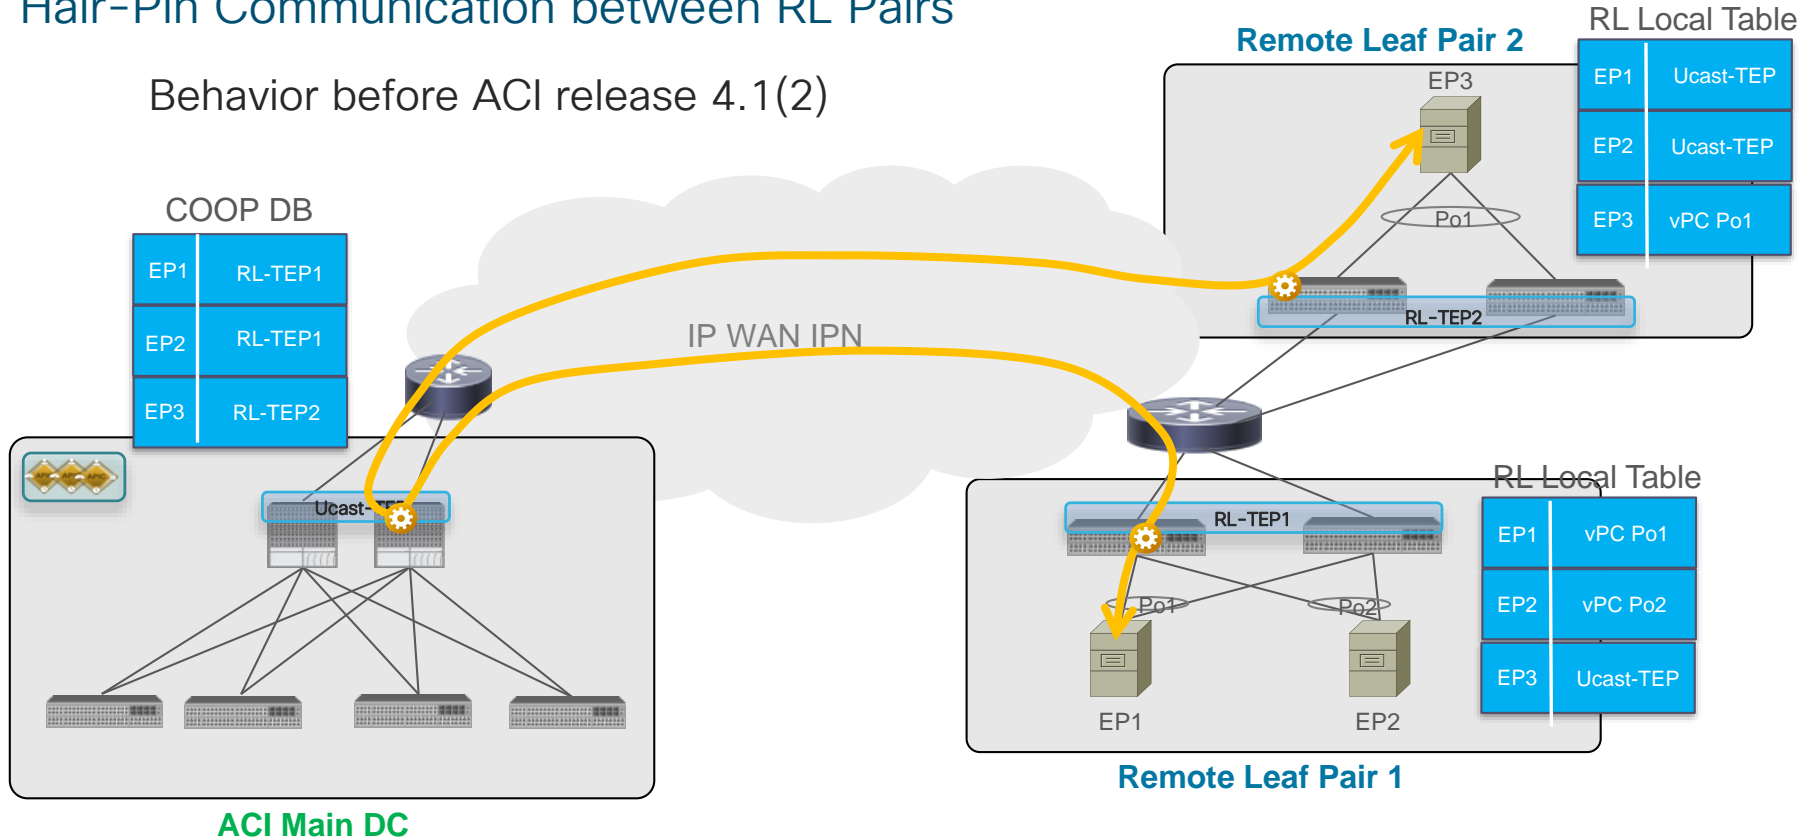
COOP for Announcing Remote Endpoint Information



ACI Remote Physical Leaf

Hair-Pin Communication between RL Pairs

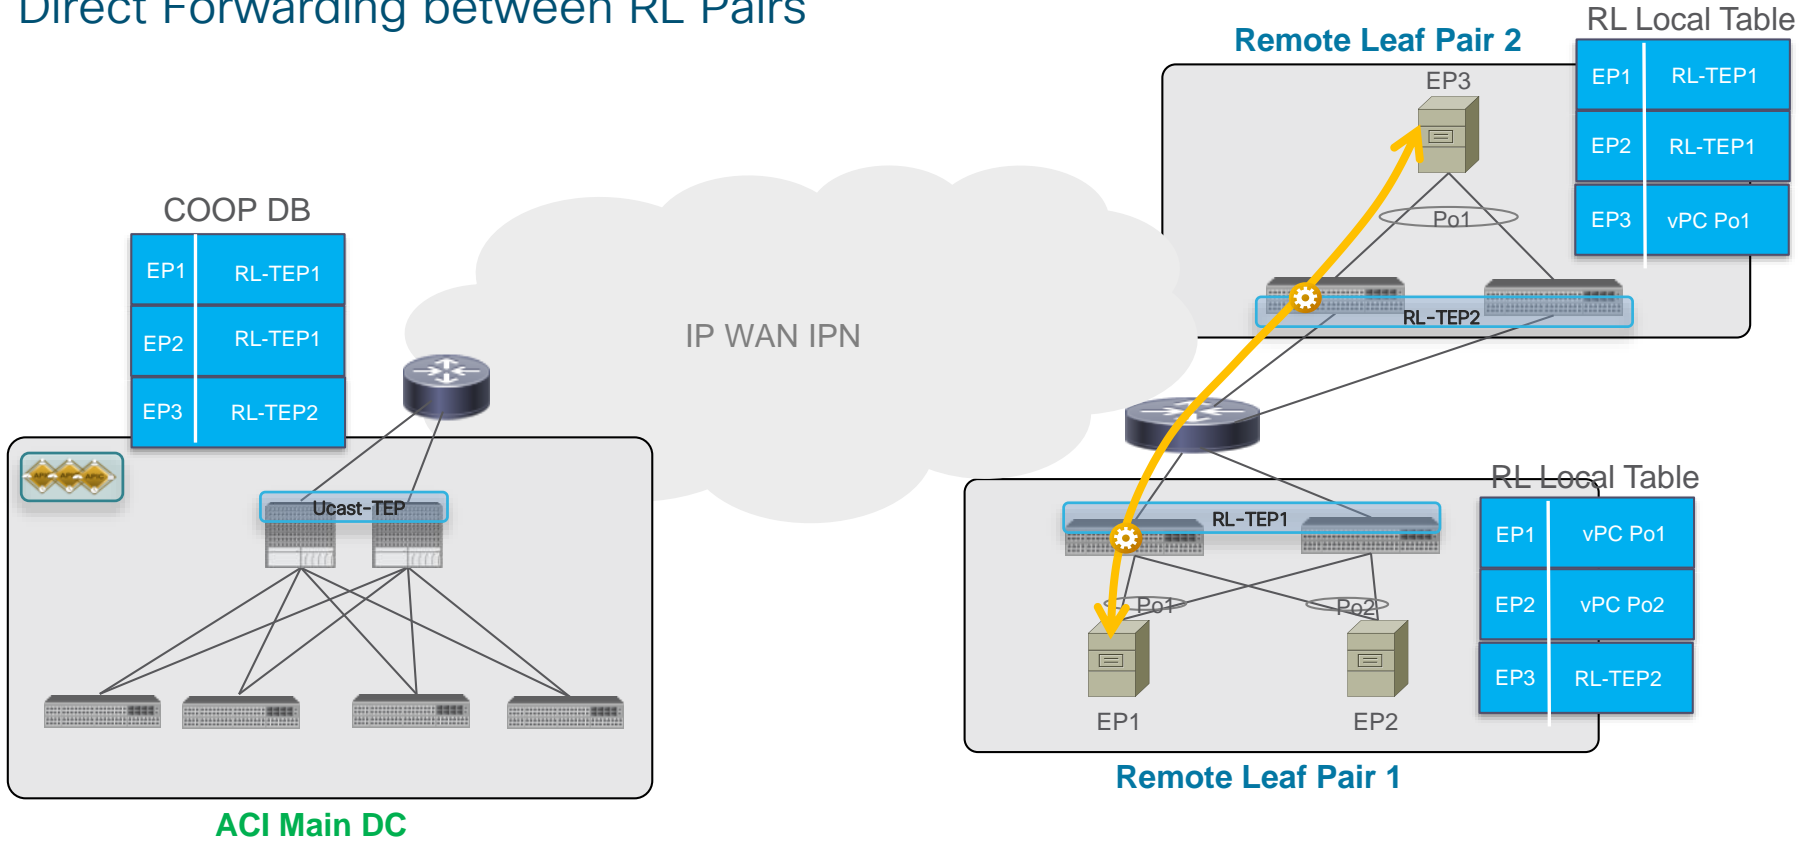
Behavior before ACI release 4.1(2)



ACI Remote Physical Leaf

Direct Forwarding between RL Pairs

ACI 4.1(2)
Release

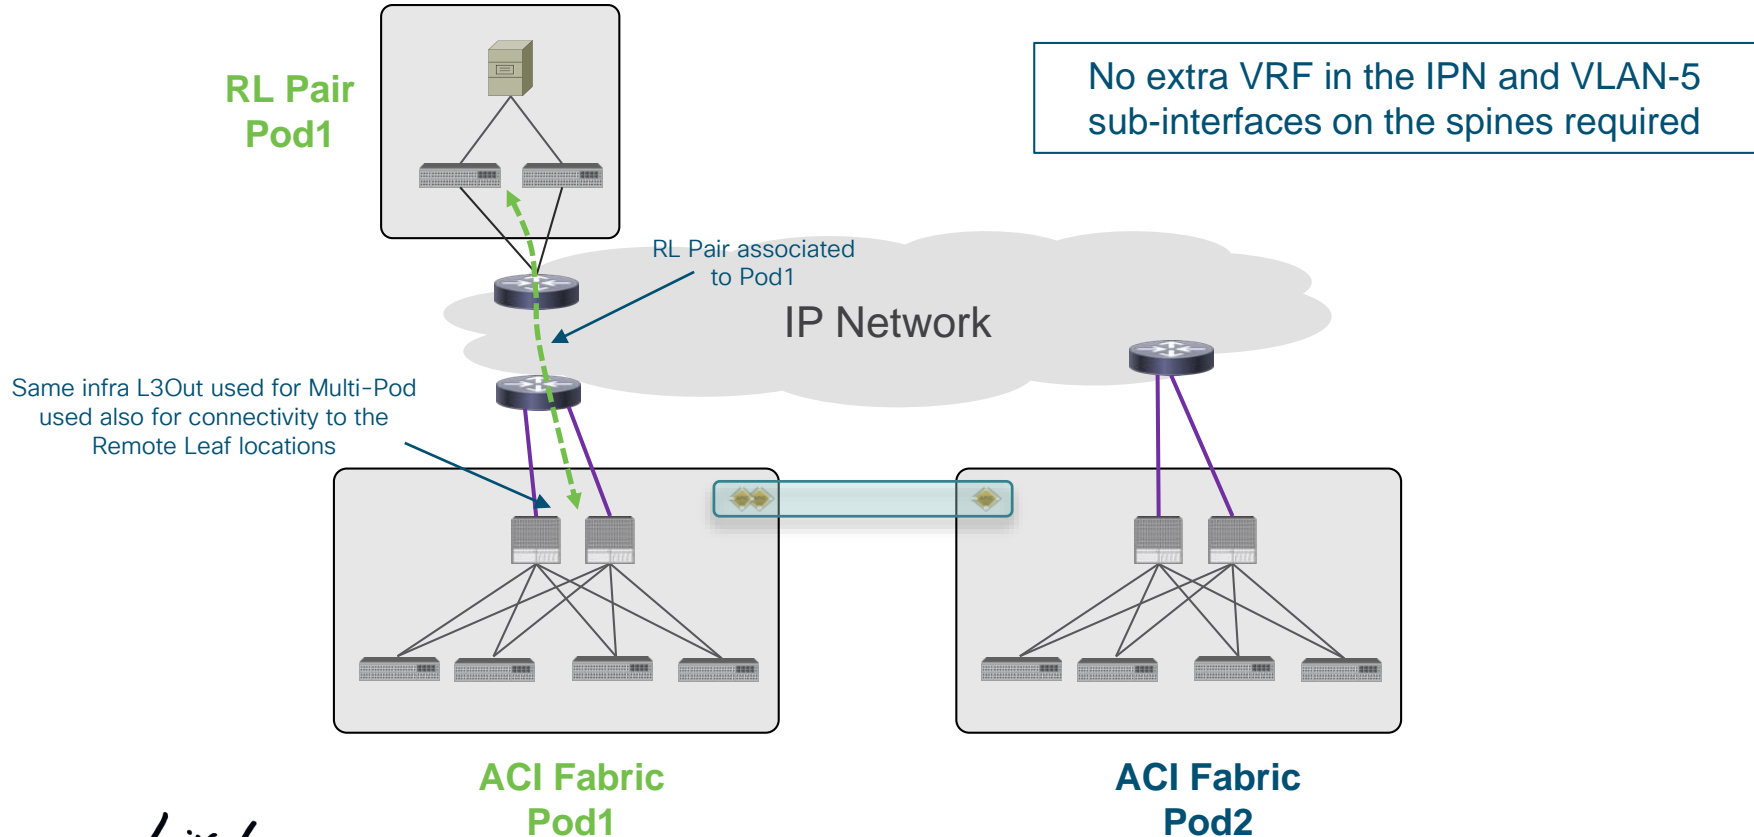


Integrating ACI Remote Leaf with Multi-Pod

ACI Remote Physical Leaf

Integration with Multi-Pod

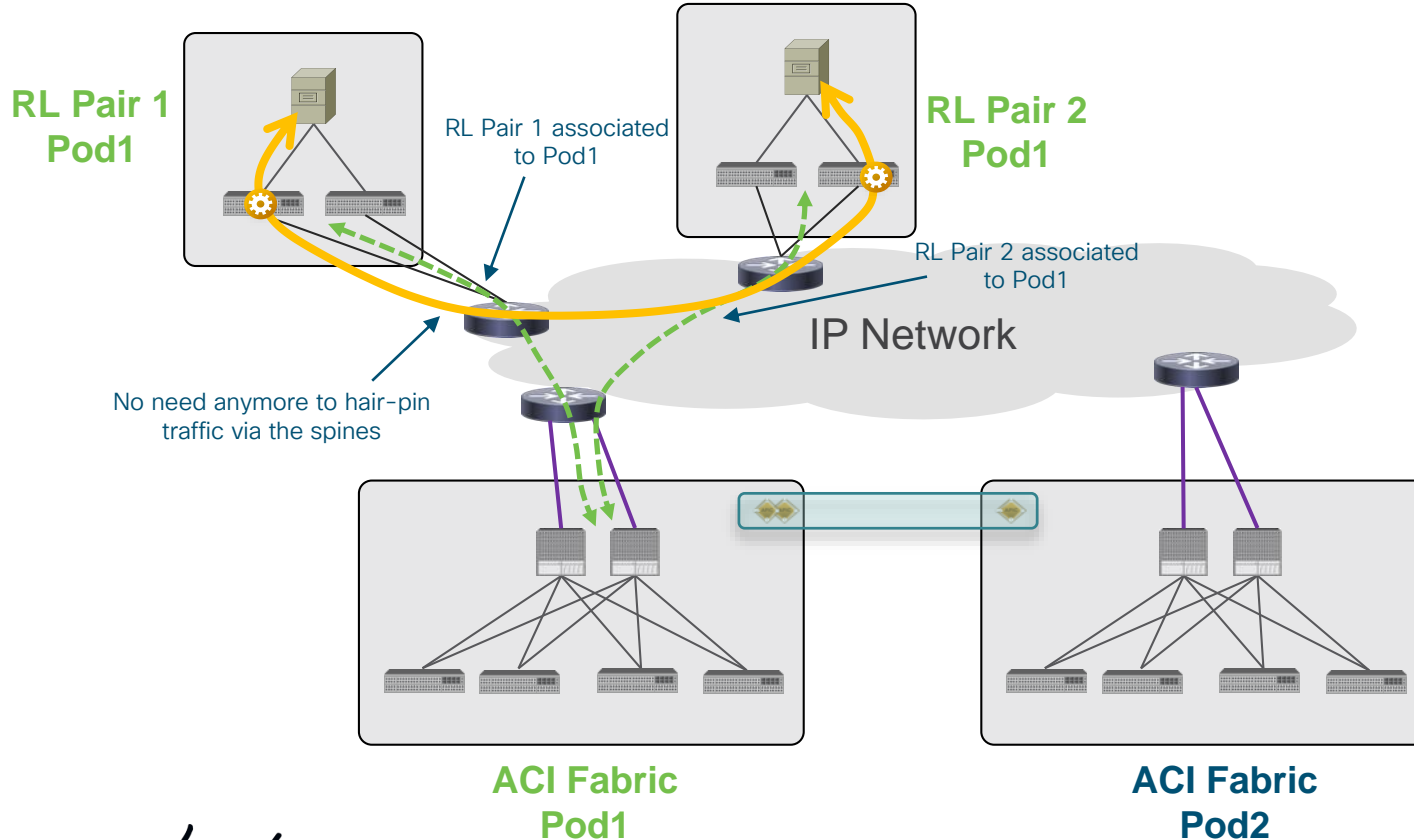
ACI 4.1(2)
Release



ACI Remote Leaf with Multi-Pod

Direct Forwarding between RL Pairs Part of the Same Pod

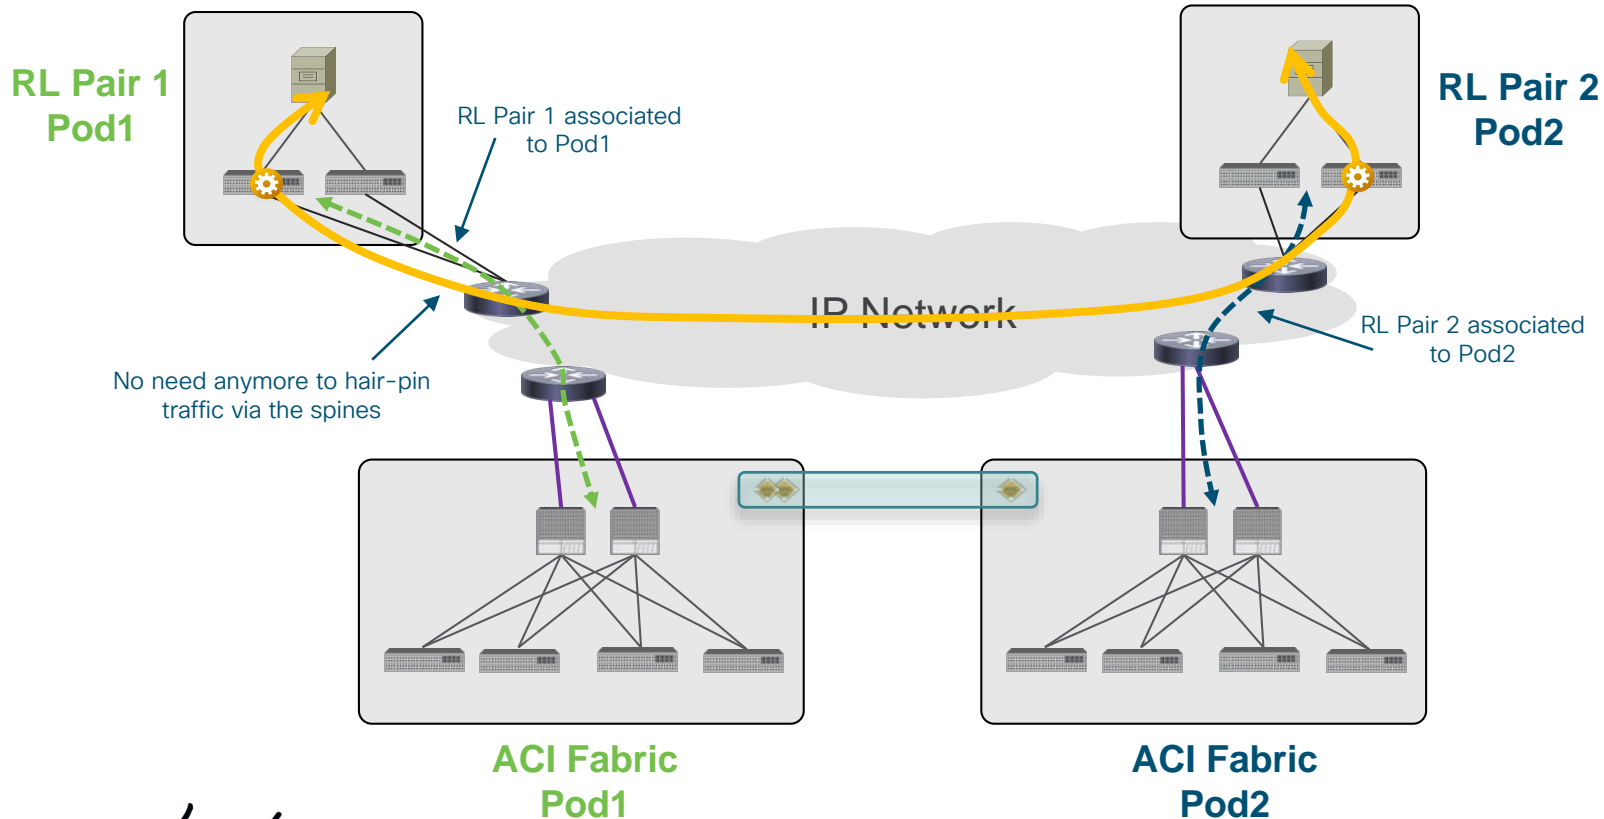
ACI 4.1(2)
Release



ACI Remote Leaf with Multi-Pod

Direct Forwarding between RL Pairs Part of Different Pods


ACI 4.1(2)
Release



ACI Remote Physical Leaf

RL Pair Resiliency in a Pod Failure Scenario

ACI 4.2(1)
Release

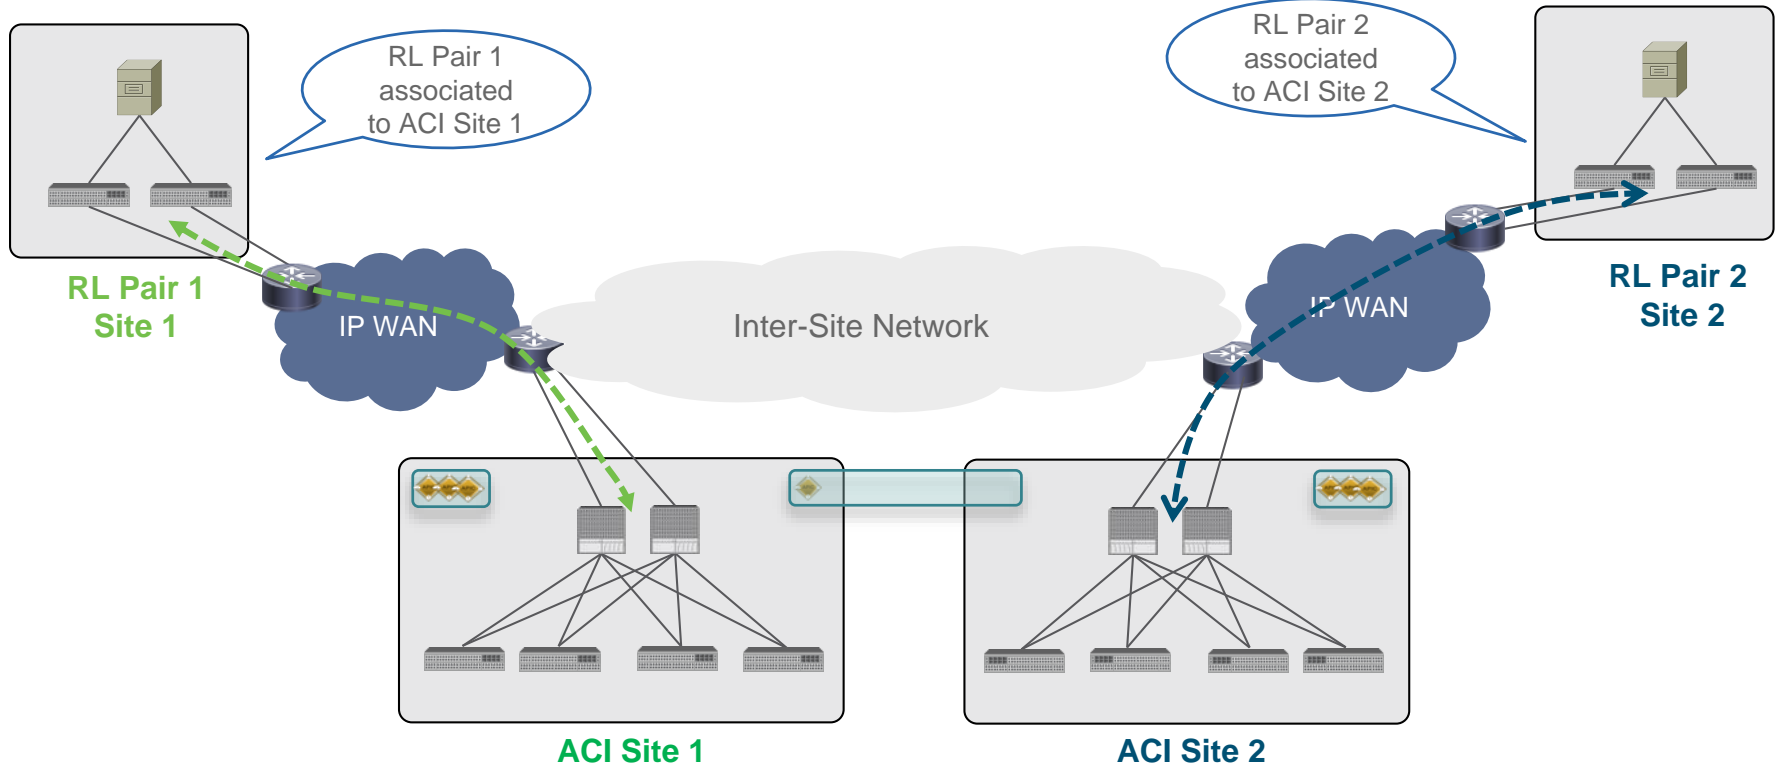


Integrating ACI Remote Leaf with Multi-Site

ACI Remote Physical Leaf and Multi-Site

RL Sites Can Be Associated to Separate Pods

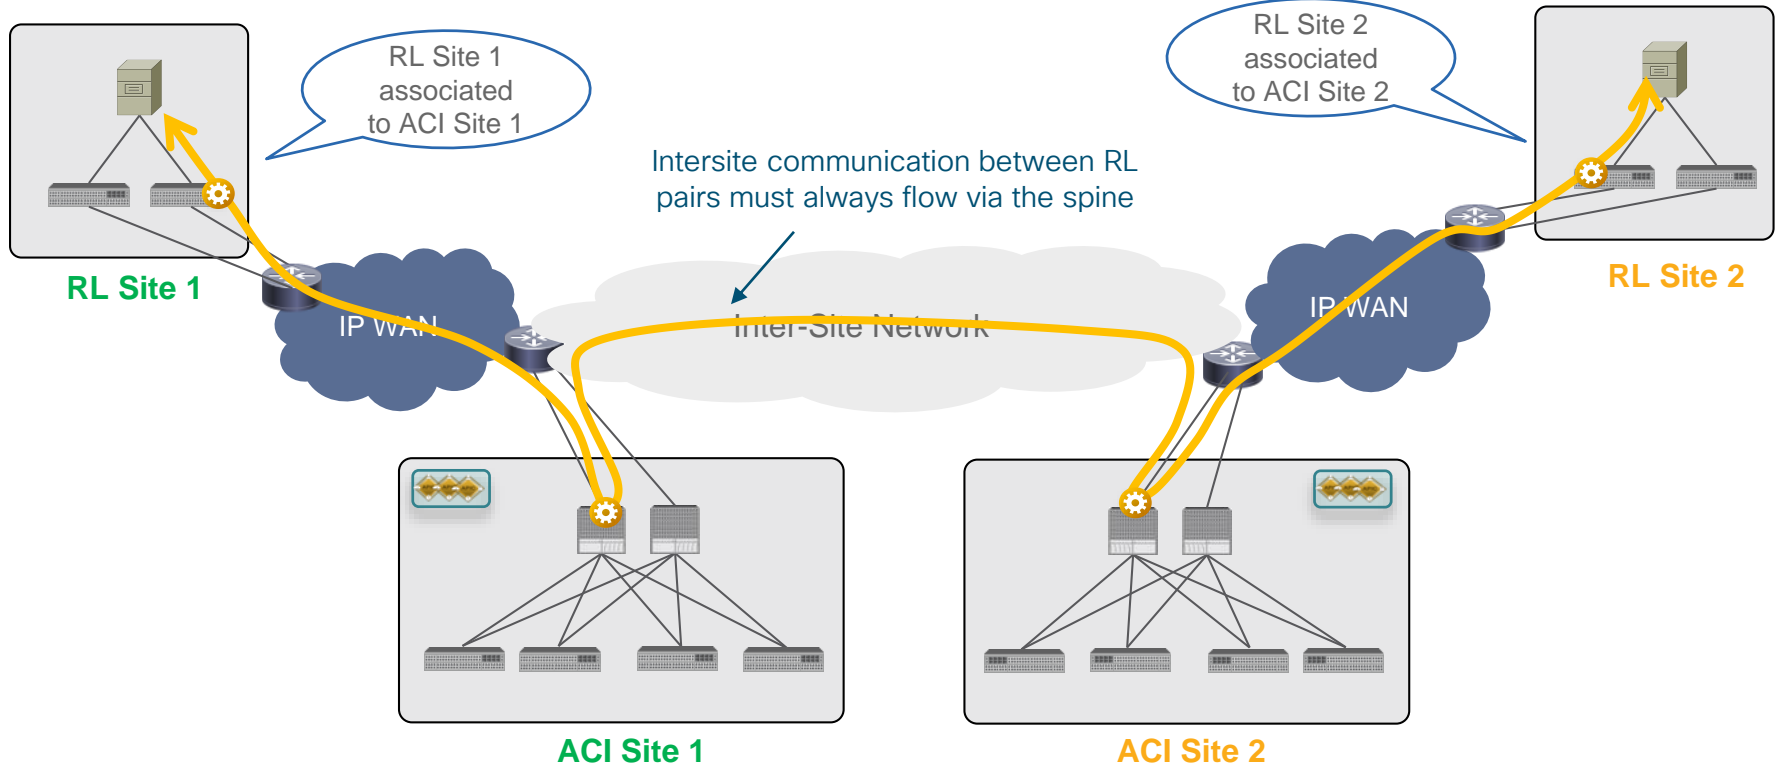
ACI 4.1(2)
Release



ACI Remote Physical Leaf and Multi-Site

RL Sites Can Be Associated to Separate Pods

ACI 4.1(2)
Release



ACI Physical Remote Leaf

Where to Go for More Information



- ✓ ACI Remote Physical Leaf White Paper

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-740861.html>


- ✓ BRKACI-2387

Agenda

- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - **Mapping use cases to the proper solutions**
 - Active/Active DC → Multi-Pod
 - Disaster Recovery → Multi-Site
 - Migration/Coexistence with Legacy DC Networks and ‘Disaggregated DCs’ Model → Physical Remote Leaf
 - **Baremetal Cloud Integration → Virtual Pod (vPod)**
 - Extending ACI to the Cloud
 - Connecting the users to the Multi-Cloud DC
 - ACI and SDA Integration
 - ACI and SDWAN Integration

ACI Virtual Pod (vPod)



Main Use Cases





On-Premises ACI Data Center




Baremetal Cloud



Brownfield




Colocation / Remote DC



ACI Virtual Pod

Extend ACI to Brownfield or Baremetal Cloud Locations



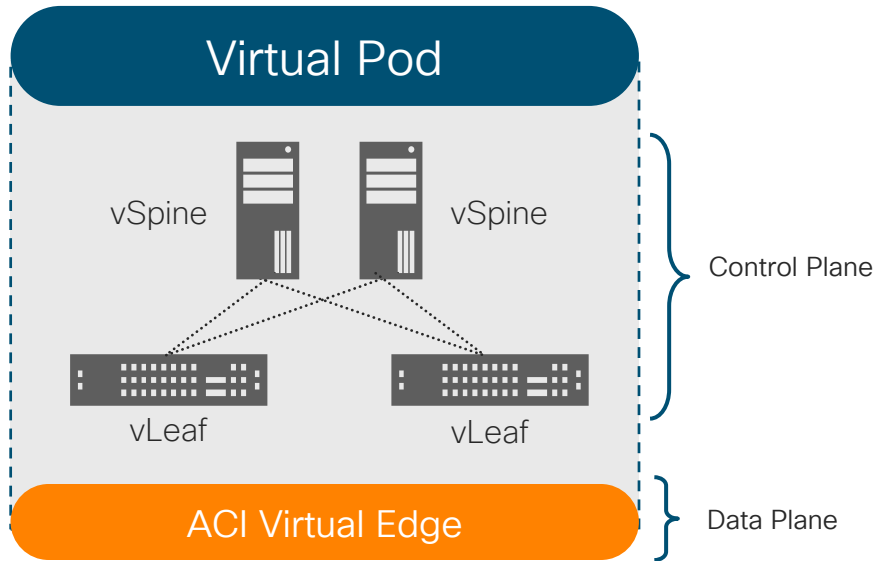
Switching/routing and policy enforcement

- vPod allows to extend ACI connectivity and policies to compute resources deployed in legacy DC networks
 - No need to deploy any ACI HW in the remote network

ACI Virtual Pod

Architectural Components

ACI 4.0(2)
Release



Management Cluster (vSpine + vLeaf)

- vSpine nodes: centralized endpoint and LPM database (COOP and BGP)
- vLeaf nodes: distribute APIC policies to ACI Virtual Edges (DME/PE on vLeaf <-> Opflex on AVE)
- vSpine and vLeaf nodes are **not used** for data-plane forwarding

ACI Virtual Edge (vPod Mode)

- Implements ACI data plane function (switching and routing) **and** policy enforcement
- iVXLAN for communication within vPod and across Pods

ACI Virtual Pod

Software and Hardware Requirements

On-Premises Datacenter

Supported Spines

Fixed Spine:

- N9364C
- N9332C

Modular Spine LC: (C9504/C9508/C95016)

- N9732C-EX with FM N9K-C950x-FM-E(2)
- N9736C-FX with FM N9K-C950x-FM-E(2)

- APIC 4.0(2) onwards

vPod Datacenter

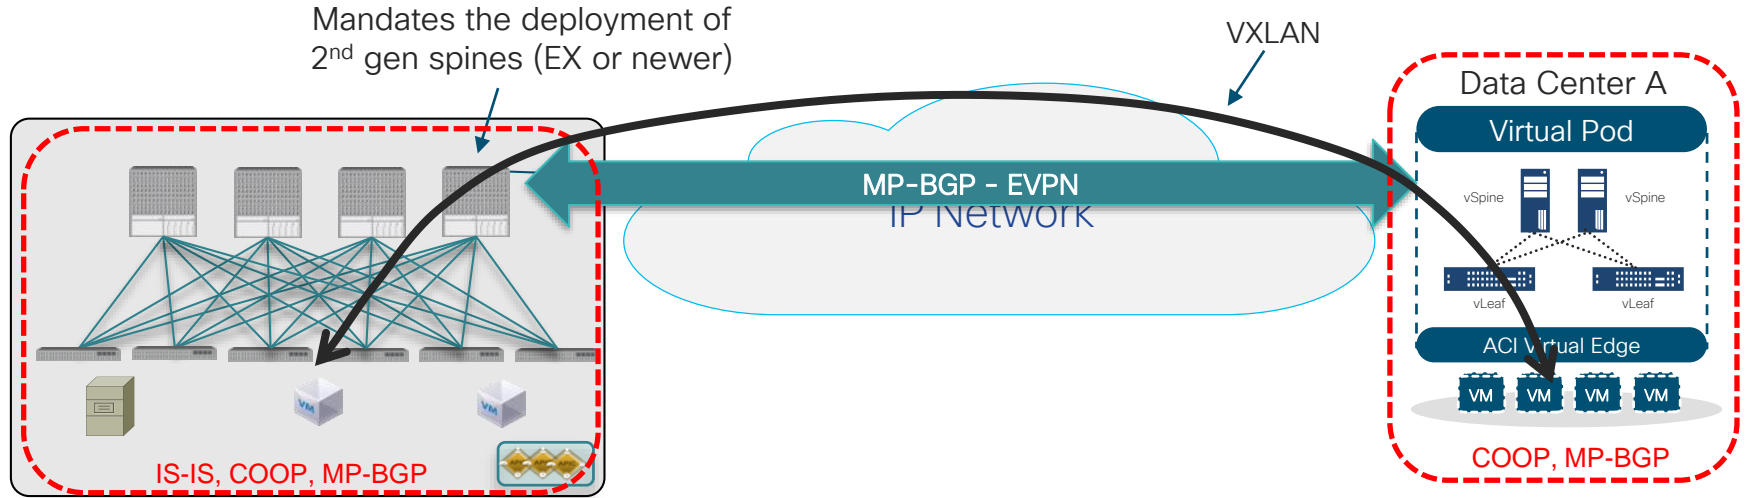
- VMware vCenter 6.0 or later
- 2 hosts for Management Cluster
- 2 hosts for Payload Cluster

- ESXi 6.0 or later

- ACI vCenter plugin or vPod python/PowerShell deployment scripts

ACI Virtual Pod

Control and Data Planes



On-Premises ACI Data Center

- Policies centrally defined on the APIC cluster deployed on-prem
- MP-BGP EVPN sessions established to exchange endpoint reachability information between Pods
- Ingress replication support on physical spines and AVEs to forward BUM traffic

vPod Management Plane

Fully Managed via APIC

Schedule Node Upgrade

Group type: Physical Virtual

Upgrade Name:
Select or create new

Target Firmware Version:

Upgrade start time: Now Schedule for later

Ignore compatibility check:

Graceful Maintenance:

Run Mode:

Node selection: Range Manual

i Only nodes that are not already in another Firmware Upgrade Group can be added in order to avoid a possible conflict in upgrades which may lead to unintentional outcome.

Selected	Node id	Node name	Model	Current Firmware	Status
Current Firmware: simsw-3.2(0.37) (5 Nodes)					
<input type="checkbox"/>	101	vleaf1	N9K-C9396PX	simsw-3.2(0.37)	Not Scheduled
<input type="checkbox"/>	102	vleaf2	N9K-C9396PX	simsw-3.2(0.37)	Not Scheduled

ACI Virtual Pod

Where to Go for More Information



- ✓ Virtual Pod White Paper

<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-742393.html>

- ✓ BRKACI-2882 (Cisco Live San Diego 2019)


<https://www.ciscolive.com/global/on-demand-library.html?search=BRKACI-2882%20#/session/1542224297572001r8Mq>

Agenda

- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - Mapping use cases to the proper solutions
 - Active/Active DC → Multi-Pod
 - Disaster Recovery → Multi-Site
 - Migration/Coexistence with Legacy DC Networks and ‘Disaggregated DCs’ Model → Physical Remote Leaf
 - Baremetal Cloud Integration → Virtual Pod (vPod)
 - **Extending ACI to the Cloud**
 - Connecting the users to the Multi-Cloud DC
 - ACI and SDA Integration
 - ACI and SDWAN Integration



ACI Anywhere



Virtual ACI




ACI



Cloud ACI



ACI Multi-POD



ACI Multisite

ACI Remote Leaf

Virtual ACI



Cloud ACI

ACI 2.0

ACI 3.0

ACI 3.1

ACI 4.0

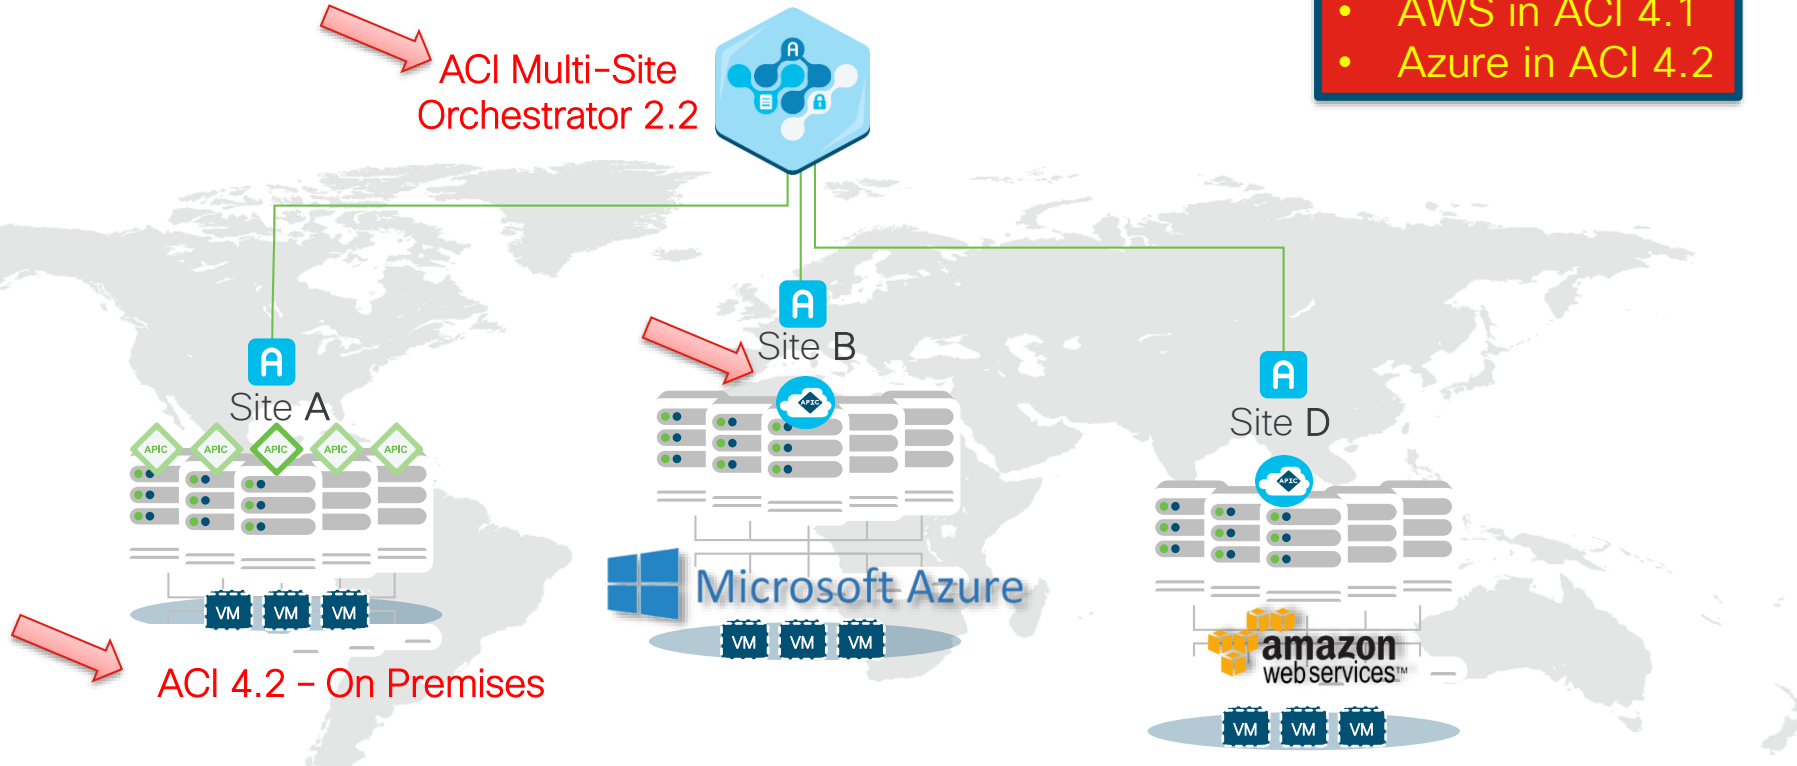
ACI 4.1 | ACI 4.2



ACI Hybrid-Cloud Deployment Model

Hybrid Cloud:

- AWS in ACI 4.1
- Azure in ACI 4.2



Common Governance

Discovery & Visibility

Policy Translation

Monitoring & Troubleshooting

Single Point Of Orchestration


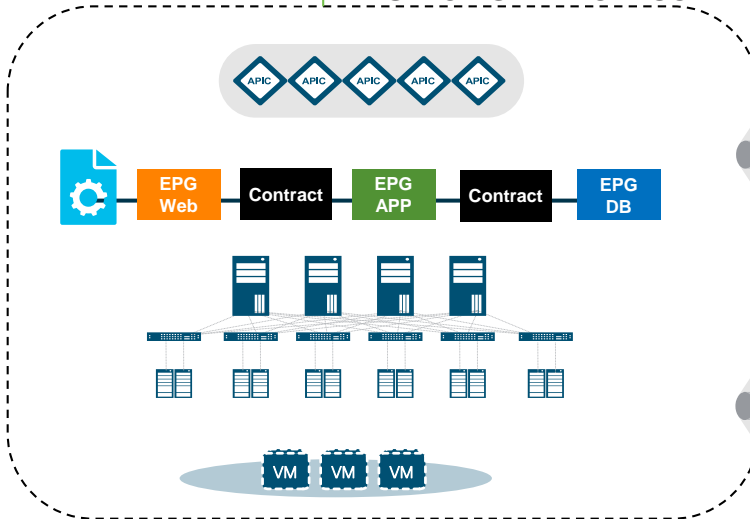
Operational Consistency

Extending ACI to the Cloud

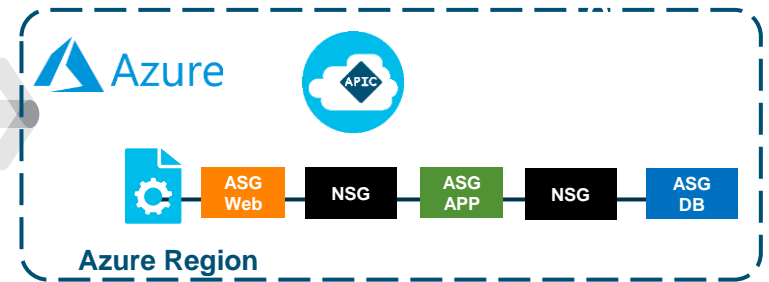


ACI Multi-Site Orchestrator 2.2

ACI for On-Premise



Cloud ACI for Public Cloud



Common Governance

Discovery & Visibility

Policy Translation

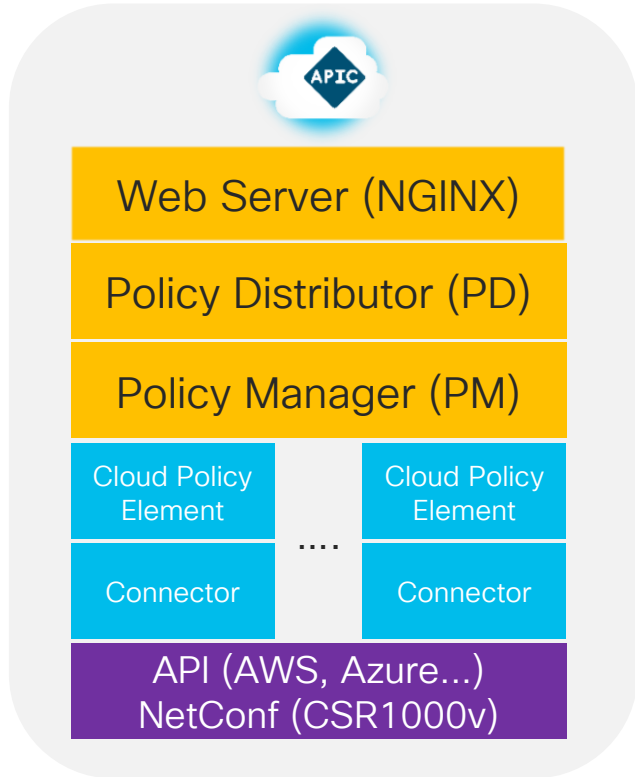
Monitoring & Troubleshooting

Single Point Of Orchestration

Operational Consistency

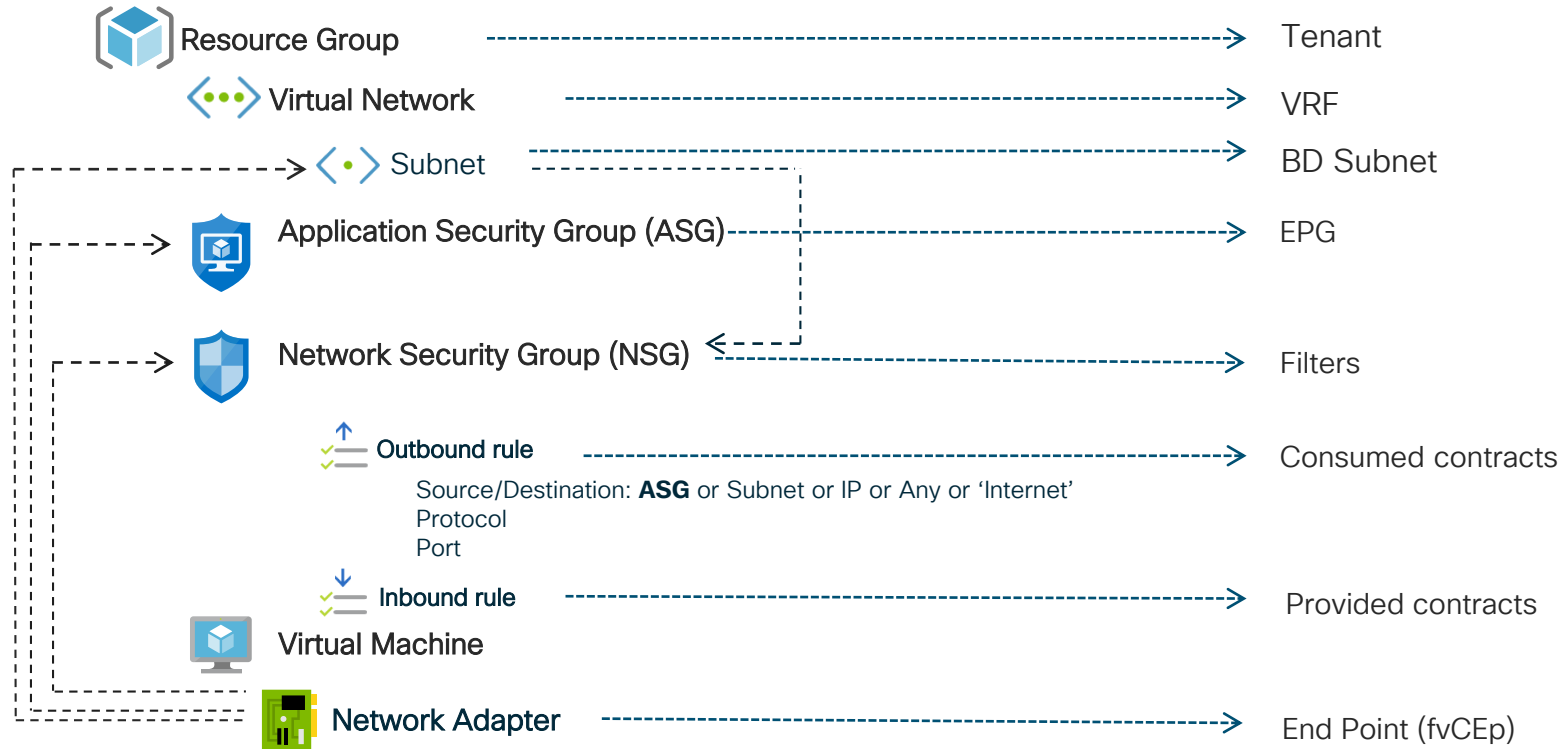
Cloud APIC (cAPIC)

Cloud APIC Architecture

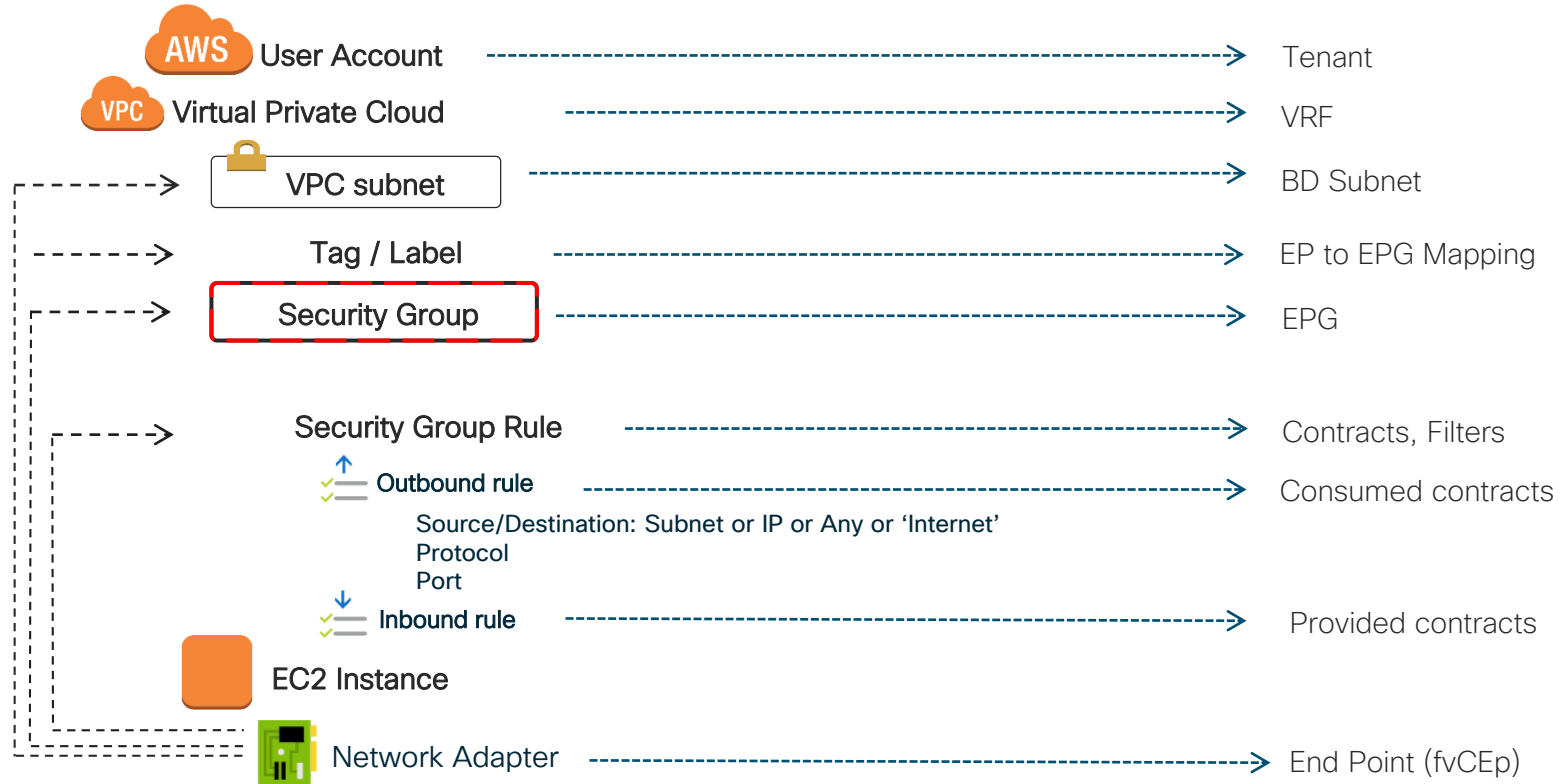


- Virtual Form Factor of APIC
- Automates / Manages Cloud Routers
- Translates ACI Policy to cloud native constructs
- Deploys cloud resources and infrastructure components
- Intuitive GUI and Similar ACI UI look and feel
- REST API North Bound Interface
- cAPIC manages 1 or more regions

Policy Mapping - Azure



Policy Mapping - AWS



Topology Health

Cloud APIC **aws**

Inter-Site Connectivity

Connectivity View

Region	ap-southeast-2
IPSec Tunnels	BGP Sessions
1 Down	0 Down
3 Up	4 Up
CRITICAL	CRITICAL
1	0
MAJOR	MAJOR
0	0
MINOR	MINOR
0	0
WARNING	WARNING
0	0

Region ap-southeast-2 Topology

ap-southeast-2

us-west-1

IPN

On-Prem Site
BarcelonaOnPrem

IPSec TUNNEL (solid line)
BGP SESSION (dashed line)

IPsec Tunnel ID 1
Region us-west-1 VPC overlay-1 CSR...

Oper State: Down

CRITICAL	MAJOR	MINOR	WARNING
1	0	0	0

Router CsrRouter1

Overview Cloud Resources ACI Relationships Statistics **Event Analytics**

Faults

Severity	Code	Cause	Affected Object	Description	Last Transition
CRITICAL	F0325	Tunnel down	acct-[Incloud-tenant-1]-[region-[us-west-1]-[vpc-overlay-1]-[ipsec-CsrRouter1/tunnel-1]	Configuration error on tunnel	1 hour ago

Page 15 of 1

- Network connectivity and Health

Endpoints in an EPGs

EPG Web 🔖 — ✕

Overview Topology **Cloud Resources** Application Management Statistics Event Analytics 🔄 ⚙️ Actions ▾

Virtual Networks

Endpoints

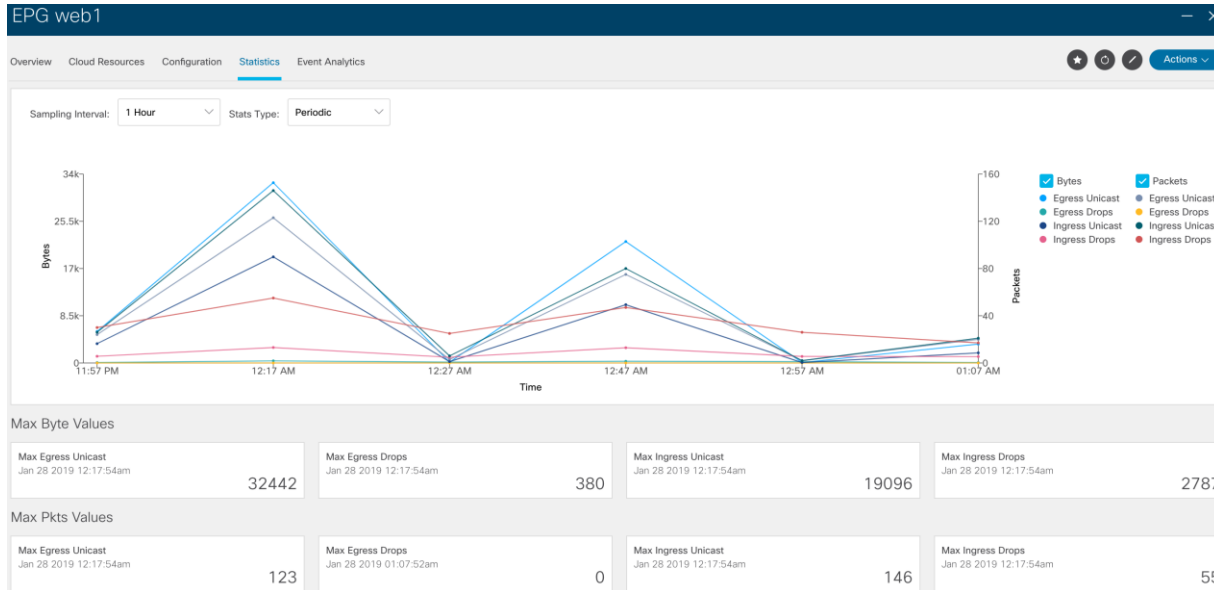
Security Groups

Name	Oper State	Private IPv4 Addr.	Public IPv4 Addr.	EP Type	Application Management		Cloud Resources	
					EPGs	Security Groups	Virtual Machines	
wos-wordpress946 WoS > westus > WoS-VRF 10.101.200.0/24 > 10.101.200.0/24 > 10.101.200.128/25	in-use	10.101.200.132	13.64.103.72	vm	1	1	1	

10 ▾ Rows Page 1 ▾ of 1 |◀◀ 1-1 of 1 ▶▶|

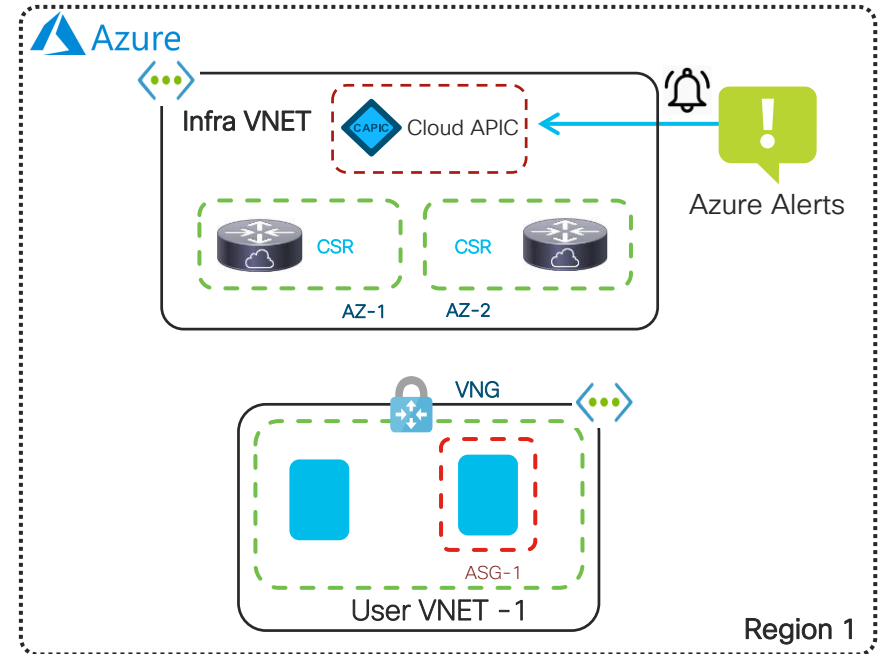
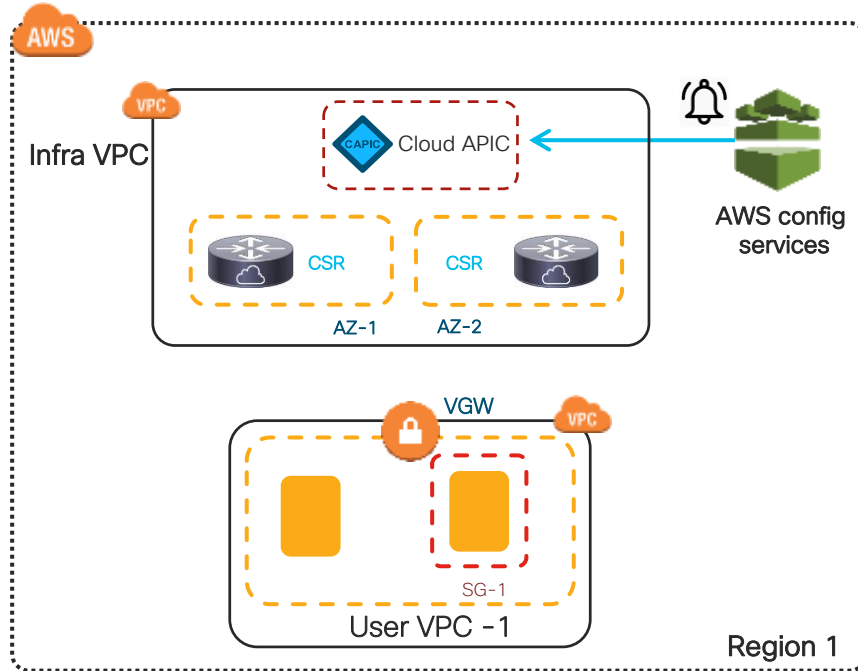


Statistics



- We will show multiple statistics:
 - Inter-site
 - Inter-region
 - Inter-VPC
 - Cloud EPG
 - Cloud Routers

End Point Learning in Cloud




Security Group (SG)

Availability Zone (AZ)

Cloud EPG

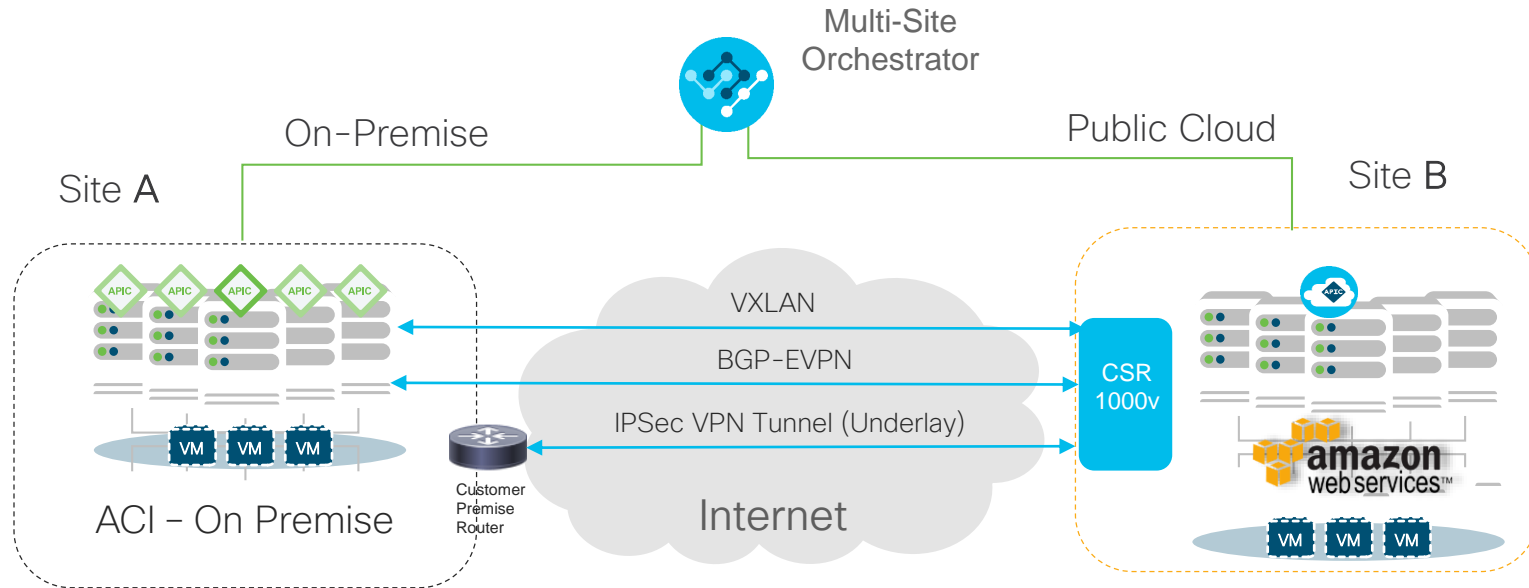
Mapping Endpoints by Tags



Cloud Connectivity (Shh, it's just Multi-Site)

ACI Multi-Site Extension to Cloud

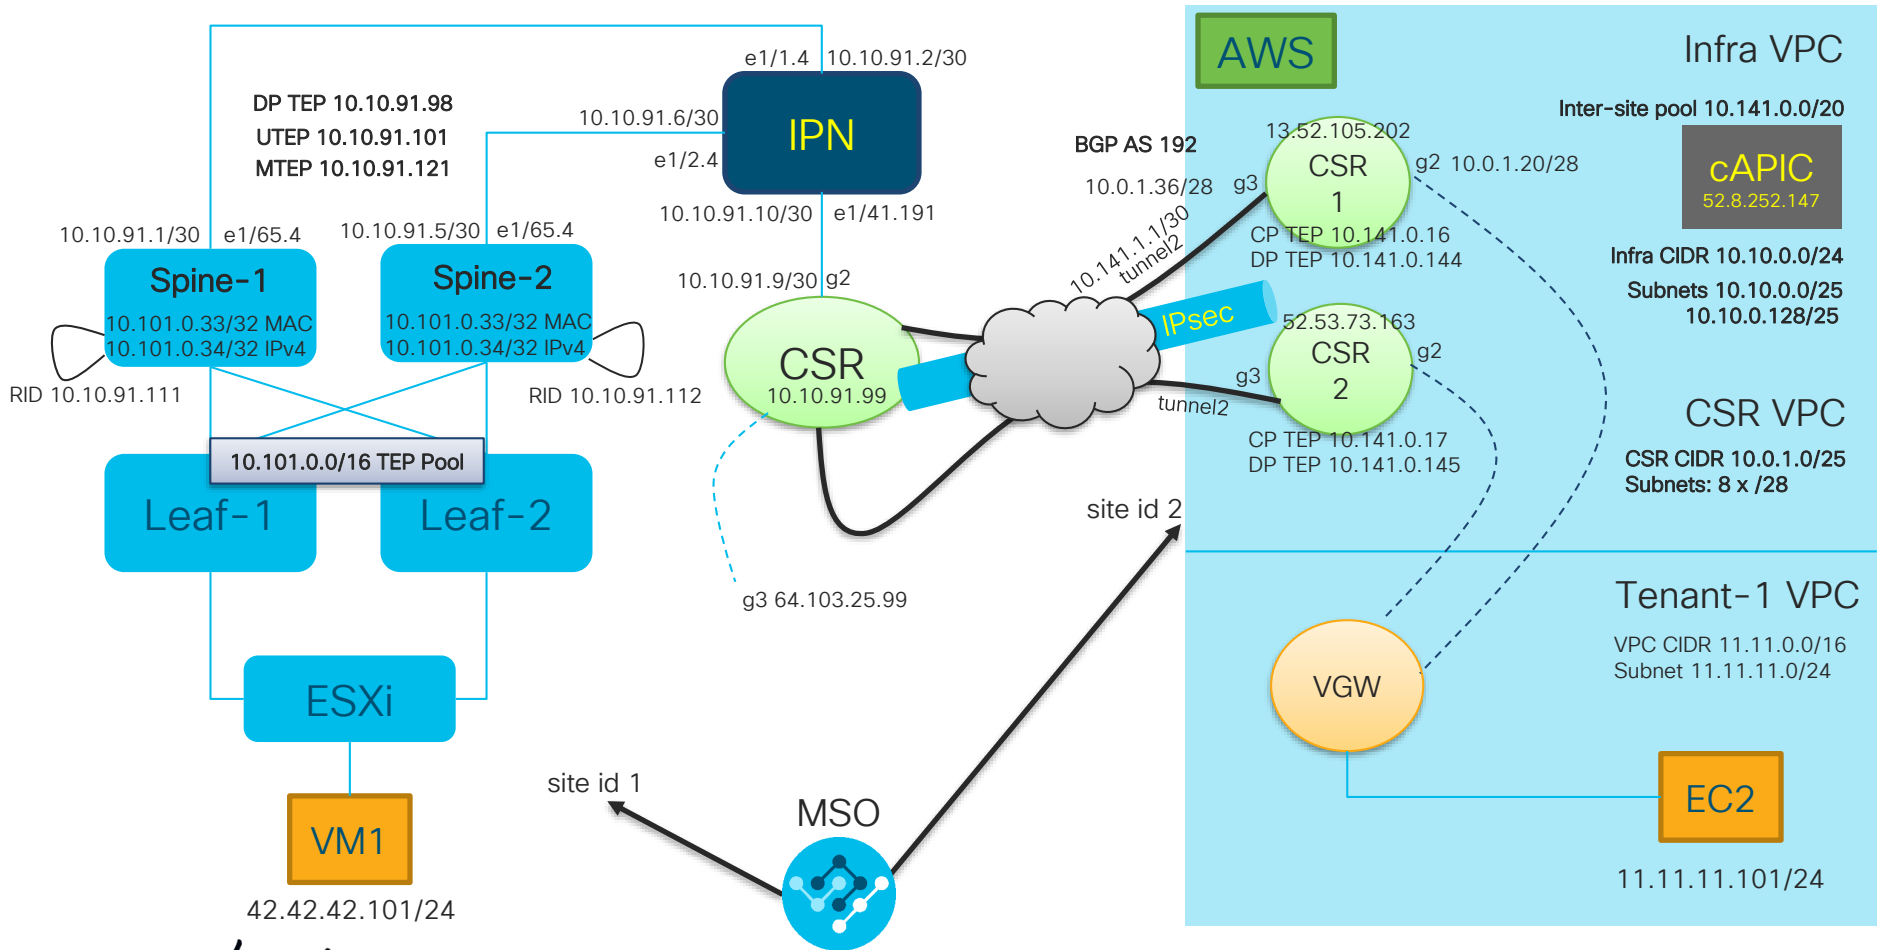
Simplest Mode through [IPSec VPN](#)

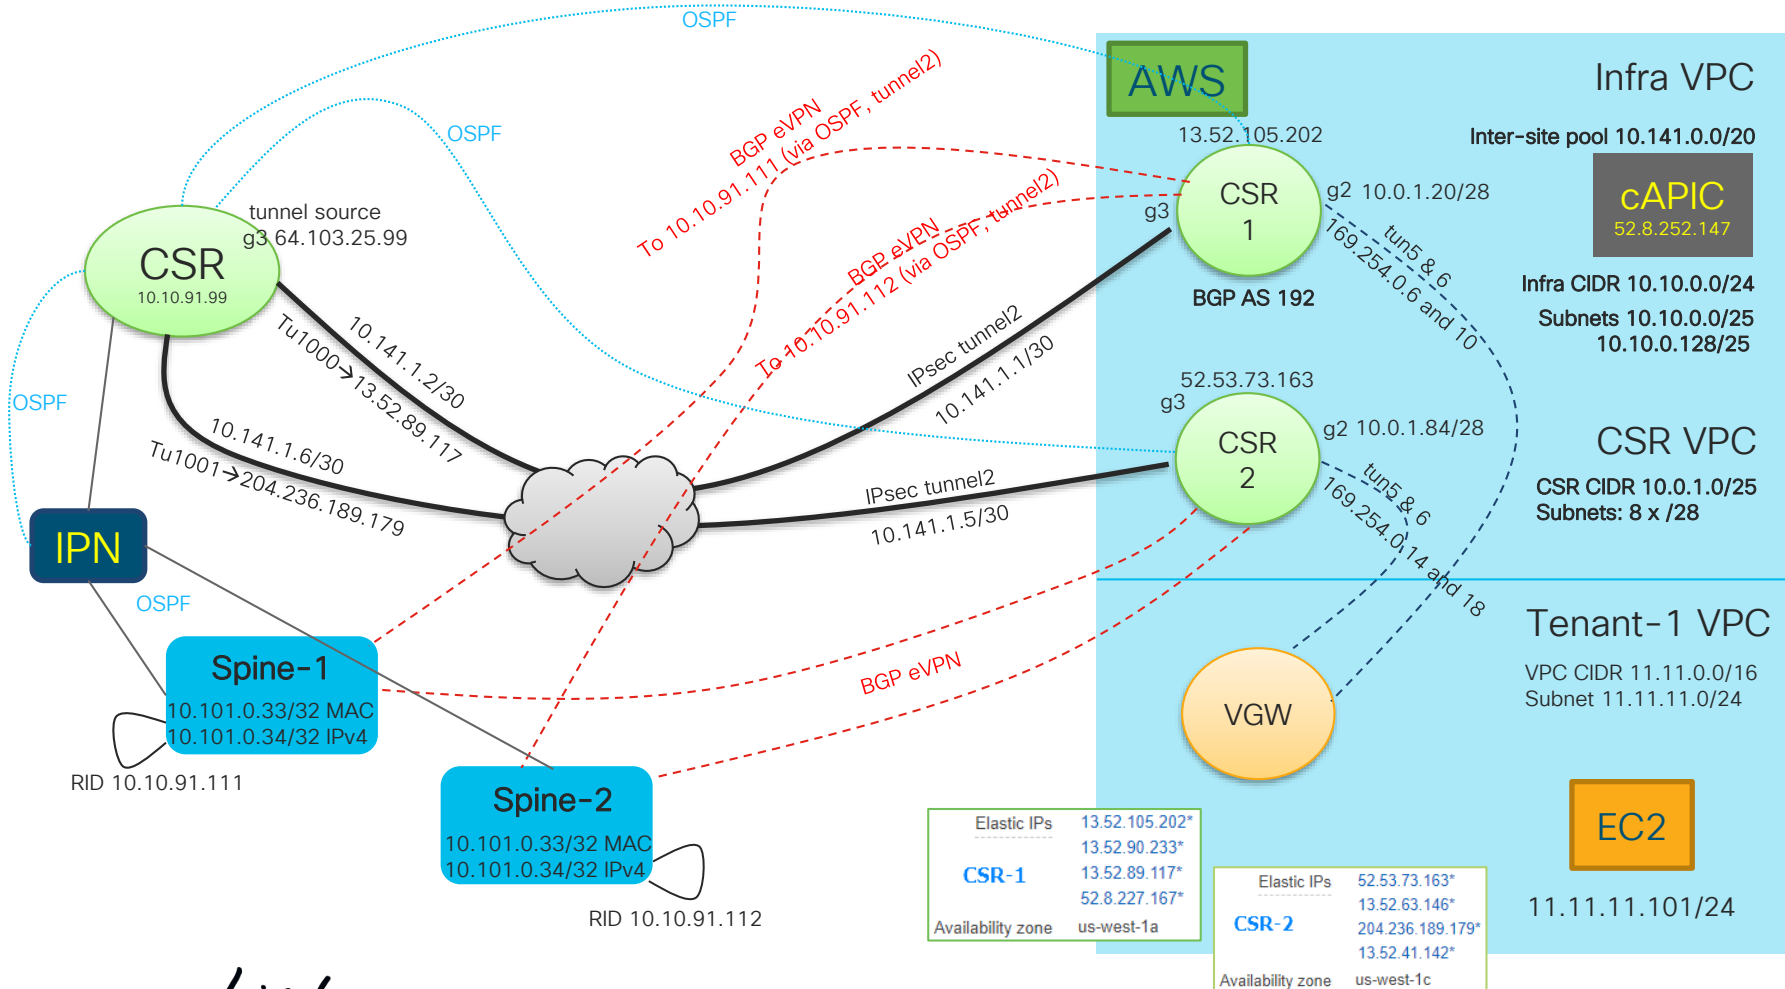


- VXLAN data-plane connects ACI fabric and Cloud site
- BGP-EVPN routing reachability between ACI fabric and Cloud Site
- IPSec VPN connection between customer Premise Router before ACI fabric and CSR1kv

Cloud Multi-Site Components

- On-prem ACI
 - AWS requires software version 4.1
 - Azure requires software version 4.2
 - 2nd generation spines required (like regular multi-site)
- MSO
 - 2.1 supports AWS (hosted on-prem)
 - 2.2 supports AWS, Azure and cloud first
- Same IPN connectivity as traditional multi-site
 - You need an IPN layer (OSPF w/Spines)
 - And an on-prem CSR1Kv or ASR to terminate IPsec
 - MP-BGP eVPN between Spines and CSRs in the cloud
- Azure and/or AWS site(s) with a cloud APIC deployed
 - AWS - Infra VPC (where cAPIC and CSRs reside) & Tenant VPC (only AWS components)
 - Azure - Infra VNET (where cAPIC and CSRs reside) & Tenant VNETs (only Azure components)





AWS

Infra VPC

Inter-site pool 10.141.0.0/20

cAPIC
52.8.252.147

Infra CIDR 10.10.0.0/24
Subnets 10.10.0.0/25
10.10.0.128/25

CSR VPC

CSR CIDR 10.0.1.0/25
Subnets: 8 x /28

Tenant-1 VPC

VPC CIDR 11.11.0.0/16
Subnet 11.11.11.0/24

EC2

11.11.11.101/24

VGW

13.52.105.202

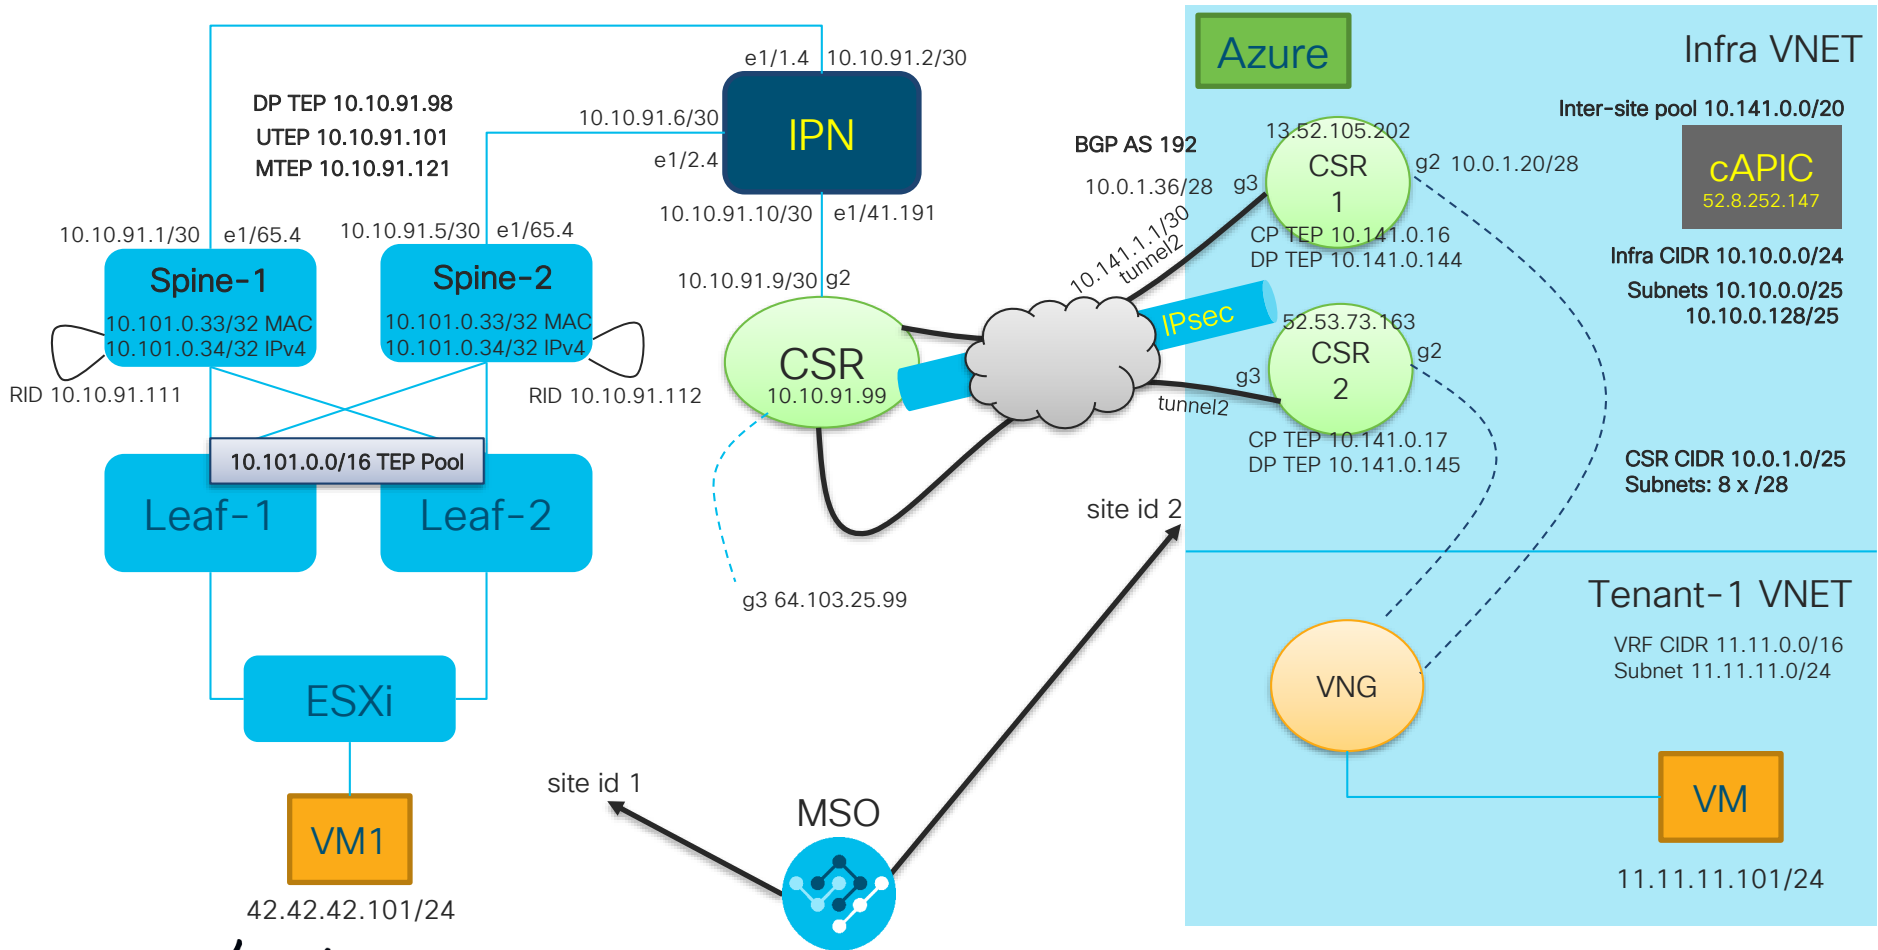
CSR 1
BGP AS 192

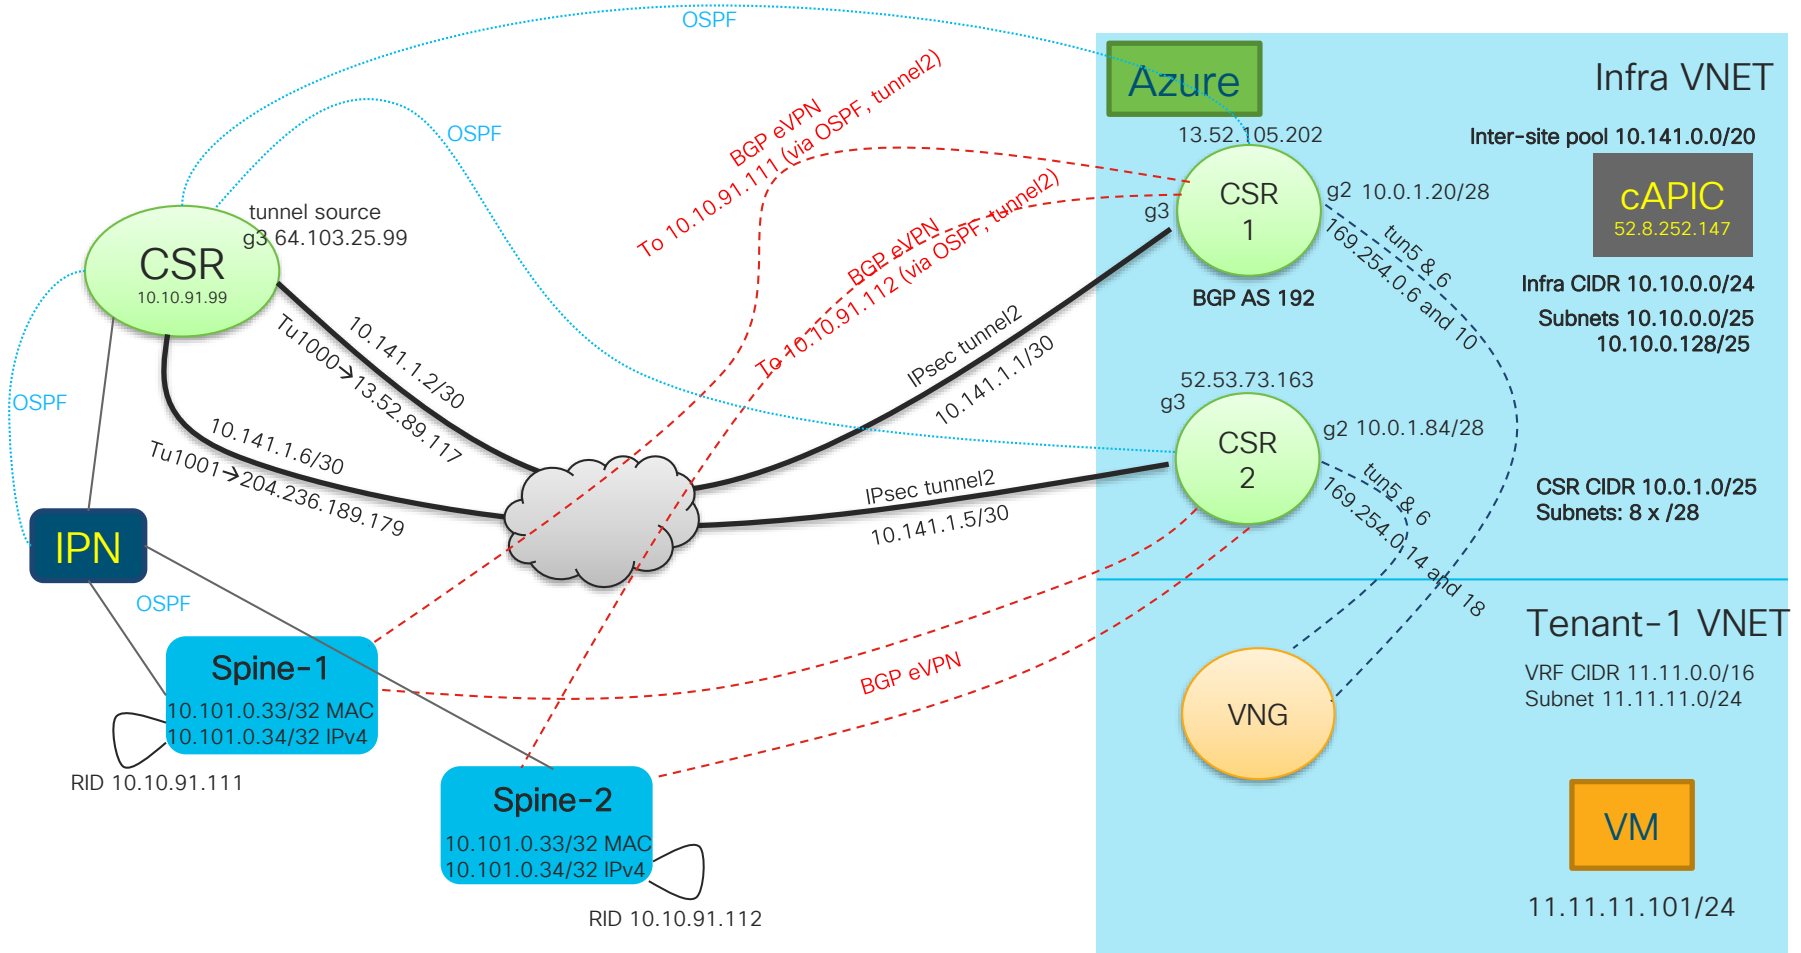
52.53.73.163

CSR 2

Elastic IPs	13.52.105.202*
	13.52.90.233*
CSR-1	13.52.89.117*
	52.8.227.167*
Availability zone	us-west-1a

Elastic IPs	52.53.73.163*
	13.52.63.146*
CSR-2	204.236.189.179*
	13.52.41.142*
Availability zone	us-west-1c



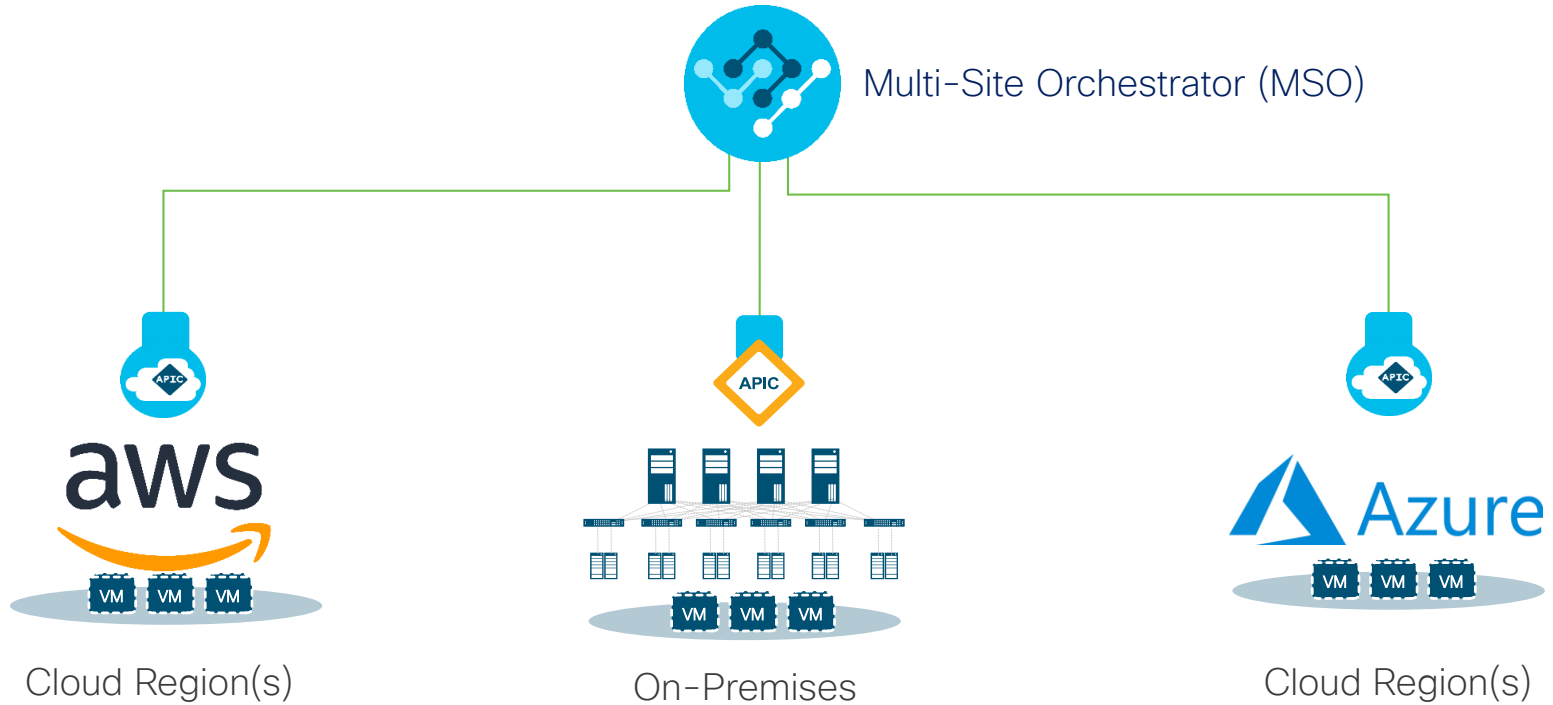


Wait, that looks quite complex

Actually, its just another instance of multi-site

- You don't configure half of what you just saw
- You spin up a cAPIC using a Cloud Formation (AWS) or ARM (Azure) Template
- Pair on-prem with your IPN just like regular multi-site
- MSO and cAPIC take care of most configuration aspects
 - You provide high-level config parameters
 - You get a ready-to-use IPsec configuration to copy/paste in your CSRs
 - All AWS configuration aspects (VGW BGP and IPsec, routing, security groups) is fully automated and abstracted
 - All Azure configuration aspects (GW BGP and IPsec, routing, ASG, NSG) is fully automated and abstracted

Cloud ACI Its Multi-Site



MSO Form Factor



Hardware Appliance
(based on SE)



VMware OVA



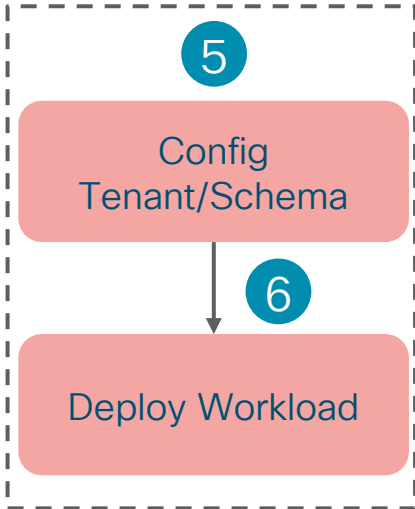
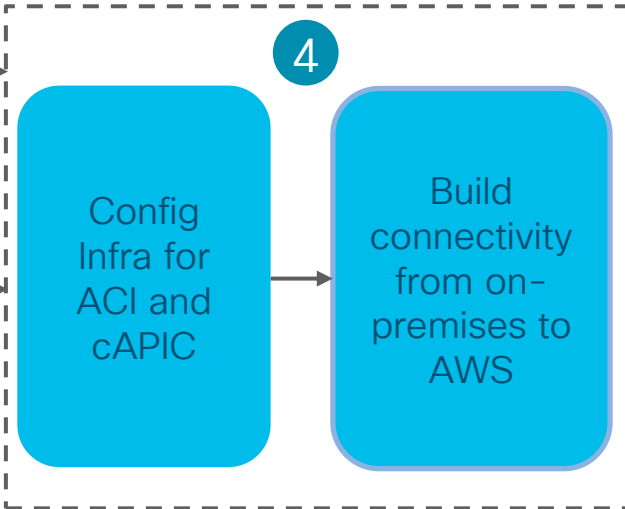
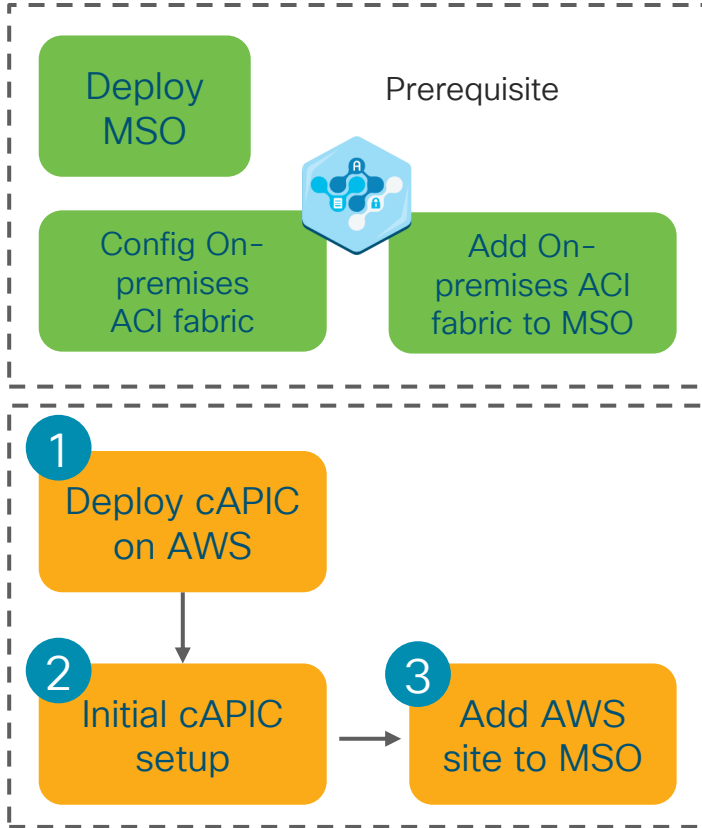
Cloud MSO for AWS

Demo 1: ACI Cloud Connectivity

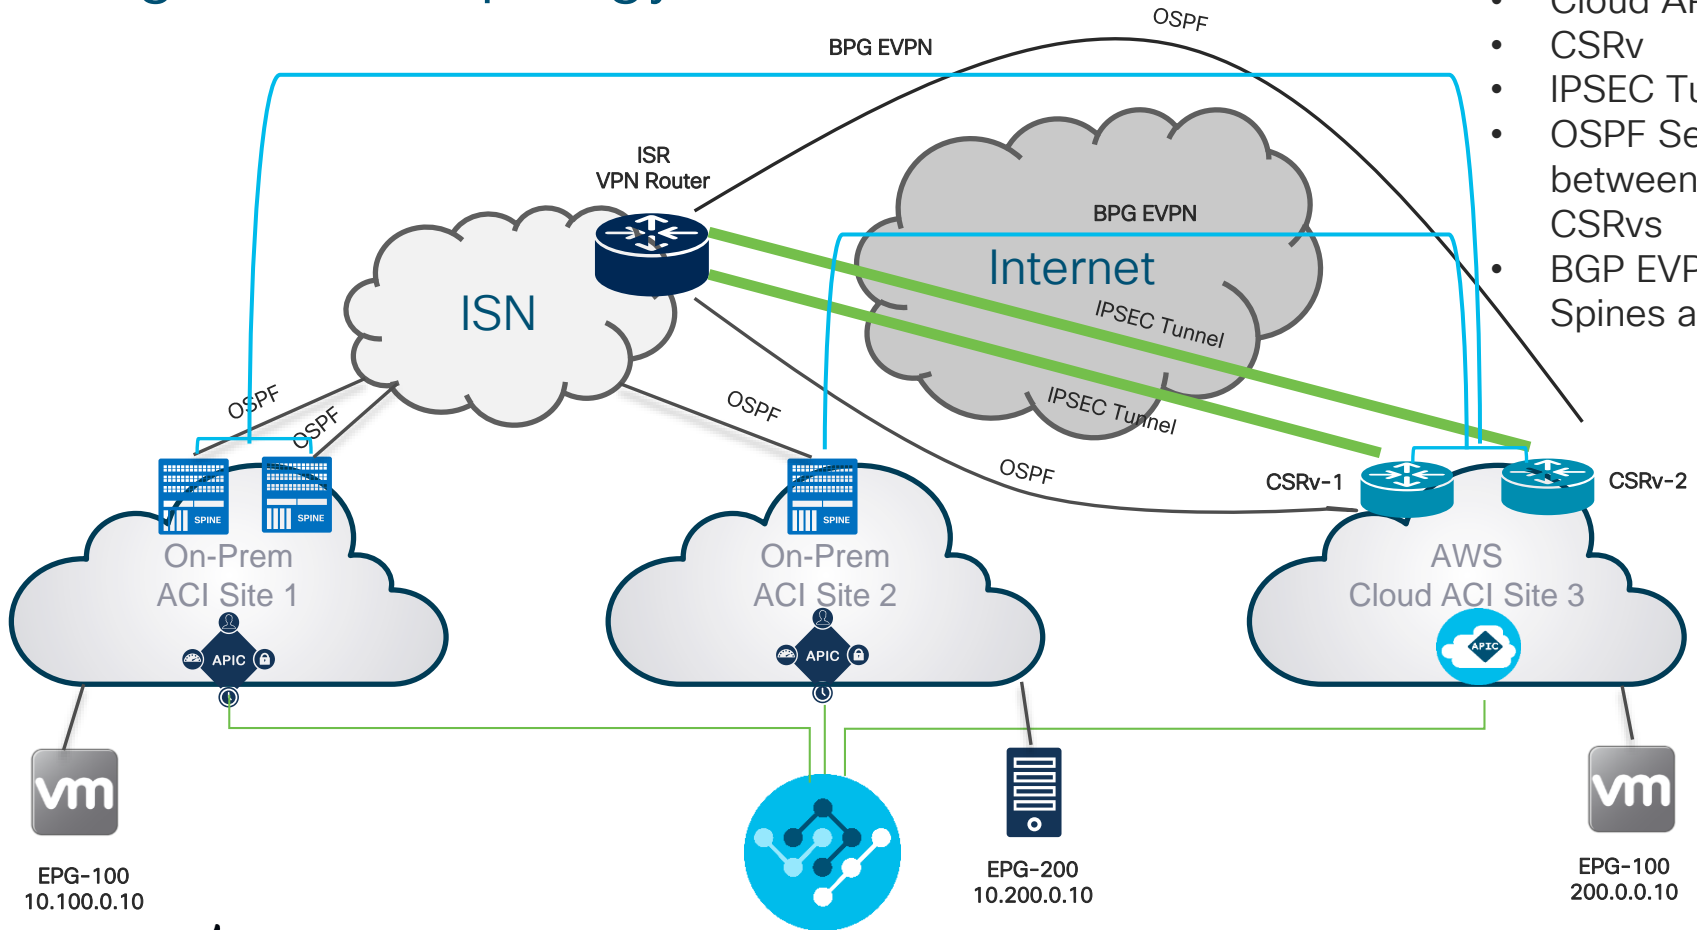


High Level steps

1. Deploy cAPIC on AWS
2. Initial cAPIC setup
3. Add the AWS site on MSO
4. Config Infra for ACI and cAPIC
Build connectivity from on-premises to AWS
5. Config Tenant/Schema
6. Deploy Workload



High Level Topology



Fully Automated:

- Cloud APIC
- CSRv
- IPSEC Tunnel
- OSPF Setup between ISR and CSRvs
- BGP EVPN between Spines and CSRvs

So I have Multi-Site to Cloud
You mean I can stretch VLANs???



No really, can I stretch VLANs to the cloud??

- **Absolutely not!**
 - Stretching an EPG does not mean stretching a VLAN or broadcast domain.
- Stretched EPG Foo uses subnet X on-prem and subnet Y in the cloud
 - You tell MSO what criteria(s) should be used to join EPG Foo
 - Could be tags, could be an IP prefix, could be a region or an AZ
 - cAPIC informs MSO when an EC2 instance matches the selector
 - EC2 instance appears as a /32 in external EPG on-prem, shadow EPG takes care of contracts to make it look like a stretched EPG
 - cAPIC programs a security group in AWS to match on-prem contracts
- Anyway, public cloud vendors do not allow broadcast or multicast.
 - There is no good use case for this. Don't let friends run NSX Cloud.

Agenda

- **ACI Anywhere, Extending the ACI Fabric**
 - Overall Design Principles (AZs and Regions)
 - Mapping use cases to the proper solutions
 - Active/Active DC → Multi-Pod
 - Disaster Recovery → Multi-Site
 - Migration/Coexistence with Legacy DC Networks and ‘Disaggregated DCs’ Model → Physical Remote Leaf
 - Baremetal Cloud Integration → Virtual Pod (vPod)
 - Extending ACI to the Cloud
 - **Connecting the users to the Multi-Cloud DC**
 - ACI and SDA Integration
 - ACI and SDWAN Integration

ACI and Multi- Domain Integration

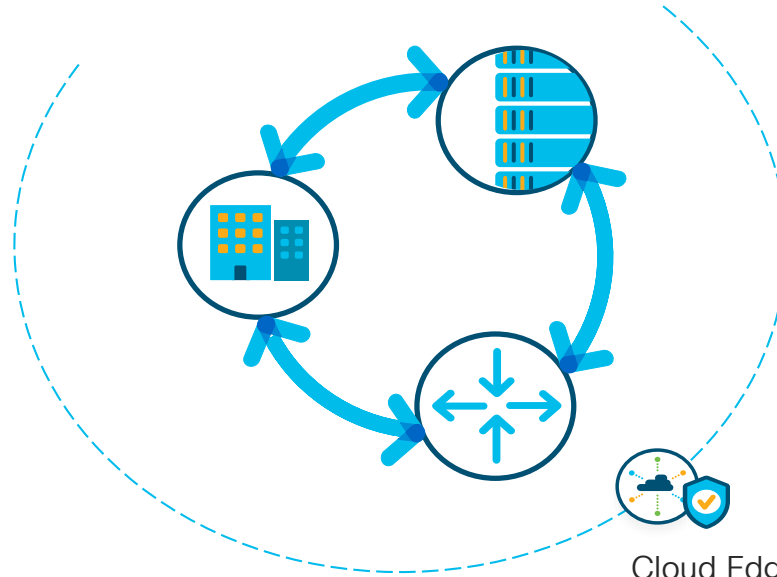
Our direction

Policy

Automation

Telemetry, Analytics and Assurance

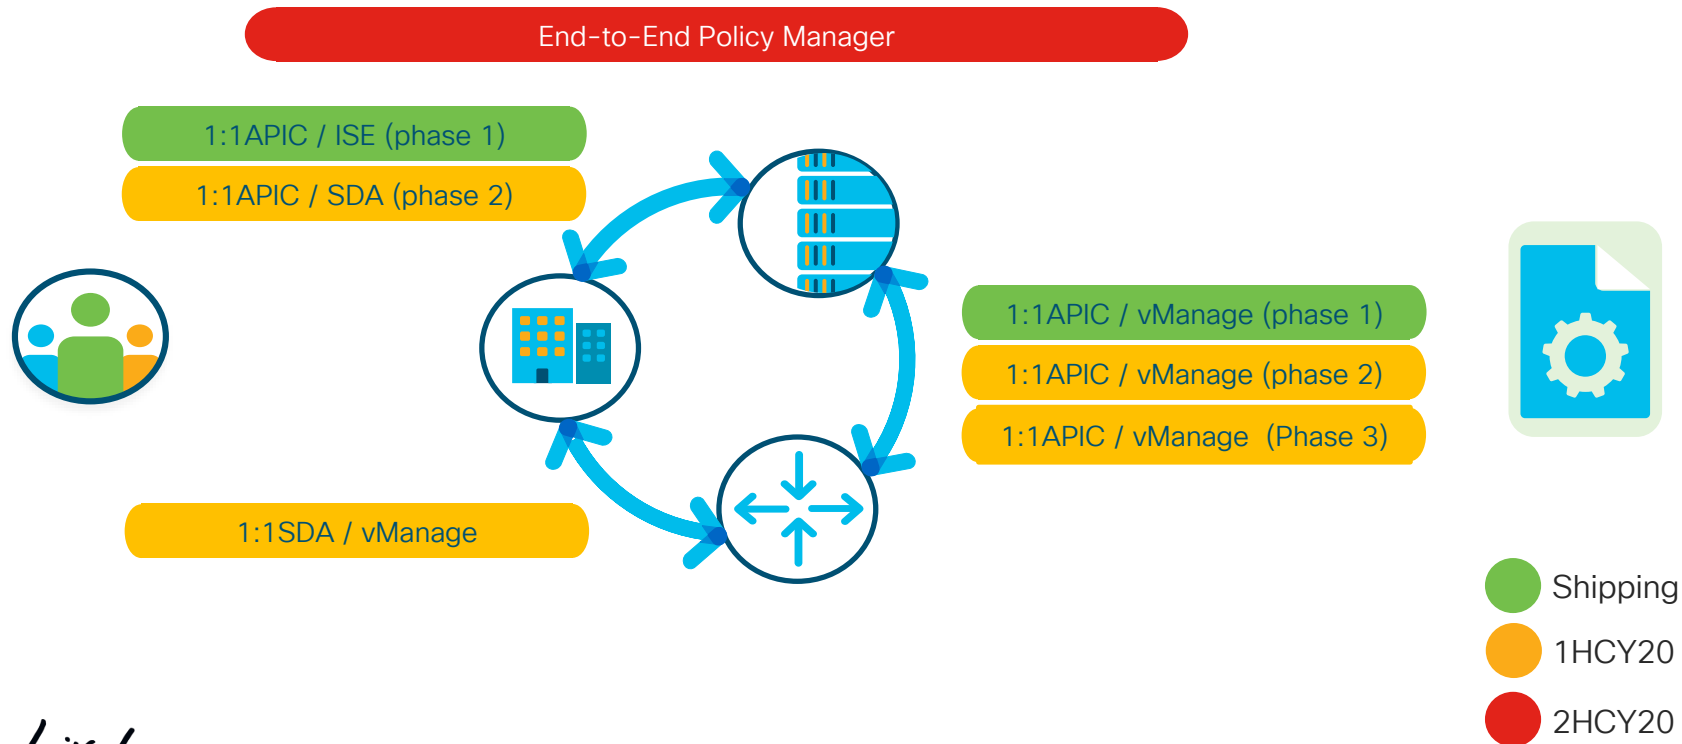
Security, Identity and Segmentation



Cloud Edge
Trust boundary

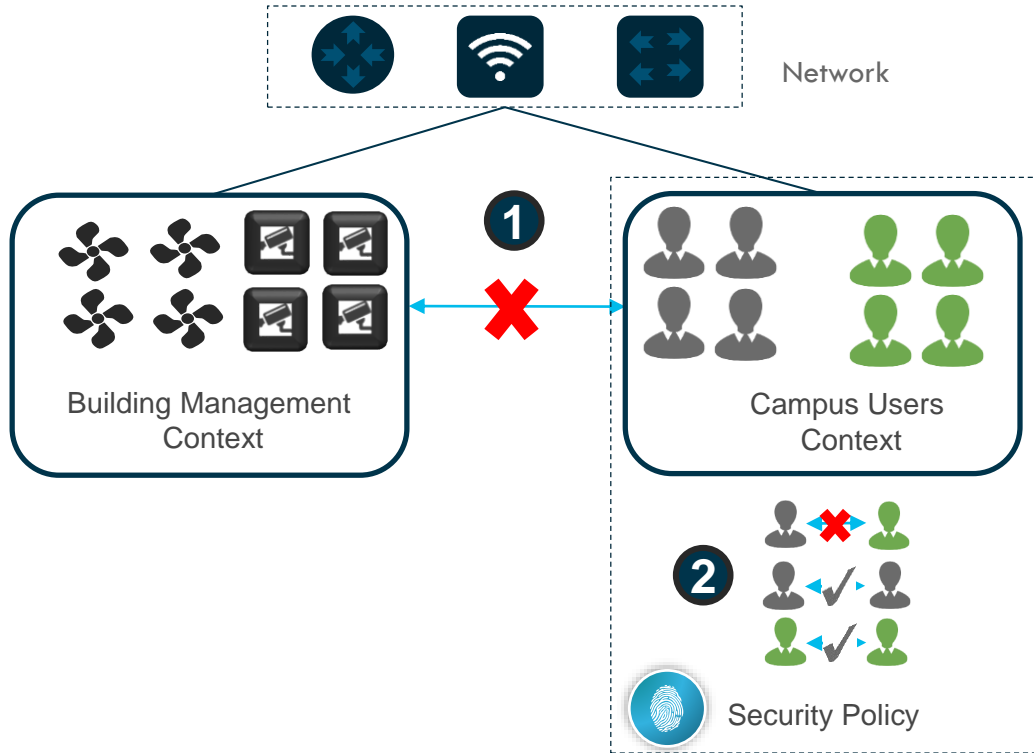
Domain specific integrations

with end-to-end policy



Two Level Segmentation

Two level segmentation/label model



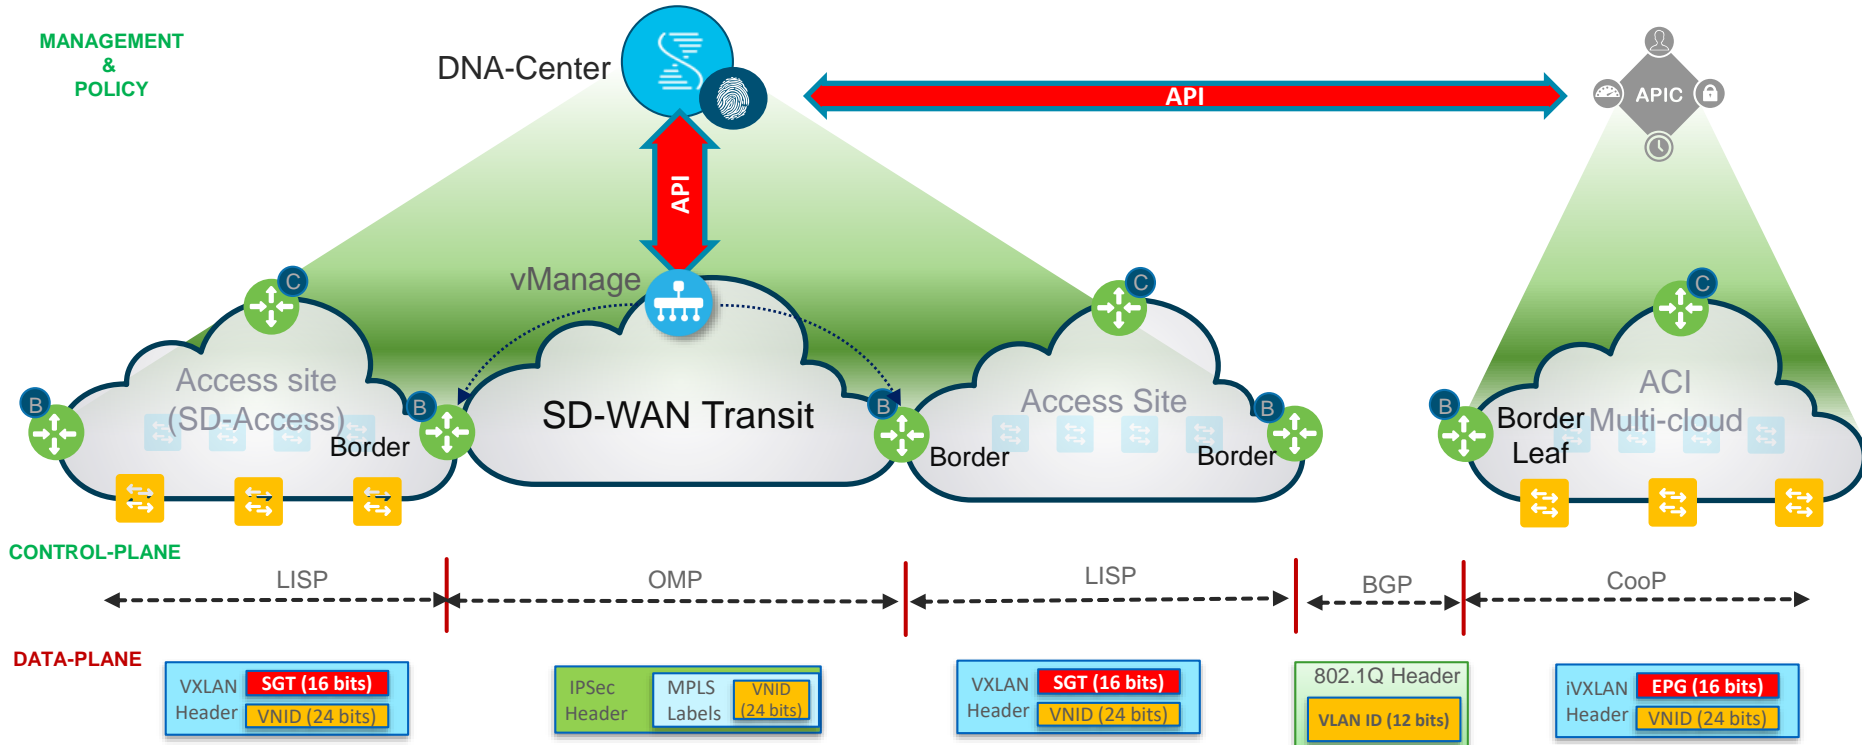
1 Virtual Network (VN)

First level Segmentation that ensures zero Communication between Building systems and Users

2 Scalable / End-point Group

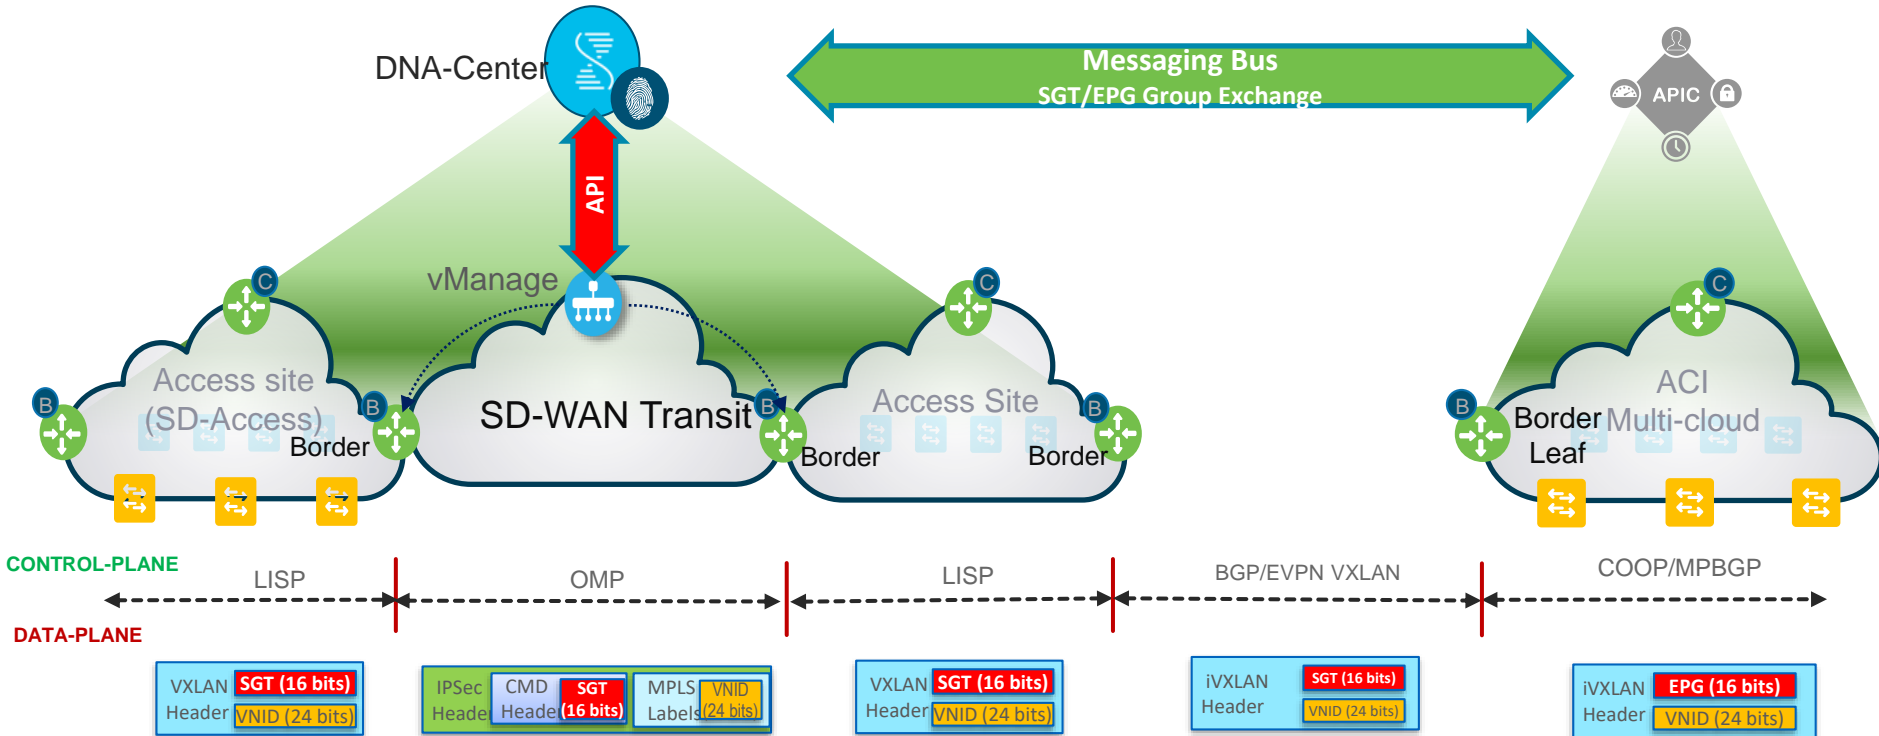
Second level Segmentation within a VN that ensures role based access control between Blue-group and Red-group

End-to-end segmentation and cross domain interworking Phase 1 (Current)



cisco Live!

End-to-end segmentation and cross domain interworking Phase 2

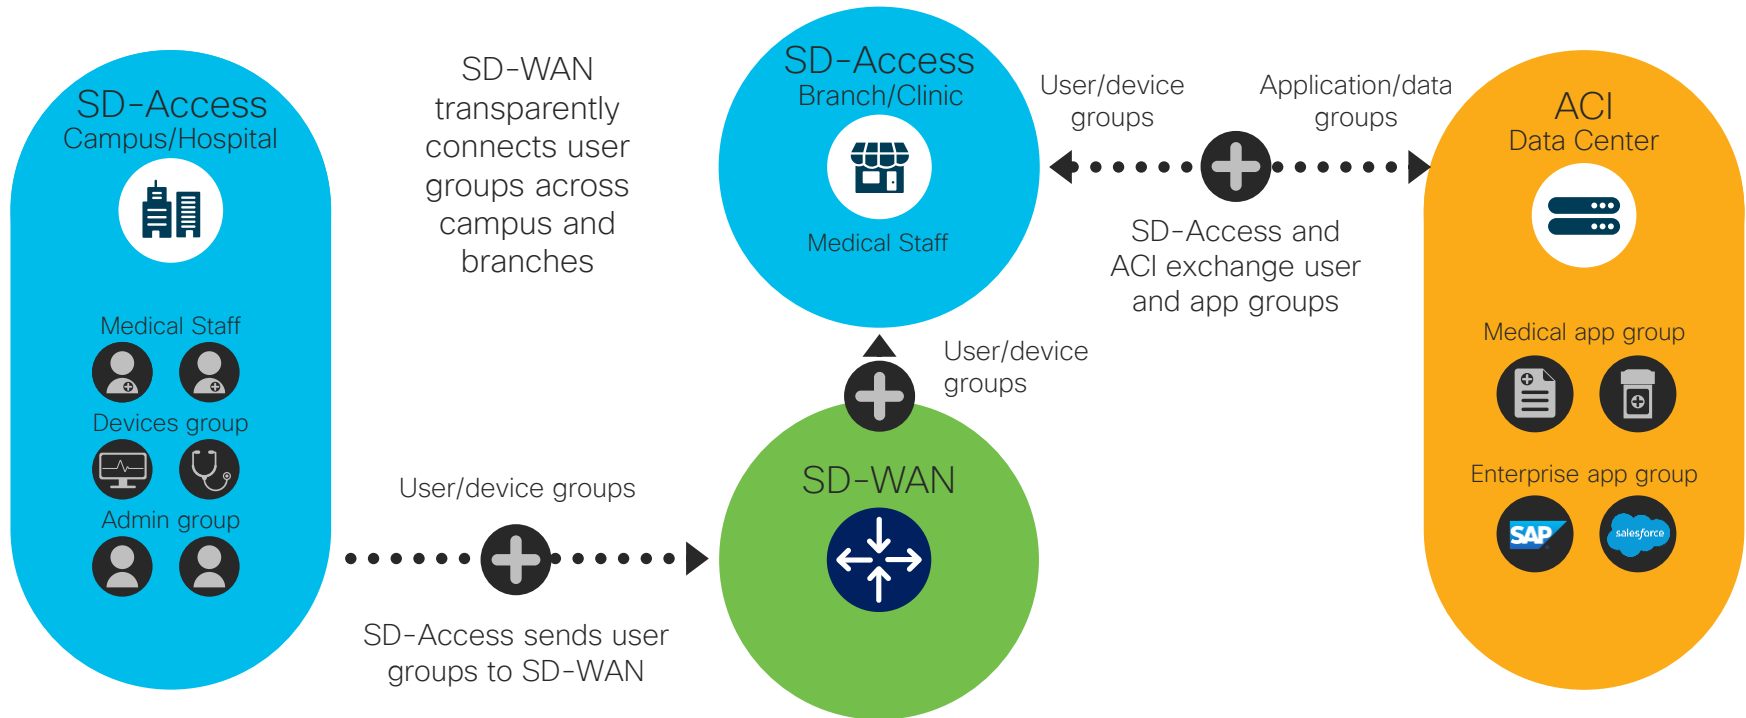


cisco Live!

SDA and SDWAN Pairwise Integration


Segmentation Policy Follows the User

No reconfiguration required if users move



Segmentation maintained between users and apps at any location

Before: Campus/WAN w/o Integration



SGT (16 bits) Micro-segments
VNID (24 bits) Macro-segments

Cisco DNA Center





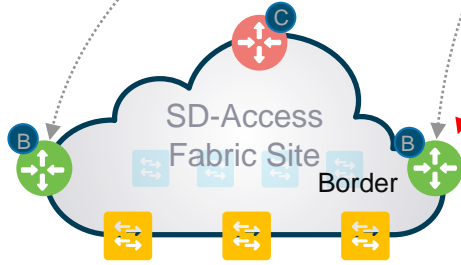
Unmanaged NNI
 Manual sync between domains
 Out of band supplementary channels

SGTs in SXP via ISE

vManage

Transit
 (Separate WAN)

MANAGEMENT
 & POLICY



LISP



BGP VRF-lite

MP-BGP / Other

BGP VRF-lite

LISP


CONTROL-PLANE



DATA-PLANE



Integrated Multi-domain SD-Access/SD-WAN




MANAGEMENT
&
POLICY

Cisco DNA Center



vManage



LISP

OMP

LISP

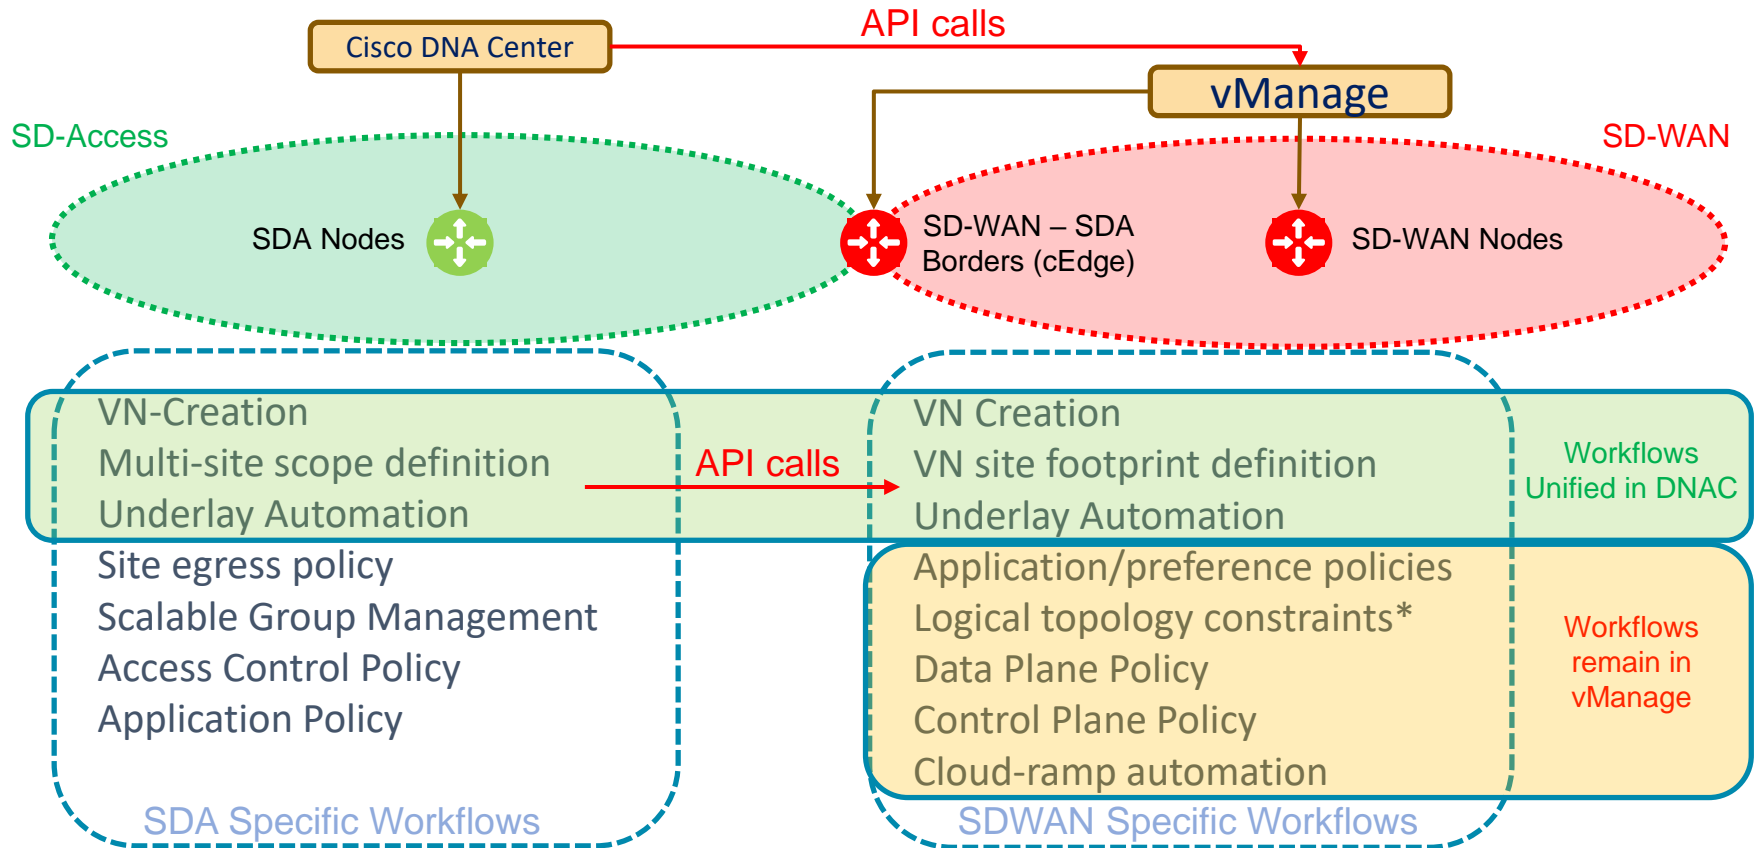
CONTROL-PLANE



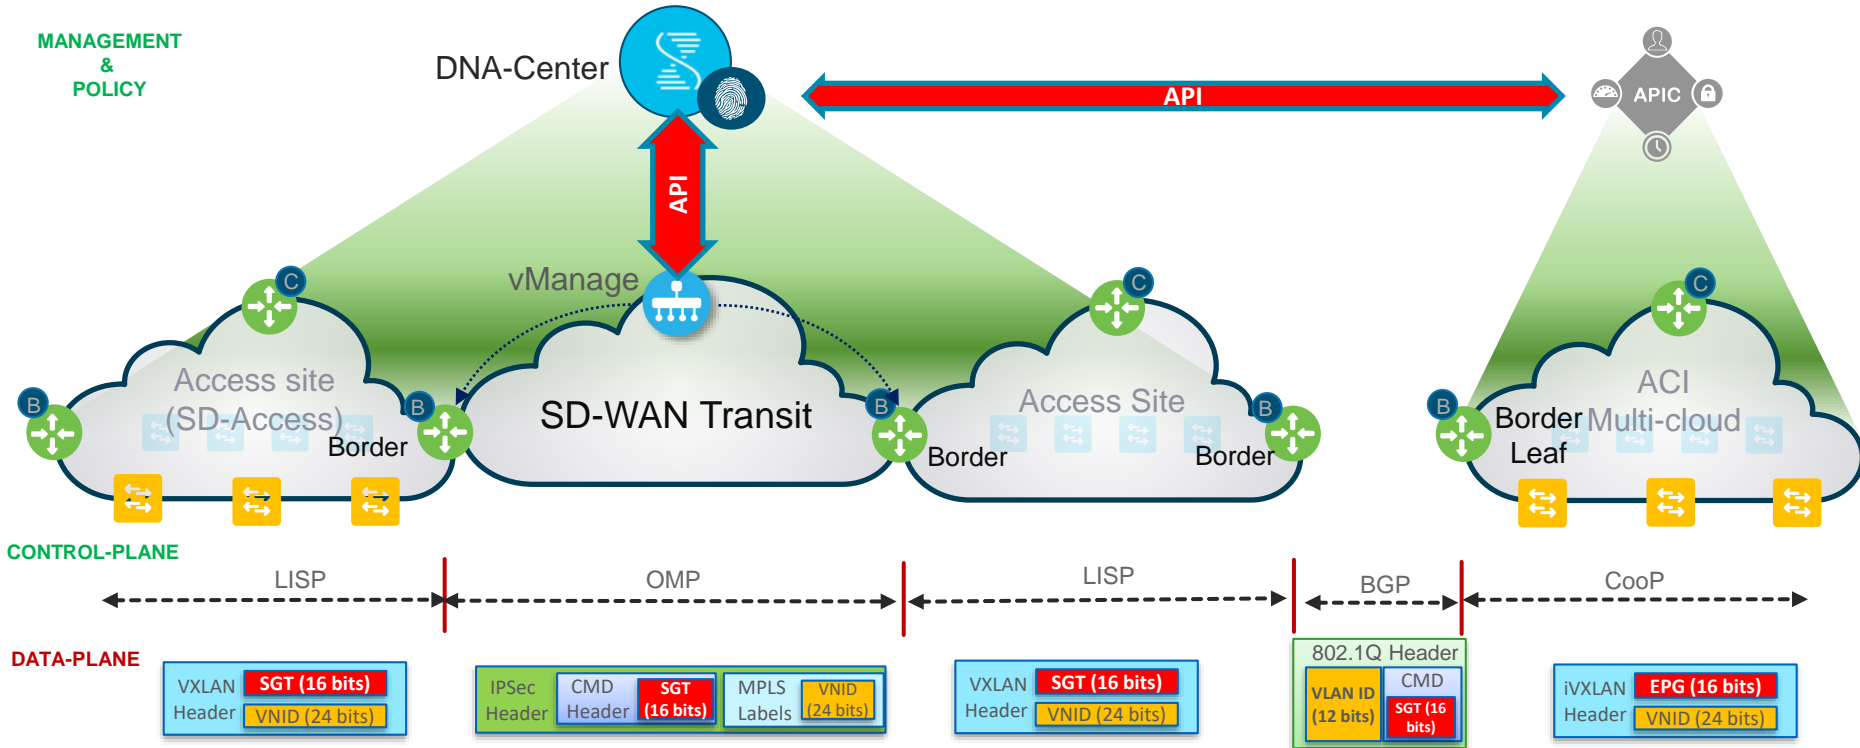
DATA-PLANE

cisco *Live!*

Management of Integrated SDA+SD-WAN



End-to-end segmentation and cross domain interworking

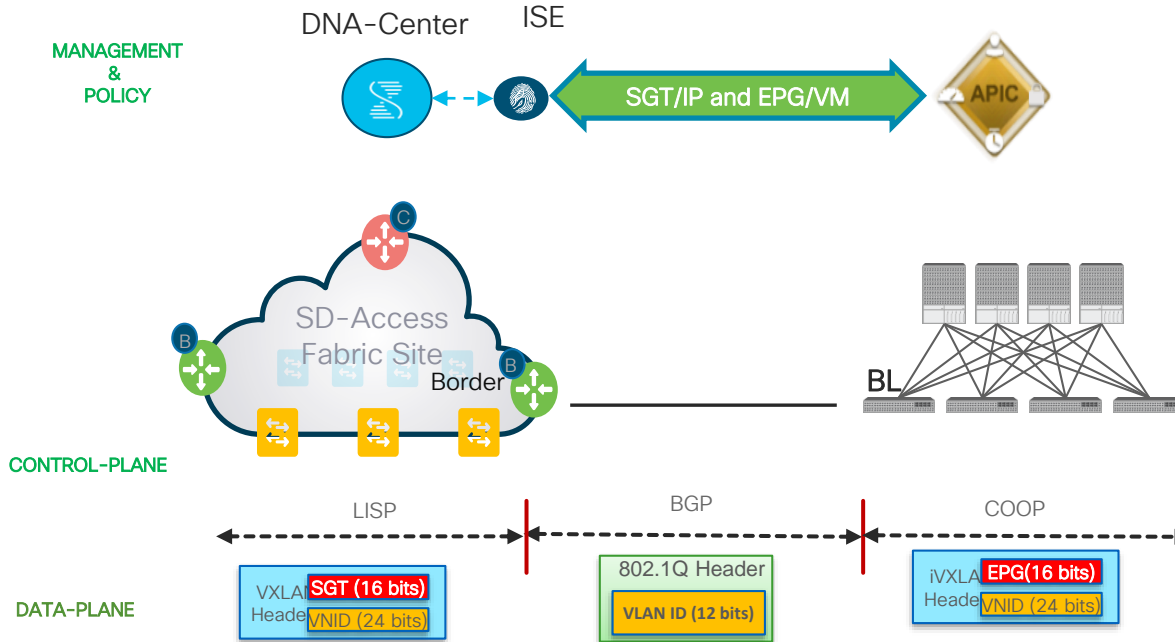


cisco *Live!*

ACI and SDA Pairwise Integration - Phase 1

ACI and SDA Pairwise Integration

Phase 1 - SDA-ACI: Group/Identity Mapping



ACI 3.2 Scale
EX, FX and FX2 Hardware

SGT -> External EPG	250
Number of Mappings	64k
Mappings per External EPG	8k
Transaction rate (target)	100/s

Phase1: Solution Testing completed and supports:

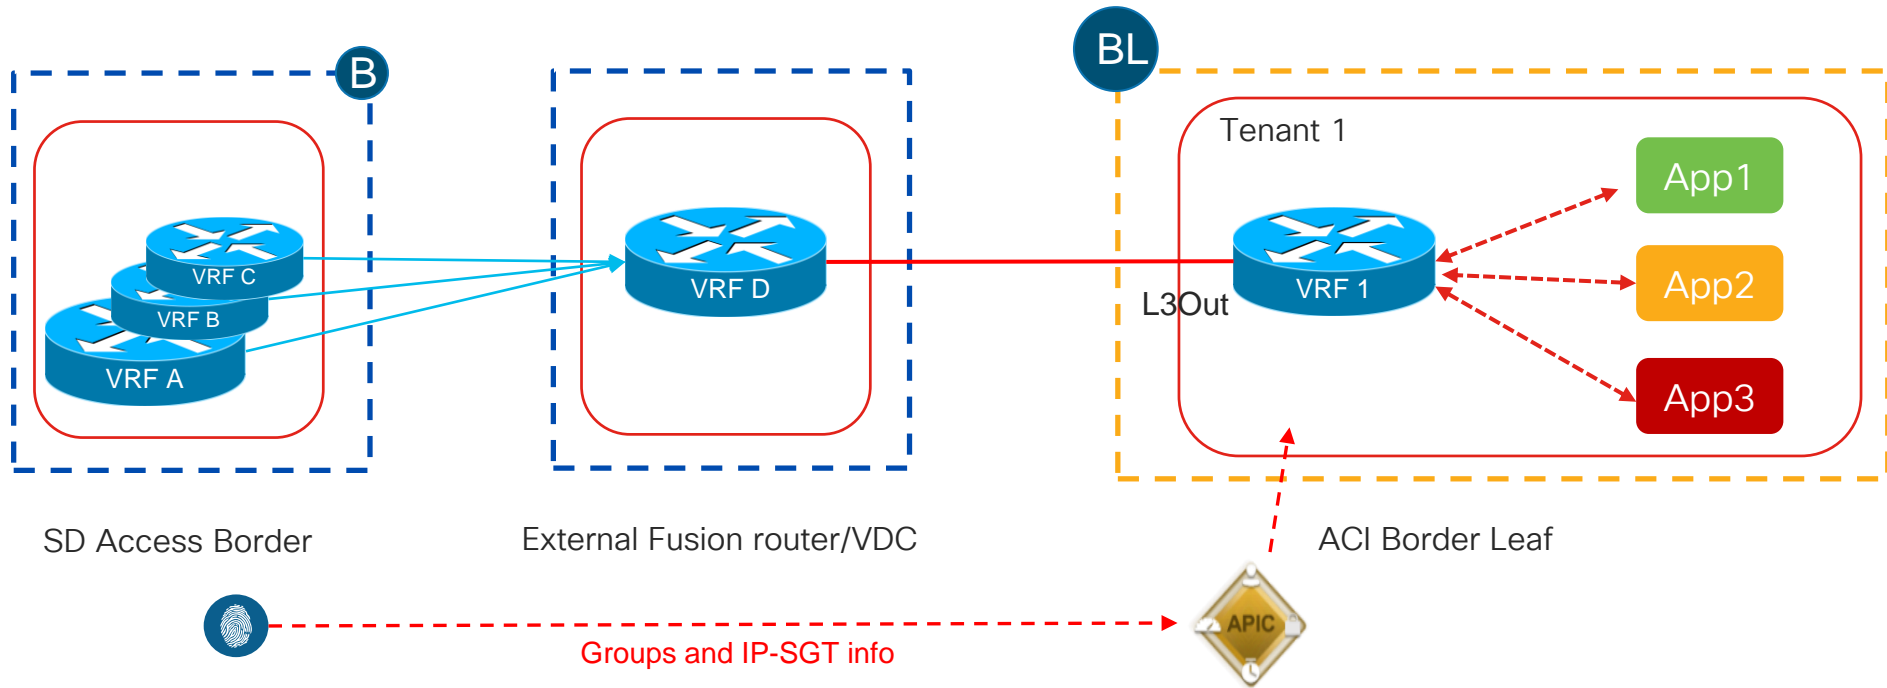
- IP Data path
- Exchange of SGT and EPG at the control plane layer
- IP-SGT/EPG bindings in both directions

- ACI 3.2.x above
- ISE Version 2.4 Patch 6
- DNA Version 1.2.10



Phase 1 SDA-ACI

Current Solution: Single VRF, Single Tenant



Note: Shared L3Out on ACI BL nodes currently not supported

Where Should the Policy Be Applied?

Policy Enforcement on the Application Domain

Groups Provisioned from SDA to ACI (by ISE)

The screenshot shows the Cisco DNA Center interface with the 'POLICY' tab selected. Under 'Group-Based Access Control', the 'Scalable Groups' sub-tab is active. A table lists various groups and their associated virtual networks.

Name	Virtual Network
Auditors	DEFAULT_VN
BYOD	DEFAULT_VN
CANADASGT	DEFAULT_VN
CBASGT	DEFAULT_VN
CloudSvrs	DEFAULT_VN
Contractors	DEFAULT_VN
Developers	DEFAULT_VN
Development_Servers	DEFAULT_VN
Employees	DEFAULT_VN

ISE provisions SGTs info to APIC via REST API*

The screenshot shows the Cisco APIC interface for Tenant Pod01. The 'Networks' section is expanded, showing a list of external EPGs associated with the L3Out. A green dashed box highlights this list.

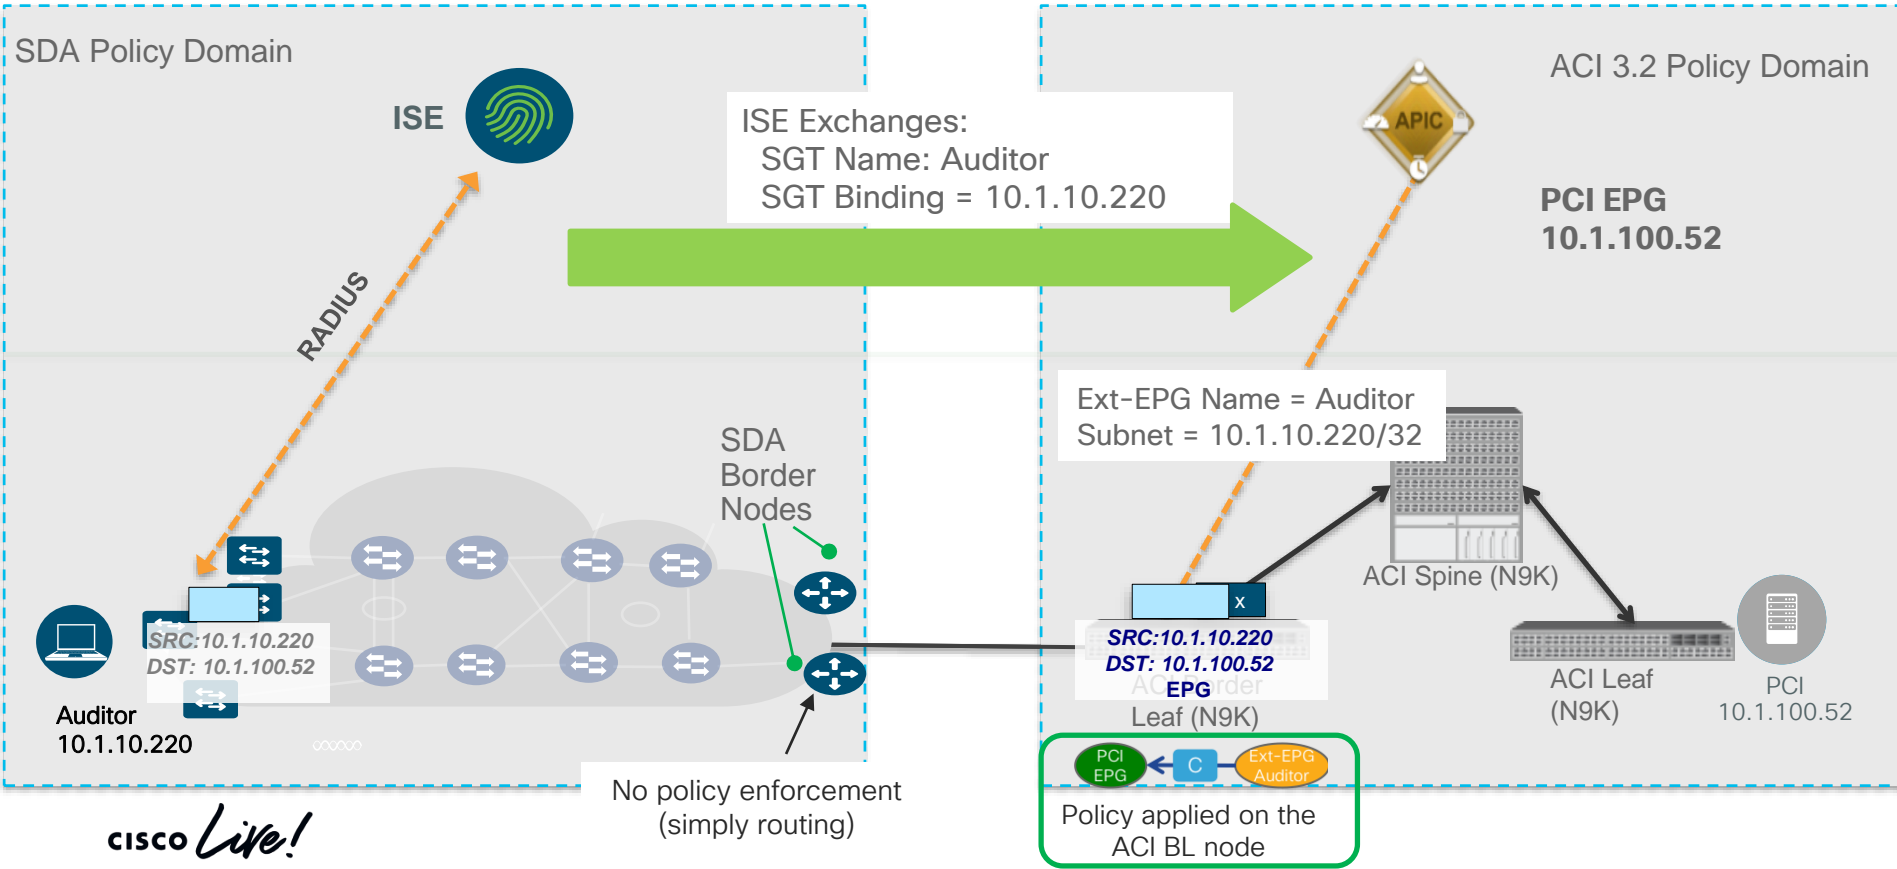
- Auditors_SGT
- BYOD_SGT
- Contractors_SGT
- Developers_SGT
- Development_Servers_SGT
- Employees_SGT
- Guests_SGT
- Network_Services_SGT
- PCI_Servers_SGT
- Point_of_Sale_Systems_SGT
- Production_Servers_SGT
- Production_Users_SGT
- Quarantined_Systems_SGT
- Test_Servers_SGT
- TrustSec_Devices_SGT
- default

External EPGs associated to the L3Out

* Provisioning of SGTs to APIC can be controlled at the SGT level

Policy Enforcement on the Application Domain

Applying Policy on the ACI BL Nodes



Policy Enforcement on the Application Domain

Enforcement Scale in ACI

	EX	FX	FX2
SGT ↔ External EPG	250	250	250
SGT IPv4 Mapping (/32)	64k	64k	64k
SGT IPv6 Mapping (/128)	24k	48k	24k
SGT IPv4+IPv6 Mapping (Dual-Stack)	24k + 24k	32k + 32k	24k + 24k
Policy	8k	128k	8k
Max. IPv4 Bindings per Ext-EPG	8k	8k	8k
Max. IPv6 Bindings per Ext-EPG	8k	8k	8k

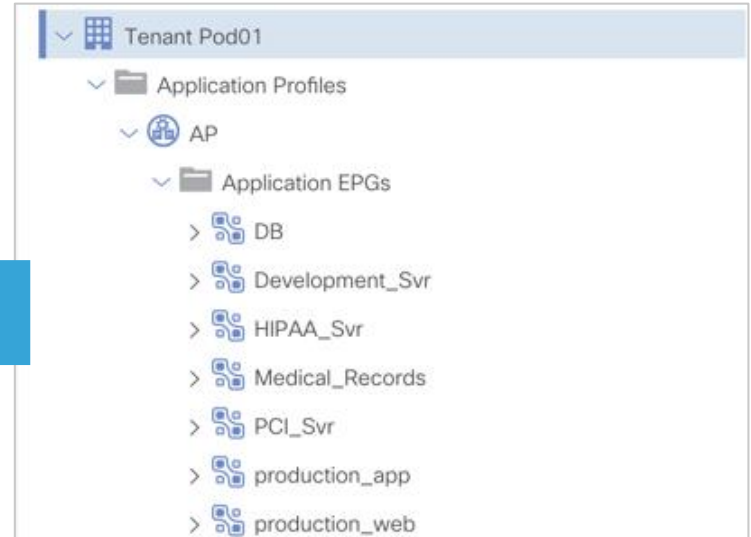
Note: “Egress Policy Enforcement” on Campus Facing ACI Border Leaf

Policy Enforcement on the SDA Domain

SDA Learning Groups from ACI

Group-Based Access Control Policies	Scalable Groups
Name	
AP_DB_EPG	
AP_Development_Svr_EPG	
AP_HIPAA_Svr_EPG	
AP_Medical_Records_EPG	
AP_PCI_Svr_EPG	
AP_production_app_EPG	
AP_production_web_EPG	

ISE retrieves Application EPGs from APIC via REST API*



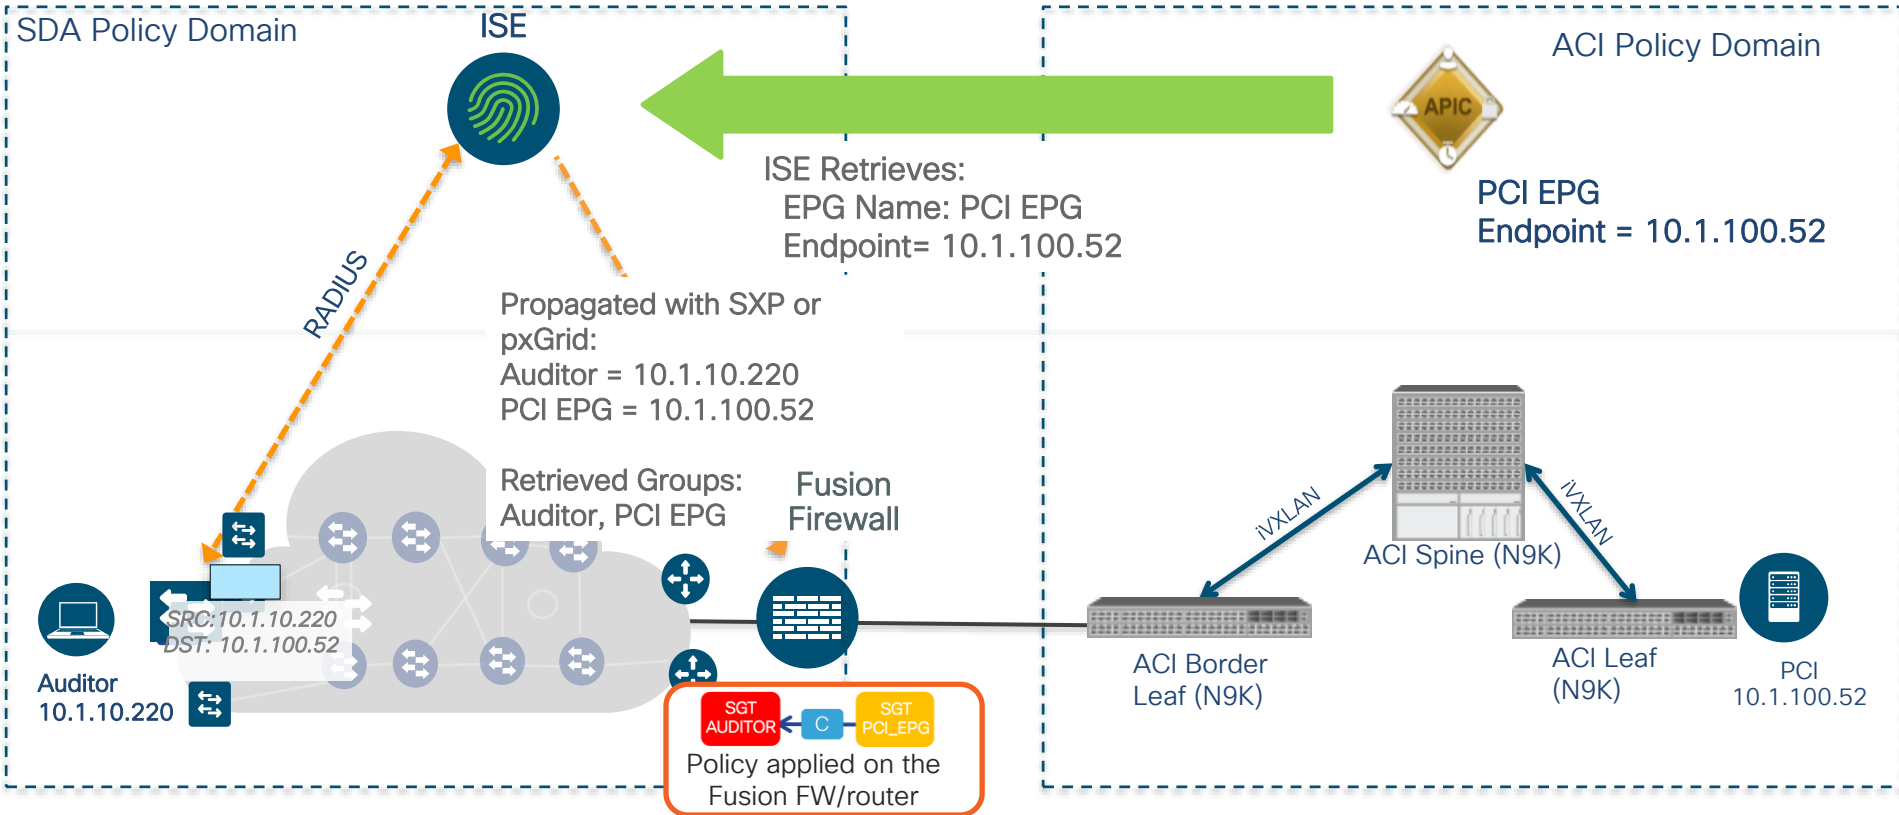
Application EPGs associated to the ACI Tenant integrated with ISE

SGTs created in DNAC

* All the EPGs defined for the specific Tenant are retrieved

Policy Enforcement on the SDA Domain

Apply Policy on the SDA Site (for Example on a Fusion Firewall)



Policy Enforcement on the SDA Domain

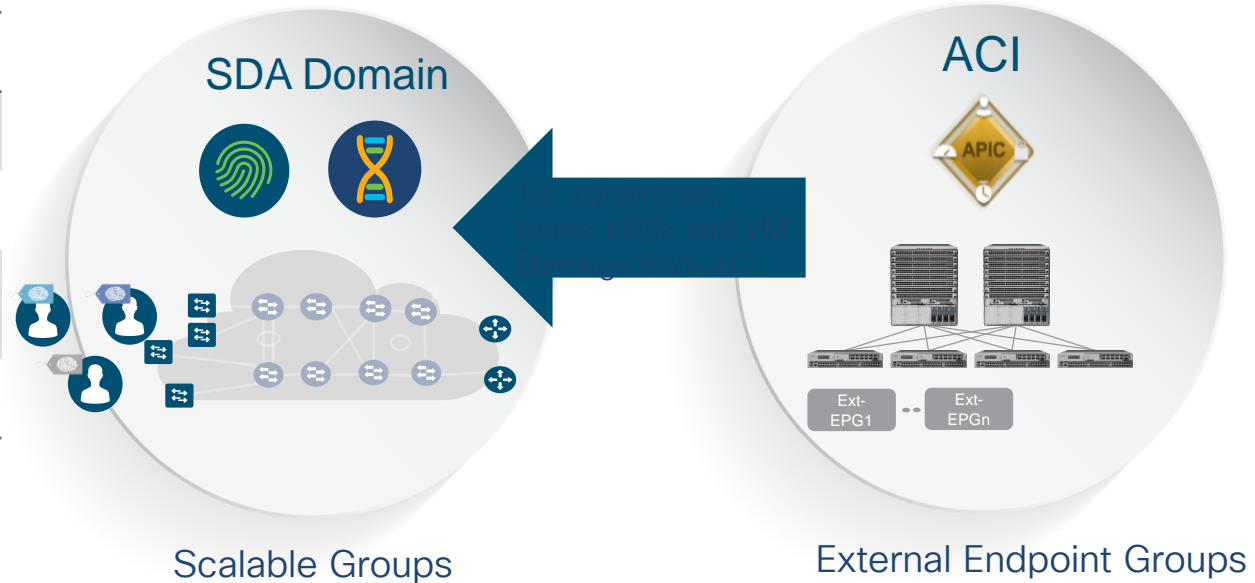
Scaling Enforcement in SDA Environment

ISE/SDA Scale

Numbers of Groups	1000
Number of Mappings*	250k
SXP Peers*	200
pxGrid Peers**	200

*Per pair of ISE SXP Nodes

** Per ISE pxGrid node

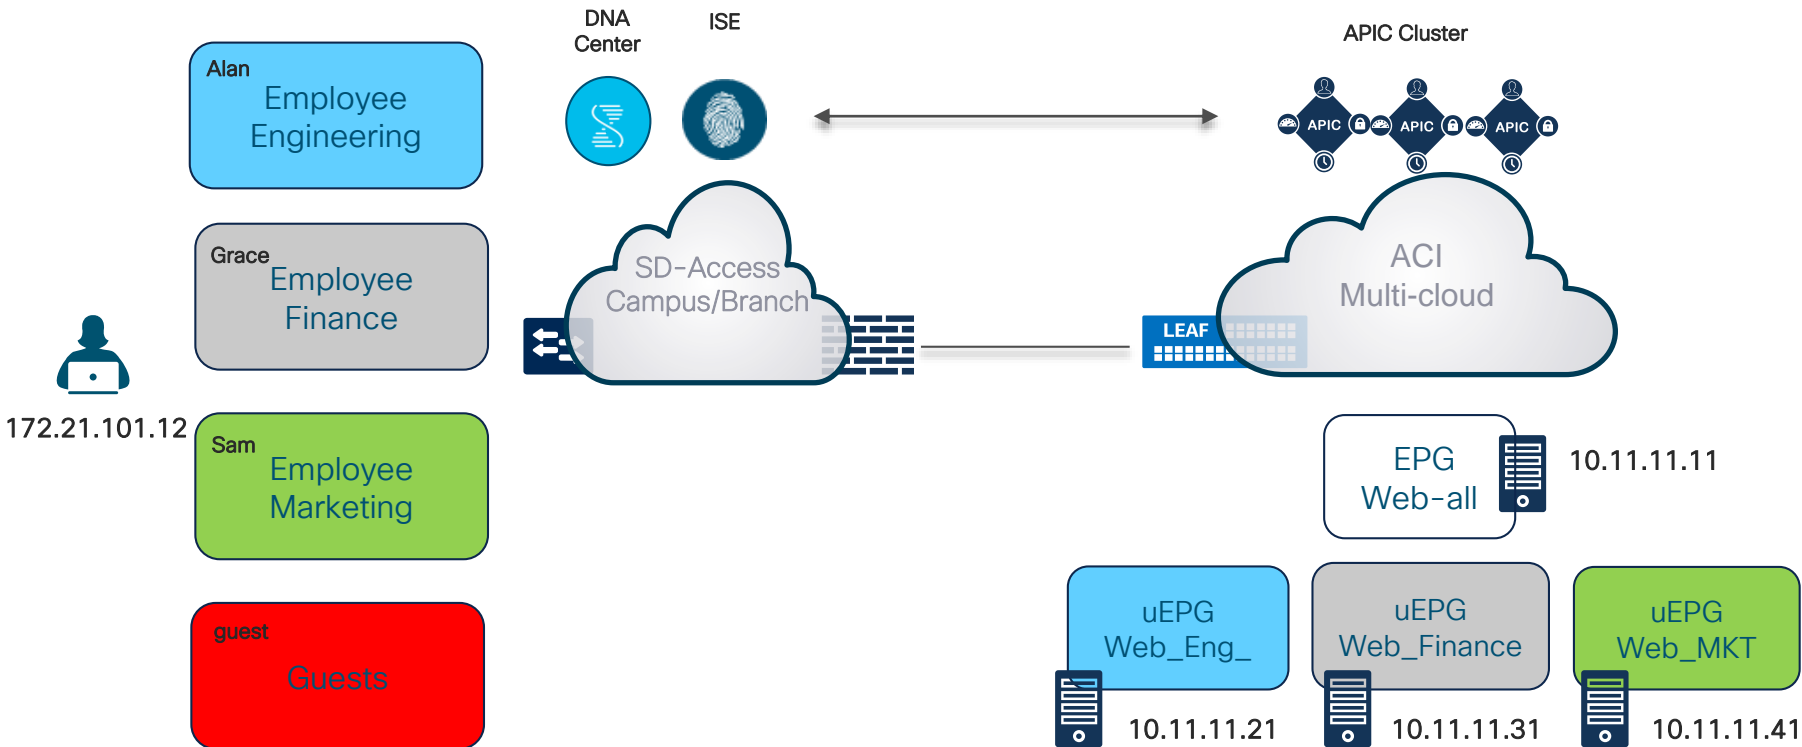


Fusion platforms capable of > 250k Mappings include:
ASR, ISR4k, C6800, ASA and FirePower Appliances

Demo 2: Multi-Domain Segmentation ACI and SDA



Topology ACI and SDA Integration - Phase 1



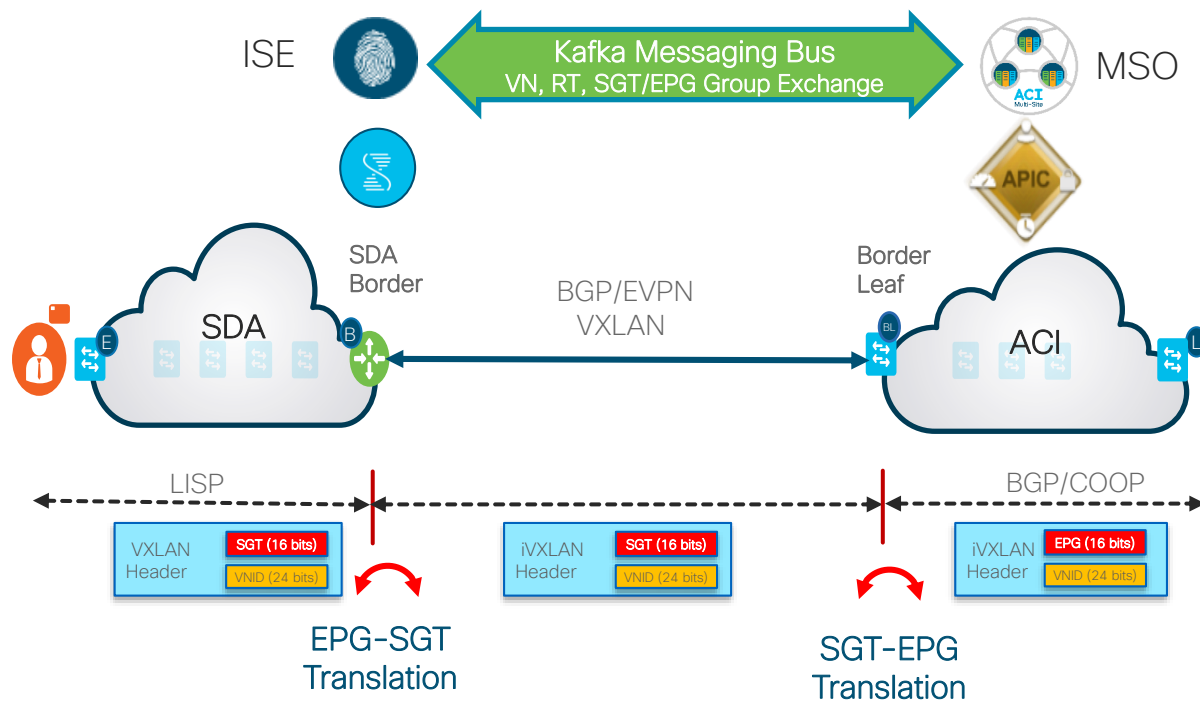
ACI and SDA Pairwise Integration - Phase 2

Disclaimer

- Phase 2 of ACI/SDA integration is currently planned for Q3CY20. As a consequence some of the specific implementation options described in the following slides may slightly change before FCS



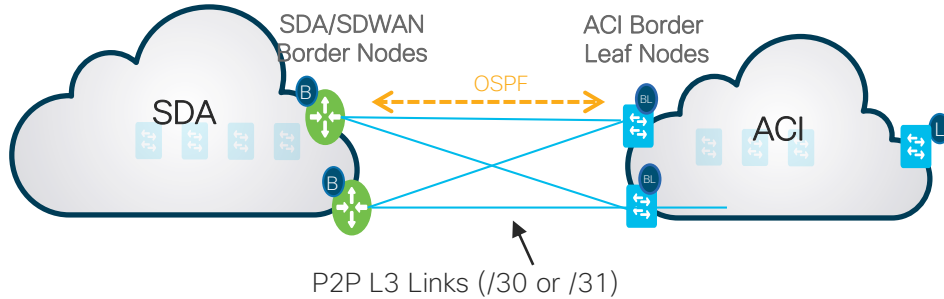
Phase 2 SDA-ACI Overall Architecture




Area	Description
Policy Plane	<ul style="list-style-type: none"> {VN, IP, group} sharing via Kafka Messaging Bus Route Targets exchange between ACI and SDA (through ISE)
Control Plane	<ul style="list-style-type: none"> BGP-EVPN with exchange of subnet information
Data Plane	<ul style="list-style-type: none"> iVXLAN with group information Endpoint data-plane learning on ACI BL nodes

Underlay Connectivity between SDA Border Nodes and ACI BL Nodes

Phase 2 SDA-ACI Connecting Border Nodes



- Direct back-to-back connectivity between SDA Border Nodes and ACI Border Leaf Nodes
 - Square and full mesh topologies both supported
 - SDA Border Nodes could be co-located in the DC location
- Use of point-to-point Layer 3 interfaces (/30 or /31 mask)
- Point-to-point OSPF peerings to exchange TEP reachability information across domains
 - Independently configured in each domain




- SDA and ACI domains can also be connected through a generic Layer 3 infrastructure
 - Underlay configuration in the L3 network must be separately handled
 - Increased MTU support required in the L3 core
- OSPF still used between the border nodes and the first Layer 3 hop device in the L3 core
 - Possible to redistribute OSPF into a different protocol used in the core

Control Plane Considerations

Control Plane Considerations


ACI-SDA BGP EVPN Peering



- BGP EVPN peerings to exchange prefix information for Campus and Application subnets
 - Full mesh BGP EVPN adjacencies between EVPN RIDs
- Information exchanged between border nodes:
 - VNIDs for the VRFs defined in each domain (downstream VNID assignment)
 - Route-Targets (RTs) value used to control import/export of prefixes into each VRF (Symmetric RT approach)
 - Router-MAC for the border node originating the prefix
 - TEP IP address to be used as next-hop
- VNID, RMAC, and TEP IP are used to construct VXLAN header for packets forwarded between SDA and ACI domains

Control Plane Considerations

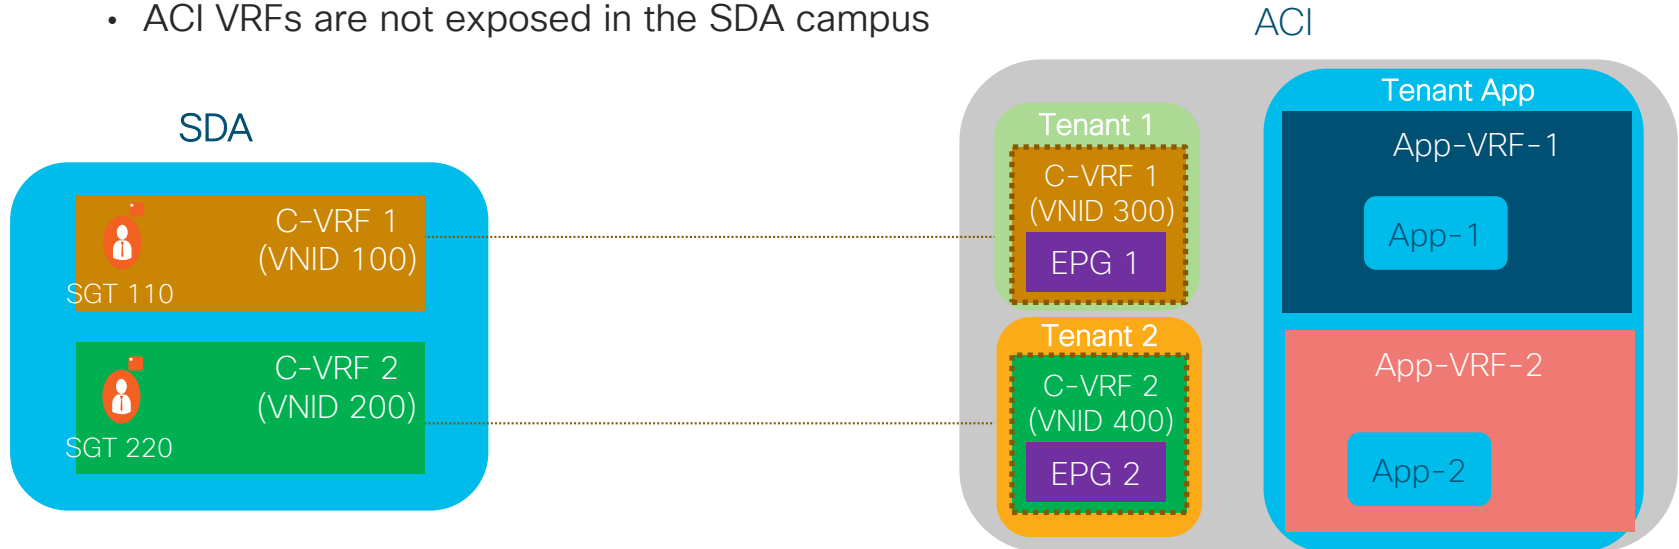
Route Exchange



- Routes are exchanged as EVPN type-5 routes with VNID of dedicated VRF
- Routes are imported within a domain based on the specific VRF RT import policies
- SDA border nodes push to ACI the subnets' routes for the users in the Campus requiring connectivity to the ACI services
- ACI BL nodes advertise the following routes to the SDA border nodes:
 - BD subnets for applications made available from ACI fabric
 - Prefix routes learnt in the ACI fabric from peers connected to other BL L3Outs
 - Specific /32 and /128 host routes for BDs that are stretched between ACI Pods/Sites (not at FCS)

Campus VRF Extension into ACI

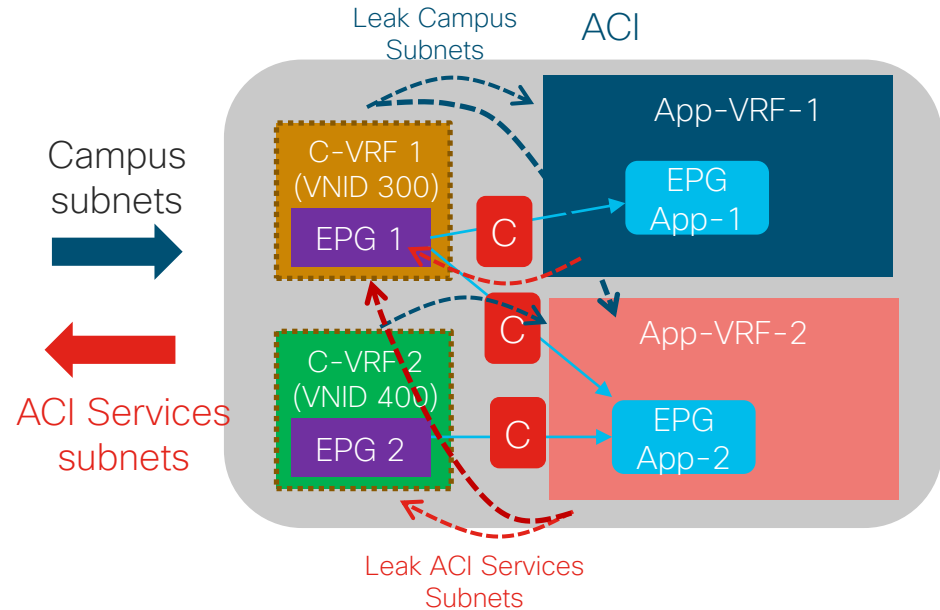
- Support multi-tenancy/multi-VRF design with minimal or no change to existing design on SDA and ACI side
- Allow campus to expose multiple VRFs to DC and ACI to expose apps from multiple VRFs to campus
 - SDA initiates a “Remote Tenant” setup in the ACI domain for each Campus VRF
 - For each defined Campus VRF (C-VRF) there is a corresponding C-VRF created on ACI
- ACI VRFs are not exposed in the SDA campus



Campus VRF Extension into ACI

Route-Leaking in ACI

- **Campus SG consuming an ACI Service:** in ACI is represented as a “shared service” contract between C-VRF and the VRF(s) of the different Application EPGs representing the ACI services
- The subnets representing the ACI services will be leaked into C-VRF on the ACI Border Leaf nodes and advertised toward the Campus through BGP EVPN
- Similarly, the campus Subnets are advertised from the SDA border nodes into the C-VRF in ACI through BGP EVPN and leaked into one or more application VRFs



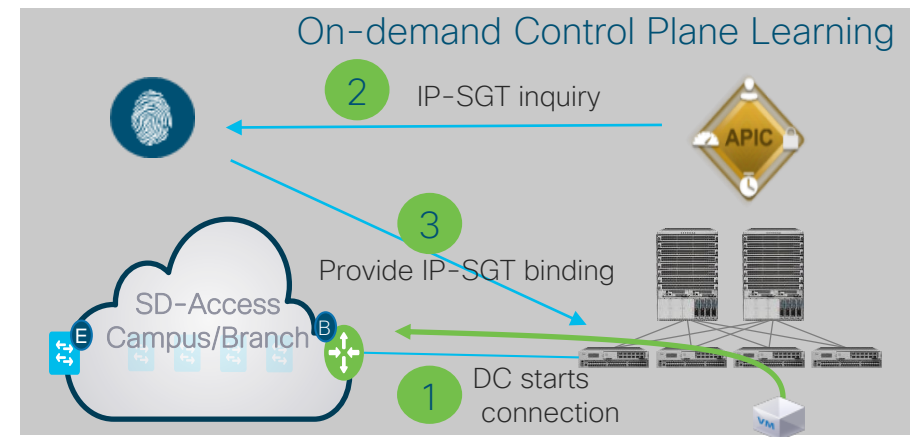
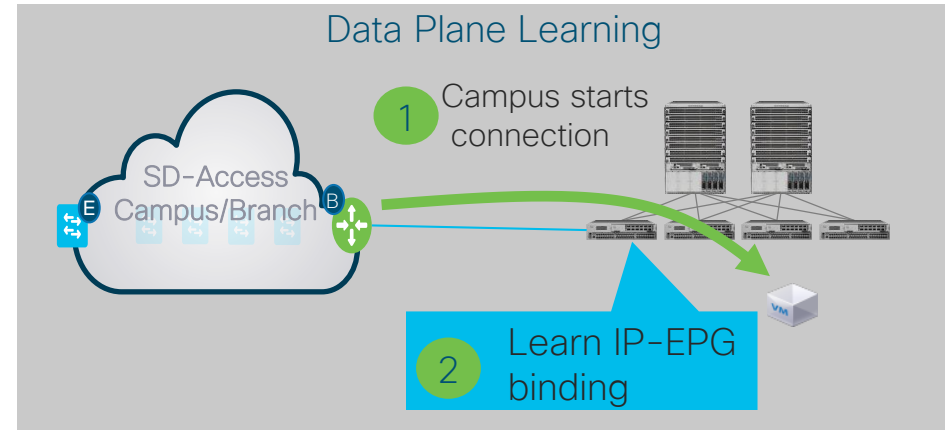
Data Plane Considerations

Phase 2 SDA-ACI

ISE-APIC/MSO Policy and Data Plane Learning

- Each campus SGT is represented as an External EPG in the ACI Border Leaf
- ACI BL learns the mapping of IP-SGT-EPG from data packet
 - Required because adjacent Campus IP addresses may be assigned to different SGTs
- Inquiry ISE for IP-SGT mapping when needed
 - E.g. when EP in ACI initiates the connection toward the campus
- Each domain can apply their policies independent of the other domain
 - On ACI the policy is always applied on the BL nodes (permit/deny/redirect to L4L7 graph, etc.)

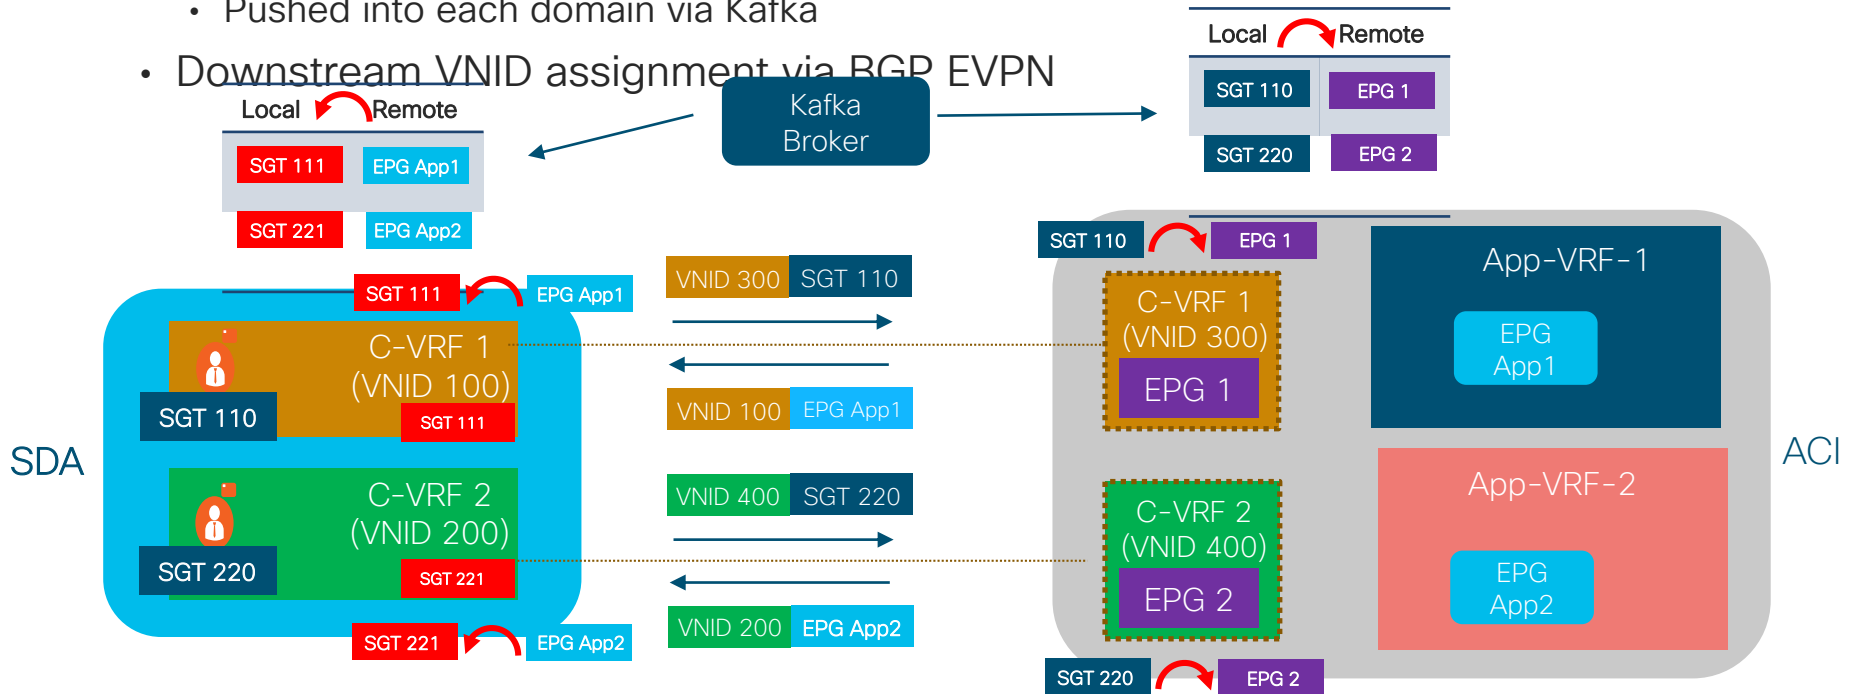
CISCO *Live!*



Campus VRF Extension into ACI

Class-ID Translation between Domains

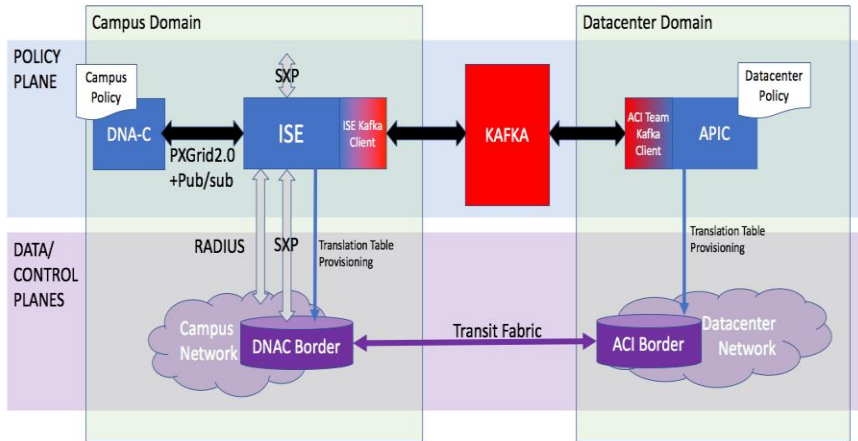
- Class-ID translations to keep SDA and ACI separated domain for resource allocation
 - Pushed into each domain via Kafka
- Downstream VNID assignment via BGP EVPN



Cross-Domains Policies

Cross-Domains Policies

Use of Kafka Communication Bus

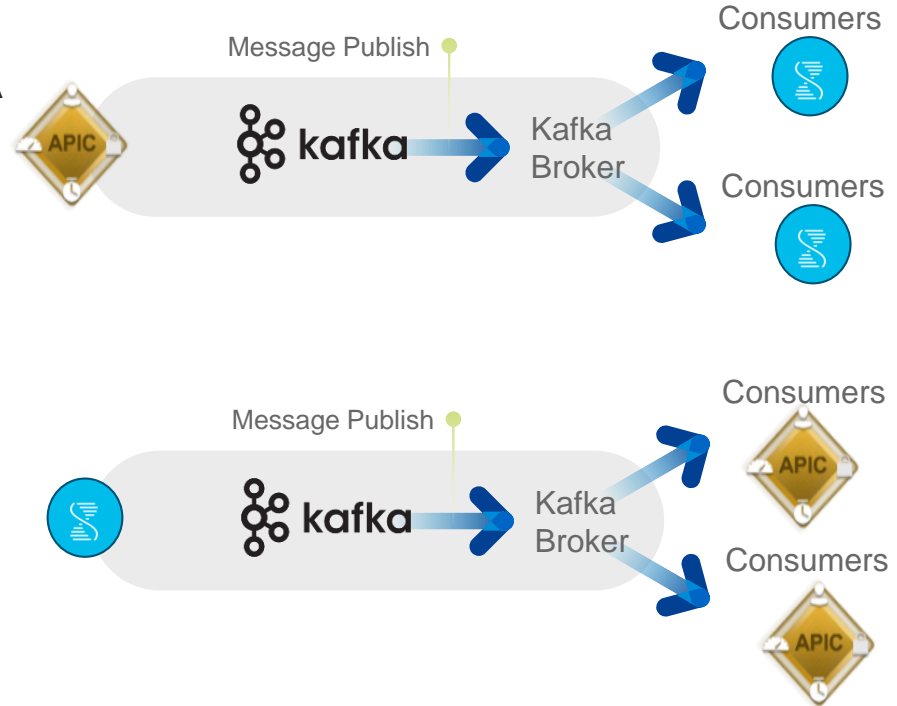


- The Kafka broker does not need to run in a specific location
 - APIC controller cluster already has a Kafka cluster running
 - This Kafka cluster will be used for inter-domain communication at FCS
 - Post-FCS this functionality may be moved to MSO (or MDP)
- Kafka clients running on APIC and ISE
- Kafka Bus is used to join Domains: ACI is one domain (Datacenter), DNAC/ISE is another domain (Campus)
- The Kafka Interchange will allow the domains to peer, and to share data (Remote Tenant, Consumer/Service Gateway and Endpoint (IP-SG)) to allow the unification of Segmentation Policy

Phase 2 SDA-ACI

Cross-Domains Policy

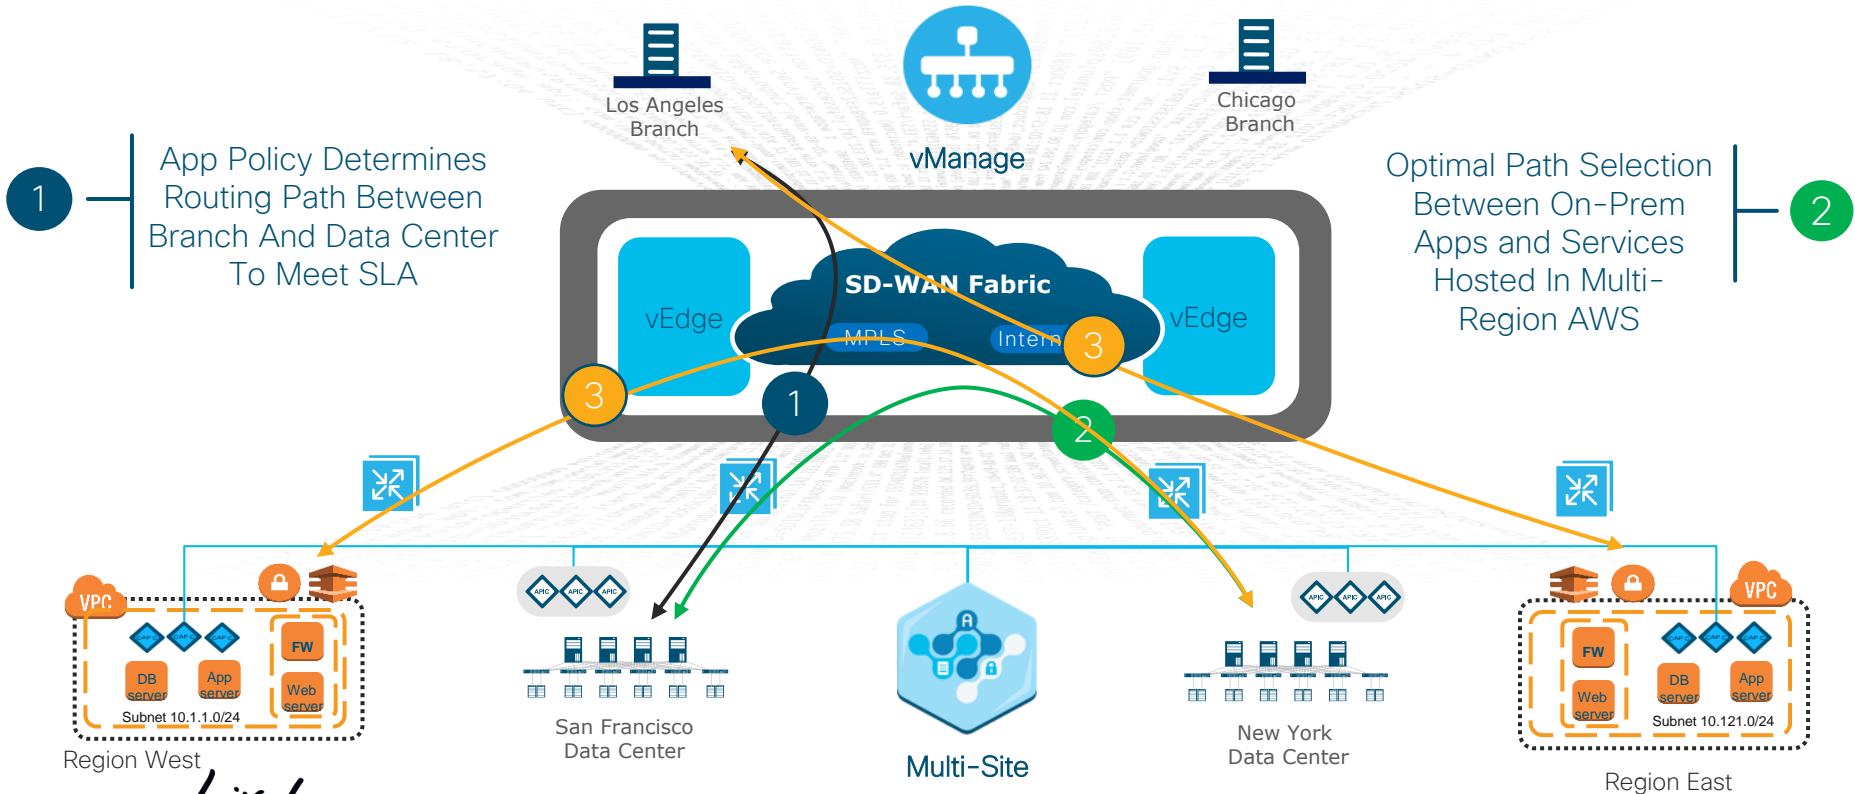
- Use messaging bus (Kafka pub sub) model to publish/consume services between ACI and SDA domains
 - Flexible to construct cross domain policy without a 'uber controller'
 - BGP-EVPN and VXLAN config automation
 - VRF/Tenant instantiation
- Publish services provided by ACI
 - Domain, service name, EPGs, bindings, contracts, provider/consumer, protocol, ports
- Campus subscribe to DC service
 - ISE publishes Campus SGTs that represent the consumer and IP to SGT bindings



ACI and SDWAN Pairwise Integration

ACI and SD-WAN Pairwise Integration

Extend Operational Domain And Policy To Branch & Public Cloud



CISCO Live!

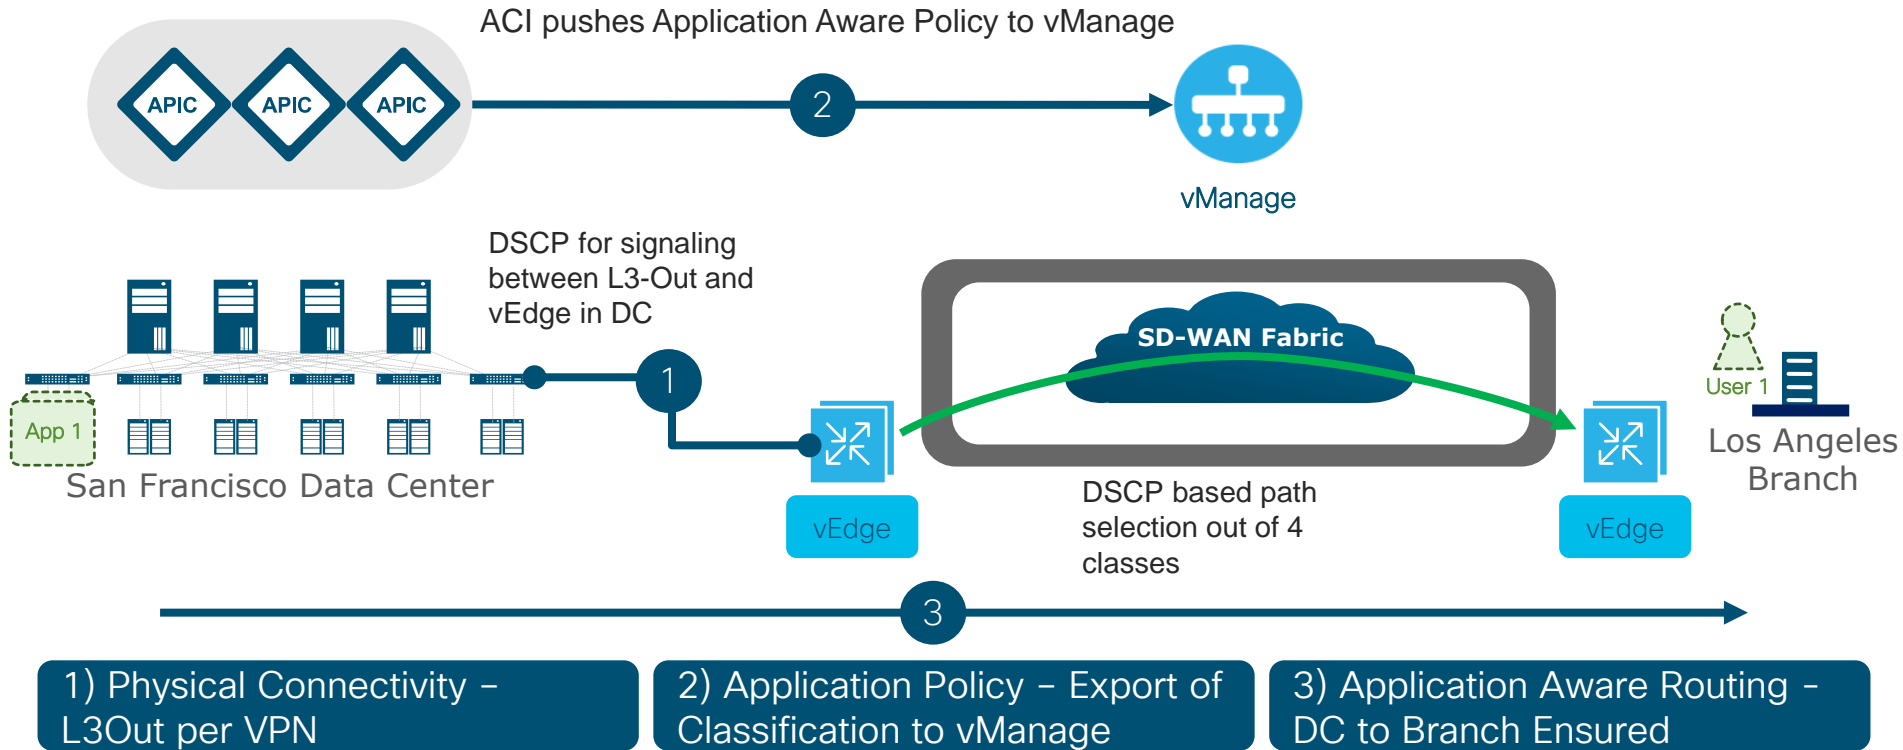
ACI and SD-WAN Pairwise Integration

Use Cases

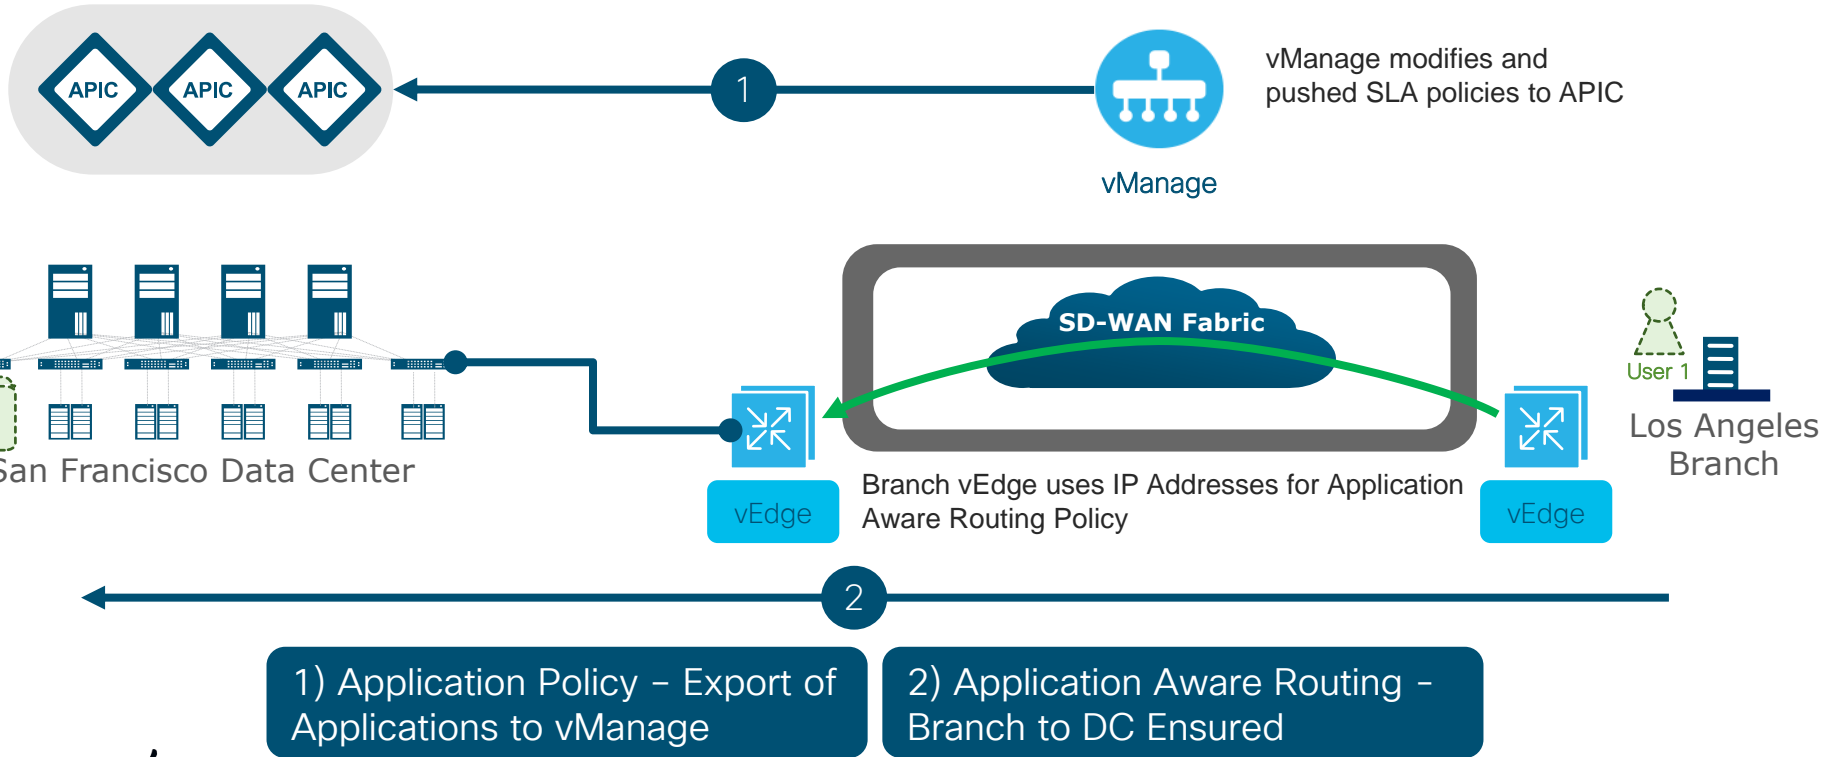
1. Application Based Routing for North South Traffic
 - a) ACI contract specifies outbound SD-WAN traffic engineering (tunnel selection, ...) to vManage and marks outbound traffic accordingly
 - b) ACI contract identifies IP/VIP to SD-WAN controller to mark inbound traffic
2. Remote Leaf, vPod and Multi-Site traffic optimization, transit of iVXLAN traffic across SDWAN
3. ACI Anywhere integration with SDWAN Cloud onRamp

ACI to SD-WAN (Viptela) Integration – Phase

ACI 4.1



ACI to SD-WAN (Viptela) Integration – Phase 1




ACI-SDWAN Release and Platform support

Release	Cisco ACI	4.2(X)/14.2(X)	Release
	vManage	19.2	
	IOS-XE SDWAN	16.12	

Platform Support	ISR4K	ISR43xx, ISR-4431, ISR-4451	Platform Support
	ASR1K	ASR1001-X, ASR1002-X, ASR 1001-HX, ASR 1002 -HX	
	ACI	EX, FX	

Configuration Details

Configuration Overview


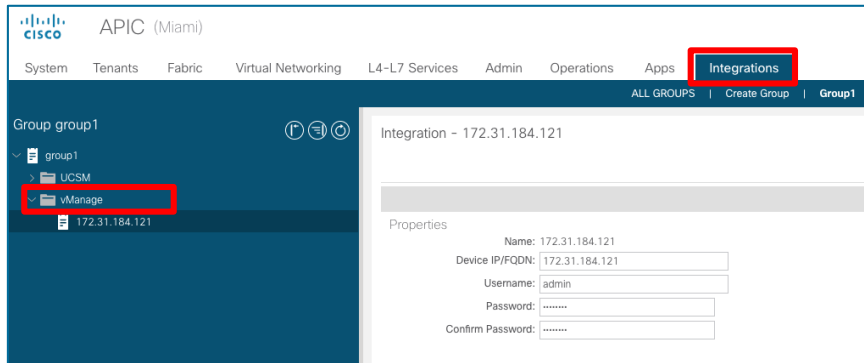


- ACI admin adds vManage integration to APIC
 - SD-WAN SLA policies are automatically pushed from vManage to APIC
 - SD-WAN VPN list is automatically pushed from vManage to APIC
- ACI admin assigns WAN SLA and VPN to the tenant VRF
- ACI admin assigns WAN SLA to L3Out contract

- SD-WAN admin provisions cEdge/vEdge devices and configures service VPNs towards ACI BL
- SDWAN admin adds a site list in vManage
- ACI partner registration appears in vManage
- SD-WAN admin attaches ACI controller to cEdge/vEdge devices
- SD-WAN admin configures central policy for SLA, VPN, and site list and pushes policy to vSmart
- vSmart pushes policy to cEdges/vEdges

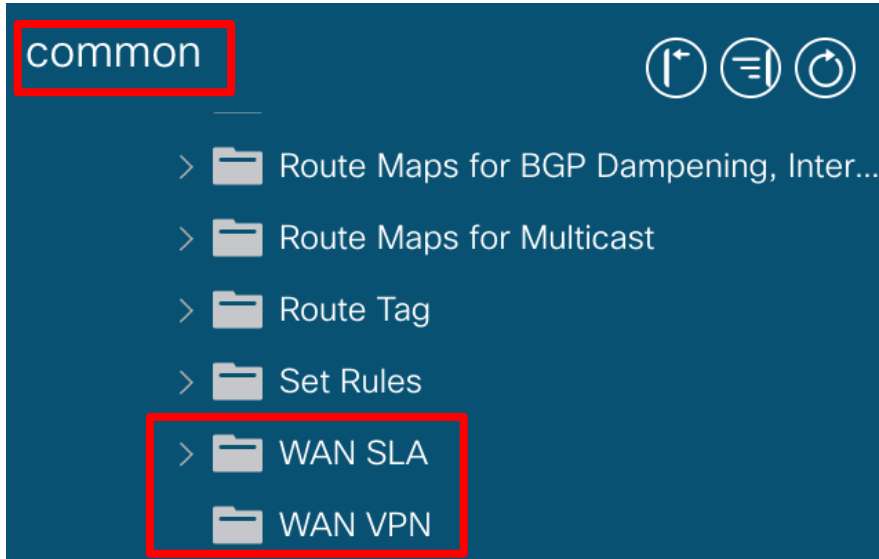
Adding vManage Integration

vManage is added to the APIC under the Integrations tab.



APIC pulls SLA policies and VPNs from vManage

ACI SD-WAN policies



From ACI 4.1 there are two new WAN folders under the common tenant protocol policies.

- **WAN SLA** displays SLA policies pulled from vManage
- **WAN VPN** displays VPNs pulled from vManage

ACI WAN SLA Policies on APIC

- The WAN SLA policies are defined as jitter, delay, and loss values of the WAN link
- SLA policy to DSCP value are randomly chosen
- When an SLA policy is added to an L3Out contract, the border leaf will mark traffic with the DSCP value mapped to that SLA class
- WAN cEdge/vEdge use the DSCP value to map to correct SLA class
- Modifying SLA policies is done on vManage
- Only 4 SLA classes are currently supported

Wan SLA Policies				
Name	DSCP	Acceptable Jitter (ms)	Acceptable Delay (ms)	Acceptable Loss (%)
Bulk-Data	AF12 medium drop	100	300	10
Default	AF13 high drop	100	300	25
Transactional-Data	AF11 low drop	100	50	5
Voice-And-Video	AF21 low drop	100	45	2

ACI WAN SLA Policies on vManage

The screenshot shows the Cisco vManage interface for configuring SLA policies. The breadcrumb navigation is 'CONFIGURATION | POLICIES Centralized Policy > Add Policy'. The main content area is titled 'Select a list type on the left and start creating your groups of interest'. On the left sidebar, 'SLA Class' is selected. The main table displays a list of SLA classes with columns: Name, Loss (%), Latency (ms), Jitter (ms), Reference Count, Updated By, Last Updated, and Action. The 'Action' column for each row contains edit, delete, and refresh icons. A red box highlights the 'Action' column, and a blue arrow points to the edit icon for the 'Bulk-Data' row.

Name	Loss (%)	Latency (ms)	Jitter (ms)	Reference Count	Updated By	Last Updated	Action
Voice-And-Video	2	45	100	0	system	25 Sep 2019 1:27:43 PM	
Transactional-Data	5	50	100	0	system	25 Sep 2019 1:27:40 PM	
Default	25	300	100	0	system	25 Sep 2019 1:27:47 PM	
Bulk-Data	10	300	100	0	system	25 Sep 2019 1:27:48 PM	

SLA parameters
modified on vManage
are exposed to APIC

WAN VPNs



WAN VPN Entries

Name
10
65528

- Service VPN segments in vManage are pulled by APIC and appear under WAN VPN
- Each VPN segment represents a routing domain in the SD-WAN network and can be mapped to a VRF on the ACI domain

Tenant Configuration


Tenant admin configuration steps:


1. Assign VPN to VRF
2. Add SLA policy to contract subject


Under the VRF select a VPN. This is read from the common tenant WAN VPN list


WAN VPN: 

Under the contract subject select the WAN SLA policy and select a QoS Priority (cannot be undefined)

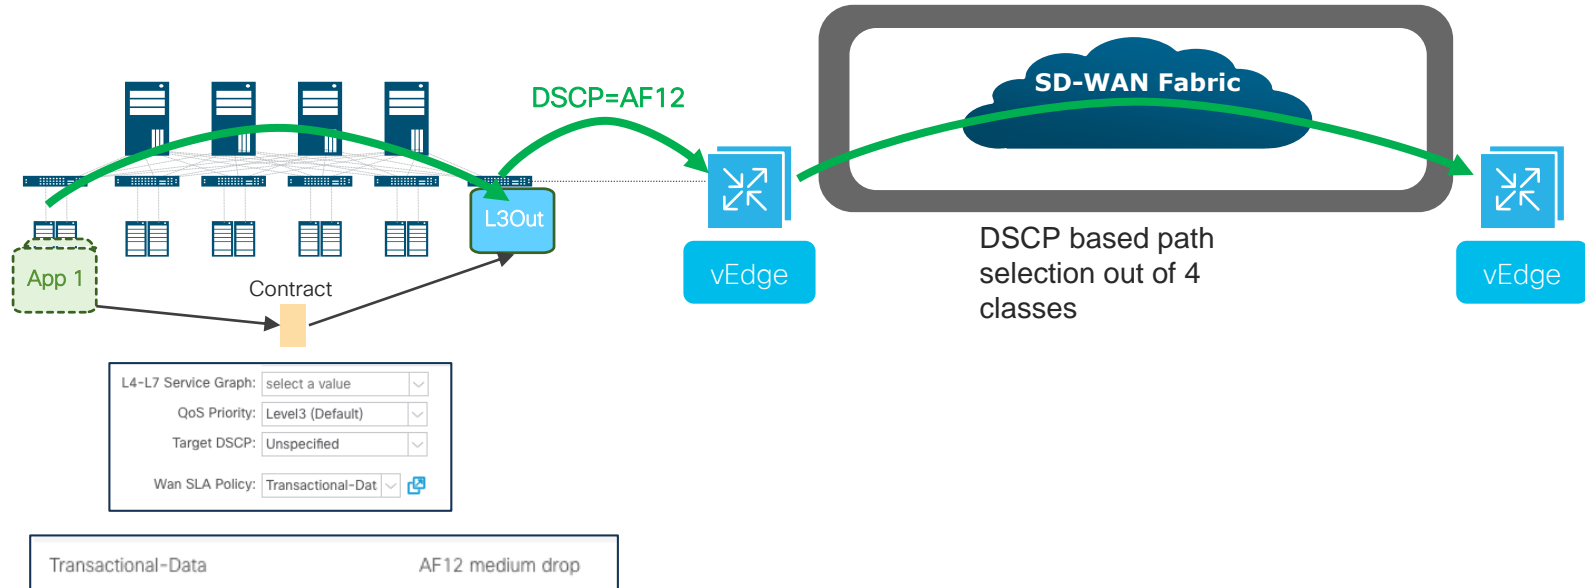
L4-L7 Service Graph: 

QoS Priority: 

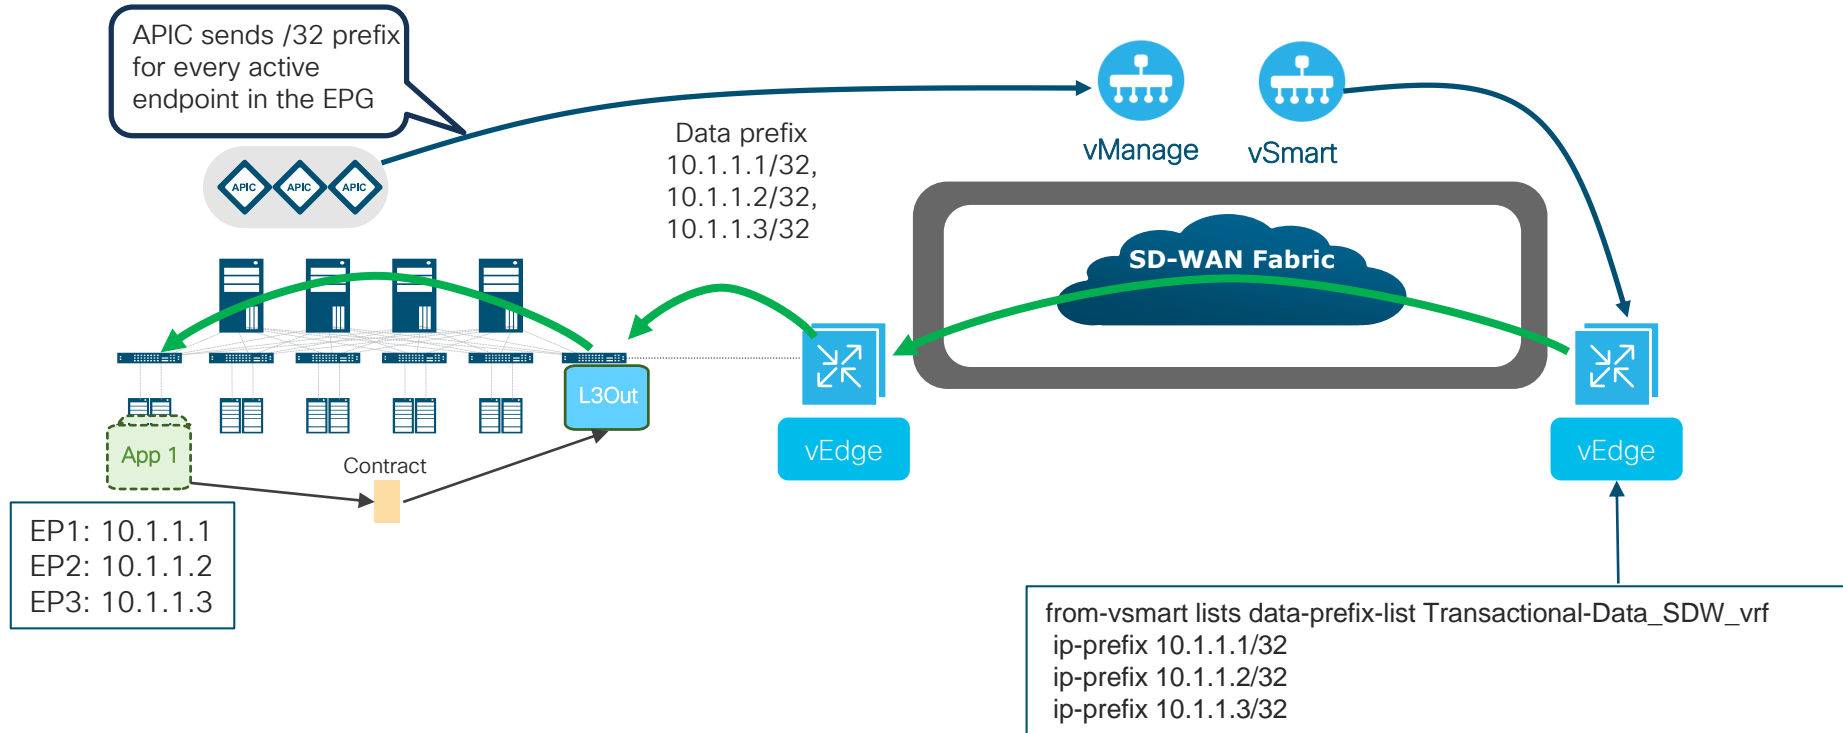
Target DSCP: 

Wan SLA Policy: 

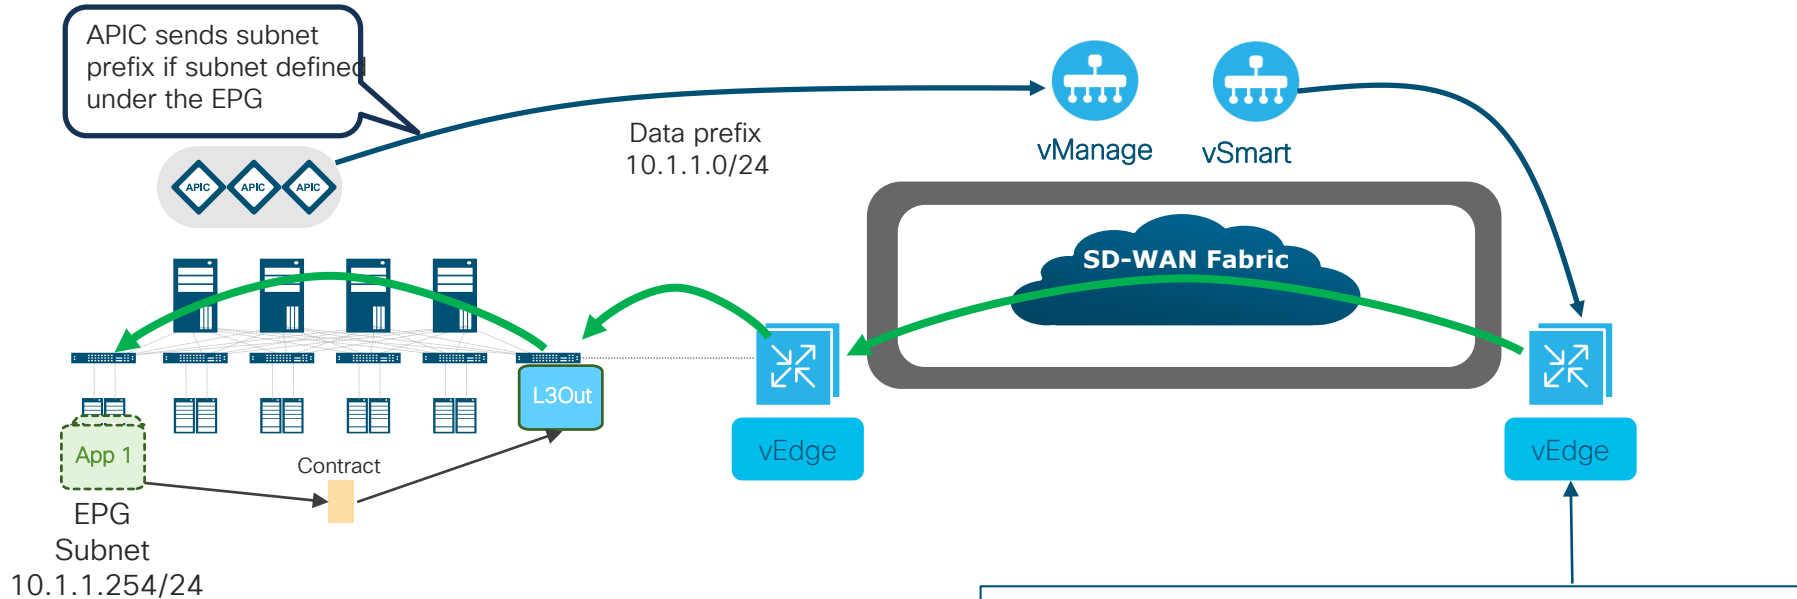
ACI Data Center to Branch



ACI Branch to Data Center



ACI Branch to Data Center



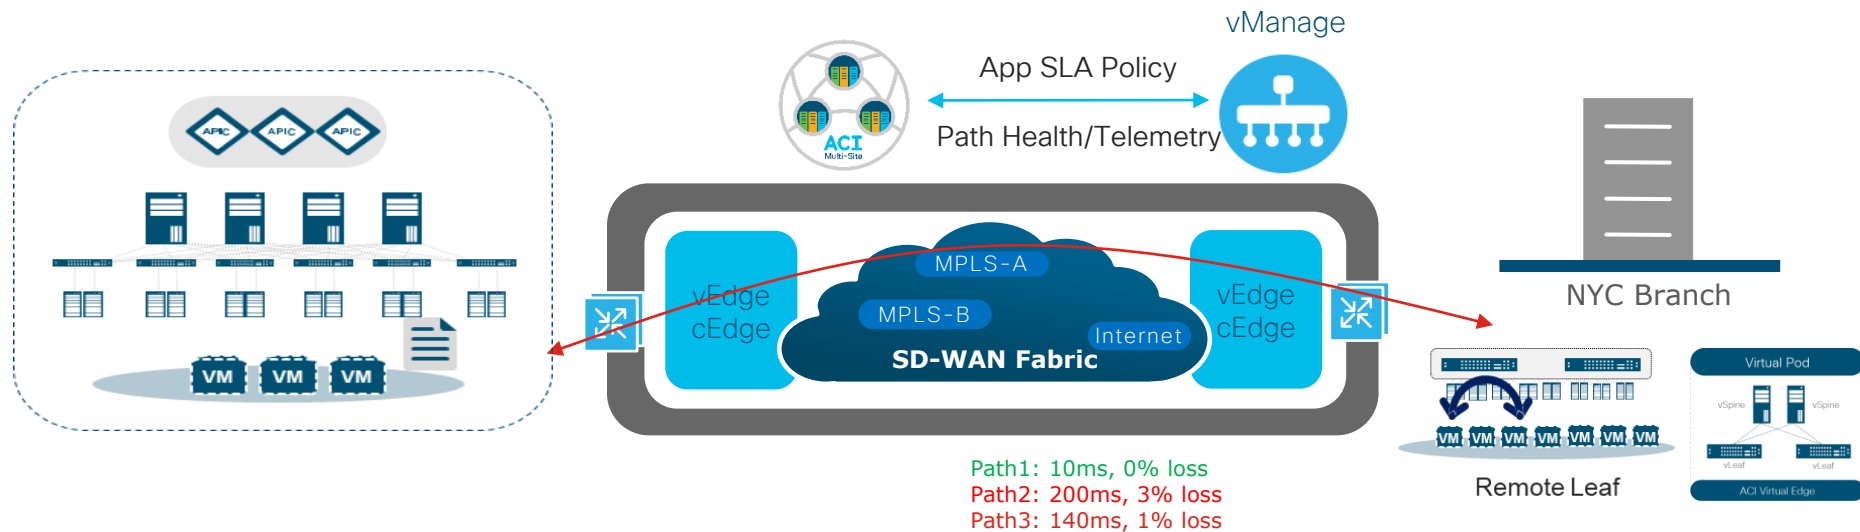
If the subnet is defined under the EPG, only the subnet prefix will be advertised to vManage (not the /32 prefixes)

Use Case 2

Differentiate Traffic Flows between ACI Domains

ACI SD-WAN – Use Case 2

Branch with Remote Leaf or vPod Use Case



- App SLA policy determines routing path selection between ACI Fabric and RL or vPod branch location to meet SLA (E-W traffic)
- iVXLAN tunnel traffic is carried over SDWAN
- SLA policy is configured in the DSCP bits of the iVXLAN header


ACI SD-WAN - Use Case 2

ACI Multi-Site + SD-WAN Integration Use Case

1

ACI MSO - Single Pane

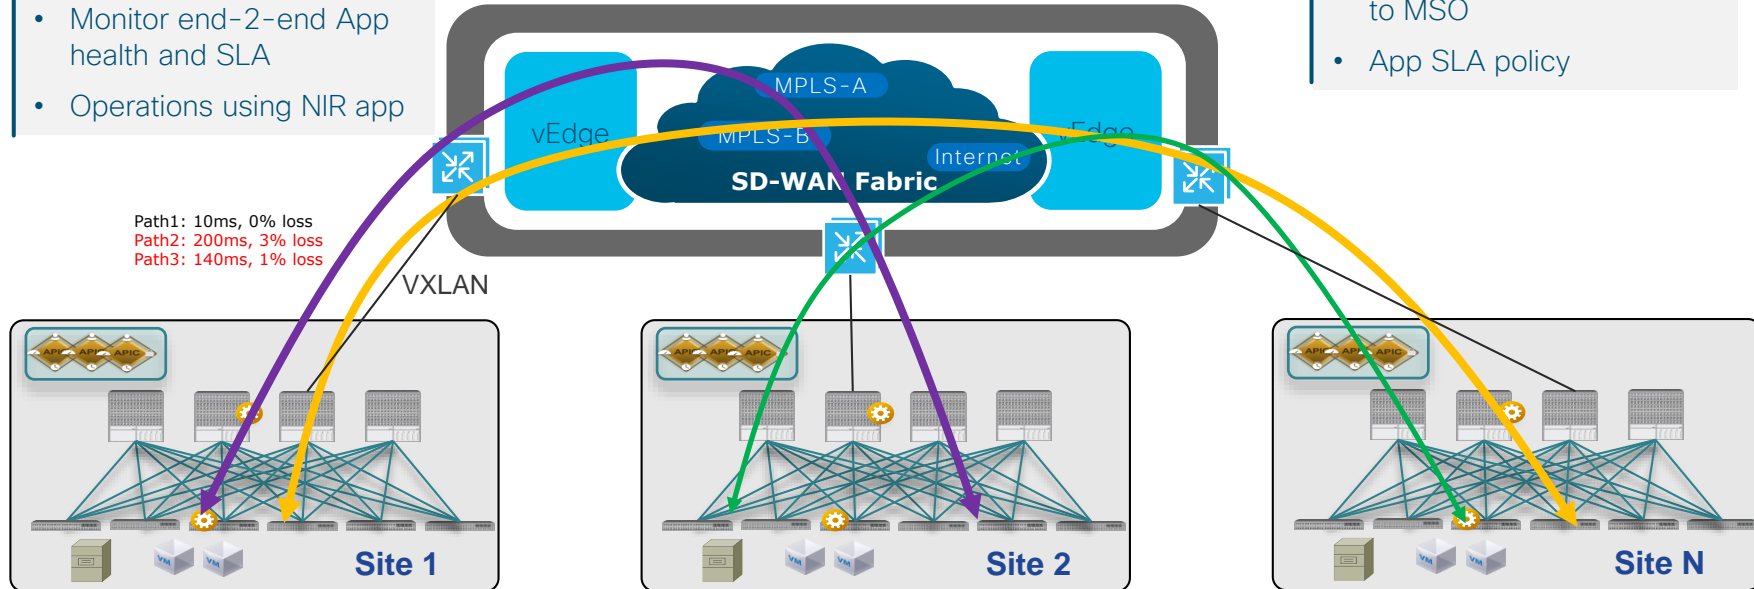
- Provision DCI overlay & Connectivity to vEdge
- Define App SLA policy
- Monitor end-2-end App health and SLA
- Operations using NIR app



2

vManage - SD-WAN

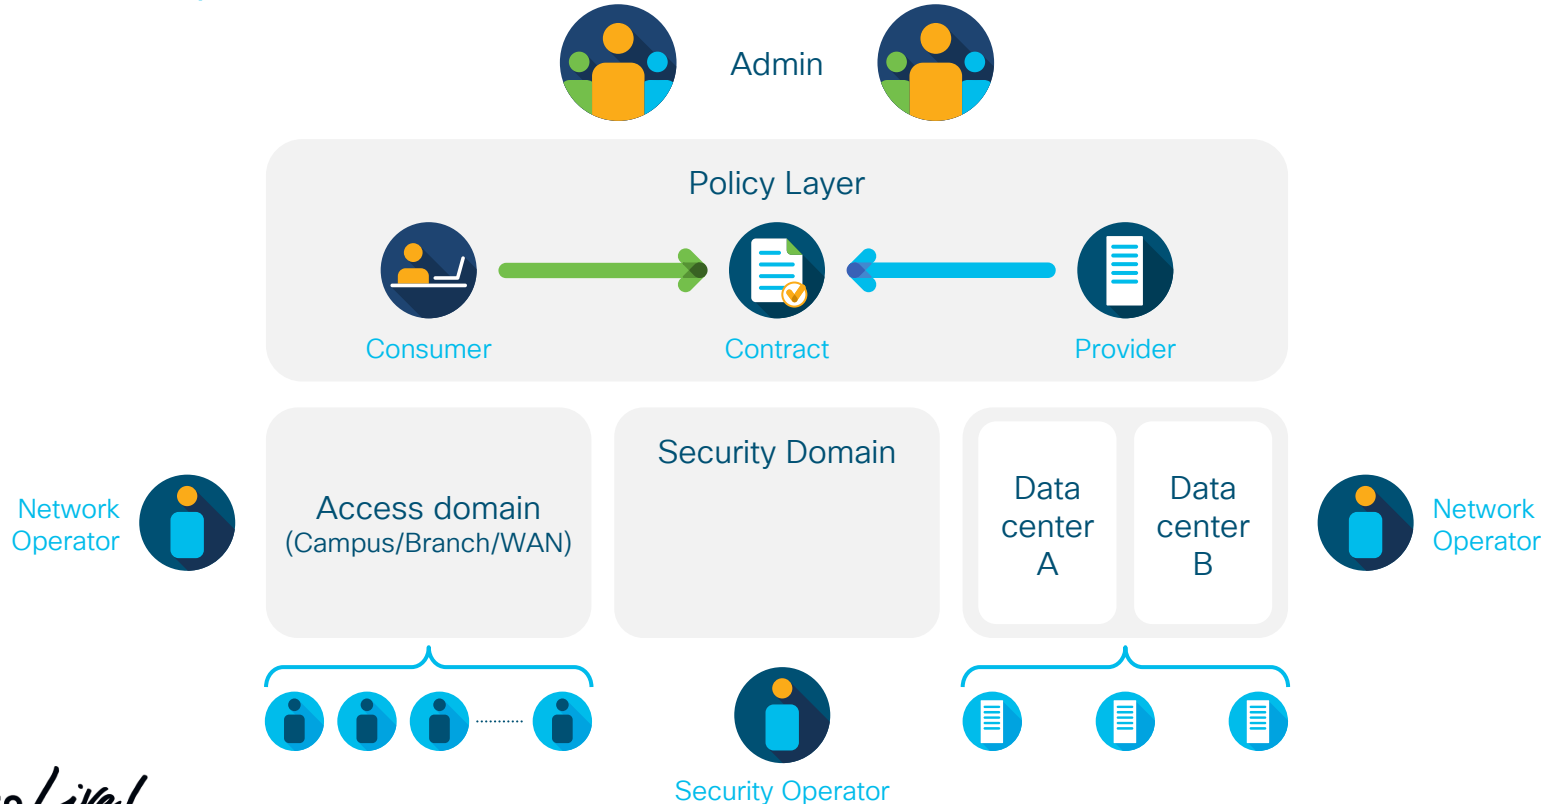
- Automated path selection using App policy
- Path Health & Telemetry to MSO
- App SLA policy



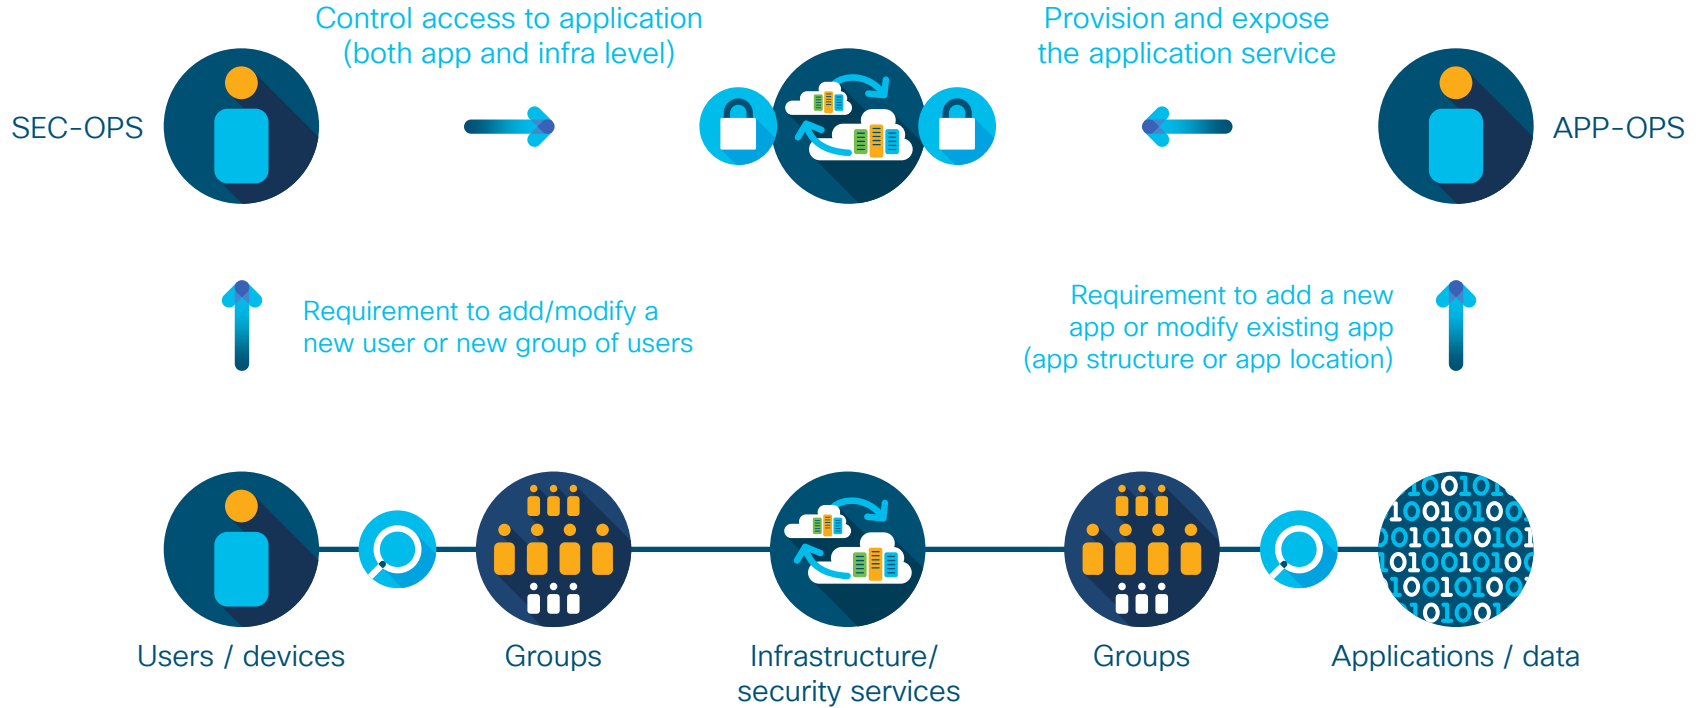
Next Steps

Unified policy language across domains

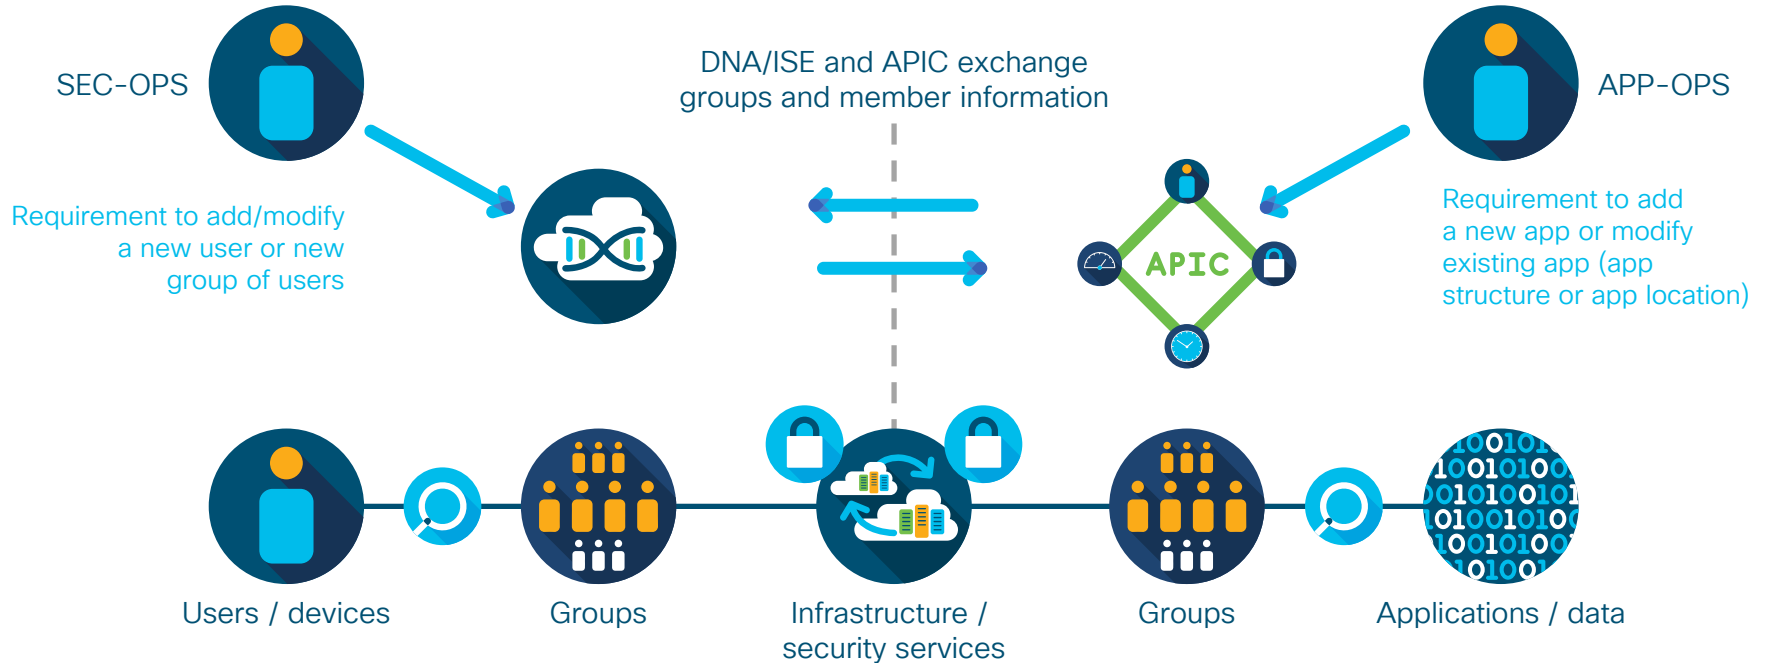
Consumer to provider



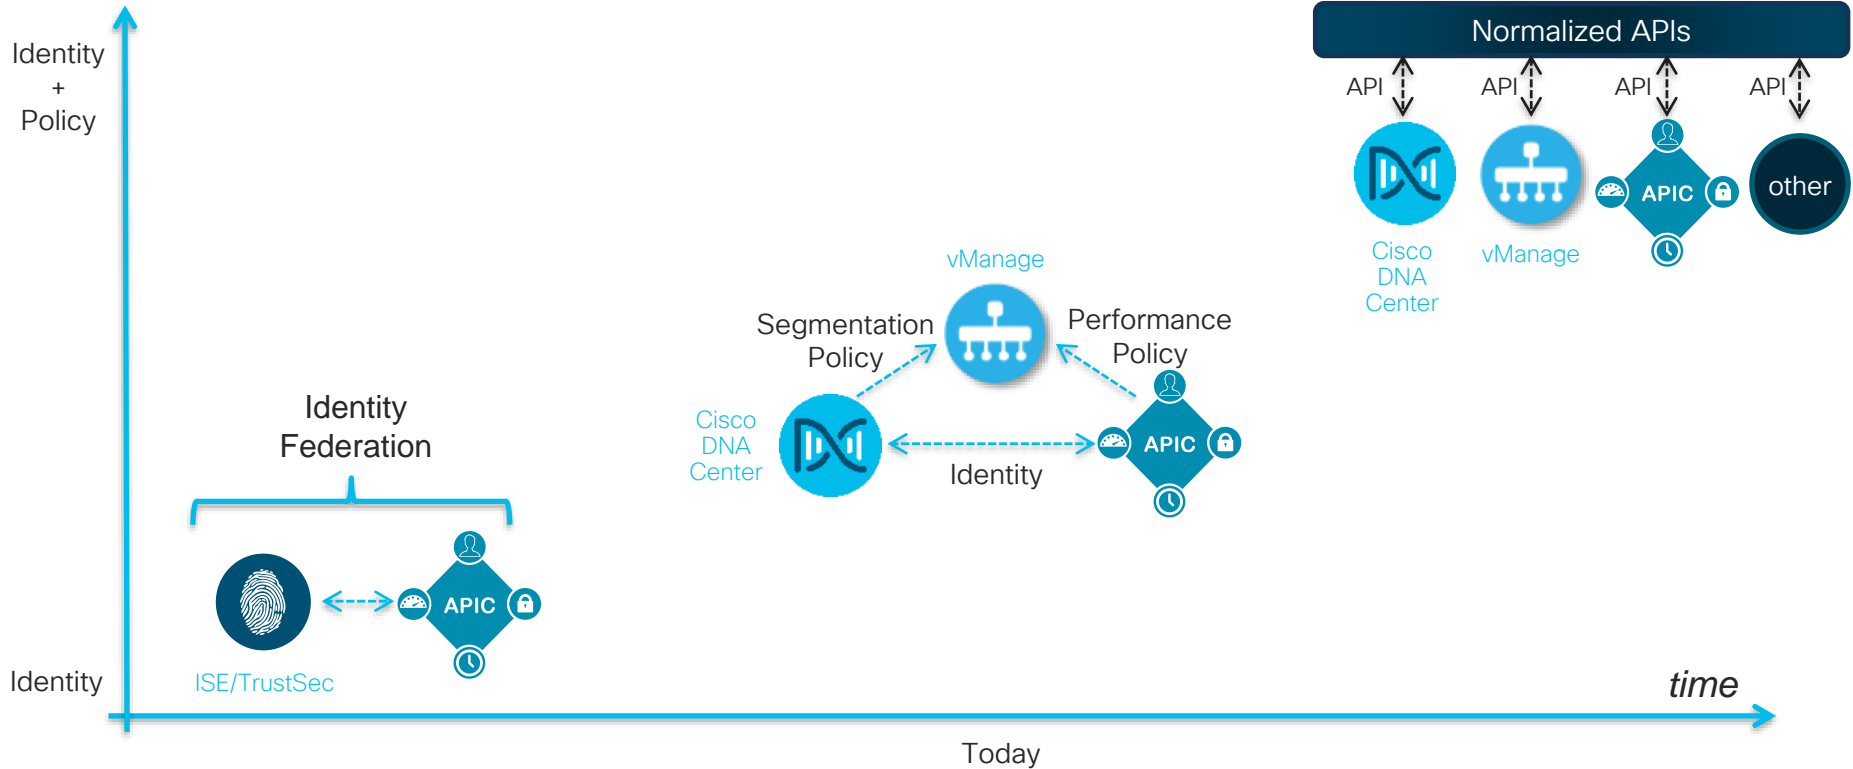
Business change view of policy



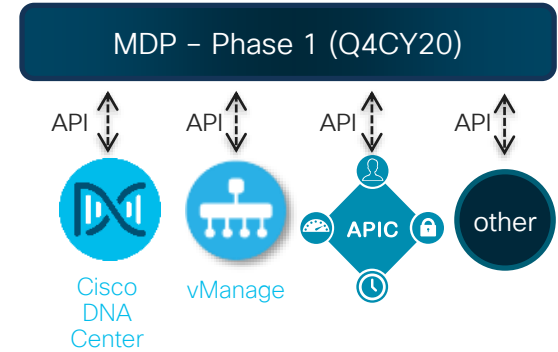
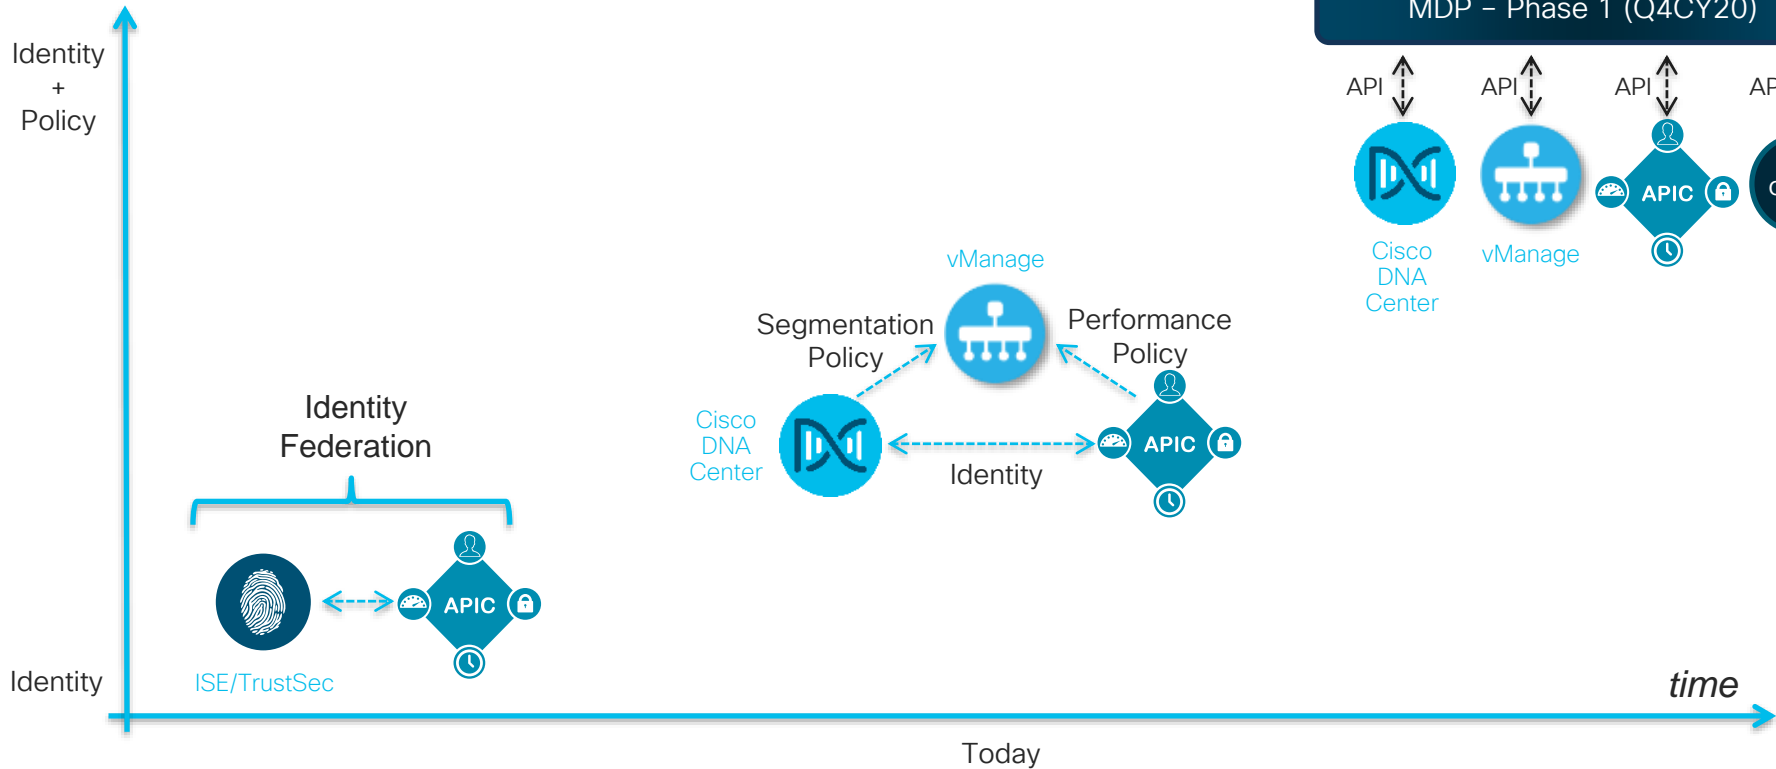
Federation of identity and policy change across domains



The Domain Federation Journey



The Domain Federation Journey



How to Operate ACI Fabric(s)

FCAPS and ITIL

FCAPS is a Network Management Framework (evolved from ISO Telecommunications Management Network)



CISCO *Live!*

ITIL v3 is a collection of (good) practices for IT as a whole. Focused on Lifecycle, Service Delivery, Support and Improvements



ITIL framework for service management and operations

IT Service Strategy - ITIL's Service Strategy provides a set of frameworks for determining what services are delivered, how their value is measured, how to measure cost and provide a measure of return on investment (ROI), and how to manage IT's relationship with its business partners.

IT Service Design - This area covers design of processes and how they relate to one another, Service-Level Agreements (SLA), capacity and availability management, business continuity management, security, and supplier management. IT Service Design also notes the need for a service catalog

IT Service Transition - Service Transition governs how services are delivered and deployed. Such areas as change management, release and deployment management, and service evaluation are typically part of the transition phase. When cloud is on scope, the actual tasks when deploying a service to the cloud change significantly. IT departments should set up a test cloud environment that mirrors the production environment in order to allow User Acceptance Testing (UAT).

IT Service Operation - Service Management covers the management and monitoring of services, and how issues are managed and resolved. Key to the Service Management component is the notion of a Service Desk

IT Continual Service Improvement - In Continual Service Improvement (CSI), IT personnel and business teams work together to ensure services can quickly meet new and emerging business requirements. CSI is heavily data-driven and relies upon operational statistics as well as business insights

ITIL framework for service management and operations



IT Service Strategy - ITIL's Service Strategy provides a set of frameworks for determining what services are delivered, how their value is measured, how to measure cost and provide a measure of return on investment (ROI), and how to manage **IT's relationship with its business partners**.


IT Service Design - This area covers design of processes and how they relate to one another, Service-Level Agreements (SLA), **capacity and availability management, business continuity management, security**, and supplier management. IT Service Design also notes the need for a service catalog

IT Service Transition - Service Transition governs how services are delivered and deployed. Such areas as **change management, release and deployment management**, and service evaluation are typically part of the transition phase. When cloud is on scope, the actual tasks when **deploying a service to the cloud** change significantly. IT departments should set up a test cloud environment that mirrors the production environment in order to allow User Acceptance Testing (UAT).

IT Service Operation - Service Management covers the **management and monitoring of services**, and how issues are managed and resolved. Key to the Service Management component is the notion of a Service Desk

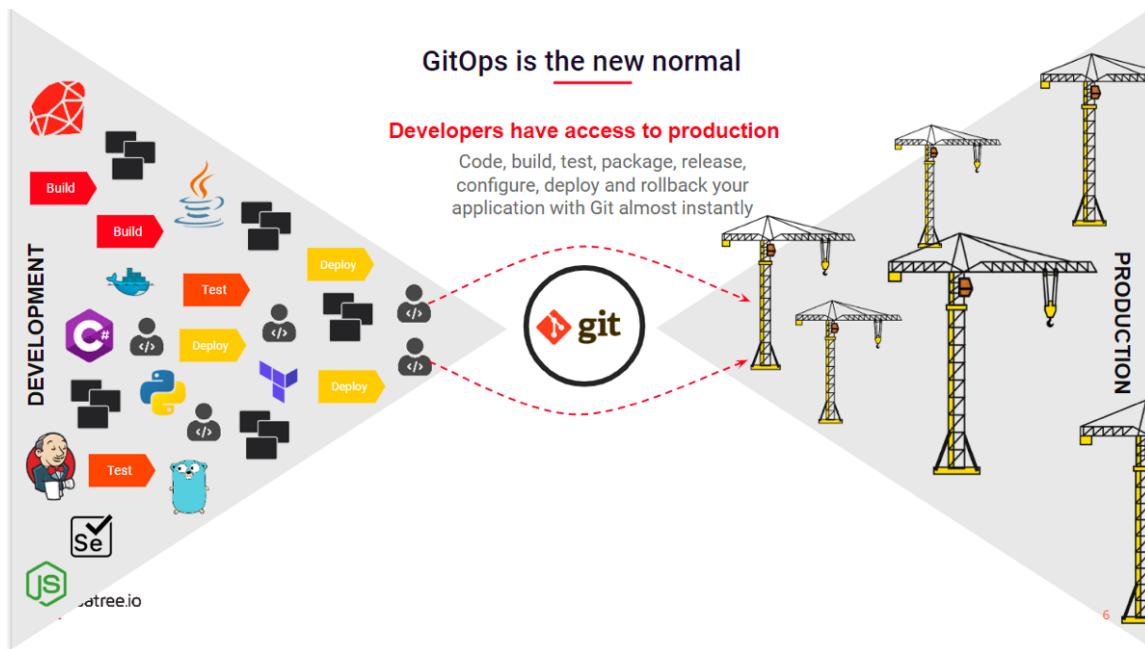
IT Continual Service Improvement - In Continual Service Improvement (CSI), IT personnel and business teams work together to ensure services can quickly meet new and emerging business requirements. CSI is heavily **data-driven and relies upon operational statistics as well as business insights**






**INTEGRATING SECURITY INTO
PRODUCT LIFECYCLE;
DEVSECOPS**





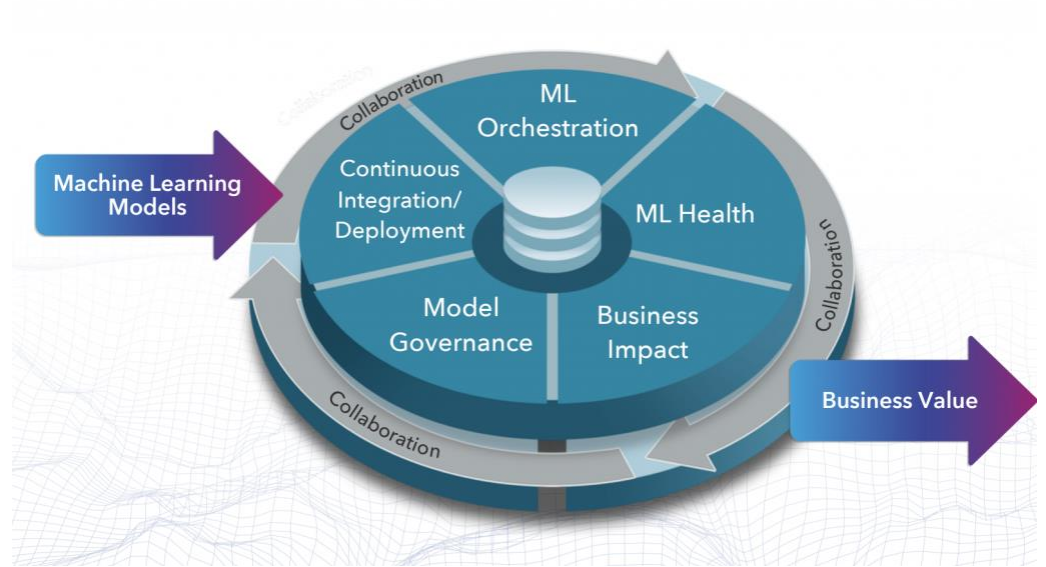


DataOps Vs DevOps Approach





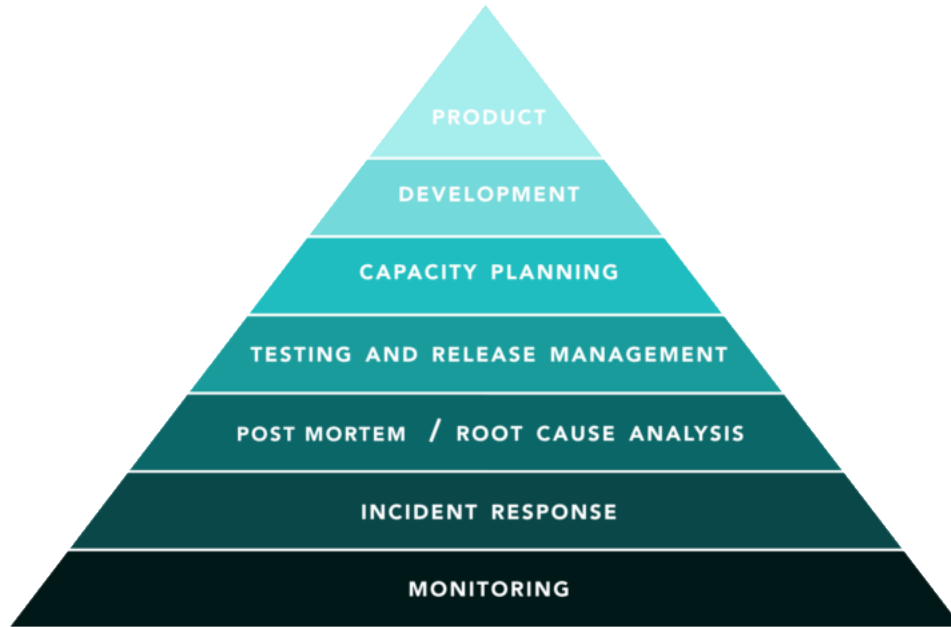
MLOps



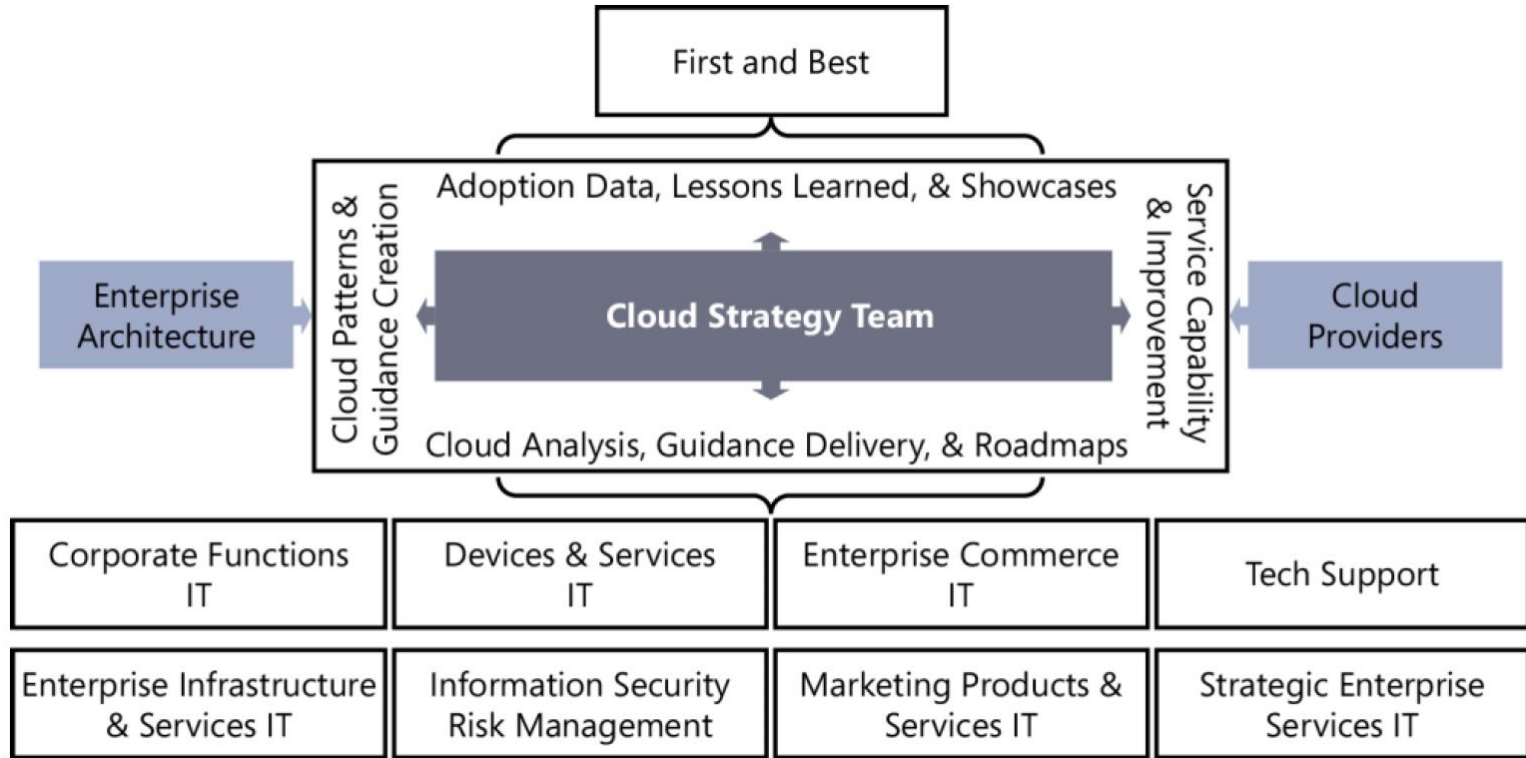
Infrastructure-as-Code (IaC) Goals

- **Unify** the view of resources
- **Support** the modern data center (IaaS, PaaS, SaaS)
- Expose a way for individuals and teams to **safely and predictably change** infrastructure
- Provide a workflow that is **technology agnostic**
- Manage anything with an **API**

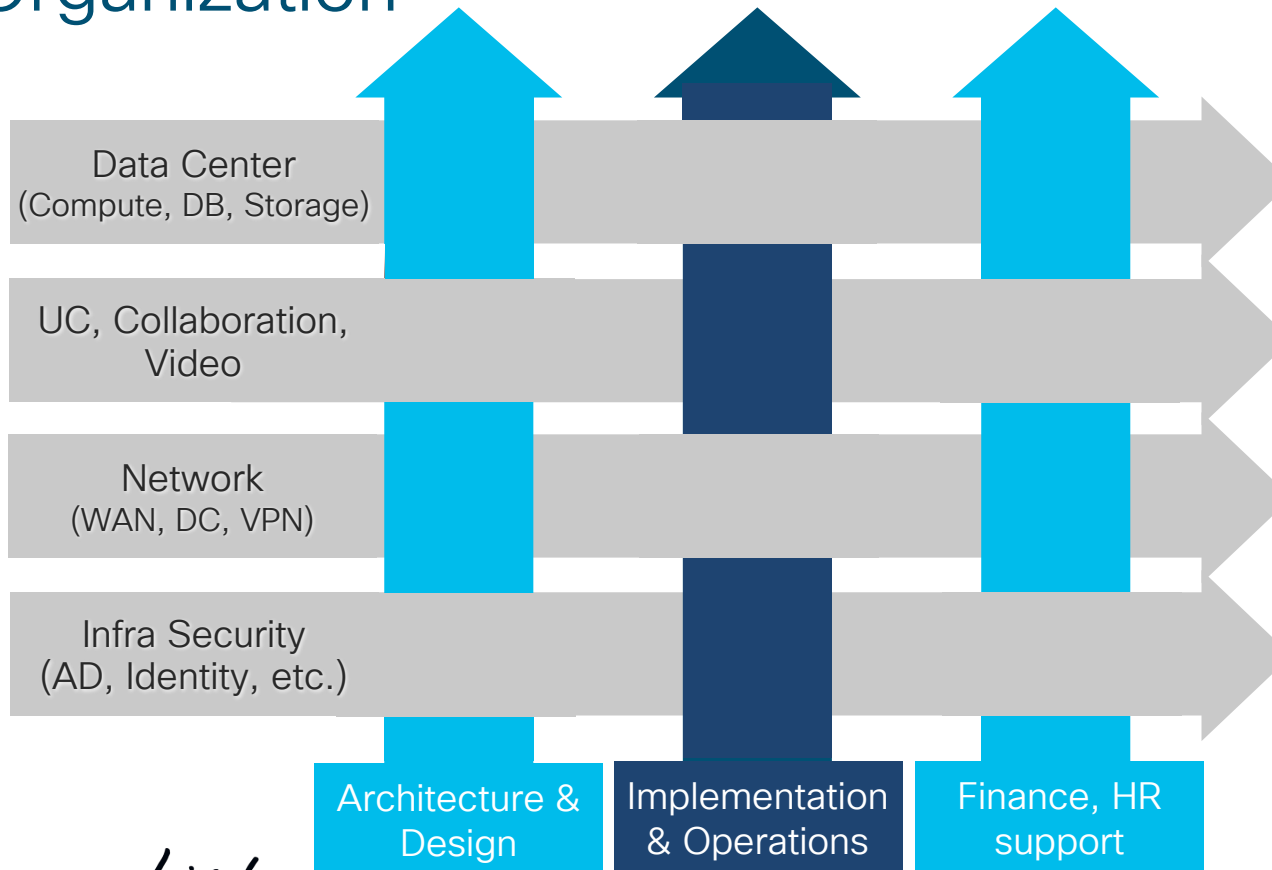
Google has defined a model called **Site Reliability Engineering (SRE)** as a collection of best practices for any digital business



Example: Cloud Strategy Team at Microsoft IT




Example: Cisco IT Service-Based Organization



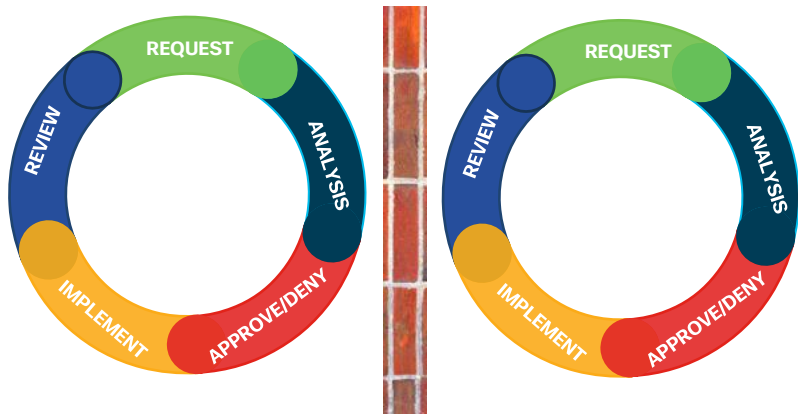
- Everything is a service
- Service owner is the GM
- Budget, roadmap, metrics, etc. are the responsibility of the service owner
- Interlaced with functions common across all services

Now What?!!



Current Culture and Mindsets within Enterprise

Traditional Mindset



Avoid failure

Change is Risky


Change is Complex

Empowered accountability

Limited Feedback systems

Manual

DevOps Mindset



Embrace failure

Change is good

Active collaboration

Empowered accountability

Feedback systems

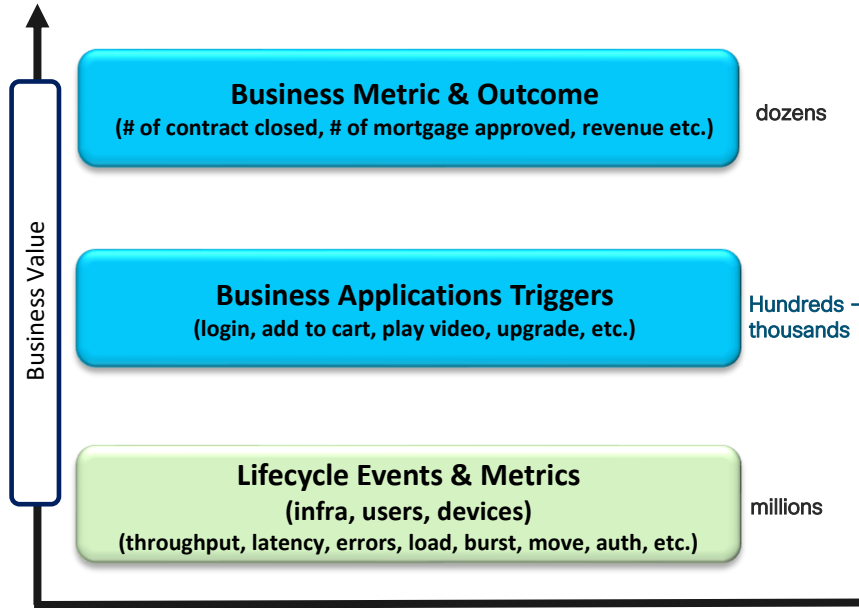
Automation

The Operations Facts

- 1) By Switching to a DevOps model of Continuous Integration and Continuous Deployment, enterprises often can realize much faster implementation of new features in their multicloud applications.
- 2) However, there are many applications for which change must be strictly controlled (e.g.: core Enterprise Resource Planning (ERP) system).
- 3) The need for agile as well as traditional frameworks (e.g.: ITIL) to control change - and the consequent risk - remains very valid.

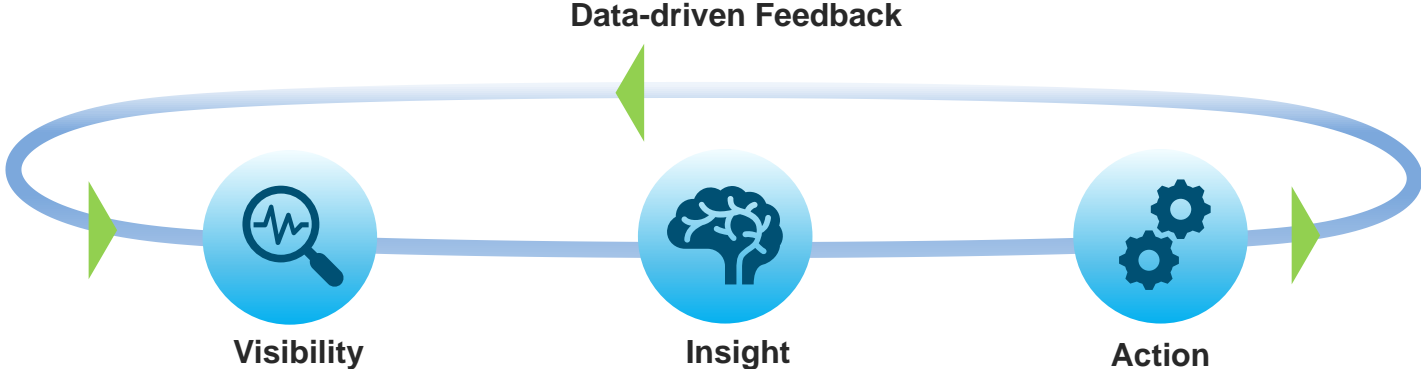
ACI Anywhere Operations address both 😊

Customers' business value chain

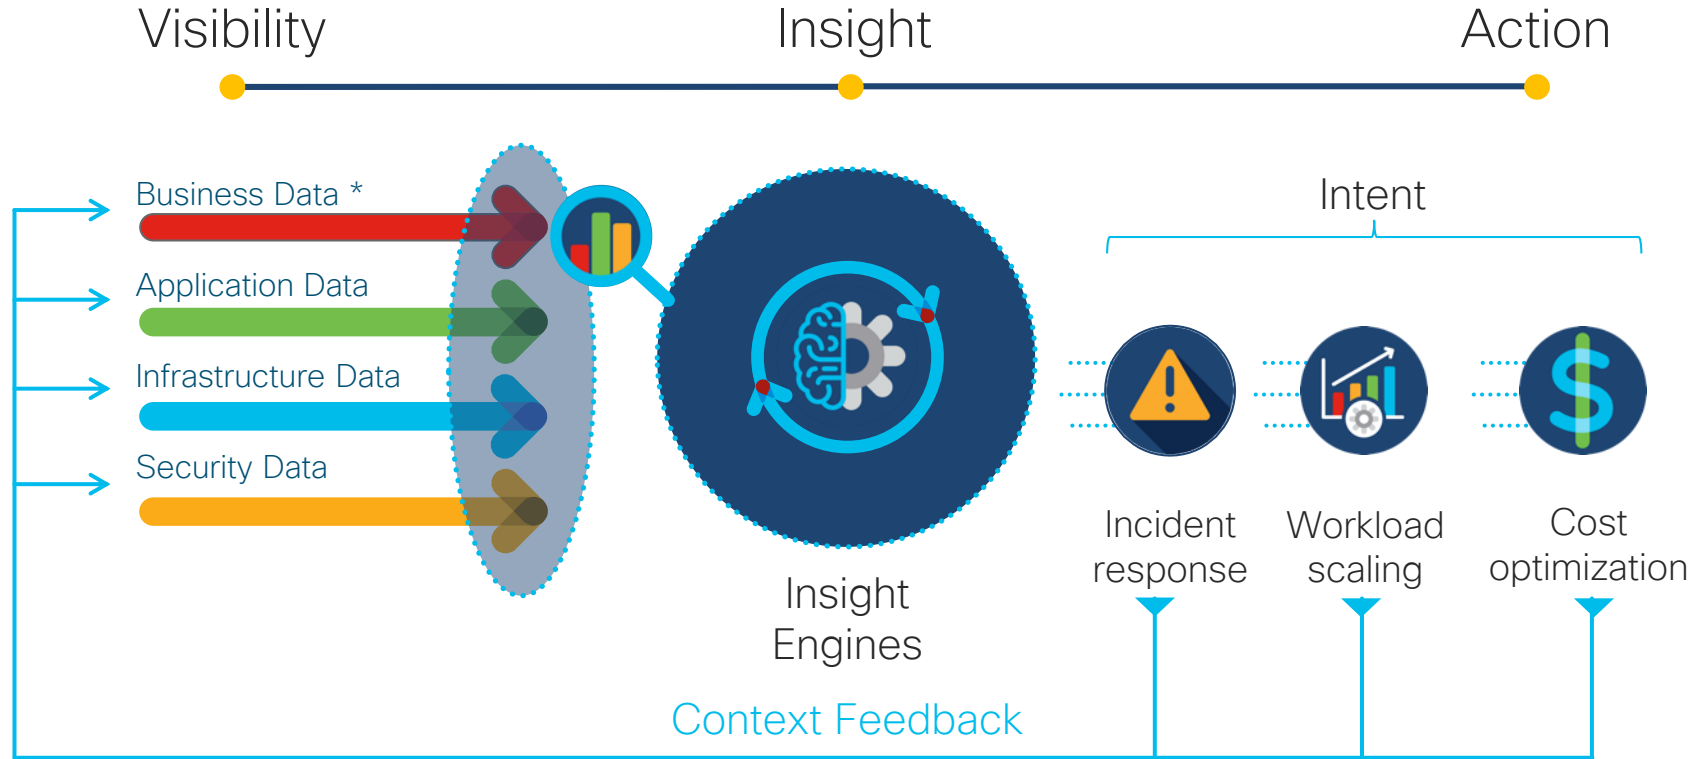


- Different parts of the organization (E.g. LOB vs IT) are focused on different metrics;
- Scale of metrics / events generated is far greater as we get closer to the infrastructure, users, devices;
- Increased evolution towards data augmented operating models (eg.: AIOps), however these must be driven by Real-Time data to be effective

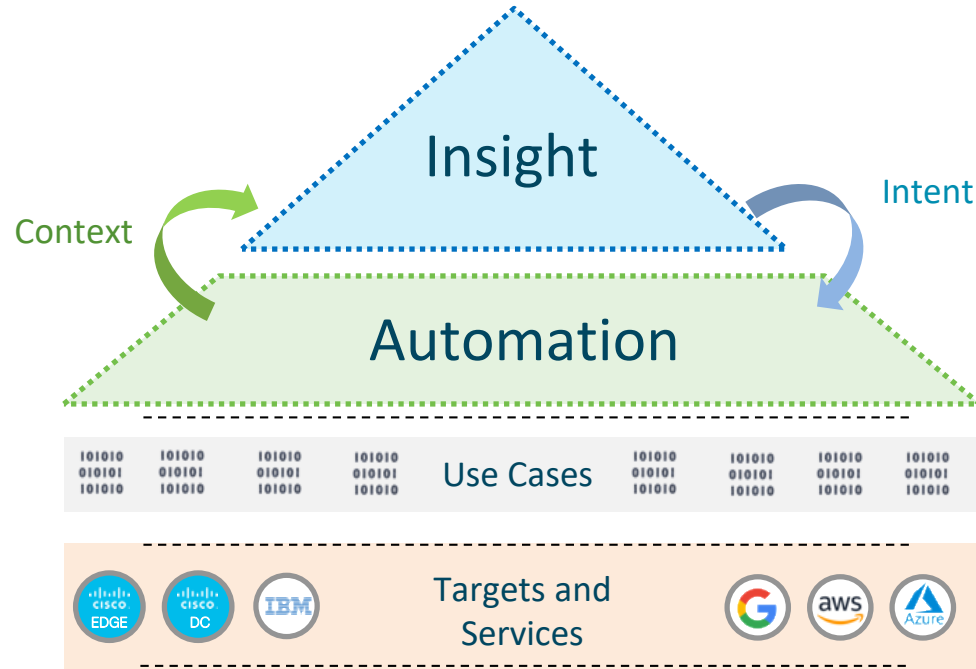
Cisco “Visibility-Insights-Actions” operations approach



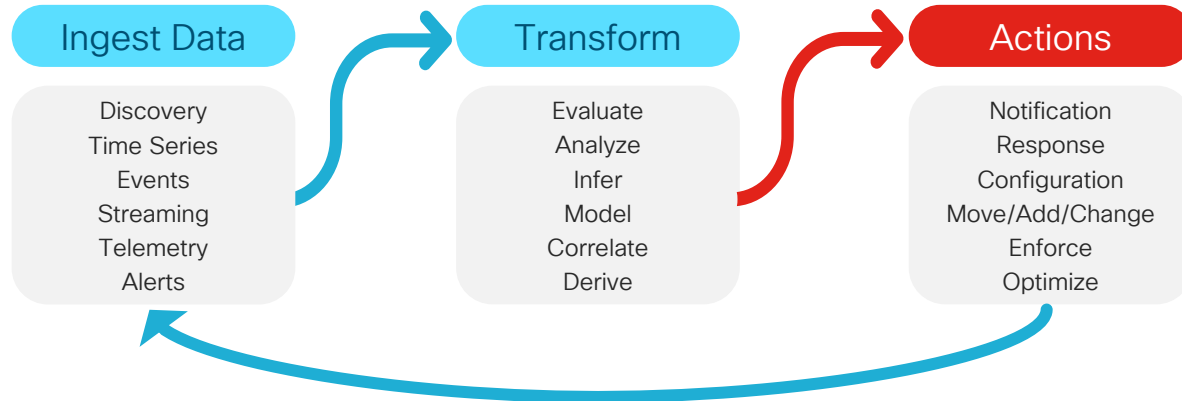
Cisco “Visibility-Insights-Actions” operations approach



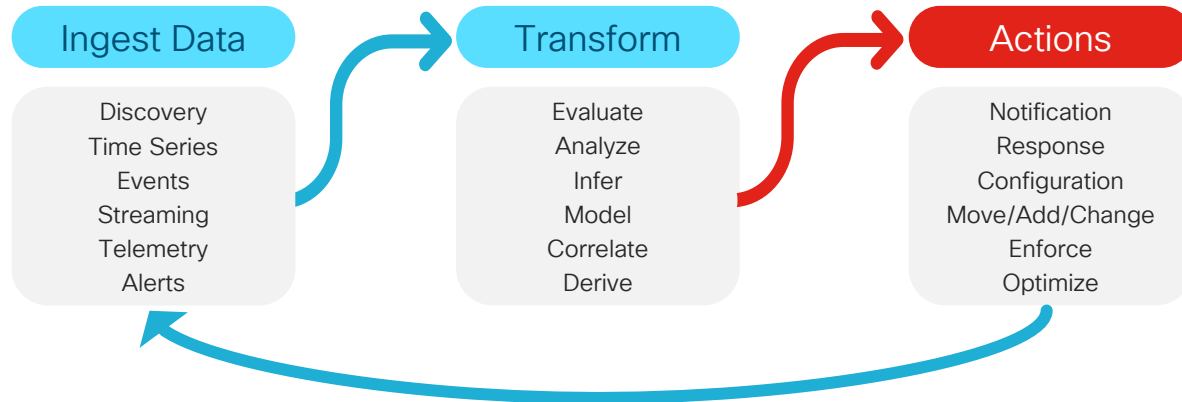
Cisco cross-architecture technology framework



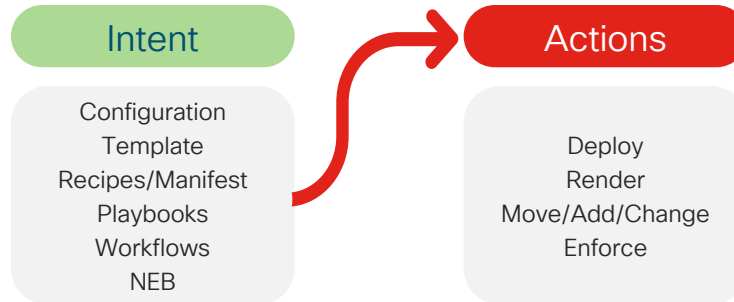
Data-driven INSIGHTS Engines



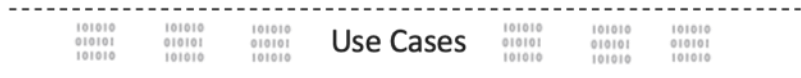
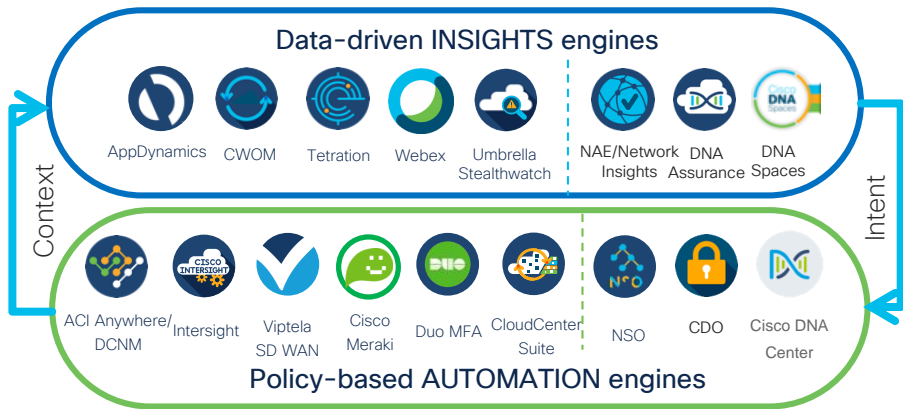

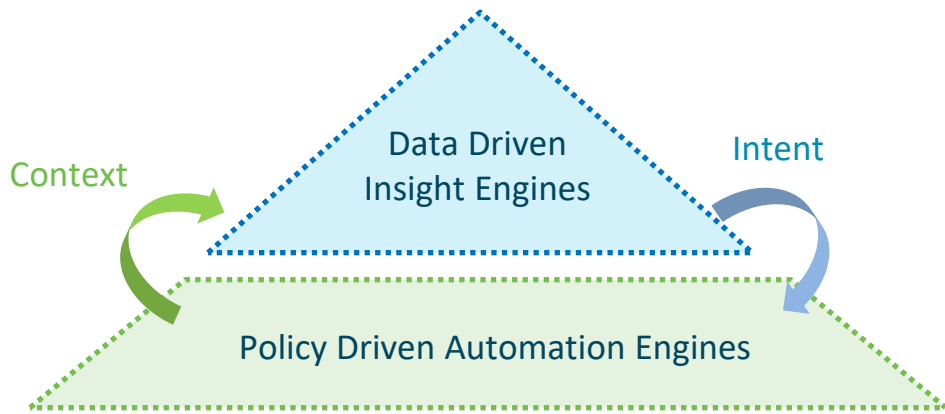
Data-driven INSIGHTS Engines



Policy-based AUTOMATION Engines

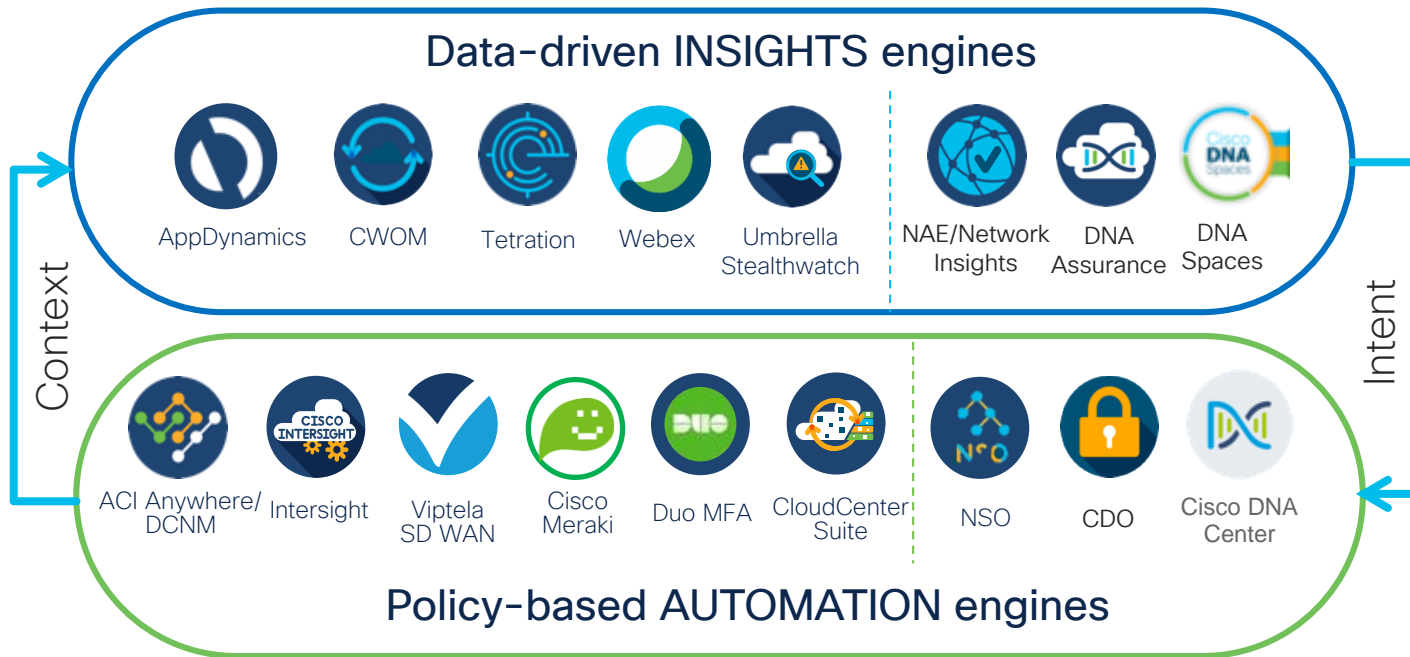


Technology framework detailed



cisco *Live!*

Technology framework detailed

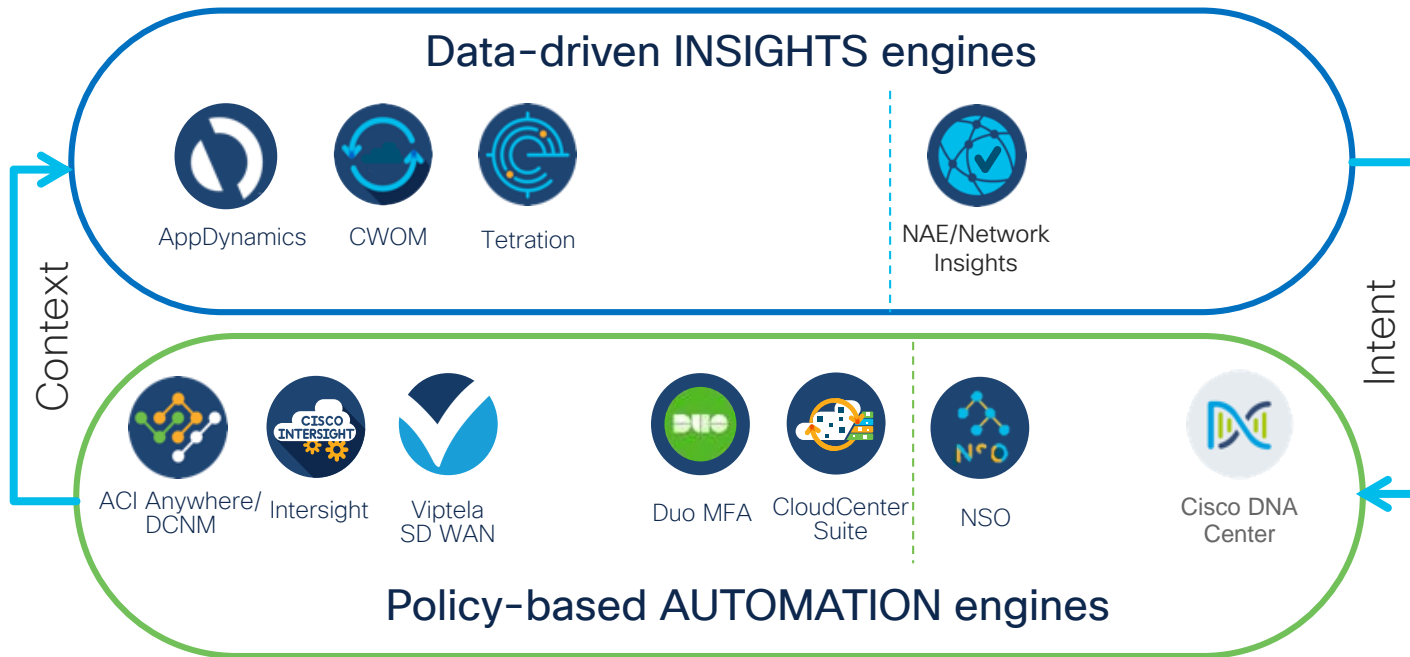


Private Cloud & Co-Lo



CISCO *Live!*

ACI integrations & technology framework

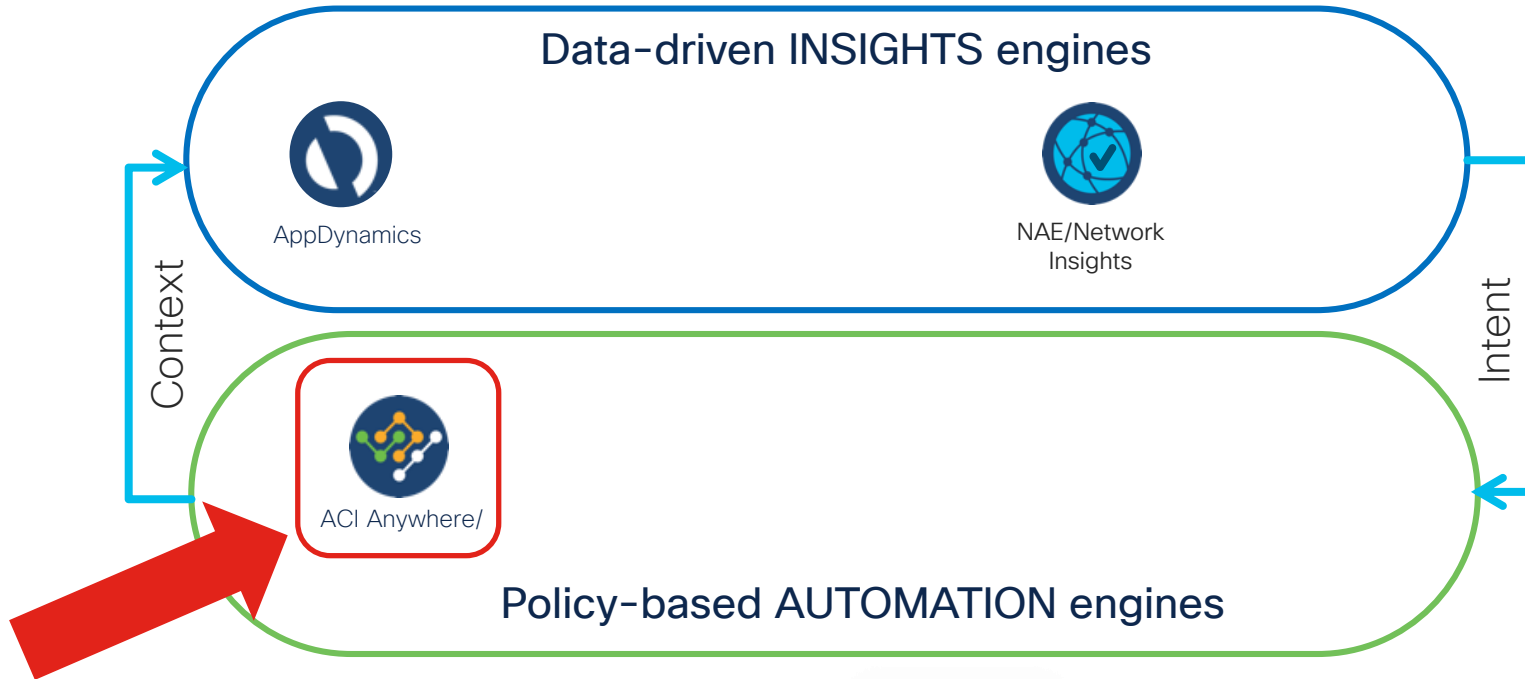


Private Cloud & Co-Lo



CISCO *Live!*

Focus of this ACI Operations section



Private Cloud & Co-Lo

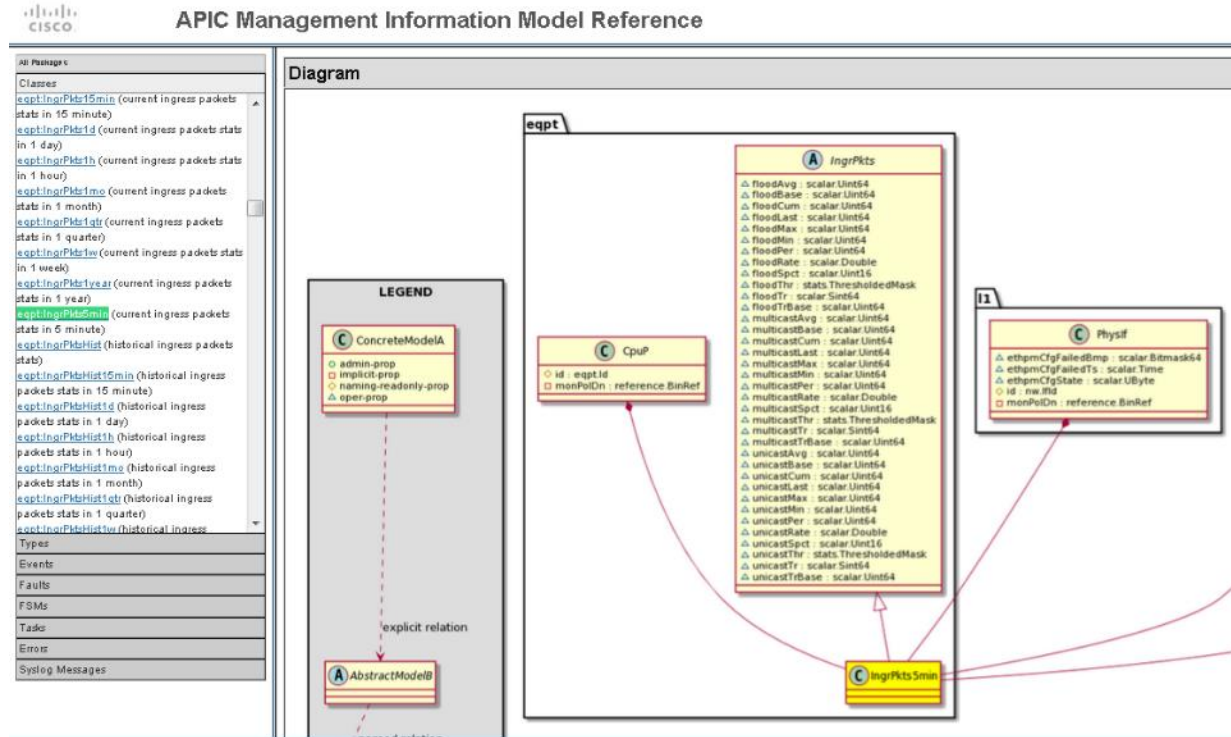
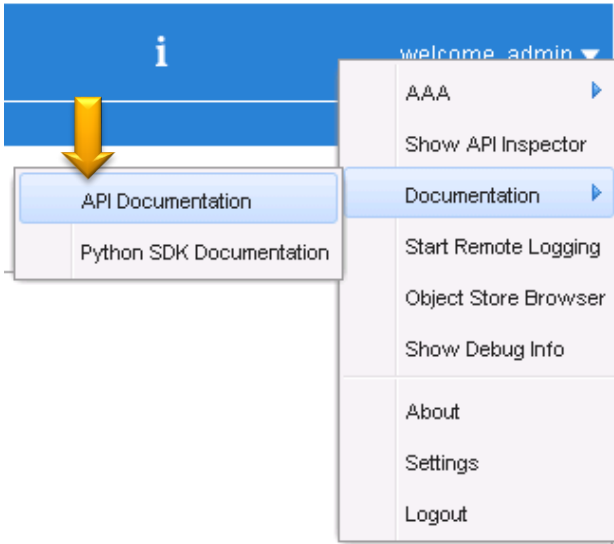


ACI Operations - Agenda

- Before getting started - setting the concepts stage

APIC Management Information Model Reference

From the WebUI



direct URL

<https://<APIC IP address>/doc/html/>



Cloud APIC Management Information Model Reference

Overview Naming Diagram Containers Contained Inheritance Stat Counters Stats Events Faults FSMs Properties Summary Properties Detail

Class cloud:EPg (CONCRETE)

<https://<APIC IP address>/doc/html>

Class ID:14644
 Class Label: cloud EPg
 Encrypted: false - Exportable: true - Persistent: true - Configurable: true - Subject to Quota: Disabled - Abstraction Layer: Logical Model - APIC NX Processing: Disabled
 Write Access: [admin, tenant-epg, tenant-network-profile]
 Read Access: [aaa, access-connectivity-11, access-connectivity-12, access-equipment, access-protocol-11, admin, fabric-equipment, fabric-protocol-12, fabric-protocol-mgmt, nw-svc-device, nw-svc-policy, tenant-connectivity-mgmt, tenant-connectivity-util, tenant-epg, tenant-network-profile, tenant-protocol-12, tenant-protocol-13, tenant-security, vmm-policy]
 Creatable/Deletable: yes (see Container Mos for details)
 Semantic Scope: EPg
 Semantic Scope Evaluation Rule: Parent
 Monitoring Policy Source: Parent
 Monitoring Flags : [!sObservable: true, HasStats: true, HasFaults: true, HasHealth: true, HasEventRules: false]

Cloud EPg

Naming Rules

```


RN FORMAT: c:cloudepg-(name)

[1] PREFIX=c:cloudepg- PROPERTY = name

DN FORMAT:

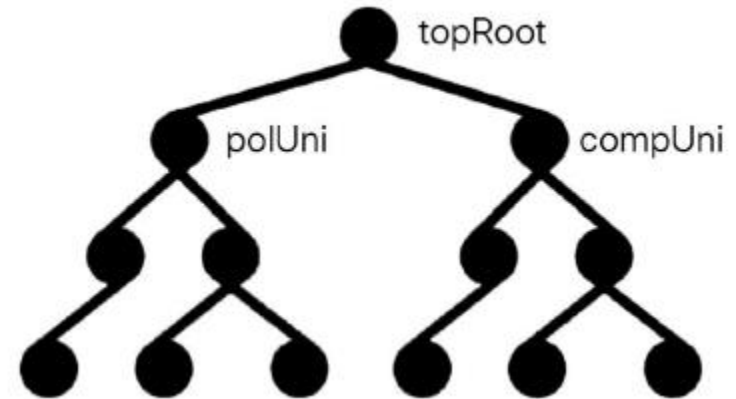
[1] un1/sn-(name)/c:cloudapp-(name)/c:cloudepg-(name)
  
```

Diagram

 Super Mo: cloud:AAEPg.
 Container Mos: cloud:App (deletable:yes).
 Contained Mos: aaa:RbacAnnotation, cloud:AEpSelector, cloud:EPSelector, cloud:RgInfoHolder, fv:CtrctCtDefCont, fv:OrchInfo, fv:RInfoHolder, fv:SharedService, fv:UpdateContract, health:NodeInst, orchs:LDevVipCfg, tag:Anst, tag:Annotation, tag:Tag, telemetry:MatchedSelector, vns:ACCfg, vns:SvcPol, vz:ConsCtrctLbl, vz:ConsLbl, vz:ConsSubjLbl, vz:ProvCtrctLbl, vz:ProvLbl, vz:ProvSubjLbl.
 Relations From: cloud:Pool, file:ARemoteHost, vns:Chassis, span:ADest, vns:ALDev, vns:DevMgr, snmp:ClientGrpP, netflow:AExporterPol, extdev:MgrP, dbgac:FromEpgCmn, infra:AFunc, vns:LifCtx, vmm:CtrlrP, datetime:NtpProv, poe:IfPol, dns:Profile, dhcp:RelayP, aaa:AProvider, span:ASrc, dnsepg:ASvrGrp, auth:Svr, vns:ATerm, dbgac:ToEpgCmn, dbgac:EpgToEpg, vns:VEpg, span:VSrcDef.
 Relations To: fv:CtX, vz:BrCP, vz:CPIf, qos:CustomPol, vz:Taboo, fv:EPg.
 Relations: cloud:RsCloudEpgCtX, cloud:RtPoolToCloudEpg, fv:RsCons, fv:RsConsIf, fv:RsCustQosPol, fv:RsIntraEpg, fv:RsProtBy, fv:RsProv, fv:RsSecInherited, fv:RtARemoteHostToEpg, fv:RtChassisEpg, fv:RtDestEpg, fv:RtDevEpg, fv:RtDevMgrEpg, fv:RtEpg, fv:RtExporterToEpg, fv:RtExtdevMgrMgmtEpg, fv:RtFromAbsEpg, fv:RtFuncToEpg, fv:RtLifCtxToInstP, fv:RtMgmtEpg, fv:RtNtpProvToEpg, fv:RtPoeEpg, fv:RtProfileToEpg, fv:RtProv, fv:RtSecInherited, fv:RtSecProvToEpg, fv:RtSrcToEpg, fv:RtSvcMgmtEpg, fv:RtSvrEpg, fv:RtSvrToMgmtEpg, fv:RtTermToEpg, fv:RtToAbsEpg, fv:RtToAbsEpgForEpgToEpg, fv:RtVConnToEpgEg, fv:RtVConnToEpgSubnet, fv:RtVSrcToEpg.

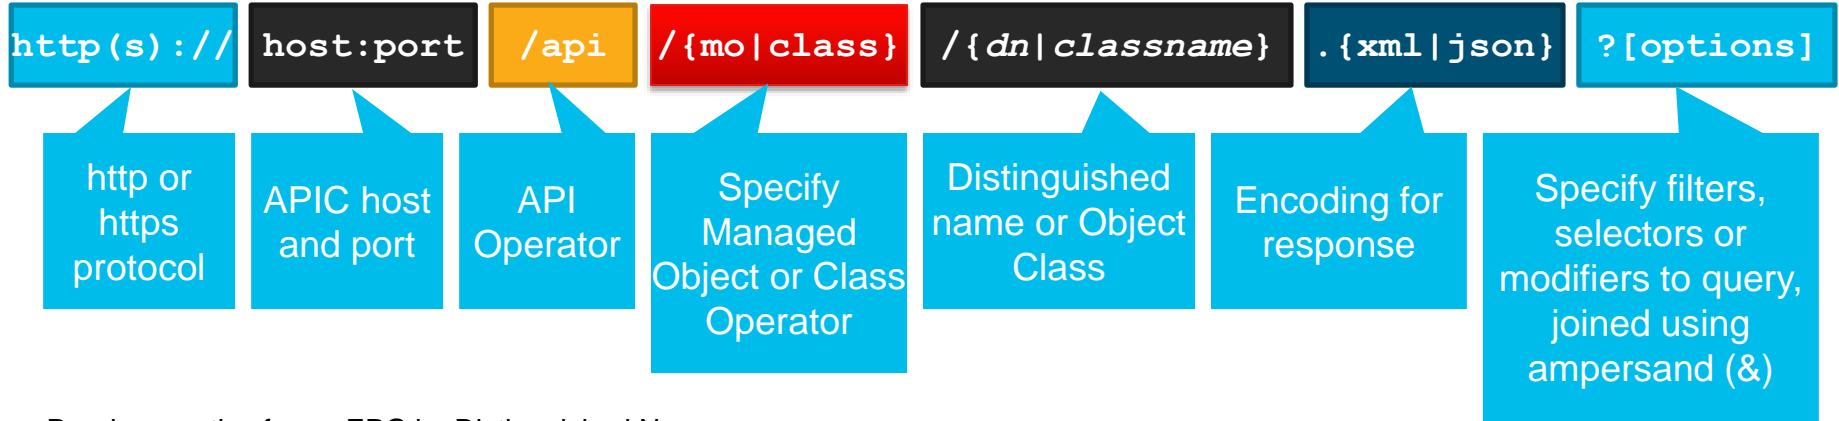
ACI Object Model Management Information Tree (MIT)

- Objects within APIC are structured in tree-based hierarchy
- Objects referred to as Managed Objects (MO)
- Every object has a parent, with exception of top:Root (top of tree)
- Relationships exist between objects



The REST API Exposes the Object Model

API Schema Follows Object Model Containment



Read properties for an EPG by Distinguished Name:

```
http://apic/api/mo/uni/tn-Cisco/ap-Software/epg-Download.xml
```

Find all 10G ports on fabric:

```
http://apic/api/class/l1PhysIf.xml?query-target-filter=eq(l1PhysIf.speed,"10G")
```

Example: Invoking the API from Python

```
import json
import requests

requests.packages.urllib3.disable_warnings()

base_url = 'https://apic-nicolas/api/'

# create credentials structure
name_pwd = {'aaaUser': {'attributes': {'name': 'admin', 'pwd': 'ins091209'}}}
json_credentials = json.dumps(name_pwd)

# log in to API
login_url = base_url + 'aaaLogin.json'
post_response = requests.post(login_url, data=json_credentials, verify=False)

# get token from login response structure
auth = json.loads(post_response.text)
login_attributes = auth['imdata'][0]['aaaLogin']['attributes']
auth_token = login_attributes['token']





# create cookie array from token
cookies = {}
cookies['APIC-Cookie'] = auth_token

# read the LED status from the chassis fan tray
sensor_url = base_url + 'mo/topology/pod-1/node-111/sys/ch/ftslot-1/ft/indled-1.json'
get_response = requests.get(sensor_url, cookies=cookies, verify=False)

# display sensor data structure
print()
print(json.dumps(get_response.json(), indent=4))
```

Object Store Browser (Visore)

- APIC has a built in object browser to navigate the object tree and inspect the state of objects.
- Point the web browser to Visore:
`http://<apic>/visore.html`
- Search for a particular object or dn (fvTenant, topSystem, topology/pod-1/node-101)

topSystem	
address	10.0.82.223
childAction	
currentTime	2014-05-14T17:53:39.944+00:00
dn	topology/pod-1/node-101/sys < >  
fabricId	1
fabricMAC	00:22:BD:F8:19:FF
id	101
inbMgmtAddr	0.0.0.0
lcOwn	local
modTs	2014-05-14T03:54:32.773+00:00
mode	unspecified
monPolDn	uni/fabric/monfab-default < >  
name	services-leaf1
oobMgmtAddr	0.0.0.0
podId	1
role	leaf
serial	SAL1733B94B
state	in-service
status	
systemUpTime	00:14:19:03.000

Cloud APIC Visore – Web Base MO Query and Browser Tool

ACI 4.1

APIC Object Store <https://<IP address>/visore.html>

Class or DN or URL: Property: Operation: Value:

All MOs of class fvTenant 6 objects found [Show URL and response of last query](#)

fvTenant ★ 🔍

dn	< uni/tn-mgmt > 📄 ▲ ▼
annotation	
childAction	
descr	
extMngdBy	
lcOwn	local
modTs	2019-01-19T23:04:02.155+00:00
monPolDn	< uni/tn-common/monepg-default >
name	mgmt
nameAlias	
ownerKey	
ownerTag	
status	
uid	0

URL and Response of last query ✕

Response Type:

URL

Response

```
{
  "totalCount": "6",
  "imdata": [
    {
      "fvTenant": {
        "attributes": {
          "annotation": "",
          "childAction": "",
          "descr": "",
          "dn": "uni/tn-mgmt",
          "extMngdBy": "",
          "lcOwn": "local",
          "modTs": "2019-01-19T23:04:02.155+00:00",
          "monPolDn": "uni/tn-common/monepg-default",
          "name": "mgmt",
          "nameAlias": ""
        }
      }
    }
  ]
}
```



APIC Managed Object Tree Browser App

The screenshot displays the APIC (Disneyland) interface. The top navigation bar includes System, Tenants, Fabric, Virtual Networking, L4-L7 Services, Admin, Operations, and Apps. The Apps tab is active, showing a sub-menu with Apps, FaultAnalytics, APIC Managed Object Browser, and APIC Postman. The APIC Managed Object Browser is selected, displaying a tree structure of managed objects. The tree is organized into two columns, with lines connecting nodes between them. The left column contains nodes such as (pt) s-ANES, (pt) s-ANES-436-dex-pool, (pt) s-LCS-HX-App, (pt) s-LCS-HX2, (pt) s-LCS-IC_com, (pt) s-csp2100-BM_pdom, (pt) s-esi-mgrt-BM_pdom, (pt) s-0out_pdom, (pt) s-01p11_0201_pdom, (pt) s-pt(s), (pt) s-mv-are-VIM_pdom, (pt) s-A-DIE, (pt) s-A-DIE-HA, and (pt) s-A-DIE-436-dex. The right column contains nodes such as (BDEQ_ADC), (BDEQ_applic_pdom), (ap-0a-rc-436), (ap-0a-0ent), (ap-templates), (rc-A-DC-HIB), (rc-A-PP-DB), (rc-AV-mgrt-L3Out), (rc-L3Out-ADC), (rc-L3Out-DC), (rc-VIB-APP), (rc-vme-L3Out), (rc-436-dex-mgrt-L3Out), and (rc-0a-grp-L3Out). A detailed view of the brc-APP-DB [vzBrCP] object is shown on the right, displaying the following configuration:

```
brc-APP-DB [vzBrCP]
MIM link: vzBrCP
annotation:
childAction:
configIssues: filter-not-present
descr:
  dn: uni/tn-ACME/brc-APP-DB
extMngdBy:
  tOwn: local
modTs: 2018-03-09T12:35:52.240+02:00
monPolDn: uni/tn-common/monepg-default
name: APP-DB
nameAlias:
ownerKey:
```

The interface also shows a 'Show me how' button on the right side and a 'Clear' button at the bottom of the detailed view. The bottom status bar indicates the last login time as 2018-06-09T22:04 UTC+02:00 and the current system time as 2018-06-09T23:36 UTC+02:00.

Easier way to access the Object Store Browser

Right click (since ACI 3.2)

The screenshot displays the Cisco ACI GUI interface. At the top, there is a navigation bar with tabs for System, Tenants, Fabric, Virtual Networking, L4-L7 Services, Admin, Operations, and Apps. Below this is a sub-navigation bar with options like ALL TENANTS, Add Tenant, and Tenant Search. The main content area is divided into a left sidebar and a main panel. The sidebar shows a tree view for Tenant Baekerei, with 'inside1' selected. A context menu is open over 'inside1', listing various actions such as 'Create EPG Subnet', 'Add VMM Domain Association', and 'Open In Object Store Browser'. The main panel shows the 'EPG - inside1' configuration page, with a 'Summary' tab selected. The summary section displays two cards: 'Static EPG Members' with a count of 1 Total, and 'Dynamic EPG Members' with a count of 4 Total. Below these are sections for 'Domains (VM and Bare Metals)' and 'Contracts'. At the bottom of the page, there is a status bar showing the last login time and current system time.

ACI Operations - Agenda

- Before getting started - setting the concepts stage

Visibility	Insights	Actions
Faults, events, stats, health, logs, trails	Application dependency	Incident Troubleshooting

Sounds familiar ?

If someone else hit this issue and its fixed, why not notify that It could impact me!

Why is support asking for access to my setup three times for this issue?

Is my configuration consistent, at the right scale and following best practices?

Are my network security patches up to date?

How many times do I need to gather these, always rolling over logs to resolve this case?

Is the code I'm running since 2016 still recommended?



Biggest Consumer of IT Time: 43% Troubleshooting




4x

Network operators spend more time collecting data than analyzing while troubleshooting



Replication challenge

Troubleshooting an issue can be impossible if IT can't replicate the issue or see the issue as is it happening real time



Slow resolution

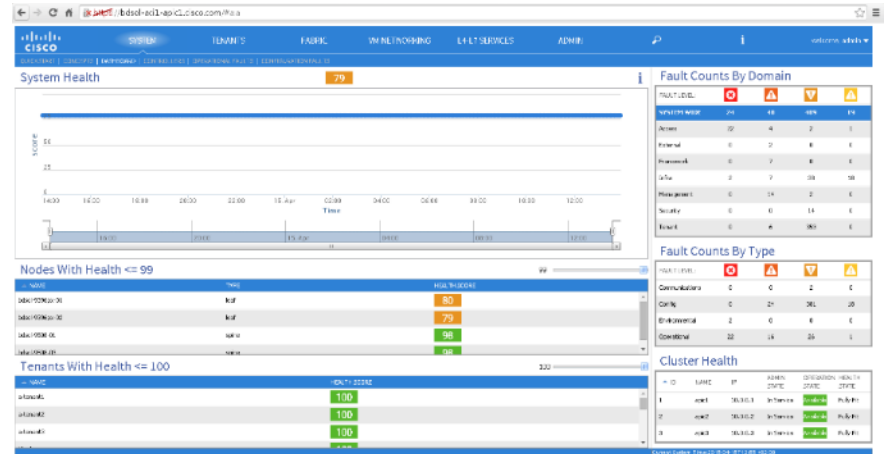
Downtime is expense; unplanned downtime cost Fortune 1000 \$1.25-2.5B annually

How do we want to troubleshoot the network?

	Switch 1	Switch 2	Switch 3	...
Hardware	✓	✓	✓	
Cabling	✓	✓	✓	
Software	✓	✓	✓	
Configuration	✓	✓	✓	
Operations	✓	✓	✓	
Switching	✓	✓	✓	
Routing	✓	✓	✓	
...	✓	✓	✓	

The ACI way:
One view for the whole Fabric!

OR



The way we're used to troubleshoot legacy ...

cisco *Live!*

End Point Tracker /Search

We can search End Point by IPv4, IPv6 or MAC address

172.16.1.0/24

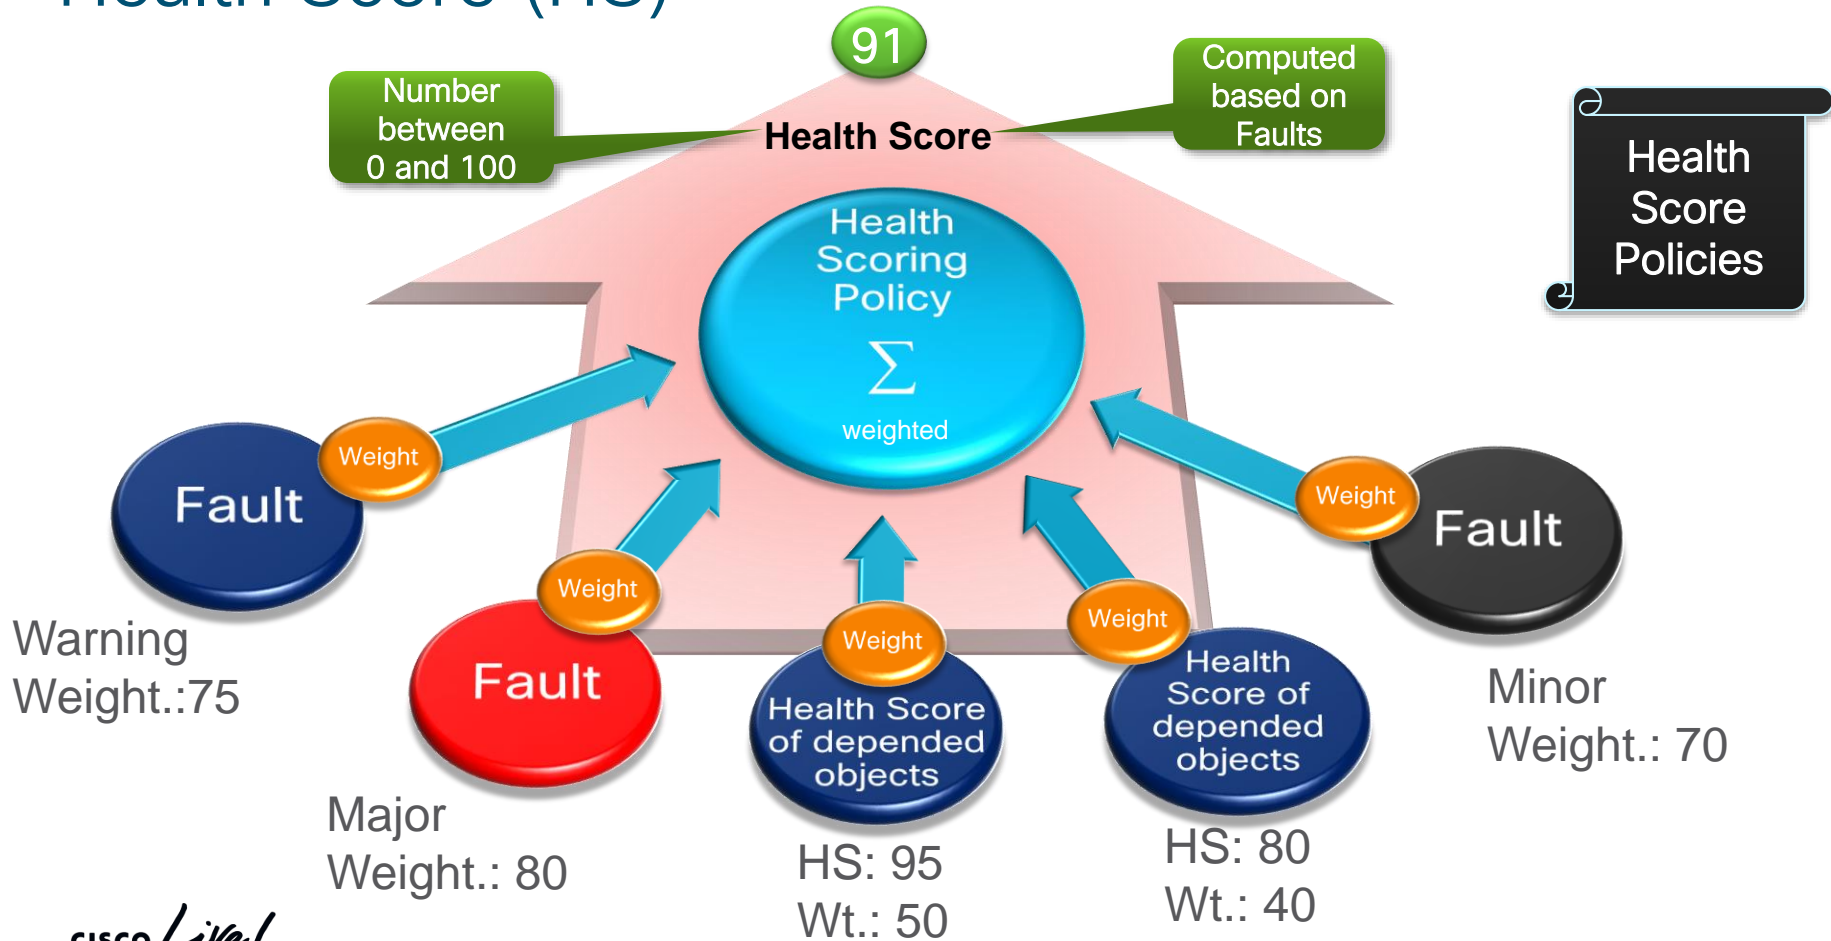
SEARCH

Learned At	Tenant	Application	EPG	IP
Leaf:101, Port:eth1/33	mio	mioAP1	mioEPG2	172.16.1.0/24

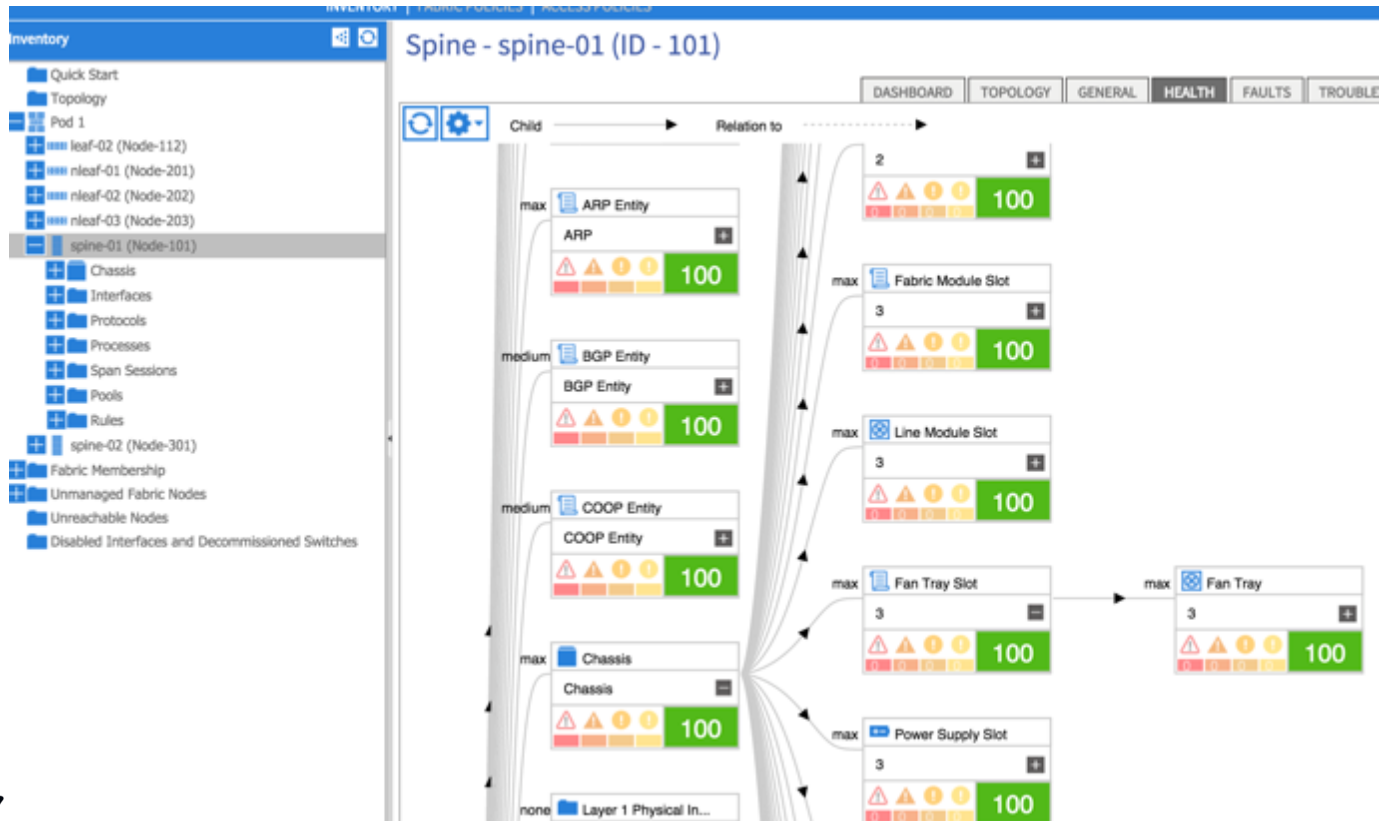
State Transitions

Date	IP	MAC	EPG	Action	Node	Interface	Encap
2015/12/31 10:42...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	detached	Node-102	eth1/33	vlan-3398
2015/11/05 15:54...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	detached	Node-101	eth1/33	vlan-3398
2015/11/05 15:53...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	attached	Node-102	eth1/33	vlan-3398
2015/10/01 16:38...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	attached	Node-101	eth1/33	vlan-3398
2015/10/01 16:38...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	detached	Node-102	eth1/33	vlan-3398
2015/10/01 16:30...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	attached	Node-102	eth1/33	vlan-3398
2015/10/01 16:06...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	detached	Node-101	eth1/33	vlan-3398
2015/10/01 15:46...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	detached	Node-102	eth1/33	vlan-3398
2015/10/01 15:40...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	attached	Node-101	eth1/33	vlan-3398
2015/10/01 15:26...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	attached	Node-102	eth1/33	vlan-3398
2015/10/01 15:25...	172.16.1.212	00:50:56:03:02:02	mio/mioAP1/mioEPG2	detached	Node-101	eth1/33	vlan-3398
2015/10/01 15:20...	0.0.0.0	00:50:56:03:02:02	mio/mioAP1/mioEPG2	attached	Node-101	eth1/33	vlan-3398


Health Score (HS)



Drill down the reason for the lowered spine Health Score



Visibility and Troubleshooting



0 Session Name: Description:

Source External IP

1

Learned At	Tenant	Application	EPG	IP
101-102, vPC: ESX01-vpc	Services			10.10.10...

Destination External IP

2

Learned At	Tenant	Application	EPG	IP
Leaf:103, Port:eth1...	DC	App	EPG2	10.201.1...

3

0 define session name

1 select end point 1

2 select end point 2

3 start

Q: Endpoints unable to communicate to each other? We're unsure where the impacted hosts and what's the datapath between them?

A: We select End Points we'd like to troubleshoot visually. The rest is done by Visibility and Troubleshooting tool

Recently Visited

- EComm-NPM-Demo - Dashboard
Applications > EComm-NPM-Demo - Dashboard
- EComm-NPM-Demo - Network Dashboard
EComm-NPM-Demo - Network Dashboard
- EComm-NPM-Demo
Applications > EComm-NPM-Demo
- Order-Tier - Network Dashboard
EComm-NPM-Demo > Tiers & Nodes > Order-Tier - Network Dashboard
- Order-Tier - Dashboard
Applications > EComm-NPM-Demo > Tiers & Nodes > Order-Tier - Das...
- EComm-NPM-Demo
Applications > EComm-NPM-Demo
- EComm-NPM-Demo
Applications > EComm-NPM-Demo
- Baselines

Applications 7

0 critical, 0 warning, 7 normal

- EComm-NPM-Demo
- NPMApplication
- NPMApplication-1
- NPMApplication-2
- NPMApplication-3
- NPMApplication-4
- NPMApplication-5

User Experience

Browser Apps 0

No Browser Apps

Get Started

Mobile Apps 0

No Mobile Apps

Get Started

Databases 0

No Databases

Get Started

Servers 0

No Servers

Get Started

Analytics

- 0 Transactions
- 0 Logs
- 0 Browser Requests

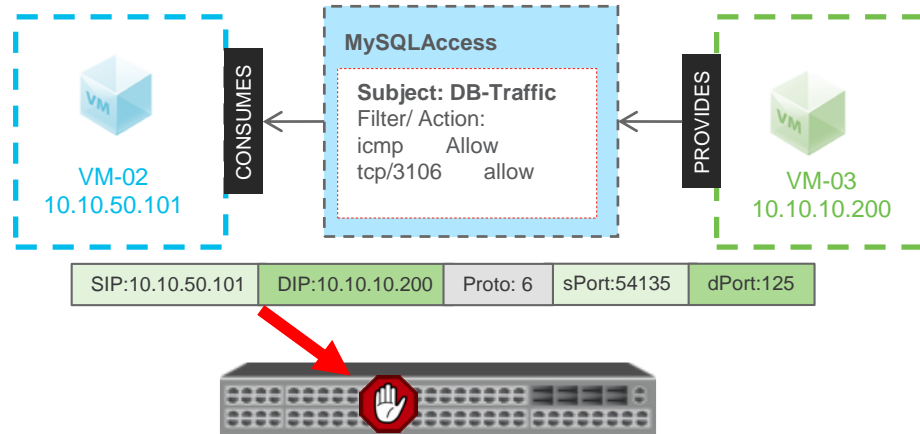
ACI Contract Logging – Denied Packets

Logging Deny

- ACI can **log implicit deny hits**
 - For Bare Metal, VMware VDS and MSFT Domains logs generated by Leaf
 - For AVS/AVE logs may be generated on Leaf or vLeaf
 - For OpenStack ML2 mode, logs configured external to the fabric at the host
- Syslog is exported according to monitoring policies and configured External Data Collectors
- Logs **include Tenant/VRF, EPG VLAN encap, ingress interfaces** and offending packet details
- **Software Dependency:** supported on all software releases
- **Hardware Dependency:** supported on all hardware models

ACL deny not logged by default:

Fabric -> Fabric Policies -> Monitoring Policies -> Common Policy -> Syslog Message Policies -> Policy for system syslog messages -> Change 'default' to 'info'

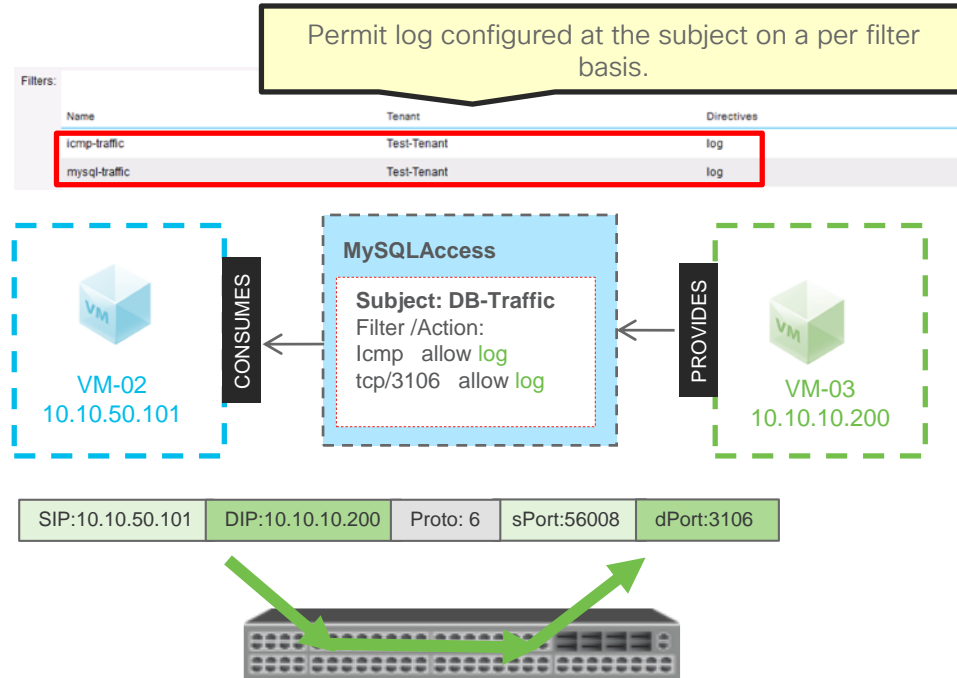


```
Feb 04 10:26:54 troy-leaf1 %LOG_LOCAL7-6-SYSTEM_MSG [E4204936][transition][info][sys] %ACLOG-5-ACLOG_PKTLOG_DENY: CName: Test-Tenant:Test-Tenant-VRF(VXLAN: 2162689), VlanType: FD_VLAN, Vlan-Id: 21, SMac: 0x00505690b43a, DMac: 0x0022bdf819ff, SIP: 10.10.50.101, DIP: 10.10.10.200, SPort: 54135, DPort: 125, Src Intf: port-channel2, Proto: 6, PktLen: 74
```

ACI Contract Logging – Permitted Packets

Logging Permit

- Permit logging is configured per Filter
 - For Bare Metal, VDS and MSFT Domains logs generated by Leaf
 - For AVS/AVE logs may be generated on Leaf or vLeaf
 - For OpenStack ML2 mode, logs configured external to the fabric at the host
- Syslog is exported according to monitoring policies and configured External Data Collectors
- Logs include Tenant/VRF, EPG VLAN encap, ingress interfaces and offending packet details
- **Software Dependency: 2.2(1n) or higher**
- **Hardware Dependency: requires EX models or newer**



Feb 04 10:14:44 troy-leaf1 %LOG_LOCAL7-6-SYSTEM_MSG [E4204936][transition][info][sys] %ACLOG-5-ACLOG_PKTLOG_PERMIT: CName: Test-Tenant:Test-Tenant-VRF(VXLAN: 2162689), VlanType: FD_VLAN, Vlan-Id: 21, SMac: 0x00505690b43a, DMac: 0x0022bdf819ff, SIP: 10.10.50.101, DIP: 10.10.10.200, SPort: 56008, DPort: 3106, Src Intf: port-channel2, Proto: 6, PktLen: 98

Audit Logs

- Any configuration changes on the system is recorded as audit logs.
- User action which could be create, modify and delete
- Objects affected by the policy that got implemented

TIME STAMP	USER	ACTION	AFFECTED OBJECT	DESCRIPTION
2015-03-04T01:03:50.288+00:00	admin	deletion	uni/tn-Customer/acEpToEp-yong_67_dst_src	EpToEp yong_67_dst_src deleted
2015-03-04T01:03:50.288+00:00	admin	deletion	uni/tn-Customer/acEpToEp-yong_67_dst_src/rstoEpForEpToEp-[uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11]	RstoEpForEpToEp uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11 deleted
2015-03-04T01:03:50.288+00:00	admin	deletion	uni/tn-Customer/acEpToEp-yong_67_dst_src/rsfromEp-[uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F]	RsFromEp uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F deleted
2015-03-04T01:03:50.228+00:00	admin	deletion	uni/tn-Customer/acEpToEp-yong_67_src_dst	EpToEp yong_67_src_dst deleted
2015-03-04T01:03:50.228+00:00	admin	deletion	uni/tn-Customer/acEpToEp-yong_67_src_dst/rstoEpForEpToEp-[uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F]	RstoEpForEpToEp uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F deleted
2015-03-04T01:03:50.228+00:00	admin	deletion	uni/tn-Customer/acEpToEp-yong_67_src_dst/rsfromEp-[uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11]	RsFromEp uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11 deleted
2015-03-04T01:01:27.905+00:00	admin	creation	uni/tn-Customer/acEpToEp-yong_67_dst_src	EpToEp yong_67_dst_src created
2015-03-04T01:01:27.905+00:00	admin	creation	uni/tn-Customer/acEpToEp-yong_67_dst_src/rstoEpForEpToEp-[uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11]	RstoEpForEpToEp uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11 created
2015-03-04T01:01:27.905+00:00	admin	creation	uni/tn-Customer/acEpToEp-yong_67_dst_src/rsfromEp-[uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F]	RsFromEp uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F created
2015-03-04T01:01:27.834+00:00	admin	creation	uni/tn-Customer/acEpToEp-yong_67_src_dst/rstoEpForEpToEp-[uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F]	RstoEpForEpToEp uni/tn-Customer/ap-Apps/epg-AppServerPool-1/cep-00:00:20:03:9D:4F created
2015-03-04T01:01:27.834+00:00	admin	creation	uni/tn-Customer/acEpToEp-yong_67_src_dst/rsfromEp-[uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11]	RsFromEp uni/tn-Customer/ap-DB/epg-SQLServerPool-1/cep-00:00:1F:FB:01:11 created
2015-03-04T01:01:27.833+00:00	admin	creation	uni/tn-Customer/acEpToEp-yong_67_src_dst	EpToEp yong_67_src_dst created
2015-03-04T01:01:20.705+00:00	admin	deletion	uni/tn-Customer/trEp-yong_67_dst_src	TrEp yong_67_dst_src deleted

Faults and Audit Logs/Events Association

Since APIC 3.2, the UI introduces an enhancement for easy association of faults and audit logs/ events. For a given fault, the APIC UI can display the audit logs and events right before the fault is raised.

To use the association for fault investigation, open up the specific fault, click on “Troubleshooting”, then view the audit logs and event logs in the pop-up window.

User can choose the time range of the correlation, such as 1, 5, 10, 15, 20 or 60 minutes prior to the fault.

Events/Audit Log Fault Association

Fault Properties

General Troubleshooting History

Fault Code: F806262
Severity: major
Last Transition: 2018-05-17T23:56:30.934+00:00
Lifecycle: Raised
Affected Object: comp/prov-VMware/ctrl-[ACI-Network]-vcenter35
Description: Fault delegate: [FSM:FAILED] Add-FSM for VM Controller: vcenter35 VM Domain: ACI-Network VM Provider: VMware Error: Cannot complete login due to an incorrect user name or password.[FSM:rc:vmmgr:CompCtrlAdd]
Type: Operational
Cause: fsm-failed
Change Set:
Created: 2018-05-17T23:56:30.934+00:00
Code: F806262
Number of Occurrences: 1
Original Severity: major
Previous Severity: major
Highest Severity: major

A fault raised: APIC failed to login to vCenter

Fault Properties

General Troubleshooting History

Audit Logs Events

Audit log 1 minutes before the fault

Time	ID	User	Action	Affected Object	Description	
2018-05-17 10:10	33.167+00:00:00	4294975329	admin	modification	uni/vmmp-VMware/dom-ACI-Network/ctrl-vcenter35/rsacc	RsAcc modified
2018-05-17 15:20	16.508+00:00:00	4294975328	admin	creation	uni/vmmp-VMware/dom-ACI-Network/usrsacc-vcenter138	UsrAccP vcenter138 created

Audit logs for the minute before the fault show that user Admin modified the VMM login credential.

Fault Properties

General Troubleshooting History

Audit Logs Events

Event log 1 minutes before the fault

Severity	Affected Object	Code	Cause	Creation Time	Description
10	comp/prov-VMware/ctrl-[ACI-Network]-vcenter35	E4204949	transition	2018-05-17T23:56:30.927+00:00	CreateListView failed: SOAP Error code: (null)/VMM Controller: non-retrievable API error
15	comp/prov-VMware/ctrl-[ACI-Network]-vcenter35	E4204949	transition	2018-05-17T23:56:30.925+00:00	Login - NOT OK: SOAP Error code: ""/ServerFaultCode: Controller: non-retrievable API error
20	comp/prov-VMware/ctrl-[ACI-Network]-vcenter35	E4204949	transition	2018-05-17T23:56:30.909+00:00	Login - NOT OK: SOAP Error code: ""/ServerFaultCode: Controller: non-retrievable API error
60	comp/prov-VMware/ctrl-[ACI-Network]-vcenter35	E4204949	transition	2018-05-17T23:56:25.836+00:00	CreateListView failed: SOAP Error code: (null)/VMM Controller: non-retrievable API error

Events during the minute before the fault

Cloud ACI - Event Analytics

ACI 4.1

The screenshot displays the Cloud APIC Event Analytics page. The left sidebar contains navigation options: Dashboard, Topology, Application Management, Cloud Resources, Operations, Active Sessions, Event Analytics, Schedulers, Infrastructure, and Admin. The main content area is titled 'Event Analytics' and includes tabs for Faults, Events, and Audit Logs. A search bar at the top allows filtering by attributes or keywords. Below the search bar is a table of events.

<input type="checkbox"/>	Severity	Code	Affected object	Description	Creation Time
<input type="checkbox"/>	Critical	F0323	topology/pod-1/node-1/lon/svc-ifc_edmgr	Lost connectivity to leader for some data subset(s) of Access Service ifc_edmgr on node 1	Jan 22 2019 11:14:58am
<input type="checkbox"/>	Critical	F0325	topology/pod-1/node-1/lon	Connectivity has been lost to the leader for some data subset(s) of a service on node 1, the service may have unexpectedly restarted or failed.	Jan 22 2019 11:14:58am
<input type="checkbox"/>	Critical	F0323	topology/pod-1/node-1/lon/svc-ifc_licensemgr	Lost connectivity to leader for some data subset(s) of Access Service ifc_licensemgr on node 1	Jan 22 2019 11:14:58am
<input type="checkbox"/>	Critical	F0323	topology/pod-1/node-1/lon/svc-ifc_domainmgr	Lost connectivity to leader for some data subset(s) of Access Service ifc_domainmgr on node 1	Jan 22 2019 11:14:58am
<input type="checkbox"/>	Critical	F0323	topology/pod-1/node-1/lon/svc-ifc_vmmmgr	Lost connectivity to leader for some data subset(s) of Access Service ifc_vmmmgr on node 1	Jan 22 2019 11:14:58am
<input type="checkbox"/>	Critical	F0323	topology/pod-1/node-1/lon/svc-ifc_topmgr	Lost connectivity to leader for some data subset(s) of Access Service ifc_topmgr on node 1	Jan 22 2019 11:14:58am
<input type="checkbox"/>	Minor	F3361	acct-[infra]/region-[us-west-1]/context-[overlay-1]/csr-[ct_routerp_us-west-1_1:0]/physical-1/oper	Operational State of Physical Interface is down due to \u0002	Jan 22 2019 11:20:15am
<input type="checkbox"/>	Major	F607450	topology/pod-1/node-1/sys	[FSM:FAILED]: Task for updating topSystem changes for chassis on Node 1 to Topology Manager(TASK:ifc:ae:TopSystemControllerChassis)	Jan 22 2019 11:28:49am
<input type="checkbox"/>	Major	F606274	topology/pod-1/node-1/sys	[FSM:FAILED]: Task for updating topSystem changes for Node 1 to Topology Manager(TASK:ifc:ae:TopSystemSendTopSystem)	Jan 22 2019 11:28:49am

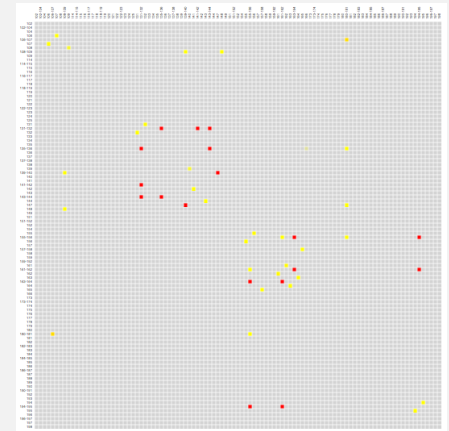


ACI Traffic Map

Help visualize and quickly spot high traffic density and underutilized nodes in the Cisco ACI™ fabric.

A grid is presented with a list of node IDs or vPC pairs on each axis. Traffic flow between a given pair of nodes or between a vPC pair is presented using color-coded cells on the heat map.

Traffic density is presented in a range of colors, from lightest (yellow), to shades of orange, to red (highest). Traffic statistics are collected using atomic counters.



- You can order by name or by traffic.
- Traffic can be seen by:
 - Sent packets
 - Received packets
 - Dropped packets
 - Excess packets

Order: **By Name** ▾

Traffic: **Sent Packets** ▾

Interval: **Cumulative** ▾

Spines: **select switches** ▾

Node Range: **a range of nodes such as 101-110,115 to filter on**

APPLY **SETTINGS**

Another way to check traffic on Fabric level

- Visualize utilization on Fabric level using APIC Apps
- We can monitor different parameters at Fabric Level
- VisuDash App:
 - Top 10 Tenants ranked by number of End-points
 - Top 20 interface by utilization



Going deeper – fTriage

- Fabric triaging tool (used by TAC)
- Python utility, runs on APIC in admin mode.
- Logs into switch nodes to capture requested data with commands/query/ELAM.
- Driven by specific user inputs which are validated first.
- Runs ELAM to determine packet data path
- Runs show commands/moquery on APIC/switch to determine control plane
- Traces a packet hop by hop until the point where it exits the fabric or gets dropped.
- Provides detailed info on interfaces + nodes where packet capture is attempted
- Drop reason is provided along with node, interface info.
- Requires consistent traffic stream flowing, as it relies on ELAM, it must have data packet flowing.
- Not a tool to detect transient and partial drops

ELAM Assistant (graphical interface of fTriage) – 1/2

The screenshot shows a web browser window with the URL `https://apic1.smalldc.ceclabs.local:1601`. The page title is "APIC (Disneyland)". The navigation menu includes "System", "Tenants", "Fabric", "Virtual Networking", "L4-L7 Services", "Admin", "Operations", and "Apps". The "Apps" menu is expanded, showing "ELAM Assistant (Beta)".

The ELAM Assistant interface has a sidebar on the left with the following items:

- Capture Packet
- node-101 (spine-01)
- node-111 (leaf-01)
- node-112 (leaf-02)
- node-113 (leaf-03)
- node-114 (leaf-04)
- node-115 (leaf-05)
- node-116 (leaf-06)

The main content area displays "Welcome to ELAM Assistant" and a "Leaf/Spine Login Password" form. The form contains two input fields: the first has the text "admin" and the second has "*****". A green "Validate & Save" button is positioned below the fields.

Below the login form is a "Supported Devices" section with the following information:

- LEAFs**: N9K-C93180YC-EX, N9K-C93108TC-EX, N9K-C93180LC-EX, N9K-C93180YC-FX, N9K-C93108TC-FX, N9K-C9348GC-FXP
- SPINE LineCards**: N9K-X9732C-EX, N9K-X9736C-EX, N9K-X9736C-FX

At the bottom of the page, the status bar shows "Last Login Time: 2018-06-09T22:55 UTC+02:00" and "Current System Time: 2018-06-09T22:59 UTC+02:00".

ELAM Assistant (graphical interface of fTriage) – 2/2

The screenshot displays the APIC (APIC (Disneyland)) interface, specifically the ELAM Assistant (Beta) section. The main heading is "Capture a packet with ELAM (Embedded Logic Analyzer Module)".

ELAM PARAMETERS

name your capture :

Status	Node	Direction	Source I/F	Parameters (Outer Header)	(Inner Header)
<input type="checkbox"/>	Not Set	node-111	from frontport	eth1/1	dst_ip 10.10.10.10

Buttons: (green), (green), (orange)

Left sidebar: Capture Packet, node-101 (spine-01), node-111 (leaf-01), node-112 (leaf-02), node-113 (leaf-03), node-114 (leaf-04), node-115 (leaf-05), node-116 (leaf-06)

Bottom status: Last Login Time: 2018-06-09T22:55 UTC+02:00, Current System Time: 2018-06-09T23:01 UTC+02:00

If you REALLY want to get deeper and love CLI

fTriage Route

```
sd-tb-99-ifcl# ftriage route -h
usage: ftriage route [-h] -ii intf [-ie encap] [-iie encap] -ei intf
                    [-ee encap] [-eie encap] -dip addr [-sip addr]

optional arguments:
  -h, --help      show this help message and exit
  -ii intf        ingress if (siteid::(node:if|VPC:<vpc-name>|PC:<pc-
                  name>)[+vtep])
  -ie encap       ingress encap (VLAN/VNID)
  -iie encap      ingress encap (VLAN/VNID)
  -ei intf        egress if (siteid::(node:if|VPC:<vpc-name>|PC:<pc-name>)[+vtep])
  -ee encap       egress encap (VLAN/VNID)
  -eie encap      egress encap (VLAN/VNID)
  -dip addr       destination IP
  -sip addr       source IP
```


If you REALLY want to get deeper and love CLI

fTriage Bridge

```
sd-tb-99-ifcl# ftrriage bridge -h
usage: ftrriage bridge [-h] -ii intf [-ie encap] [-iie encap] -ei intf
                        [-ee encap] [-eie encap] -dmac addr [-smac addr]

optional arguments:
  -h, --help      show this help message and exit
  -ii intf        ingress if (siteid::(node:if|VPC:<vpc-name>|PC:<pc-
                  name>)[+vtep])
  -ie encap       ingress encap (VLAN/VNID)
  -iie encap      ingress encap (VLAN/VNID)
  -ei intf        egress if (siteid::(node:if|VPC:<vpc-name>|PC:<pc-name>)[+vtep])
  -ee encap       egress encap (VLAN/VNID)
  -eie encap      egress encap (VLAN/VNID)
  -dmac addr      destination MAC
  -smac addr      source MAC
```

If you REALLY want to get deeper and love CLI


fTriage Logs

```
sd-tb-99-afc1# ftriage route -ii sd-tb-99-leaf1:Eth1/47 -ie 101 -ei sd-tb-99-leaf2:Eth1/47 -ee 203 -sip 101.1.1.10 -dip 101.1.2.10
ftriage: info : 2018-03-13 09:06:33.286: Building egress BD(s), Ctx
ftriage: info : 2018-03-13 09:06:34.675: Egress BD(s) {sd-tb-99-leaf2: 'bd-[vxlan-16154554]'}
ftriage: info : 2018-03-13 09:06:34.675: Egress Ctx ctx-[vxlan-2916352]
ftriage: info : 2018-03-13 09:06:34.675: Building ingress BD(s), Ctx
ftriage: info : 2018-03-13 09:06:36.297: Ingress BD(s) {sd-tb-99-leaf1: 'bd-[vxlan-16613251]'}
ftriage: info : 2018-03-13 09:06:36.297: Ingress Ctx ctx-[vxlan-2916352]
ftriage: info : 2018-03-13 09:06:36.297: Capturing L3 packet on [sd-tb-99-leaf1]
ftriage: info : 2018-03-13 09:07:09.560: L3 packet seen on sd-tb-99-leaf1:Eth1/47
ftriage: info : 2018-03-13 09:07:16.149: SIP 101.1.1.10 DIP 101.1.2.10
ftriage: info : 2018-03-13 09:07:16.149: sd-tb-99-leaf1: Ingress function
ftriage: info : 2018-03-13 09:07:27.851: sd-tb-99-leaf1: Dst EP is remote
ftriage: info : 2018-03-13 09:07:33.464: sd-tb-99-leaf1: DMAC(00:22:BD:F8:19:FF) same as RMAC(00:22:BD:F8:19:FF)
ftriage: info : 2018-03-13 09:07:33.464: sd-tb-99-leaf1: L3 packet getting routed in SUG
ftriage: info : 2018-03-13 09:07:37.586: sd-tb-99-leaf1: Dst IP is present in SUG L3 tbl
ftriage: info : 2018-03-13 09:07:42.223: sd-tb-99-leaf1: RwdMAC DIPo(10.0.232.66) is one of dst TEPs ['10.0.232.66']
ftriage: info : 2018-03-13 09:07:46.732: Computing next set of nodes
ftriage: info : 2018-03-13 09:07:56.872: Capturing L3 packet on [sd-tb-99-spine1]
ftriage: info : 2018-03-13 09:08:31.547: L3 packet seen on sd-tb-99-spine1:Eth1/1
ftriage: info : 2018-03-13 09:08:40.749: sd-tb-99-spine1: Transit function
ftriage: info : 2018-03-13 09:08:40.749: sd-tb-99-spine1: Capturing L3 packet on egress LC
ftriage: info : 2018-03-13 09:08:44.656: sd-tb-99-spine1: Infra route 10.0.232.66 present in RIB
ftriage: info : 2018-03-13 09:09:21.082: sd-tb-99-spine1: L3 packet seen on egress LC1
ftriage: info : 2018-03-13 09:09:24.989: Computing next set of nodes
ftriage: info : 2018-03-13 09:09:36.779: Capturing L3 packet on [sd-tb-99-leaf2]
ftriage: info : 2018-03-13 09:10:09.364: L3 packet seen on sd-tb-99-leaf2:Eth1/53
ftriage: info : 2018-03-13 09:10:15.207: sd-tb-99-leaf2: Egress function
ftriage: info : 2018-03-13 09:10:23.674: sd-tb-99-leaf2: Dst EP is local
ftriage: info : 2018-03-13 09:10:23.674: sd-tb-99-leaf2: EP if(Eth1/47) same as egr if(Eth1/47)
ftriage: info : 2018-03-13 09:10:27.359: sd-tb-99-leaf2: Dst IP is present in HOM L3 tbl
ftriage: info : 2018-03-13 09:10:29.223: sd-tb-99-leaf2: RW seg_id in HOM same as EP fd seg_id
ftriage: hunch: None
sd-tb-99-afc1#
```

ACI Multi-Site

Since ACI 3.2

Day-2 Operations: Full-Stack Consistency Checker



- Multi-Site Infra: Unicast, Multicast, BGP TEPs and Tunnel state
- Multi-Site Tenant and EPG granularity:
 - Inspect and validate full-stack programming: Multi-Site Controller (MSC), APICs and Spine translations
 - Validate the consistency of local and remote inter-site EPGs, BD, VRF, External EPG, policies, etc.
 - Root cause configuration programming issues without calling TAC*
- GUI and APIs are both supported

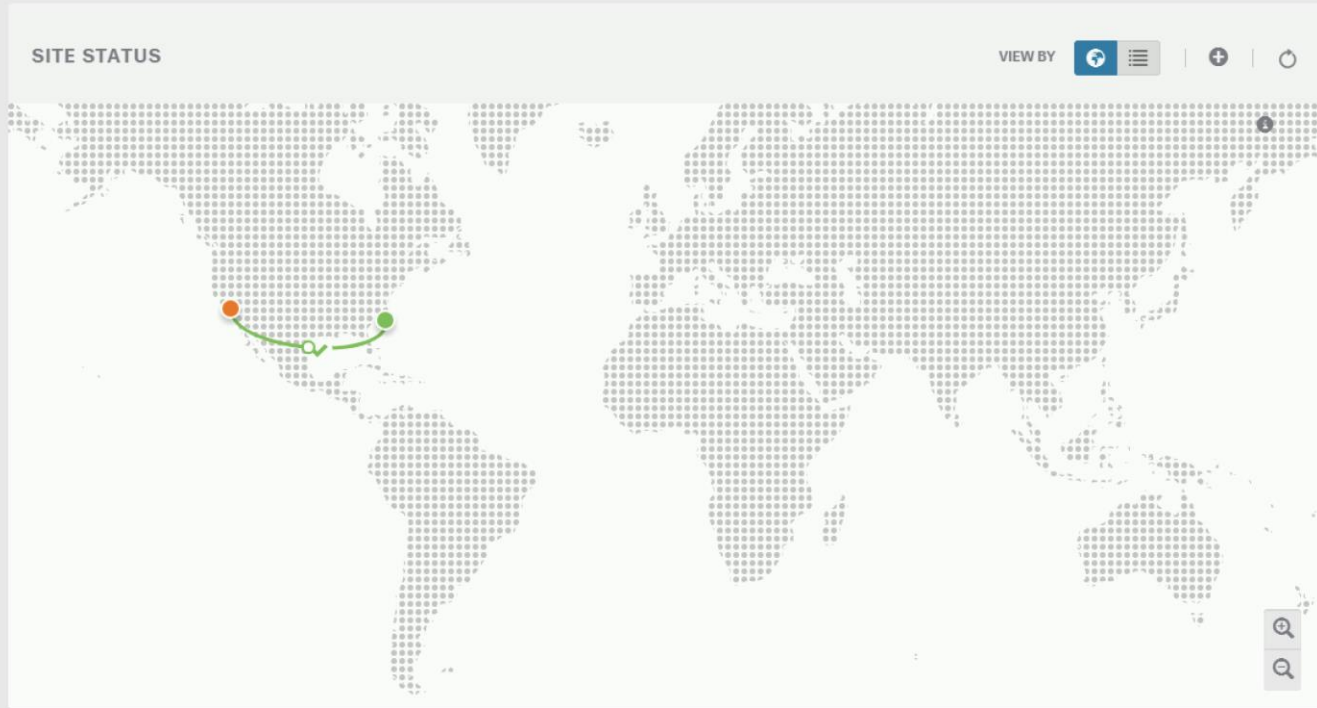
* Restrictions apply. View NIR/NIA

Eg.: Multi-Site Orchestrator - Dashboard

- Dashboard
- Sites
- Schemas
- Tenants
- Users
- Admin

 Multi Site Orchestrator

Cluster Status 3/3



Eg.: Multi-Site Orchestrator – Schema details

The screenshot displays the Hybrid-App Multi-Site Orchestrator interface. On the left, a sidebar lists 'TEMPLATES' (Hybrid-App) and 'SITES' (AWS aws, Hybrid-App, On-Prem Hybrid-App). The main area shows the 'AWS Hybrid-App' configuration for a 'HybridTenant'. Under the 'AP' section, 'WebShop' is listed with a green status dot. Below it, 'WebServer' and 'Application' components are shown as 'CONNECTED'. The 'CONTRACT' section contains 'cp-test-01'. The 'VRF' section shows 'WebShop'. On the right, a detailed view for 'VRF WebShop' is shown, including a progress bar (0/0/0/0), 'TEMPLATE PROPERTIES' (Display Name: WebShop), and 'SITE LOCAL PROPERTIES' (Regions: us-west-1, CIDR: 200.200.0.0/16 Primary).

Eg.: Multi-Site Orchestrator – Schema deployed

ACI 4.1

The screenshot displays the Cisco Multi Site Orchestrator interface. A modal window titled "Schema Details" is open, showing the configuration for a schema named "Hybrid-App". The modal includes a table with columns for severity levels: CRITICAL, MAJOR, MINOR, and WARNING. The table lists five entries, each with a type (ANP, EPB, EPB, CONTRACT, VRF) and a name (WebShop, WebServer, Application, cp-test-01, WebShop). All entries show zero counts for all severity levels. The background interface shows a navigation menu on the left and a world map in the main area.

		CRITICAL	MAJOR	MINOR	WARNING
ANP	WebShop	0	0	0	0
EPB	WebServer	0	0	0	0
EPB	Application	0	0	0	0
CONTRACT	cp-test-01	0	0	0	0
VRF	WebShop	0	0	0	0

Eg.: Multi-Site Orchestrator 2.2.3 – diffs

Deploy to sites ✕

+ Created
 ↙ Modified
 Deleted

Object Type	Name	Site 1	Site 2
Application Profile	anp_template_1	+ Created	
EPG	upeg_1_ctx_1_bd_1	↙ Modified ⚠	+ Created ⚠
Contract	contract_1_ctx_1_bd_1	↙ Modified	
VRF	vrf_1_ctx_1		Deleted
Bridge Domain	ctx_1_bd_1		

Modified Properties

Scope

Tenant → Global

5 more

[Deploy](#)

Eg.: Multi-Site Orchestrator Audit Logs

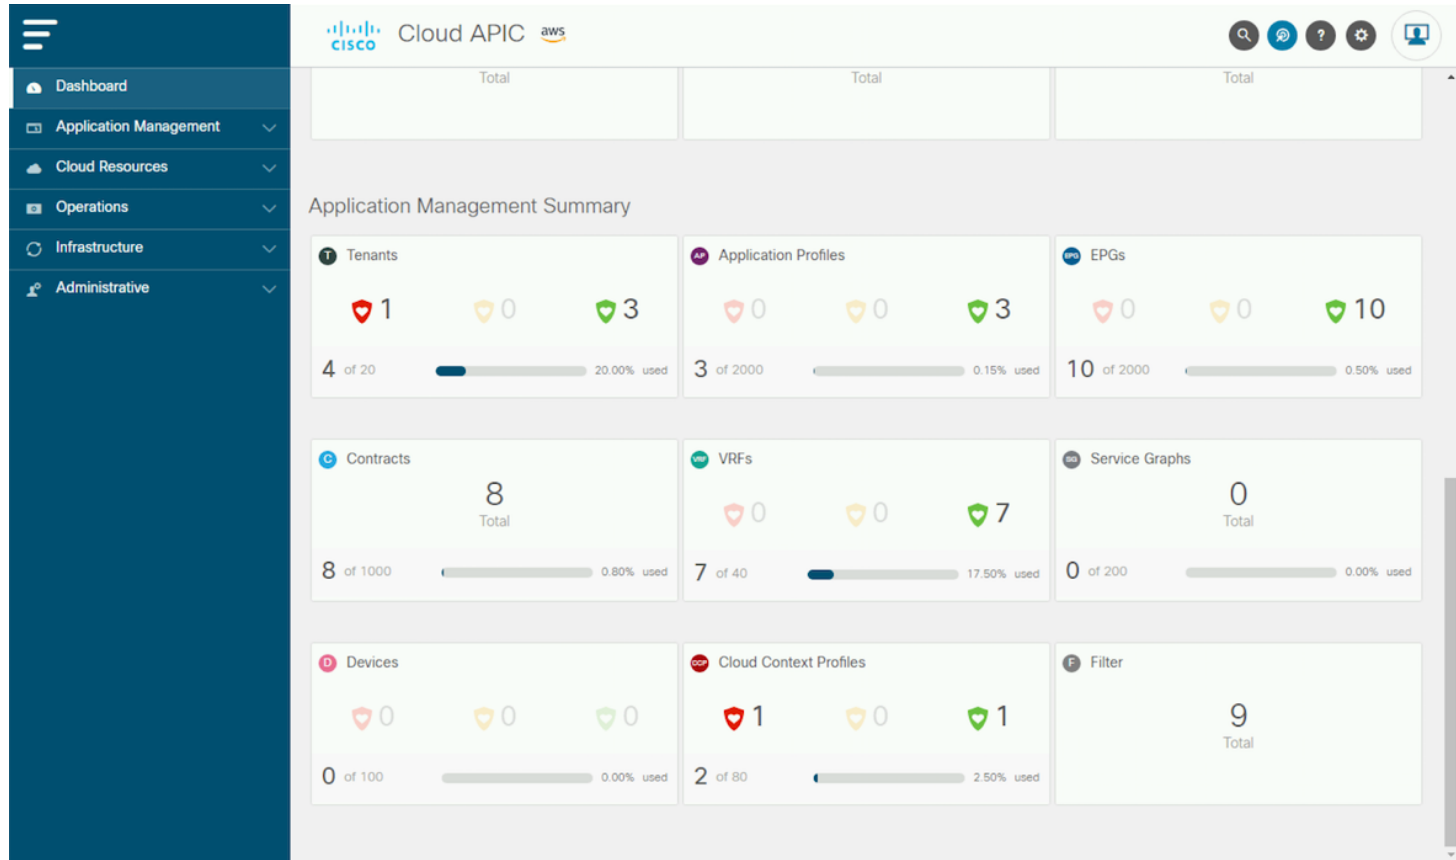
Multi Site Orchestrator Cluster Status 1/1

Audit Logs (592 Logs)

Most Recent

DATE	ACTION	TYPE	DETAILS	USER
Aug 2, 2018 7:39:19 AM	Logged In	Authentication	User admin has successfully logged in	admin (Admin User) Local
Aug 1, 2018 11:12:01 PM	Logged In	Authentication	User admin has successfully logged in	admin (Admin User) Local
Aug 1, 2018 10:43:24 PM	Logged In	Authentication	User admin has successfully logged in	admin (Admin User) Local
Aug 1, 2018 10:11:41 PM	Logged In	Authentication	User admin has successfully logged in	admin (Admin User) Local
Aug 1, 2018 10:11:25 PM	Login Failed	Authentication	Login failed for admin	-
Jul 29, 2018 8:44:45 PM	Logged In	Authentication	User admin has successfully logged in	admin (Admin User) Local
Jul 27, 2018 11:15:33 AM	Deployed	Schema Site	Template Hybrid-App of Hybrid-App was deployed to On-Prem	admin (Admin User) Local
Jul 27, 2018 11:15:33 AM	Deployed	Schema Site	Template Hybrid-App of Hybrid-App was deployed to AWS	admin (Admin User) Local
Jul 27, 2018 11:15:31 AM	Updated	Schema	Schema Hybrid-App was updated	admin (Admin User) Local
Jul 27, 2018 11:15:31 AM	Updated	Template	Template Hybrid-App on Schema Hybrid-App was updated	admin (Admin User) Local
Jul 27, 2018 11:15:31 AM	Updated	EPG	EPG Application on Template Hybrid-App was updated	admin (Admin User) Local

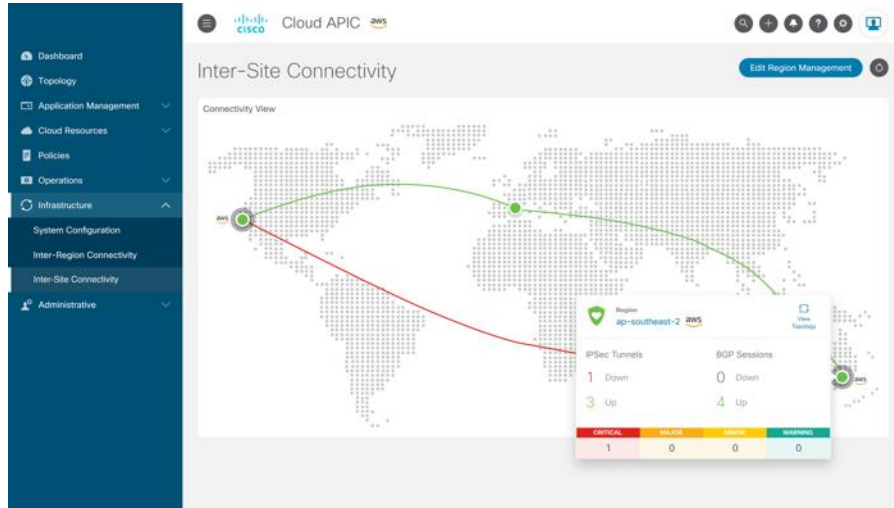
Cloud ACI – Dashboard (AWS example)



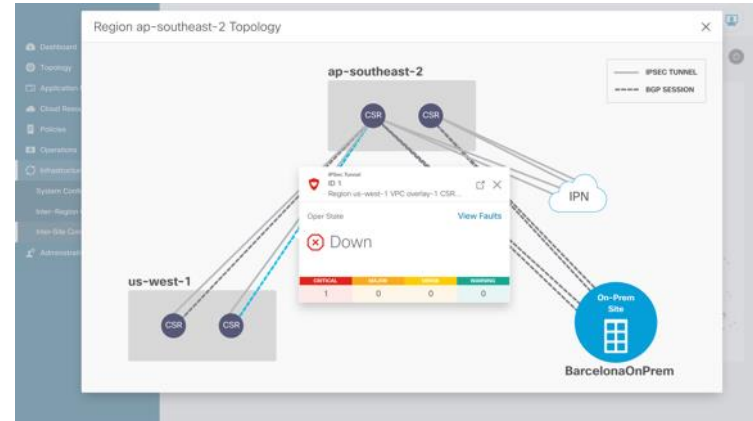
Cloud Native Events, Faults, Stats and Audit Trails

- Cloud events and stats are captured through native services like:
 - Flow Logs – VPC, NIC Flow Stats.
 - Cloud Trail – Audit History of changes to various cloud resources.
 - AWS CloudWatch – Health Monitoring metrics and Log monitoring for events.
 - AWS Config service – Configuration changes to the resources.
 - Additional statistics and events are collected from CSR1000V, like tunnel health, BGP session health etc.
- The cloud events and stats are converted into the ACI Object Model to provide consistent events, faults, stats, audit like on-premises ACI deployment.
- APIC manages cloud inventory objects per account. Aggregation and reporting can be handled at logical objects like CloudEP, CloudEPG , VRF (or) at cloud resource objects such as NIC, host instance, VPC.

Topology Health



- Network connectivity and Health



Router CsrRouter1

Event Analytics

Severity	Code	Cause	Affected Object	Description	Last Transition
CRITICAL	F0325	Tunnel down	acct-[Redacted-tenant-1]-[region-us-west-1]-[account-1]-[vpc-overlay-1]-[vpc-car-router1]-[tunnel-1]	Configuration error on tunnel	1 hour ago

Cloud ACI - Inter-Site Connectivity

Cloud APIC

Inter-Site Connectivity

Connectivity Status

Operational Configuration

US West (N. California) (cAPIC Deployed)

CSRs

Health	Name	Management IP Addr...	Connected VPCs	IPsec Tunnels	BGP Sessions
Healthy	ct_routerp_us-west-1...	13.57.119.15	1 ↑1 ↓0 Connected VPCs	2 ↑2 ↓0 IPsec Tunnels	2 ↑2 ↓0 BGP Sessions
Healthy	ct_routerp_us-west-1...	184.169.218.172	1 ↑1 ↓0 Connected VPCs	2 ↑2 ↓0 IPsec Tunnels	2 ↑2 ↓0 BGP Sessions

IPN

Cloud ACI – Events analytics

ACI 4.1

Tenant HybridTenant

Overview Cloud Resources Application Management Statistics **Event Analytics**

⊙ ✎ Actions ▾

Faults

Description contains **clus** ✕

Events

Audit Logs

ID	Category	Affected object	Description	Action	Local User	Creation Time
4294969932	transition	subj-[uni/tn-HybridTenant/cloudapp-WebShop/cloudepg-S01/rscons---msc--WebShop-S01]	RsCons --msc--WebShop-CLUS01 created	creation	admin	Jun 08 2019 07:47:36pm
4294969929	transition	subj-[uni/tn-HybridTenant/cloudapp-WebShop/cloudepg-CLUS01/rsprov---msc--WebShop-CLUS01]	RsProv --msc--WebShop-CLUS01 created	creation	admin	Jun 08 2019 07:47:36pm
4294969935	transition	subj-[uni/tn-HybridTenant/cloudapp-WebShop/cloudepg-CLUS01]	EPg CLUS01 created	creation	admin	Jun 08 2019 07:47:36pm
4294969923	transition	subj-[uni/tn-HybridTenant/cloudapp-WebShop/cloudextepg---msc--CLUS01/rscons---msc--WebShop-CLUS01]	RsCons --msc--WebShop-CLUS01 created	creation	admin	Jun 08 2019 07:47:36pm
4294969922	transition	subj-[uni/tn-HybridTenant/cloudapp-WebShop/cloudextepg---msc--CLUS01/rsprov---msc--WebShop-CLUS01]	RsProv --msc--WebShop-CLUS01 created	creation	admin	Jun 08 2019 07:47:36pm



Cloud ACI – Infra CSR1000v Routes

ACI 4.1

context overlay-1

Overview Cloud Resources ACI Relationships Event Analytics

General

Account
infra

Region
us-west-1

Tags
No tags set

Cloud Resources

2 Availability Zones 2 Routers

Configuration Relationships

3 EPGs 1 Cloud Context Profiles

Settings

Encap: 16777199 Encap Type: Vxlan

Cloud Context Profile: ct_ctxprofile_us-west-1
Tenant infra

CIDRs

Address	Subnets	Primary
10.0.1.0/25	10.0.1.0/28 7 more	yes

Route Tables

destinationAddress	Subnets
0.0.0.0/0 1 more	10.0.1.96/28
0.0.0.0/0 1 more	10.0.1.32/28
0.0.0.0/0 1 more	10.0.1.0/28
0.0.0.0/0 1 more	10.0.1.64/28
52.53.52.97 1 more	10.0.1.48/28
52.53.52.97	10.0.1.16/28

Health: N/A

Faults

CRITICAL	MAJOR	MINOR	WARNING
0	0	1	0

Audit Logs

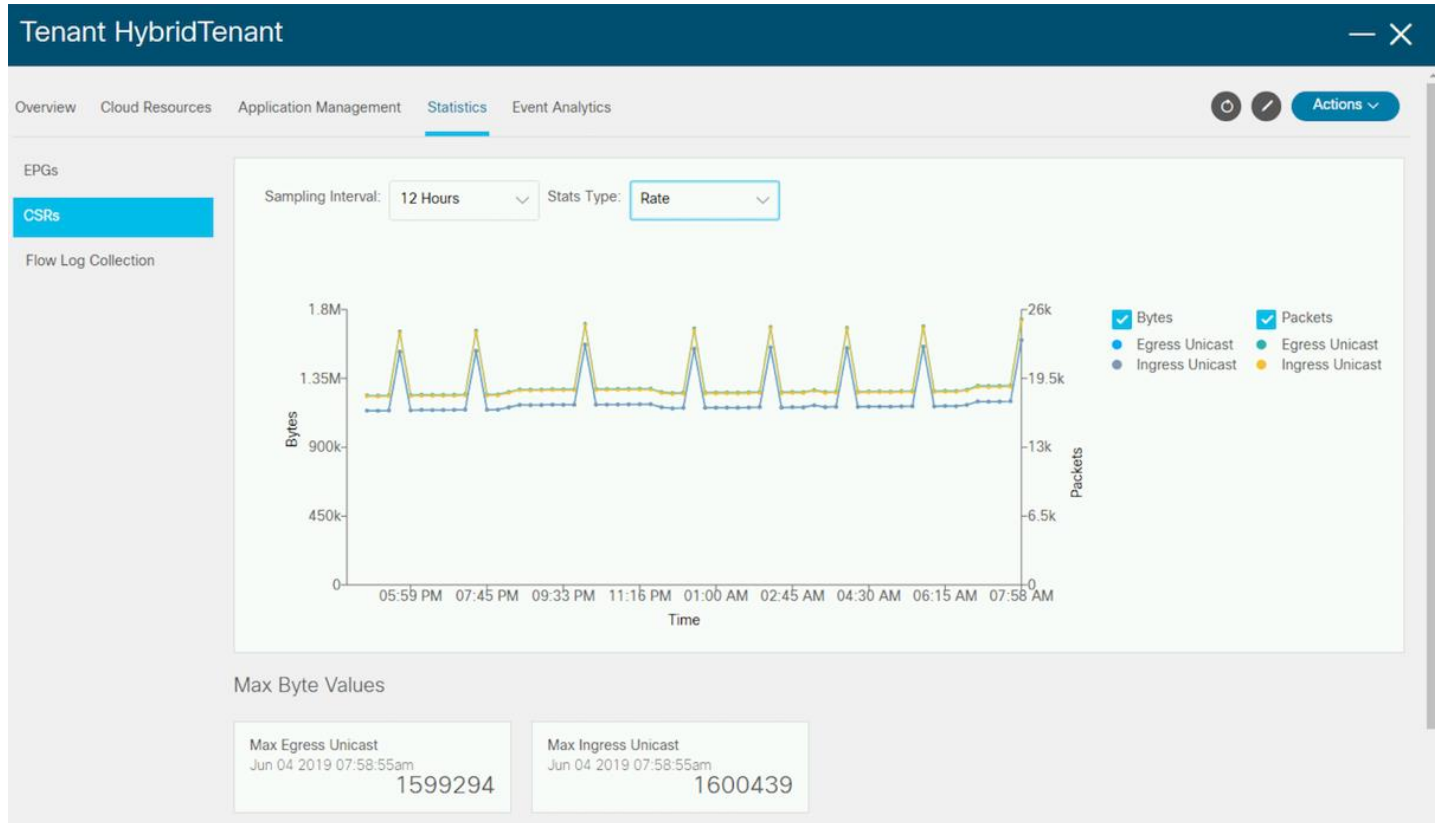
0 Deletion 0 Creation 0 Modification

Events

0 Critical 0 Major 0 Minor 0 Other



Cloud ACI – Statistics example



Cloud ACI – Events Analytics

ACI 4.1

The screenshot displays the Cisco Cloud APIC Event Analytics interface. The top navigation bar includes 'Dashboard', 'Application Management', and 'Cloud Resources'. The main header shows 'Event Analytics' with sub-tabs for 'Faults', 'Events', and 'Audit Logs'. A modal window titled 'Fault F609093' is open, showing the following details:

Details	
Severity	Code
Major	F609093
Cause	Affected object
fsm-failed	uni/controller/setuppol/setup-1
Description	Last Transition
[FSM:FAILED]: Remove fabricPod after deleting the pod policy. (TASK:ifc.policymgr.FabricSetupPPodCleanup)	May 01 2019 12:59:34pm

Below the details, there is a 'Change Set' section and a table of related events:

Severity	Code	Affected Object	Description	Time
Major	F606666	uni/fabric/macprotp-default/macexpg-default	[FSM:FAILED]: Assign Virtual IP address fr Protection Group Name(TASK:ifc.policymgr.FabricAProtGE	May 01 2019 12:59:40pm

ACI Operations - Agenda

- Before getting started – setting the concepts stage

Visibility	Insights	Actions
Faults, events, stats, health, logs, trails	Application dependency	Incident Troubleshooting
Configuration	Containers Integration	Change Management

ACI Configuration Rollback



APIC (ACI-SJC-1)

What are you looking for?

admin



System | Tenants | Fabric | Virtual Networking | L4-L7 Services | **Admin** | Operations | Apps | Integrations

AAA | Schedulers | Historical Record Policies | Firmware | External Data Collectors | **Config Rollbacks** | Import/Export | Downloads

Config Rollbacks for:

Snapshots	File Name	Description	File Size (bytes)
2018-03-08 15:25:1...	ce2_defaultOneTime-2018-03-08T13-25-13.ta...	before tenant Rivella ...	54667
2018-06-10 22:50:2...	ce2_defaultOneTime-2018-06-10T22-50-13.ta...	CLUS18-SNAP1	580416
2018-06-10 22:52:1...	ce2_defaultOneTime-2018-06-10T22-52-15.ta...	CLUS18-SNAP2	580698
2018-06-10 22:59:3...	ce2_defaultOneTime-2018-06-10T22-59-29.ta...	CLUS18-SNAP3	578894
2018-12-12 18:33:4...	ce2_defaultOneTime-2018-12-12T17-33-42.ta...	pre-CCP-2.1	823781
2018-12-14 17:05:4...	ce2_defaultOneTime-2018-12-14T16-05-39.ta...	Pre-HX2	817899
2018-12-21 12:27:4...	ce2_defaultOneTime-2018-12-21T11-27-38.ta...	chrischtchindli	818092
2019-05-16 10:53:4...	ce2_config_backup-2019-05-16T10-53-38.tar...	Backups taken before...	850052
2019-05-16 10:57:5...	ce2_config_backup-2019-05-16T10-57-44.tar...	Backups taken before...	848866
2019-05-16 11:00:5...	ce2_config_backup-2019-05-16T11-00-49.tar...	Backups taken before...	848304
2019-05-16 11:03:5...	ce2_config_backup-2019-05-16T11-03-50.tar...	Backups taken before...	850067
2019-05-16 11:05:5...	ce2_config_backup-2019-05-16T11-05-49.tar...	Backups taken before...	850481
2019-05-16 11:10:5...	ce2_config_backup-2019-05-16T11-10-49.tar...	Backups taken before...	849511
2019-05-16 11:16:0...	ce2_config_backup-2019-05-16T11-15-55.tar...	Backups taken before...	850002
2019-05-16 11:17:4...	ce2_config_backup-2019-05-16T11-17-39.tar...	Backups taken before...	850175
2019-05-16 12:42:2...	ce2_config_backup-2019-05-16T12-42-20.tar...	Backups taken before...	850673

Actions

Rollback

Select any one snapshot on left to start.

Take a snapshot

Location:

Description:

Import export file to snapshot

Click icon on top

Modify import/export security settings

Click icon on top

Create recurring snapshots

Click icon on top

Automatically create snapshot

Snapshots taken every

Mon Tue Wed Thu Fri Sat Sun

at Hour: Minute:

Save to Remote Location instead: Create a remote location:



ACI Configuration Rollback

- Snapshot for whole fabric / per tenant basis. You can create snapshots manually or periodic
- Example shows the difference between 2 snapshots at fabric level

The screenshot displays the APIC (ACI-SJC-1) configuration rollback interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'Virtual Networking', 'L4-L7 Services', 'Admin', 'Operations', 'Apps', and 'Integrations'. The 'Admin' tab is selected, and the 'Config Rollbacks' section is active. A search bar at the top right contains the text 'What are you looking for?'. The main content area is divided into two panels. The left panel, titled 'Config Rollbacks for: Tenant', shows a table of snapshots with columns for 'Snapshots', 'File Name', 'Description', and 'File Size (bytes)'. The right panel shows a detailed view of the configuration differences between two snapshots, with a 'Go Back' button and an 'Undo these changes' button.

Snapshots	File Name	Description	File Size (bytes)
2018-03-08 15:25:1...	ce2_defaultOneTime-2018-03-08T13-25-13.ta...	before tenant Rivella ...	54667
2018-06-10 22:50:2...	ce2_defaultOneTime-2018-06-10T22-50-13.ta...	CLUS18-SNAP1	580416
2018-06-10 22:52:1...	ce2_defaultOneTime-2018-06-10T22-52-15.ta...	CLUS18-SNAP2	580698
2018-06-10 22:59:3...	ce2_defaultOneTime-2018-06-10T22-59-29.ta...	CLUS18-SNAP3	578894
2018-12-12 18:33:4...	ce2_defaultOneTime-2018-12-12T17-33-42.ta...	pre-CCP-2.1	823781
2018-12-14 17:05:4...	ce2_defaultOneTime-2018-12-14T16-05-39.ta...	Pre-HX2	817899
2018-12-21 12:27:4...	ce2_defaultOneTime-2018-12-21T11-27-38.ta...	chrischtchindli	818092
2019-05-16 10:53:4...	ce2_config_backup-2019-05-16T10-53-38.tar...	Backups taken before...	850052
2019-05-16 10:57:5...	ce2_config_backup-2019-05-16T10-57-44.tar...	Backups taken before...	848866
2019-05-16 11:00:5...	ce2_config_backup-2019-05-16T11-00-49.tar...	Backups taken before...	848304
2019-05-16 11:03:5...	ce2_config_backup-2019-05-16T11-03-50.tar...	Backups taken before...	850067
2019-05-16 11:05:5...	ce2_config_backup-2019-05-16T11-05-49.tar...	Backups taken before...	850481
2019-05-16 11:10:5...	ce2_config_backup-2019-05-16T11-10-49.tar...	Backups taken before...	849511
2019-05-16 11:16:0...	ce2_config_backup-2019-05-16T11-15-55.tar...	Backups taken before...	850002
2019-05-16 11:17:4...	ce2_config_backup-2019-05-16T11-17-39.tar...	Backups taken before...	850175
2019-05-16 12:42:2...	ce2_config_backup-2019-05-16T12-42-20.tar...	Backups taken before...	850673
2019-05-16 12:47:4...	ce2_config_backup-2019-05-16T12-47-40.tar...	Backups taken before...	850290
2019-06-08 09:00:1...	ce2_DailyAutoBackup-2019-06-08T09-00-09.t...		852364
2019-06-08 17:00:2...	ce2_DailyAutoBackup-2019-06-08T17-00-20.t...		851513
2019-06-09 01:00:0...	ce2_DailyAutoBackup-2019-06-09T01-00-01.t...		851677

```
<po!uni
  rn="uni"
  >
  <aaaUserEp
    rn="userext"
    pwdStrengthCheck="no"
  >
  <aaaPreloginBanner
    rn="preloginbanner"
    guiTextMessage="ACI-SJC-1"
    guiTextMessage="ACI-MLSN-1"
    message="Application Policy Infrastructure Controller"
  >
  </aaaPreloginBanner>
  </aaaUserEp>
  <fabricInst
    rn="fabric"
  >
  <fileRemotePath
    name="dn-7c"
    rn="path-dn-7c"
    remotePort="21"
    identityPrivateKeyPassphrase="SE51gD4tB1rUAY5GHSmYkHFq71rTAxo6HmY8XhHnzF3joYnrK3pE2ndBK66hSAH1b..."
    protocol="ftp"
    identityPrivateKeyContents="SE51gD4tB1rUAY5GHSmYkHFq71rTAxo6HmY8XhHnzF3jqG9DDcFP3K4p1y7H9LZ..."
    authType="usePassword"
    identityPublicKeyContent="CF51eD4tB1rUAY5GHSmYkHFq71rTAxo6HmY8XhHnzF3jqG9DDcFP3K4p1y7H9LZ..."
  >
  </fileRemotePath>
  </fabricInst>
  </po!uni>
```

Duplicate IP

The screenshot shows the APIC (Disneyland) System Health dashboard. A central alert list is displayed, containing two alerts:

- Same IP Address used for multiple MAC's** (Critical): One or more duplicate IP addresses are being used by multiple MAC addresses. Please resolve this to avoid conflicts. [View all Duplicate IP Usage](#)
- Smart Licensing is not configured.** (Warning): The evaluation period has 71 day(s) remaining. There will be no impact on the functionality of the ACI fabric at the end of evaluation period. [Go to Smart Licensing](#)

The dashboard also features a System Health score of 92, a Nodes With Health ≤ 99 table, and Fault Counts By Domain and By Type sections.

Name	Pod ID	Health Score
leaf-01	1	78
leaf-02	1	98
leaf-03	1	98
leaf-05	1	98
leaf-06	1	98

Domain	Count	Count	Count	Count
2	106	52	74	
2	0	0	8	
0	24	5	1	
0	10	0	0	
0	70	7	61	
0	0	0	0	
0	0	0	0	
0	2	40	4	

Type	Count	Count	Count	Count
Communications	2	18	0	8
Config	0	56	43	64
Environmental	0	0	7	0
Operational	0	32	2	2

Cloud APIC - Upgrade

ACI 4.1

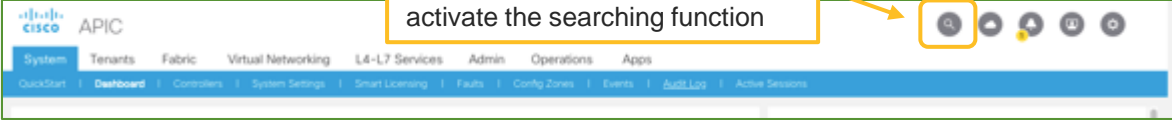
The screenshot displays the Cloud APIC Firmware Management page. The left sidebar contains navigation options: Dashboard, Application Management, Cloud Resources, Operations, Event Analytics, Active Sessions, Backup & Restore, Tech Support, Firmware Management (selected), Schedulers, Remote Locations, Infrastructure, and Administrative. The main content area is titled 'Firmware Management' and includes a 'Schedule Upgrade' button. Below this, there are sections for 'Current Firmware Version' (4.1(1i)) and 'Upgrade Settings' (Target Firmware Version, Ignore Compatibility Check: no, Upgrade Start Time: 2019-04-25T09:47:36.662+00:00). A table below shows the upgrade progress for controller ACI-Cloud-Fabric-1, which is at 100% completion with the message 'Upgrade Successful at Apr 25 2019 11:44:25am'. The table has columns for ID, Controller Name, Current Firmware Version, and Upgrade Status. The bottom of the page shows a pagination control for 5 rows, Page 1 of 1.

ID	Controller Name	Current Firmware Version	Upgrade Status
1	ACI-Cloud-Fabric-1	4.1(1i)	100% Upgrade Successful at Apr 25 2019 11:44:25am


- Similar steps as APIC
- Under Firmware Management select image location
- Schedule a time to upgrade
- Once done, it will show upgrade got completed

Improved Native Searching Function

Since the APIC 3.2, the UI provides a more flexible and easy-to-use native searching function with the support of search keywords and wild card.



Click on the Search icon to activate the searching function

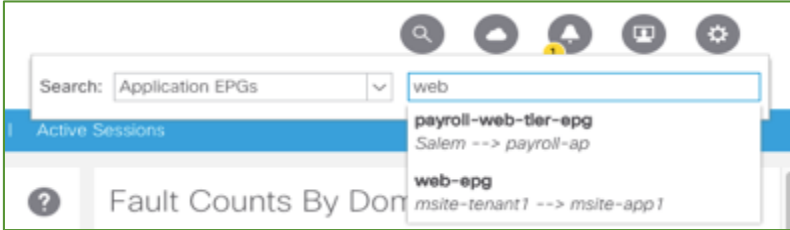


Select search category. Can choose from the dropdown menu, or search with keywords, such as "epg", etc.

Type in searching keyword, or use * as wild card.

The diagram illustrates the steps to activate and use the search function in the APIC UI. It shows the search icon being clicked, the search category dropdown menu being opened, and the search input field being used to enter a keyword or wild card.

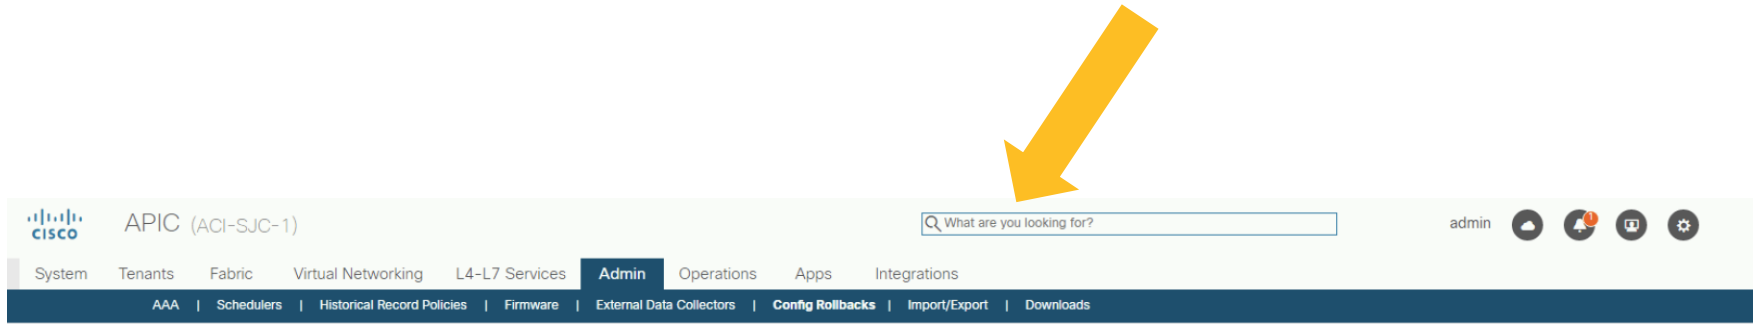
Search Output Example:



The screenshot shows the search results for the keyword "web". The search category is set to "Application EPGs". The results list includes:

- web
- payroll-web-tier-epg
Salem --> payroll-ap
- web-epg
msite-tenant1 --> msite-app1

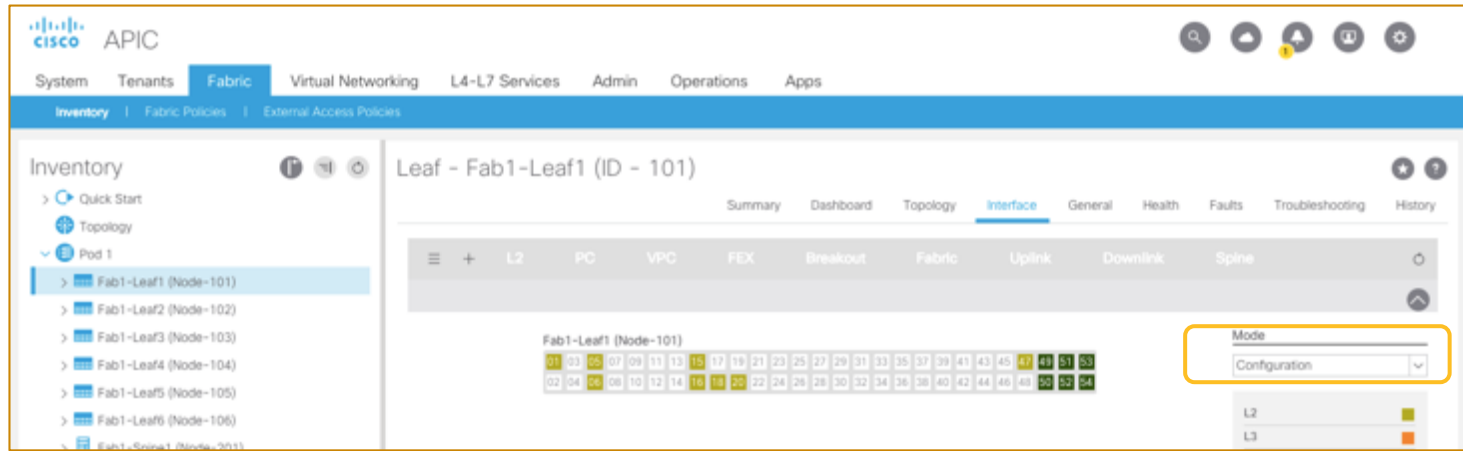
Since ACI 4.1 – search to rescue ...



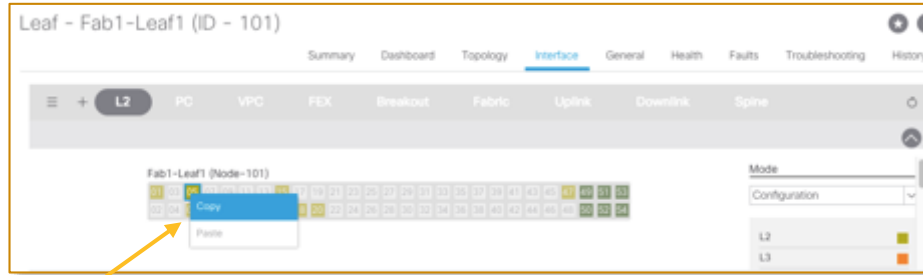
Copy/Paste L2 Port Configuration

Since the APIC 3.2, the UI adds the copy/paste function to replicate L2 port configuration on leaf switches. This is an alternative method to quickly configure L2 ports individually. The system will automatically create the associated objects behind the scenes.

Navigate to the function via Fabric → PoD → Leaf Switch → Interface, then toggle the mode to Configuration



Copy/Paste L2 Port Configuration

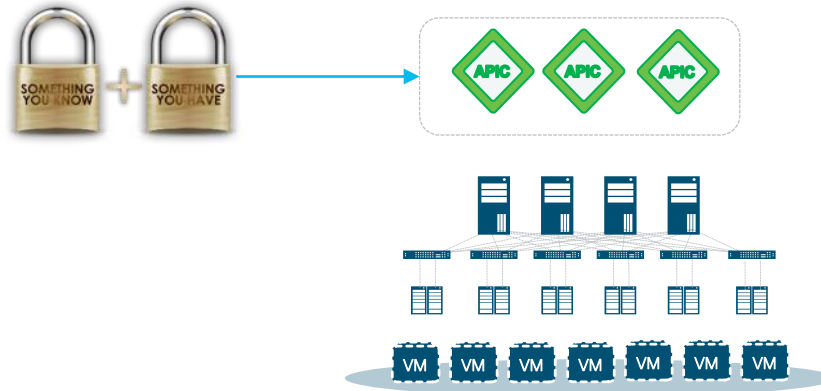


Select the copy-from port



Select and confirm the copy-to port

ACI 2-Factor Authentication Options



ACI 3.0

External Authentication via SAML and IDPs supported Okta & MSFT ADFS

ACI 3.0

Local Authentication TOTP using Google Authenticator for 2nd factor pin/barcode



ACI 3.1

RSA SecurID



ACI 3.2

PingFederate SSO
PingID 2-FA



ACI 4.0

Federal Common Access Card (CAC)



Culture change @ network operations ...

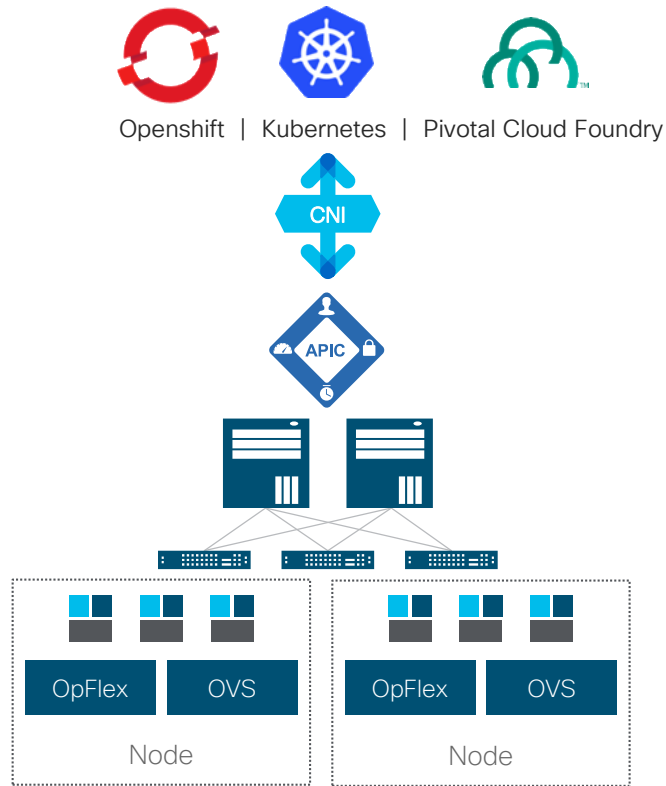


“ACI is my SDN solution and I only need access to its APIs as it must be part of my CI/CD pipeline for upgrades so I can have the latest integrations with the tools I need.”



“ACI is the production network of my key business processes, so any upgrade goes through change management and as/if needed, must be done at a maintenance window”

ACI Integration with Container Application Platforms



ACI and Containers



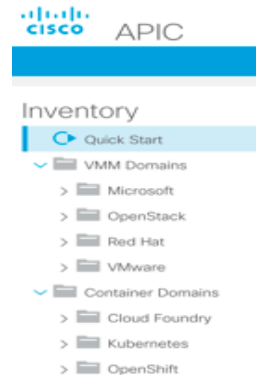
Unified networking:
Containers, VMs, and bare-metal



Integration of containers network
policies and ACI policies



Visibility: Live statistics in
APIC per container and
health metrics



VMM domain helps to bridge the gap between Kubernetes admin and network operations

The screenshot displays the Cisco APIC interface, which is used for managing network infrastructure. The interface is divided into several sections:

- Navigation:** The top navigation bar includes 'System', 'Tenants', 'Fabric', 'VM Networking', 'L4-L7 Services', 'Admin', 'Operations', and 'Apps'. The 'VM Networking' section is currently selected.
- Inventory:** The 'Inventory' section on the left shows a tree view of the network configuration. It includes 'Kubernetes', 'kube-p1', 'Nodes', 'Namespaces', and 'Services'. The 'sock-shop' namespace is highlighted with a red box.
- Services:** The 'Services' section on the left lists various services in the 'sock-shop' namespace, such as 'user-db', 'user', 'shipping', 'orders-db', 'front-end', 'orders', 'catalogue-db', 'payment', 'queue-master', 'rabbitmq', 'catalogue', 'carts-db', and 'carts'. A red box highlights the 'sock-shop' namespace in the left sidebar, with a red arrow pointing to the 'sock-shop' entry in the 'Inventory' section.
- Namespace - sock-shop:** The right-hand side of the interface shows the 'Namespace - sock-shop' details. It includes a 'Properties' table with columns for 'Name', 'Interface', 'IP', 'MAC', 'Encap', 'Labels', and 'EPG'. The table lists various pods and their associated network details.

Name	Interface	IP	MAC	Encap	Labels	EPG
carts-2469883...	pi-vethc02386f0	11.2.0.234	4E:62:A8:96:5...	vxlan-7995398	interface-name:vethc02386f0,pod-template-hash:246...	kube-p1 kubernetes kube-default
carts-db-1721...	pi-veth671d26...	11.2.0.227	0E:4F:68:AD:F...	vxlan-7995398	pod-template-hash:1721187500,name:carts-db,interf...	kube-p1 kubernetes kube-default
catalogue-429...	pi-vethd049db...	11.2.0.223	82:4E:50:87:0...	vxlan-7995398	pod-template-hash:4293036822,name:catalogue,inter...	kube-p1 kubernetes kube-default
catalogue-db-...	pi-veth635cd3...	11.2.0.230	FE:05:BC:1B:6...	vxlan-7995398	interface-name:veth635cd3c,pod-template-hash:18...	kube-p1 kubernetes kube-default
front-end-233...	pi-vethd29910...	11.2.0.233	66:14:54:18:4...	vxlan-7995398	interface-name:vethd29910d7,pod-template-hash:23...	kube-p1 kubernetes kube-default
orders-73348...	pi-veth65193d...	11.2.0.224	FE:DD:A2:02:8...	vxlan-7995398	interface-name:veth65193d94,pod-template-hash:73...	kube-p1 kubernetes kube-default
orders-db-372...	pi-veth00740d...	11.2.0.226	7A:7D:4D:95:8...	vxlan-7995398	interface-name:veth0074d0bd,pod-template-hash:37...	kube-p1 kubernetes kube-default
payment-3050...	pi-veth8fdb208a	11.2.0.225	96:28:83:86:6...	vxlan-7995398	interface-name:veth8fdb208a,pod-template-hash:305...	kube-p1 kubernetes kube-default
queue-master...	pi-vethdea7049f	11.2.0.231	BA:4B:58:BF:6...	vxlan-7995398	interface-name:vethdea7049f,pod-template-hash:206...	kube-p1 kubernetes kube-default
rabbitmq-2416...	pi-vethf8d9724f	11.2.0.232	72:93:5E:8B:0...	vxlan-7995398	interface-name:vethf8d9724f,pod-template-hash:241...	kube-p1 kubernetes kube-default
shipping-2463...	pi-veth6c1673...	11.2.0.228	36:5B:DB:60:3...	vxlan-7995398	interface-name:veth6c167323,pod-template-hash:24...	kube-p1 kubernetes kube-default
user-1574605...	pi-veth93e31b...	11.2.0.222	32:69:AA:4D:3...	vxlan-7995398	pod-template-hash:1574605338,name:user,interface-...	kube-p1 kubernetes kube-default
user-db-3152...	pi-veth1c3917...	11.2.0.229	C2:BA:6B:91:C...	vxlan-7995398	interface-name:veth1c39177d,pod-template-hash:31...	kube-p1 kubernetes kube-default

Fabric Administrator has inventory of Kubernetes objects – simplify operations

Fabric admin can search APIC for k8s nodes, masters, pods, services ...

The screenshot displays the Cisco APIC Fabric Administrator interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'VM Networking', 'L4-L7 Services', 'Admin', 'Operations', and 'Apps'. The 'Inventory' section is active, showing a tree view on the left with 'kubernetes' expanded to 'Nodes', where 'fab10-compute-2' is selected. The main panel shows the 'Hypervisor - fab10-compute-2' details, including a '100' health indicator and a 'Properties' section with fields for Name, Hostname, IP Address, and Status. Below this is a table of pods with columns for Name, Interface, IP, MAC, Encap, Namespace, Labels, and EPG.

Pods:	Name	Interface	IP	MAC	Encap	Namespace	Labels	EPG
kubedn...	pi-veth6ca705f9		10.1.2.2	B2:5C:...	vxlan-7634944	kube-system	kubernetes.io/cluster-service...	k8s k8s-app kube...
kubedn...	pi-vethf97eca53		10.1.2.4	0A:95:...	vxlan-7634944	kube-system	interface-name:vethf97eca53...	k8s k8s-app kube...

APIC keeps inventory of pods and their metadata (labels, annotations), deployments, replicasets, etc.

View pods per node, map to encapsulation, physical point in the fabric.

Supported Container Application Platforms

	Baremetal	ESXi	KVM (Openstack)
Open source Kubernetes 1.6 - 1.15	✓	✓	Future
Cisco Container Platform	Future	✓	✓
Openshift 3.6, 3.9, 3.11	✓	✓	✓
Pivotal Cloud Foundry (PCF) 2.1.1*	✓	✓	Future
Docker EE 2.1 (only with Kubernetes and/or Openshift)	✓	✓	Future
Mesosphere	Not currently planned		





Demo 3: ACI Operations



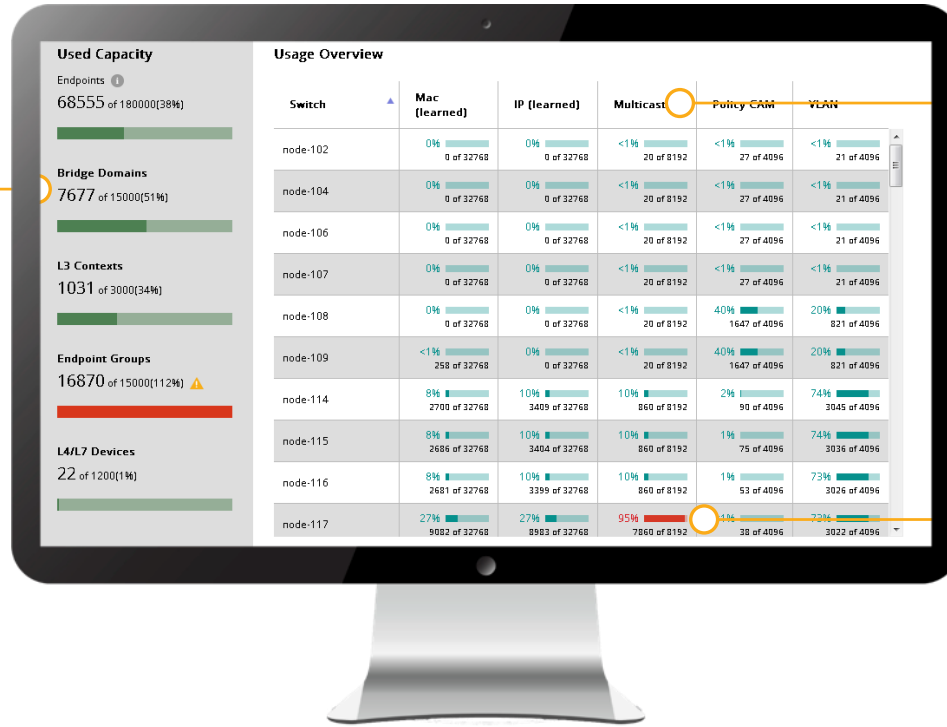
ACI Operations - Agenda

- Before getting started – setting the concepts stage

Visibility	Insights	Actions
Faults, events, stats, health, logs, trails	Application dependency	Incident Troubleshooting
Configuration	Containers Integration	Change Management
Capacity, Fabric metrics (utilization, flows, states, environmental, etc.), Telemetry	Anomalies detection (via SW & HW correlation) Trends	Increase Performance, Availability & Reliability Prevent Outages

ACI Capacity Dashboard

Comprehensive
Management and
Automation



Understand easily ALL
major aspects of
capacity (MAC tables,
Policy CAM etc.)

Instantly identify
capacity issues for the
ENTIRE fabric with
clear visuals


System Info and Environmentals



TCAM



Memory




Power



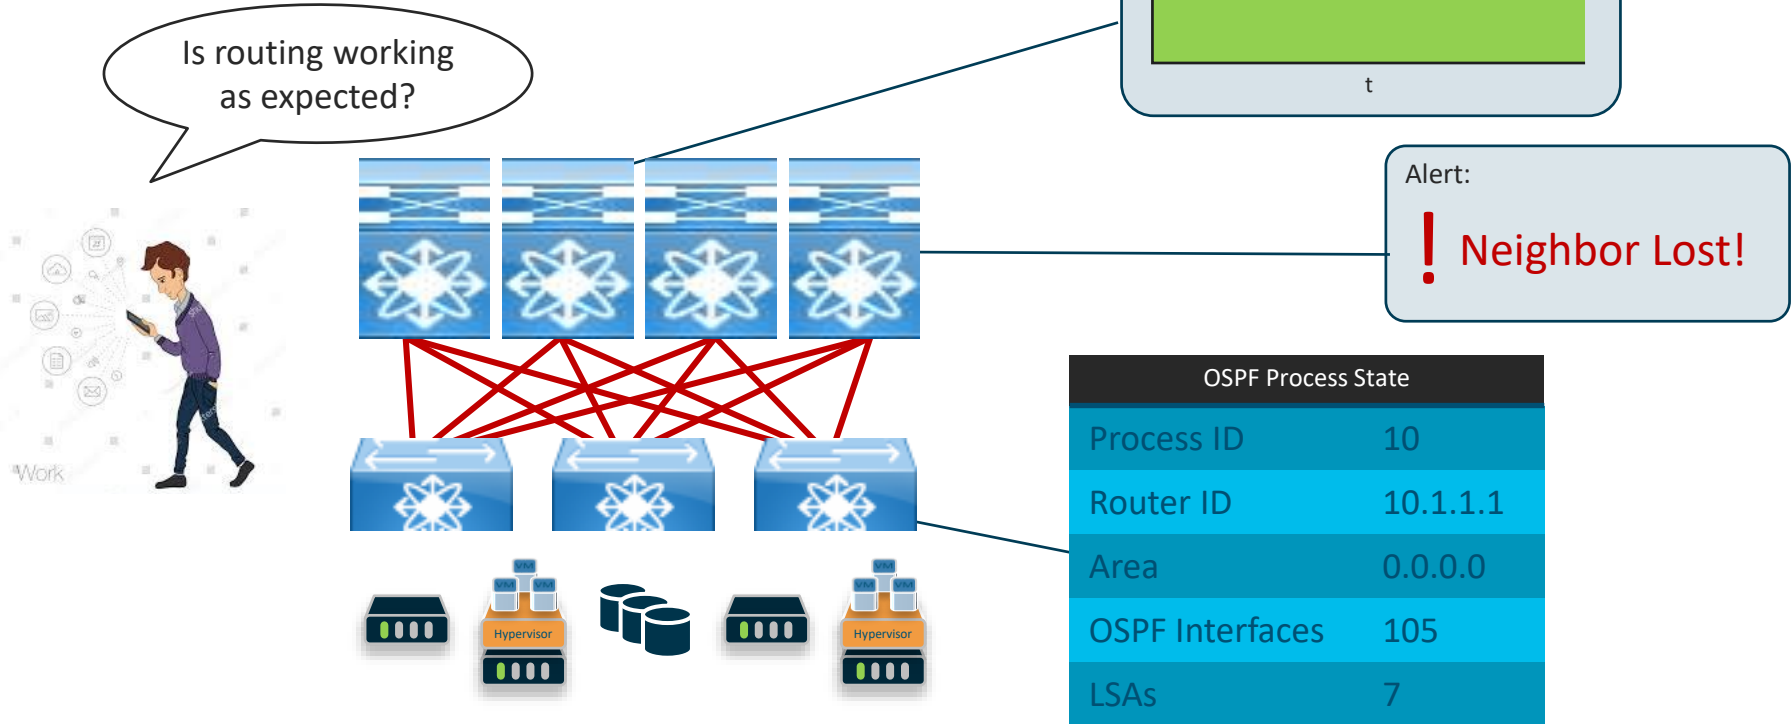
CPU




Temperature



Protocol State and Events




Monitoring Buffer Utilization and Drops

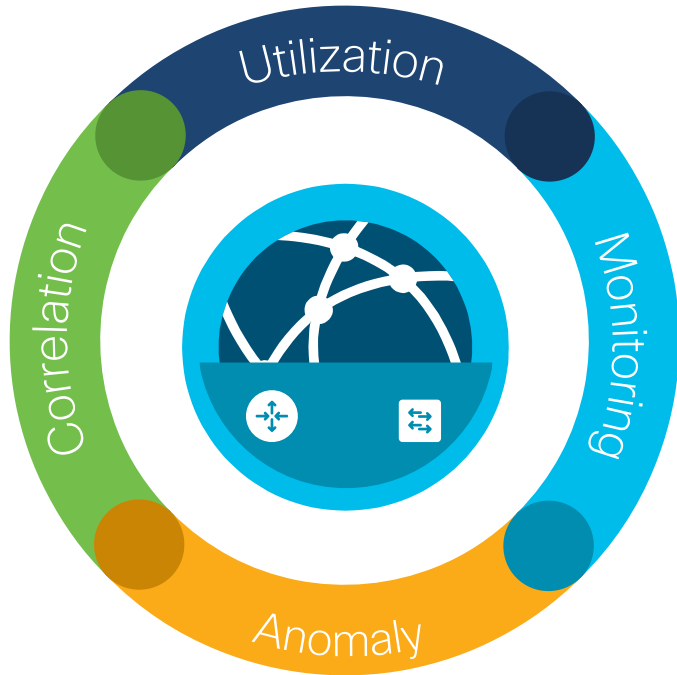


Network Path and Latency Measurement

Application performance is slow between Server A & Server B!






ACI Telemetry to rescue



Telemetry is the only true way of seeing data that represents what the network is experiencing at any point in time...


Key Telemetry Characteristics



Encoding Options

Google Protocol Buffers (GPB)

- Designed for simplicity, performance
- Intended for machine consumption, not human consumption



JavaScript Object Notation (JSON)

- Human-readable, self-describing, text-based encoding format
- Open-standard
- Not designed with performance or extensibility in mind



Transport Options

gRPC

- Modern, open source RPC framework
- Low latency, scalable, distributed
- Enables extension such as authentication, load balancing, logging and monitoring etc.



HTTP

- Ubiquitous transport option
- Many available open source stacks on multiple operating systems
- Well understood in industry



UDP

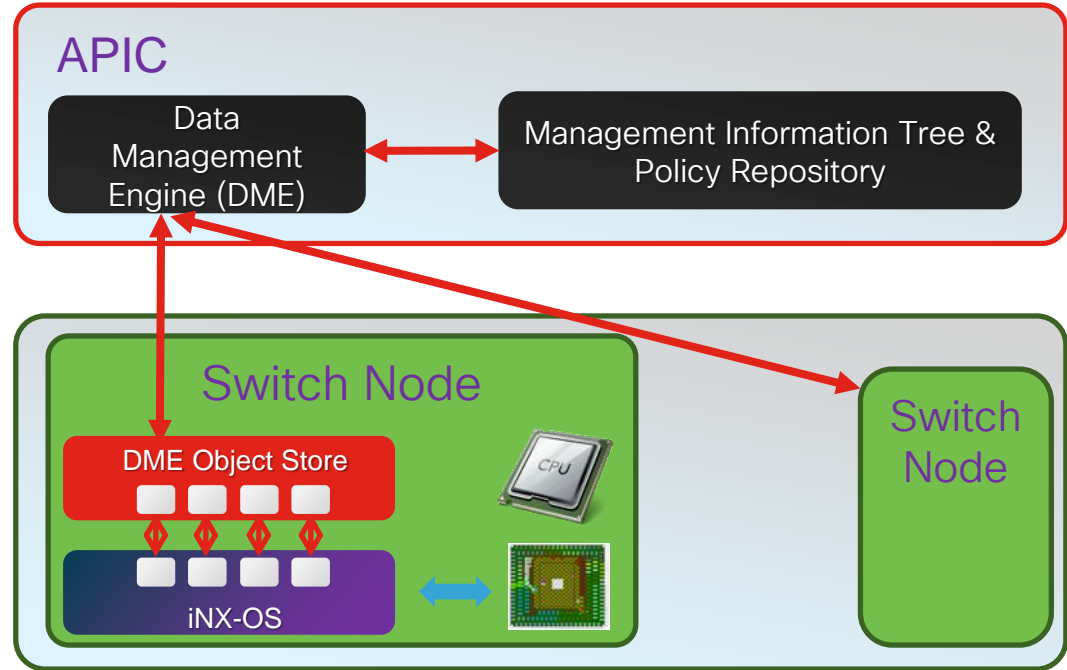
- Connectionless transport
- No upper-layer overhead



ACI Model Based Telemetry

Embedded in the Architecture

- ACI leverages a distributed DME (distributed database)
- Provisioning State is durably stored in fault-resistant manner by keeping multiple in-synch replicas
- Active Hardware and Software state (counters, faults, logs, ...) is distributed via DME as well
- Centralized Reporting and Correlation
- Application-level visibility



Cloud Scale Streaming Hardware Telemetry

Flow Table (FT)

Captures full data-plane packet flow information, plus metadata

Flow Table Events (FTE)

Triggers notifications based on thresholds / criteria met by data-plane packet flows

Streaming Statistics Export (SSX)

Streams ASIC statistics based on user configuration

Data-Plane Flow Data

ASIC State

Available only on [Cisco Cloud Scale](#) platforms!

Telemetry Data Source

Flow Table (FT)

- Captures full data-plane packet flow information, plus metadata
 - 5-tuple flow info
 - Interface/queue info
 - Flow start/stop time
 - Flow latency

Direct hardware export with low flush times (100 milliseconds)

Streaming Statistics Export (SSX)

- Streams statistics and other ASIC-level state data based on user config
 - Interface statistics (packets/bytes/drops)
 - Buffer depth
 - Queue-level microburst stats

Direct hardware export with very low collection intervals (10's of microseconds)

Flow Table Events (FTE)

- Triggers notifications based on thresholds / criteria met by data-plane packet flows
 - 5-tuple flow info
 - Interface/queue info
 - Forwarding drop indication
 - Buffer drop indication
 - Latency/burst threshold indication

Direct hardware export with flow-level and global throttling

9300-FX / 9300-FX2 platforms support triggered flow table events

Flow Table

- Collects full flow information plus metadata
 - 5-tuple flow info
 - Interface/queue info
 - Flow start/stop time
 - Flow latency
- 32K flow table entries per ASIC slice
- Direct hardware export
- EX / FX / FX2 based 9k platforms support hardware flow table

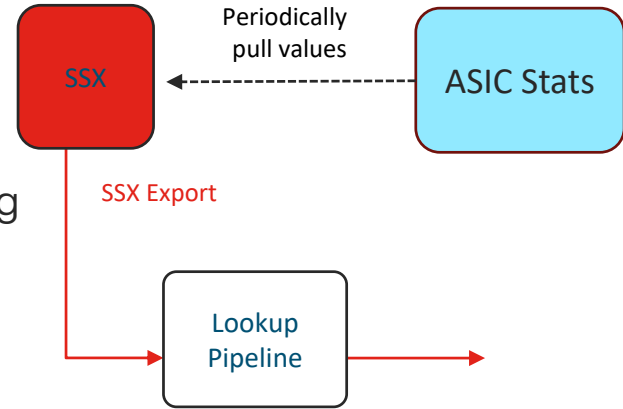


Flow Table Events

- Triggers notifications based on criteria / thresholds met by data-plane packet flows
- Collects full flow information plus metadata
 - 5-tuple flow info with timestamp
 - Interface/queue info
 - Buffer drop indication
 - Forwarding drop, ACL drop, policer drop indication
 - Latency/burst threshold exceeded indication
- Direct hardware export, with flow-level and global throttling
- FX / FX2 based 9k platforms support triggered flow table events



Streaming Statistics Export (SSX)




- Streams ASIC statistics at rapid cadence based on user config
 - Interface counters (packets/bytes/drops)
 - Ingress/Egress queue depth
 - Ingress/Egress queue drops
 - Egress queue microbursts
 - Buffer depth
- User defines streaming parameters – which statistics, how often, and to which collector
- Direct export from ASIC to front-panel port – no switch CPU involvement
- Hardware support in 9364C / 9300-FX2/GX / 9500-FX2/GX

Hardware Telemetry Platform Support



Platform	FT	FTE (roadmap)	SSX (roadmap)
9300/9500-EX	✓	✗	✗
9300/9500-FX	✓	✓	✗
9364C	✗	✗	✓
9300-FX2	✓	✓	✓
9k GX platforms	✓	✓	✓

Next step: Build your own Telemetry Platform ?



Investing in a software development team

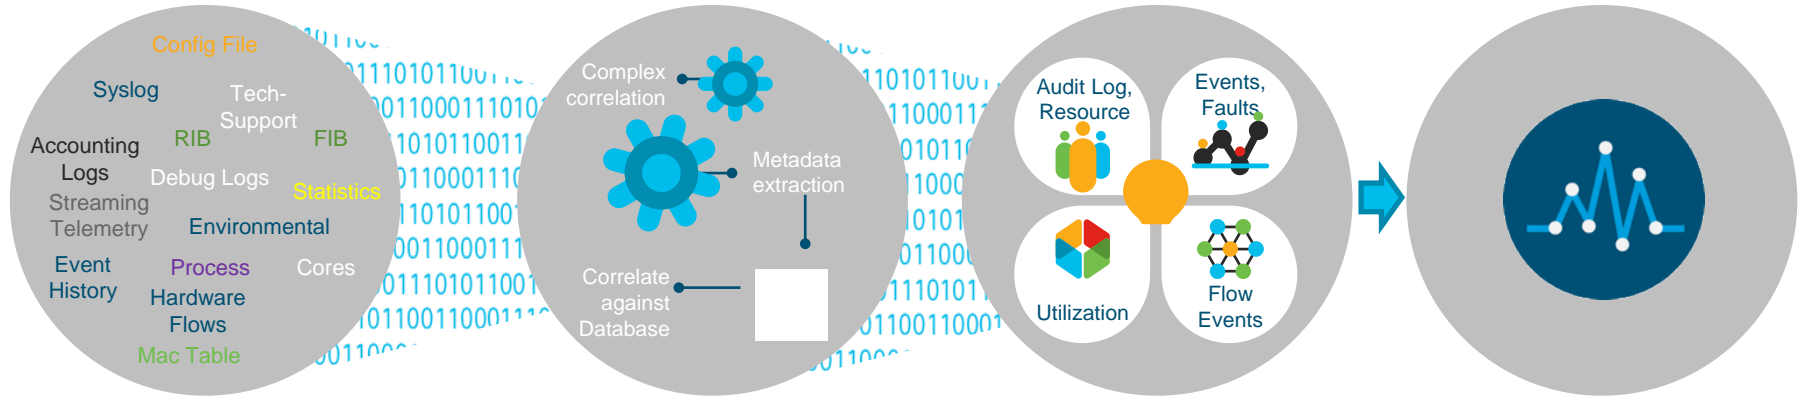
Network Insights Framework – Enabling proactive action

Sources of Telemetry Data

Ingest and Process

Derive Insights

Suggest Action



Increase Availability, Performance and Visibility

Leverage Knowledge Base

ACI | NX-OS

Network Insight Telemetry Applications

Providing Network Health Visibility & Enabling Proactive Insights

ACI App

Network Availability



Advisor

Proactive Software Recommendations/Notifications
Issue Vulnerability Detection & Remediation

Network Health



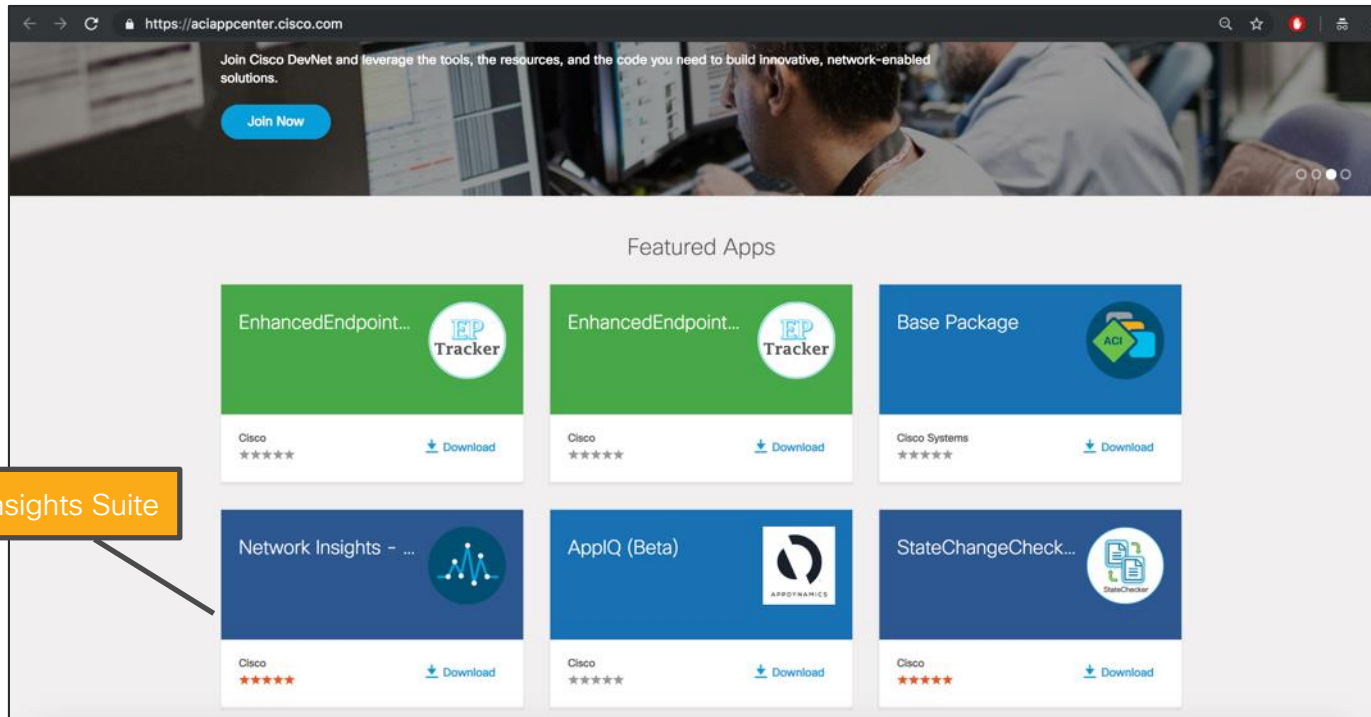
Resource Analysis

Physical/Logical Network Capacity & Utilization
Data & Control Plane & Environmental Health

Enhance Availability, Uptime & Network Wide Visibility

Download application from the App Store

Common app store for ACI and NXOS - <https://aciappcenter.cisco.com/>



Network Insights Suite


Accessing the application from APIC

The screenshot displays the APIC user interface. At the top, the navigation bar includes 'System', 'Tenants', 'Fabric', 'Virtual Networking', 'L4-L7 Services', 'Admin', 'Operations', and 'Apps'. The 'Apps' tab is selected, and a sub-menu shows 'Apps' and 'Facts'. The main content area features a grid of application tiles:

- APIC Postman** by Cisco: A graph based application to create json/xml configuration for post operation to APIC. [Open](#)
- ELAM Assistant (Beta)** by Cisco: Help you perform ELAM(Embedded Logic Analyzer Module) on ACI nodes to capture a single packet at a time and analyze where the packet goes. [Open](#)
- NAE Policy Explorer** by Cisco: Cisco Network Assurance Engine Policy Explorer provides capability to explore policy configuration and connectivity in ACI networks. [Open](#)
- Network Insights - Resources** by Cisco: Network Insights - Resources is a platform for predictive analytics, correlation and alerting using streaming telemetry data for networking fabrics. [Open](#)
- Network Insights Advisor** by Cisco: The Network Insights Advisor (NIA) Application constantly scans your Nexus datacenter environment and provides proactive advice with a focus on maintaining availability and alerting customers about. [Open](#)
- StateChangeChecke** by Cisco: Create and compare snapshots of switch state for interesting objects across all switches in the fabric. [Open](#)

An arrow points from the 'Open' button of the 'NAE Policy Explorer' tile to a callout box containing the text: **Browse to Application in APIC**

Network Insights App Architecture



- GPB over Kafka
- Flow tables over UDP
- Buffer and queue stats over UDP

Collect and Normalize
Telemetry data

Store correlated data for
use by GUI

Analyze

Visualize



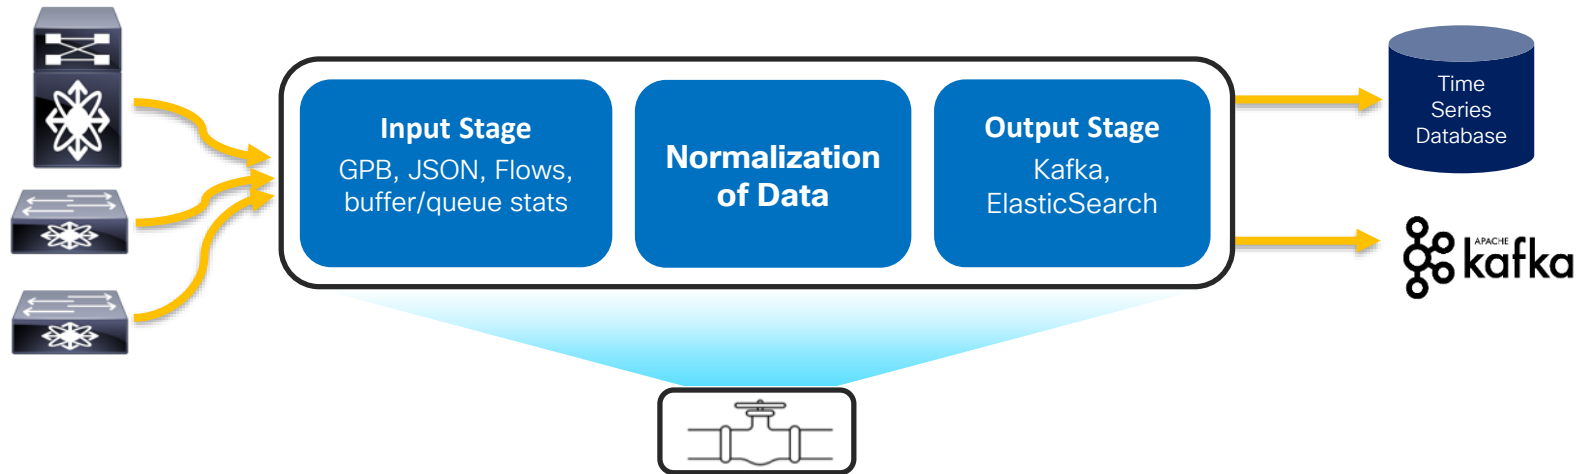
Microservices



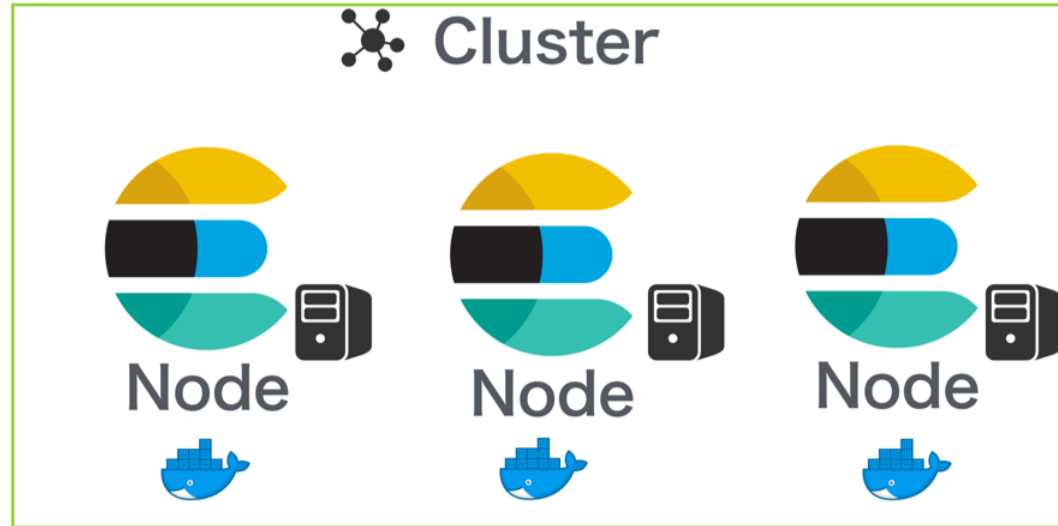
CISCO Live!

Transformation Receiver

- Telemetry Gateway Platform
- Pipeline normalizes disparate telemetry inputs to a common output
- Scalable architecture leveraging container-based scale-out model



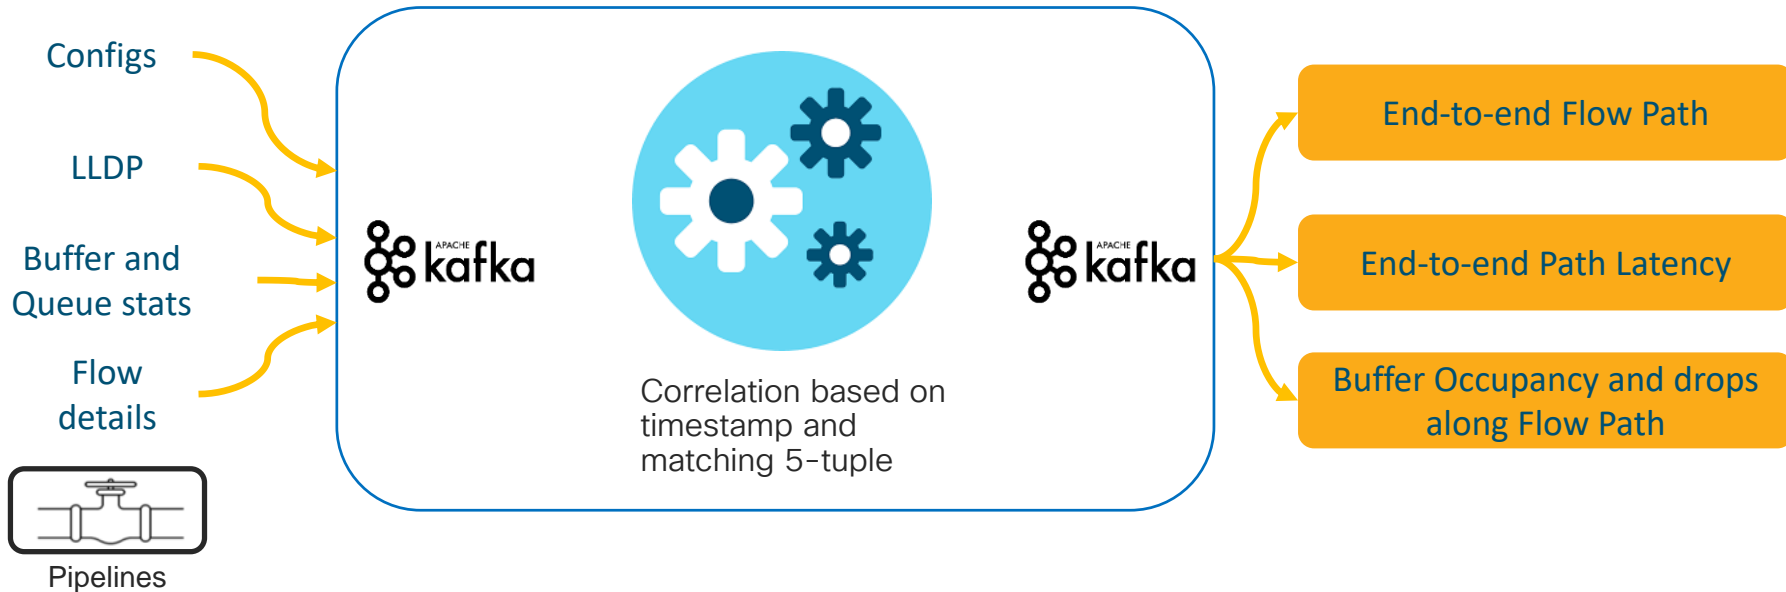
Data Lake



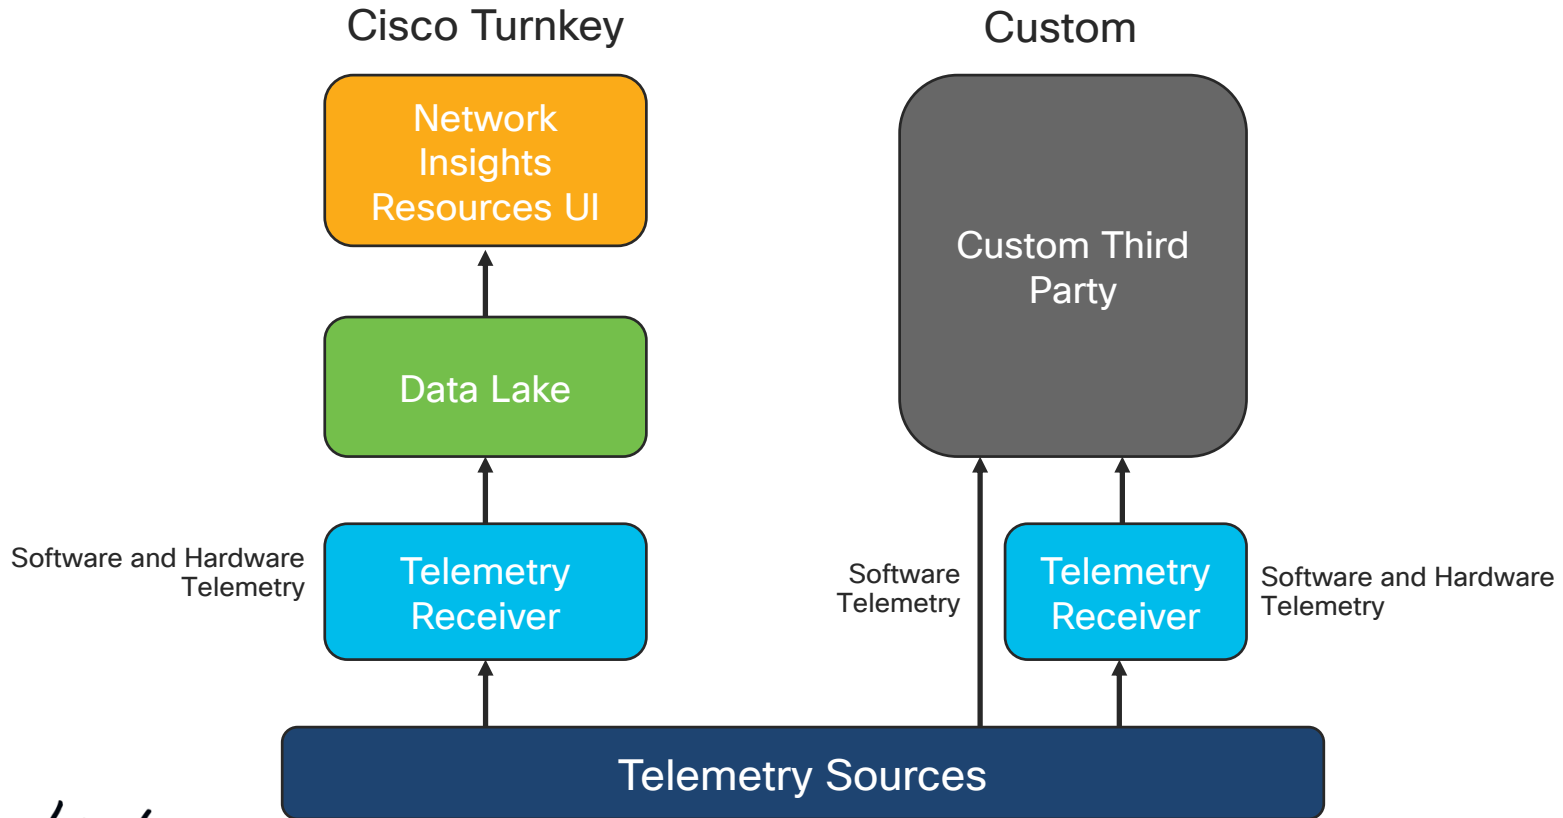
- Horizontally scalable Elasticsearch Cluster used for storing correlated data
- Up to 2 TB of storage to store 3-5 days of flow telemetry (30 days for s/w telemetry)

Correlation Engine


Correlate normalized telemetry data streams from Transformation Receiver




Telemetry and Analytics Deployment Models



Telemetry and Analytics Deployment Options




- 1 Use current APIC and store whatever is available




- 2 Use additional 3 or more Application servers (aka. Service Engines) to host dedicated Telemetry data

Introducing the Services Engine

Application Hosting Platform for APIC / DCNM



+




- Network Insights
- Network Assurance Engine (roadmap)
- 3rd Party Apps (Splunk, F5, etc.)

Dual Boot Option

| Cluster for Redundancy

What Is Services Engine and When Is It Needed?

- Provides additional compute & storage to NIA/NIR for both APIC and DCNM
- Required to support large-scale fabrics, and for flow data collection due to increased ingest rate, correlation workload, and data consumption on disk
- Operates as a cluster of 3 compute nodes
- Can boot in “ACI mode” or “DCNM mode”



2 x 10-core 2.2GHz CPUs

256GB memory

9.6TB storage

Network Insights Resources



- Analysis and correlation of software and hardware telemetry data with focus on **network operations use-cases**
- Integrates directly into APIC with common visualizations
- Focus on **anomalies** and **quick drill-down** to specific issues

Pre-requisites for using Network Insights on APIC

- NTP needs to be configured and working for all nodes to enable software telemetry
- Inband management needs to be setup on both APIC and switch nodes
- PTP needs to be enabled in the fabric. PTP is used to sync the Leaves and Spines for Flow Analytics. GrandMaster clock is not required
- Flow Analytics must be enabled and flow rules must be configured

Integration with External System

NIR has REST-APIs exposed to pull the data available within the UI to interface with 3rd party tools

REST-APIs are available to interact with:

- Anomalies
- Resources
- Events
- Nodes

REST-API examples

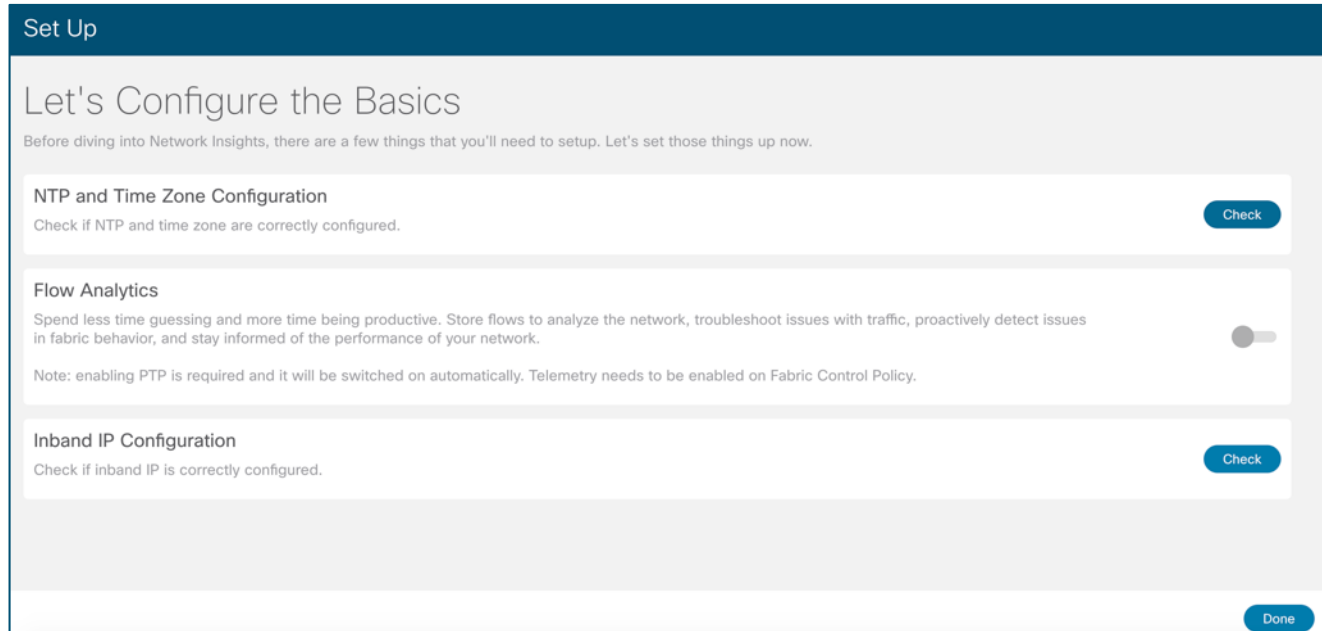
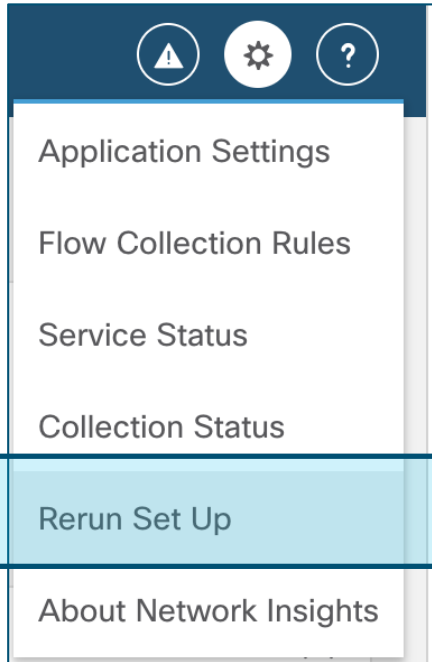
The screenshot displays the 'Network Insights - Resources' application interface. At the top, there is a navigation bar with a 'Time Range' dropdown set to 'Apr 10th 2019, 4:50 PM - Apr 10th 2019, 5:50 PM'. A sidebar on the left contains menu items: 'Dashboard', 'System', 'Resource Utilization', 'Environmental', and 'Operations'. The main content area is titled 'Dashboard' and shows 'Total Nodes' with '0 Spines' and '0 Leaves'. A search bar is located at the top right of the main area.

Three overlapping windows are shown, illustrating the REST API examples:

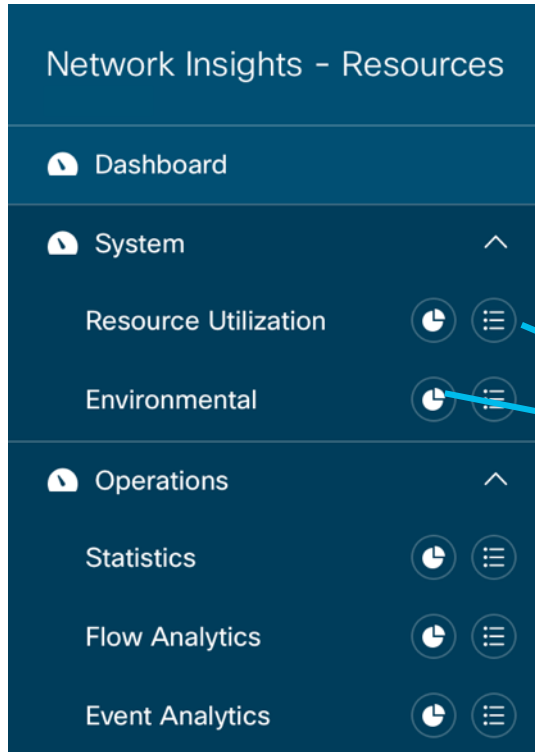
- Window 1 (Left):** Shows the 'Contents' page for 'Network Insight - Resources Application'. The 'NIR REST API Examples' link is circled in blue. Below the list of sections, the 'NIR REST API Examples' link is also circled in blue.
- Window 2 (Middle):** Shows the 'NIR REST API Examples' page. The title 'NIR REST API Examples' is circled in blue. A list of API endpoints is displayed, including 'all_resources()', 'anomalies_details()', 'anomalies_summary()', 'anomalies_top_flows()', 'anomalies_top_nodes()', 'anomalies_top_processes()', 'events_buckets()', 'events_details()', 'events_summary()', 'fabric_summary()', 'get_nodes_list()', 'get_processes_list()', 'get_stats_details()', 'get_stats_nodes()', 'health_diagnostics()', 'service_health()', 'set_logging()', 'symmetrics_detail_stats()', 'symmetrics_get_top_nodes()', 'symmetrics_get_top_processes()', 'symmetrics_summary()', 'utilization_node_details()', and 'utilization_top_nodes()'.
- Window 3 (Right):** Shows the 'all_resources()' API endpoint. The title 'all_resources()' is circled in blue. The REST URL is '/api/telemetry/utilization/resources.json'. The parameters are 'None'. The example curl command is 'curl -k -i -XGET 'https://ip:port/api/telemetry/utilization/resources.json''. The response is a JSON object with 'totalResultsCount': 5, 'totalItemsCount': 5, and an 'entries' array containing one entry with 'categoryName': '' and 'resourceName': 'EndPoints'.

APIC has a wizard to help you get started!

Follow all the steps till the setup is finished



Network Insights Resources - Tabs



Dashboard : Summary of anomalies or unusual behavior

System : Software telemetry data

- Utilization, trends and anomalies around operational, config and hardware resources
- Utilization, trends and anomalies around environmental

Browse view for details

Dashboard view for quick view

Statistics : Interface counters, LLDP, CDP and ISIS errors

Flow analytics : End to end flow details, drop reasons – directed flow monitoring of 10k flows/sec

Event analytics : Audit logs, Events, Faults

Network Insight Resources – What you get

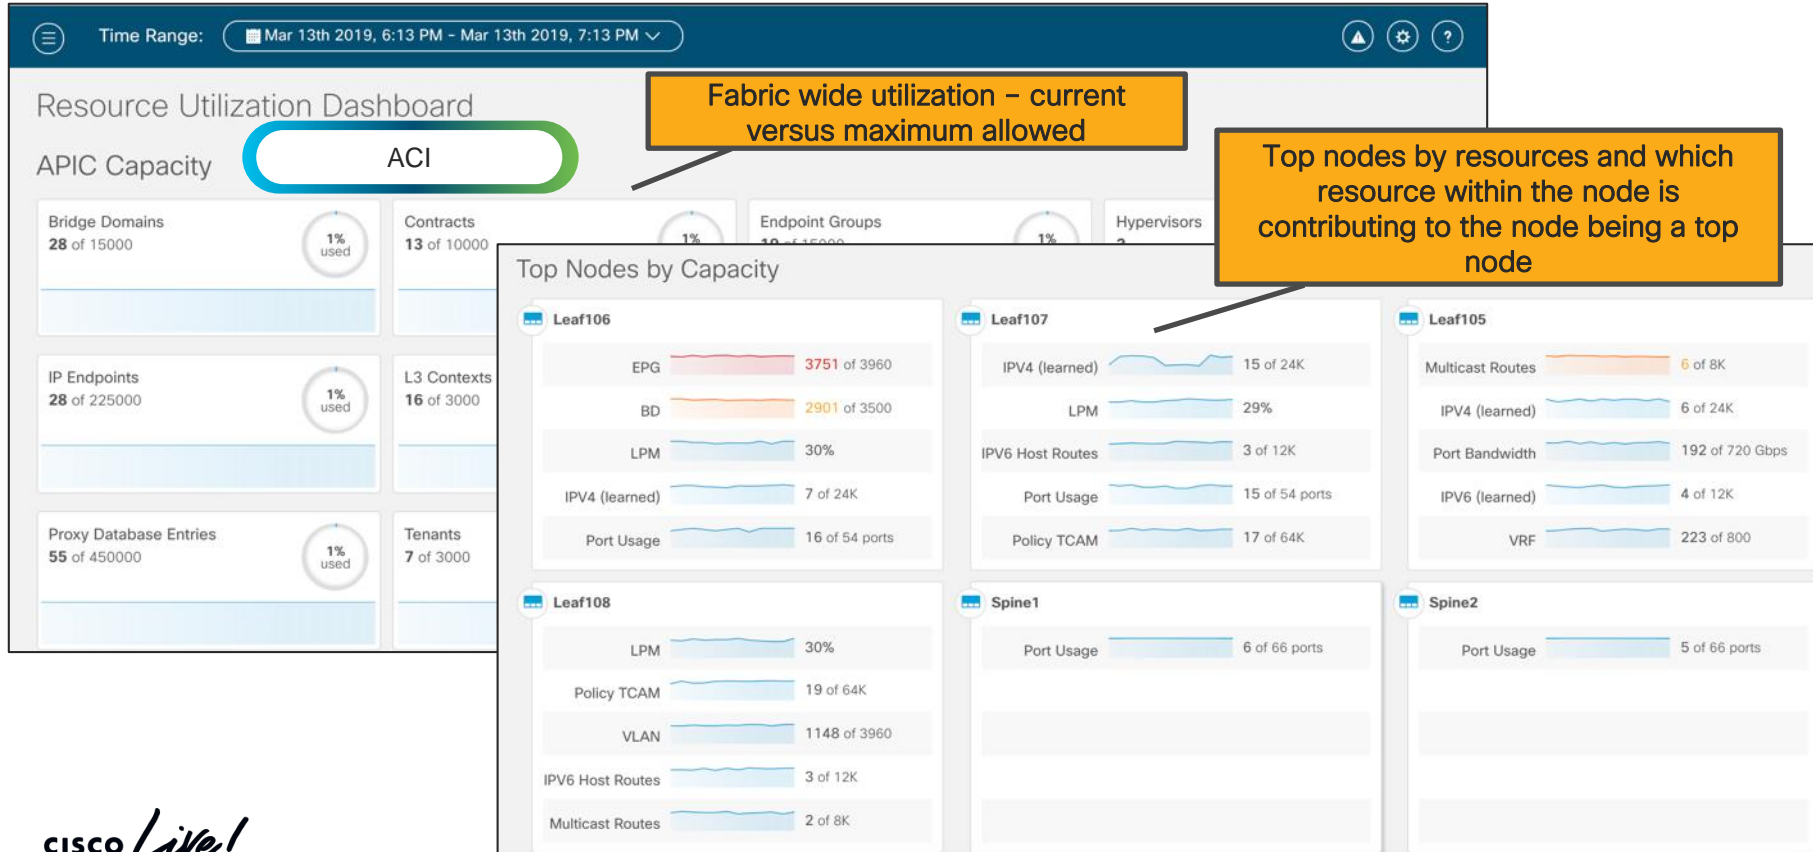
System (shipping)

- Resource Utilization [Fabric Wide]
 - Trend Monitoring (rising/falling), anomaly detection, prediction and alerting
 - Fabric Capacity
 - Endpoints, Bridge Domains, VNIs, Virtualization ratio.
 - Per device stats.
- Environmental
 - CPU Power, Storage, Fan, Memory.
 - Side-by-side analysis

Operations

- Statistics (**shipping**)
 - Protocol Stats (Errors, ingress/egress..)
 - Interface Stats (Utilization, CRC, FW..)
 - Error Trends, anomaly detection, prediction and alerting
 - Bandwidth
- Flow Analytics (shipping)
 - Anomaly detection
 - Path Tracing & Latency
 - MAC moves
- Flow Analytics (**Phase 2**)
 - Microburst Detection
- Event Analytics (**shipping**)
 - Audit and Correlation
 - Hot-Spot / Congestion Monitoring

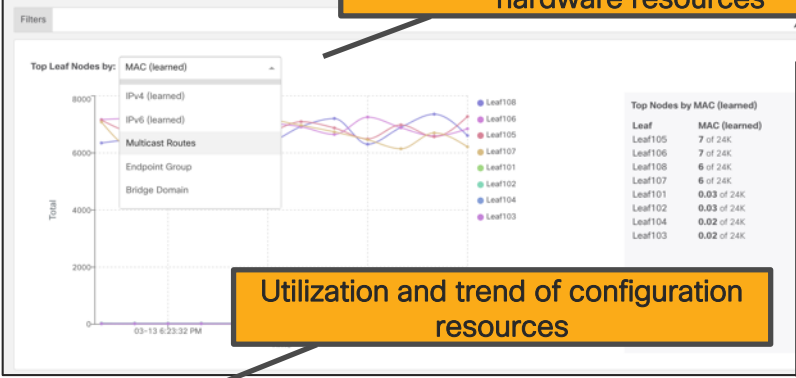
Resource Utilization Dashboard



Resource Utilization Browse View

Anomaly score accounts for anomalies as seen in dashboard

Browse Resource Utilization



Sort nodes by operational, config and hardware resources

Utilization and trend of operational resources

Utilization and trend of configuration resources

Operational Resources		Configuration Resources	Hardware Resources					
Anomaly Score ^	Node	MAC (learned)	IPv4 (learned)	IPv6 (learned)	IPv4 Host Routes	IPv6 Host Routes	Multicast Routes	Policy TCAM
🔴	Leaf106	7 of 24K	7 of 24K	3 of 12K	13 of 48K	3 of 12K	2 of 8K	18 of 64K
🟢	Leaf105	7 of 24K	6 of 24K	4 of 12K	14 of 48K	4 of 12K	6 of 8K	19 of 64K
🟢	Leaf107	6 of 24K	15 of 24K	3 of 12K	13 of 48K	3 of 12K	2 of 8K	17 of 64K
🟡	Leaf108	6 of 24K	7 of 24K	3 of 12K	13 of 48K	3 of 12K	2 of 8K	19 of 64K


Operational Resources		Configuration Resources	Hardware Resources			
Anomaly Score ^	Node	VRF	BD	EPG	VLAN	LPM
🔴	Leaf106	213 of 800	2901 of 3500	3751 of 3960	997 of 3960	322
🟢	Leaf105	223 of 800	1049 of 3500	1013 of 3960	1018 of 3960	
🟢	Leaf107	207 of 800	936 of 3500	1026 of 3960	1028 of 3960	
🟡	Leaf108	214 of 800	933 of 3500	1067 of 3960	1148 of 3960	

Utilization and trend of hardware resources – how many ports are admin up and how much bandwidth is being used

Operational Resources		Configuration Resources	Hardware Resources	
Anomaly Score ^	Node	Port Usage	Port Bandwidth	
🔴	Leaf106	16 of 54 ports	182 of 720 Gbps	
🟢	Leaf105	14 of 54 ports	192 of 720 Gbps	
🟢	Leaf107	15 of 54 ports	190 of 720 Gbps	
🟢	Spine2	5 of 66 ports	0 of 6240 Gbps	

On double clicking a line item..

Applicable to all the resources




Takes you to the Resource Trends page

The icon represents a time graph

All resources and corresponding anomalies over time can be observed and correlated on this page

Trends across resources can be seen here

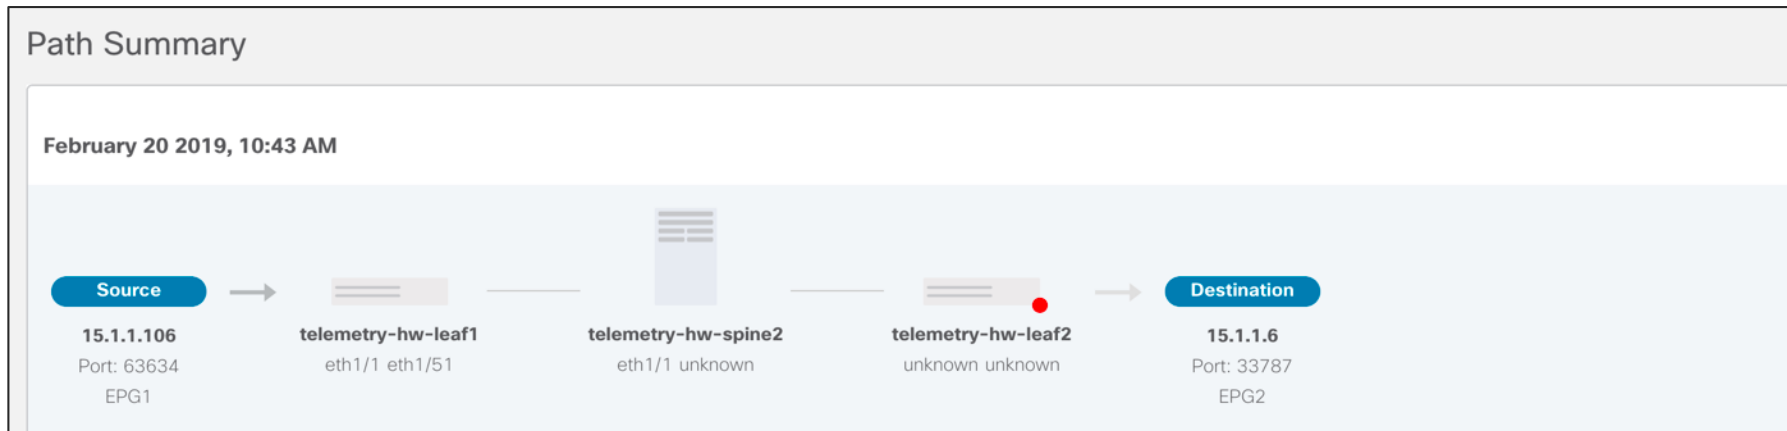


Resource Analysis – Flow Analytics

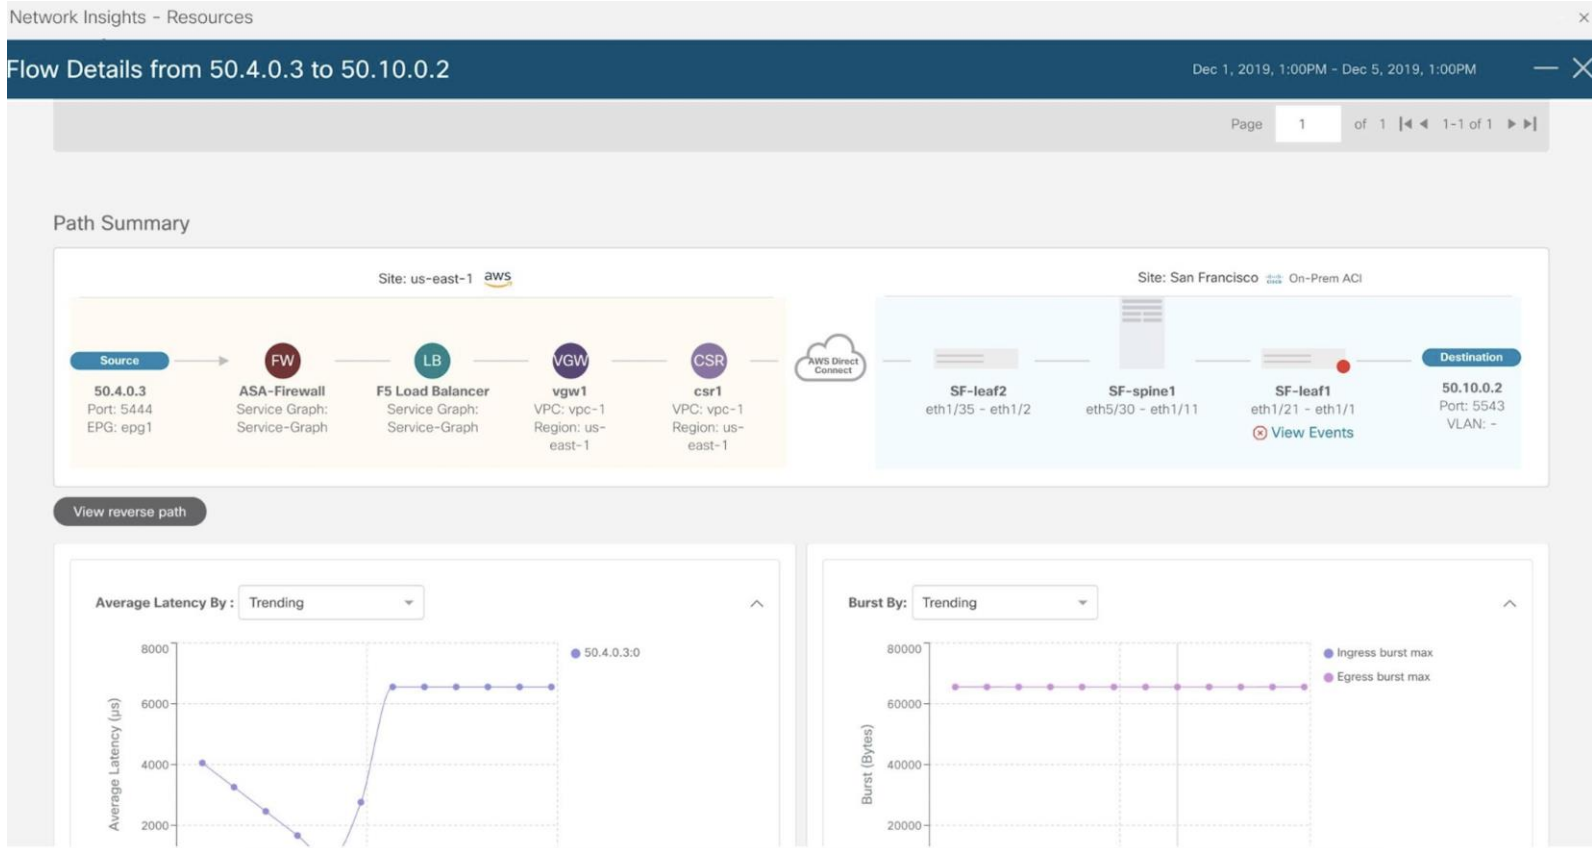
Proactive Anomaly Detection for ACI Deployments

Targeted Flow Monitoring Use Cases –

- Application Performance Issues:
 - Forwarding/policy Drops indicating congestion
 - High end to end application latency
- Application Downtime Event –
 - Policy misconfiguration due to ACL's

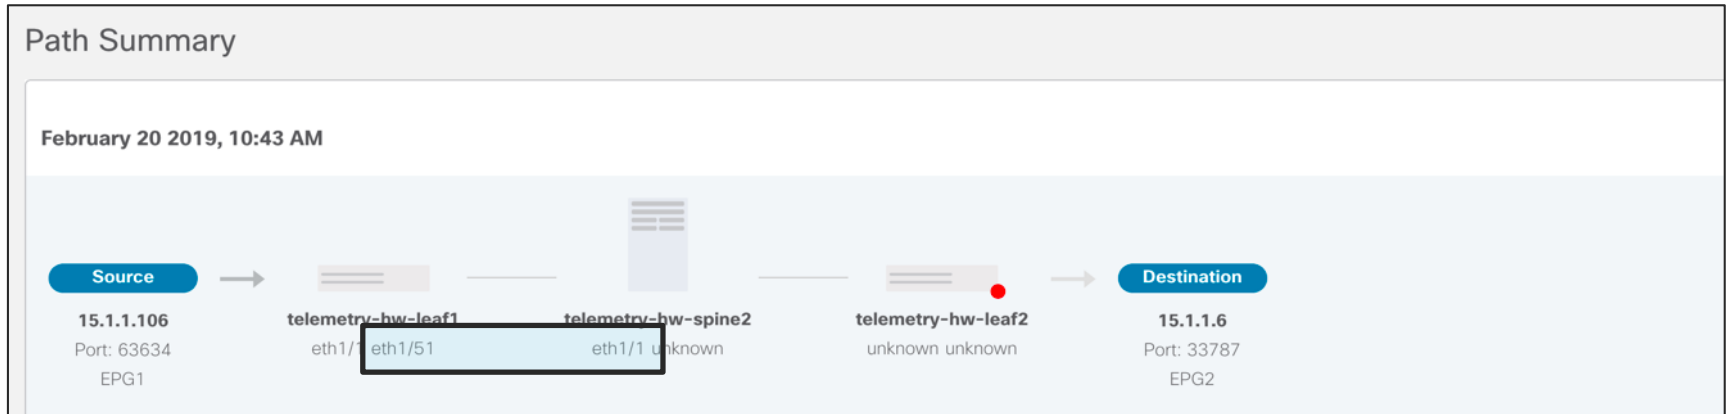


The intent - Network Insights and Cloud ACI



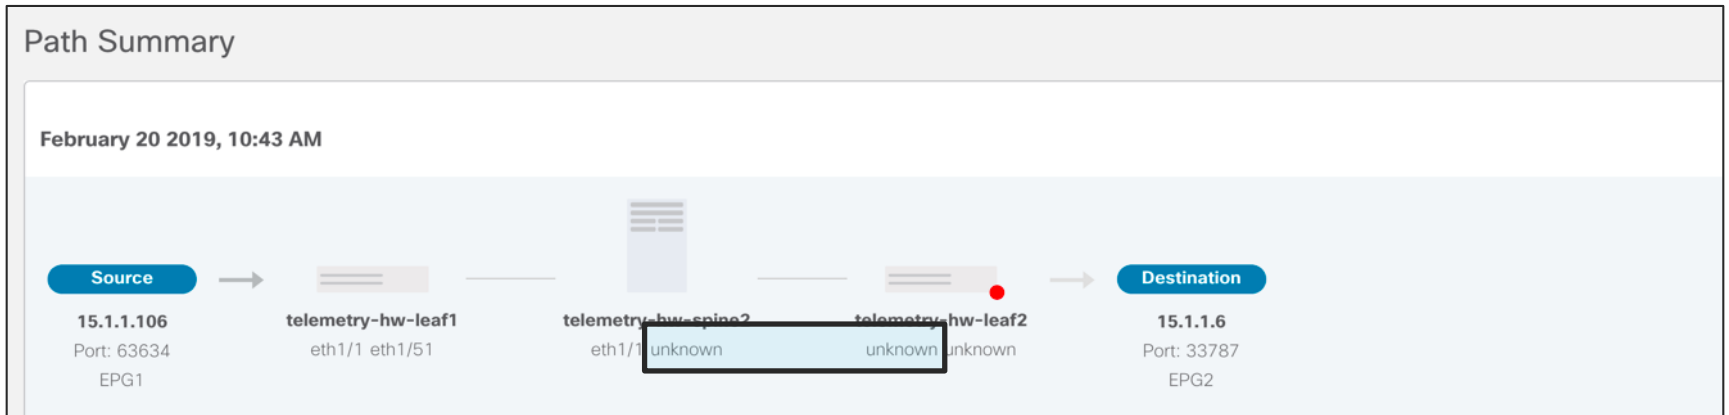
When Spine doesn't support FT

Even if Spine doesn't support FT (for example Nexus 9364 as Spine), Ingress Interface is gathered via LLDP between Ingress Leaf and Spine. The flow will look like the below when exported to NIR -



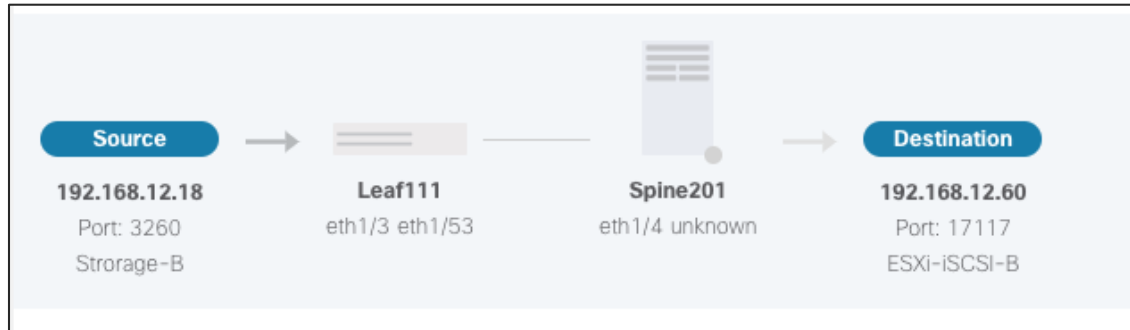
Limitation of FX/FX2

FX/FX2 can not determine the ingress port, so it is expected to see 'unknown' ingress port on egress leaf. The flow will look like the below when exported to NIR -



Mixed Fabric with FT and non-FT capable Devices

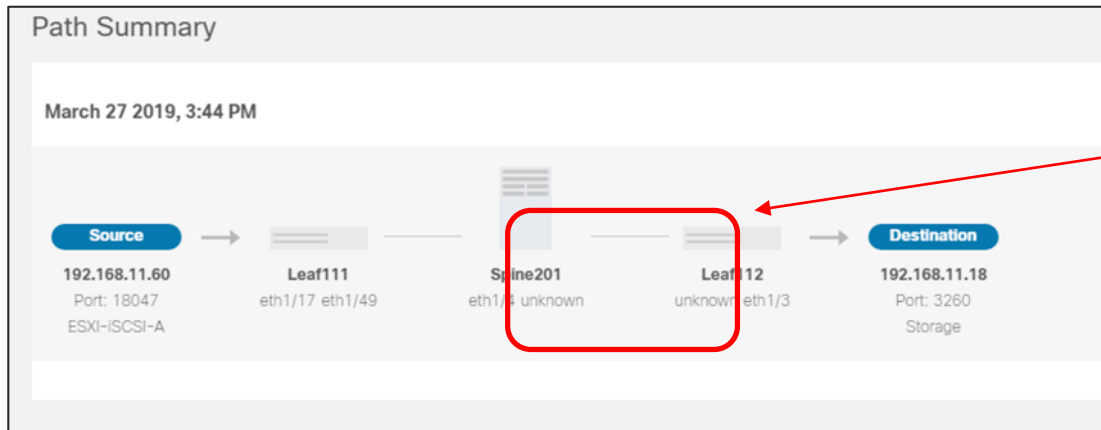
EP 192.168.12.60 is attached to non-FT capable leaf. The flow will look like the below when exported to NIR -



Multipod Support

Multipod is supported. Following limitations and guidelines apply:

- IPN is not visible in path trace
- Spine egress interface to IPN is only visible if spine is FT capable
- Spine in the other POD is only visible in pathtrace if it is FT capable



Both spines are not FT capable in this MPOD scenario. So no egress interface information nor spine in other POD is visible

Directed flow monitoring – Supported in Nexus 9K-FX/EX and FX2 in NIR 2.1

- Phase1 will support 10k flows/sec across the fabric
- User must define flow rules in APIC after which only those flows will be pushed to NIR for flow analysis

User defined flow rules

ACI


Settings -> Flow Collection Rules

Time Range: Mar 10th 2019, 12:40 PM - Mar 17th 2019, 12:40 PM

Flow Analytics Dashboard (Limited Availability)

- Application Settings
- Flow Collection Rules
- Service Status

Top 10 Nodes by Average Latency



VRF Based Rules

Add a rule

Flow Collection Rules

Define parameters - Name, Tenant, VRF, Subnets

VRF Based Rules

Name	Tenant	VRF	
Demo			+
NewTestRule	mgmt	inb	🗑
vEPC	vEPC	VRF	🗑

Demo

Rule will be applied to all relevant switches.

Subnets

Subnet: 10.10.10.0/24

Add Subnet

Cancel

Save

Flow Analytics Dashboard

Each chart is clickable to drill down on details



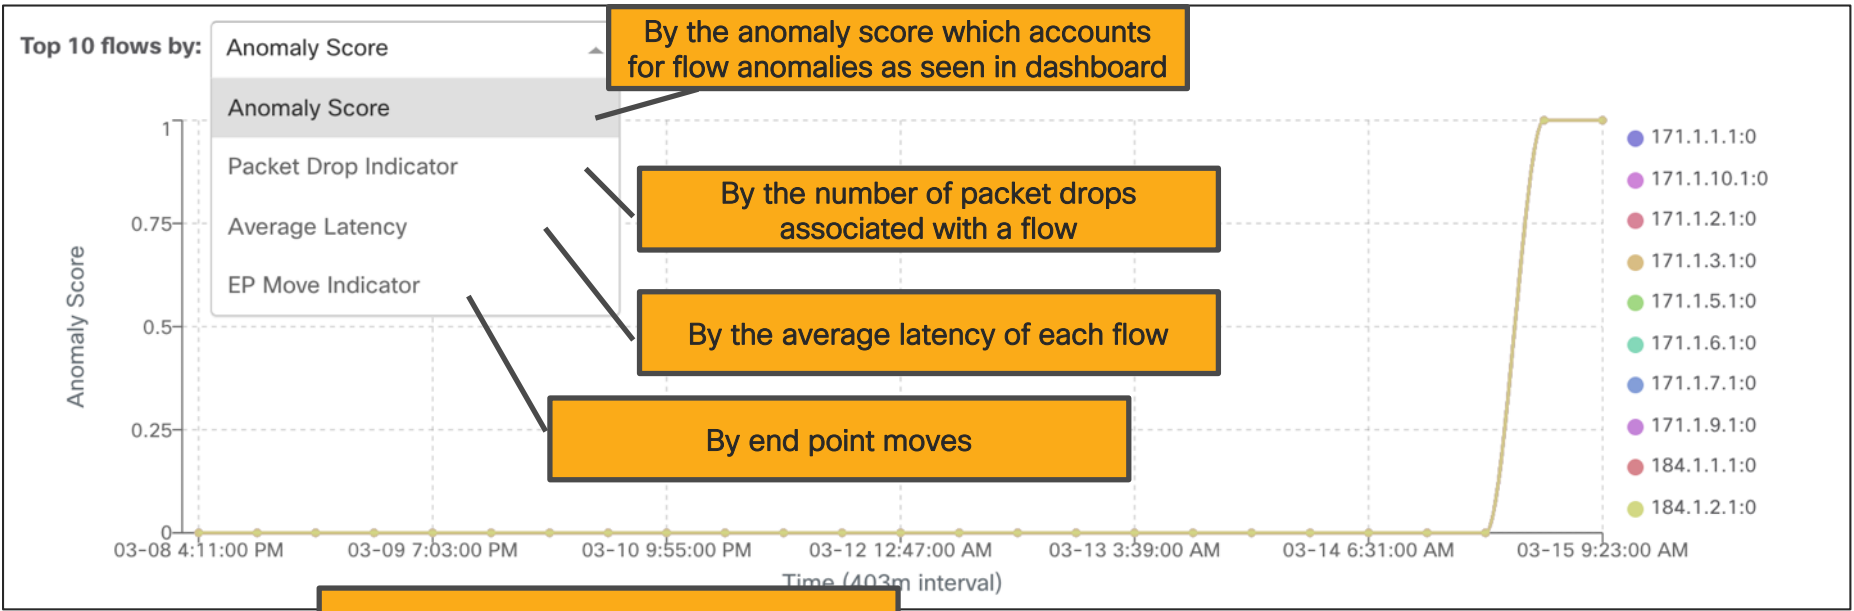
Top 10 nodes by Average Latency of flows - useful to detect unusual latency

Top 10 nodes by packet drops - Quick drill down on packet drops per node

Top 10 nodes by End Point Moves - track EP moves

Flow Analytics Browse View

Sort top 10 flows over a selected period of time




Populates all the flows seen by the app


Anomaly Score	Origination Timestamp ^	Nodes	Ingress Nodes	Egress Nodes	Ingress Tenant	Egress Tenant	Source			Destination	
							EPG	Address	Port	EPG	Address
	Mar 15 2019 11:33:31am	Leaf101, Leaf104, Spine1	Leaf101	Leaf104	vEPC	vEPC	PGW	171.1.1.1	60	Ext-EPG	174.1.1.1

On double clicking a flow in the browse view

Packet Drop Indicator




EP Move Indicator




Related flow details - Packet drops, EP moves, Traffic pattern

Traffic by Packets



Traffic by Bytes

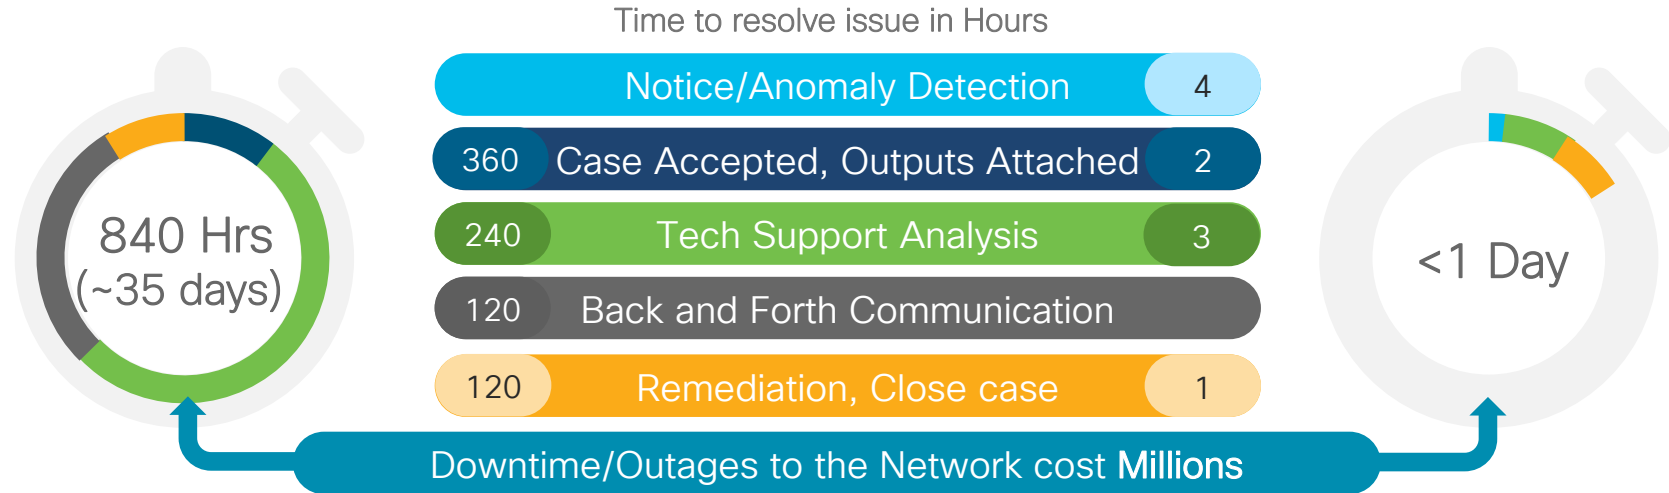


Network Insights Advisor (NIA)

Reduce Downtime/Outages

Before Network Insights Advisor

After Network Insights Advisor



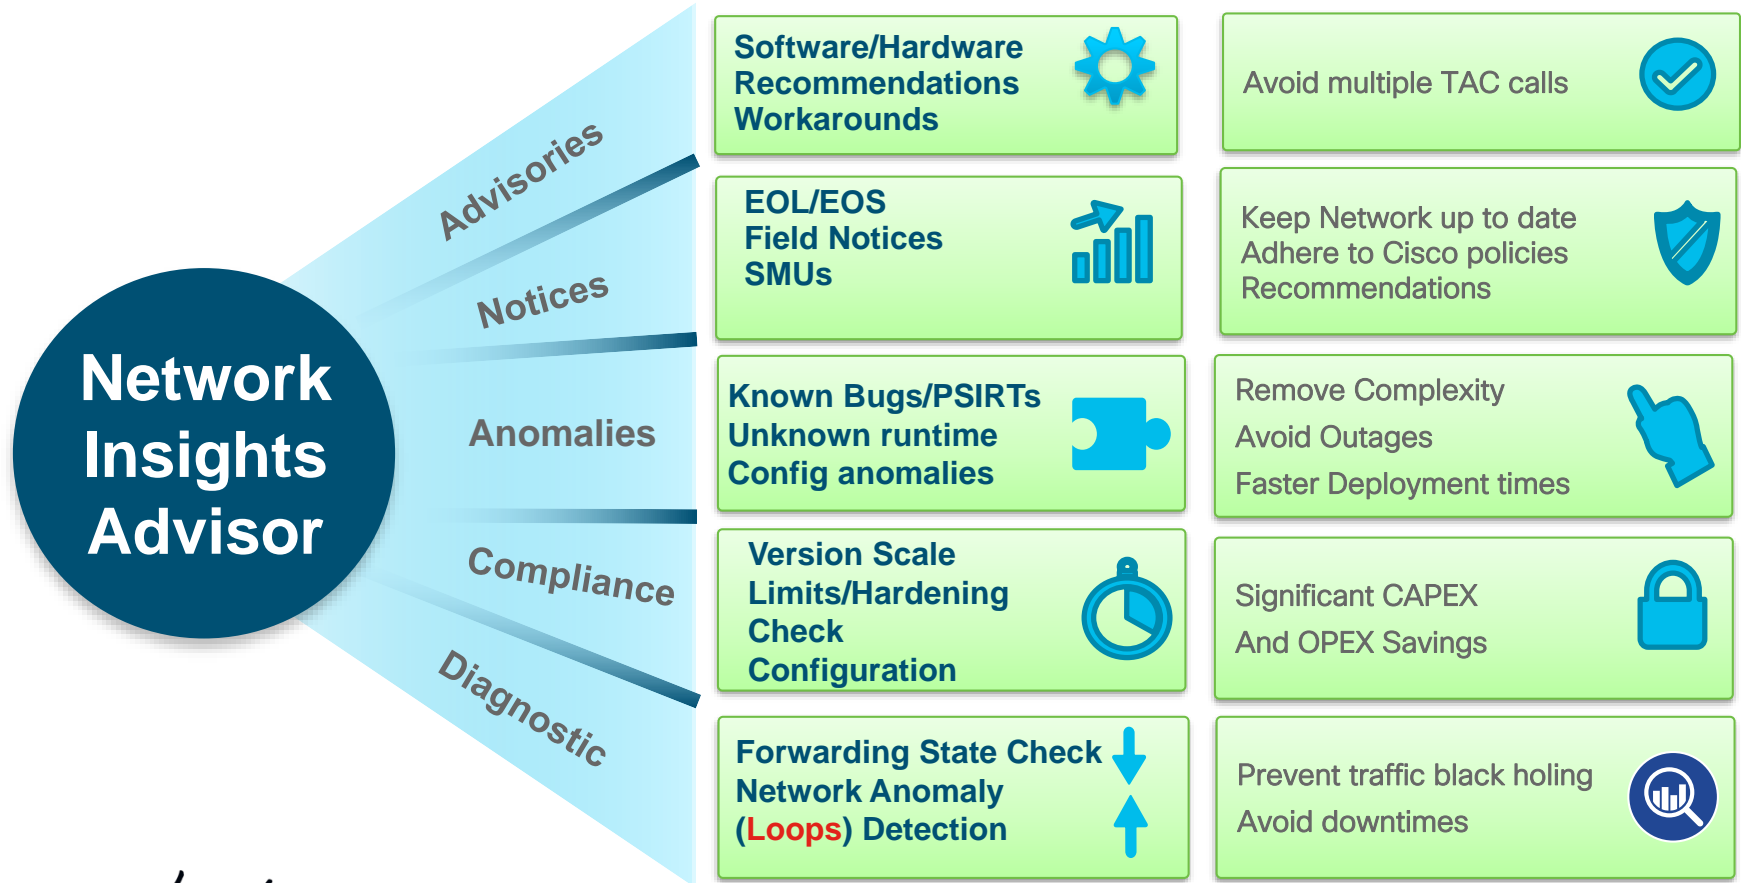
Success metrics

NIA immediately flags anomalies and optimizes your network

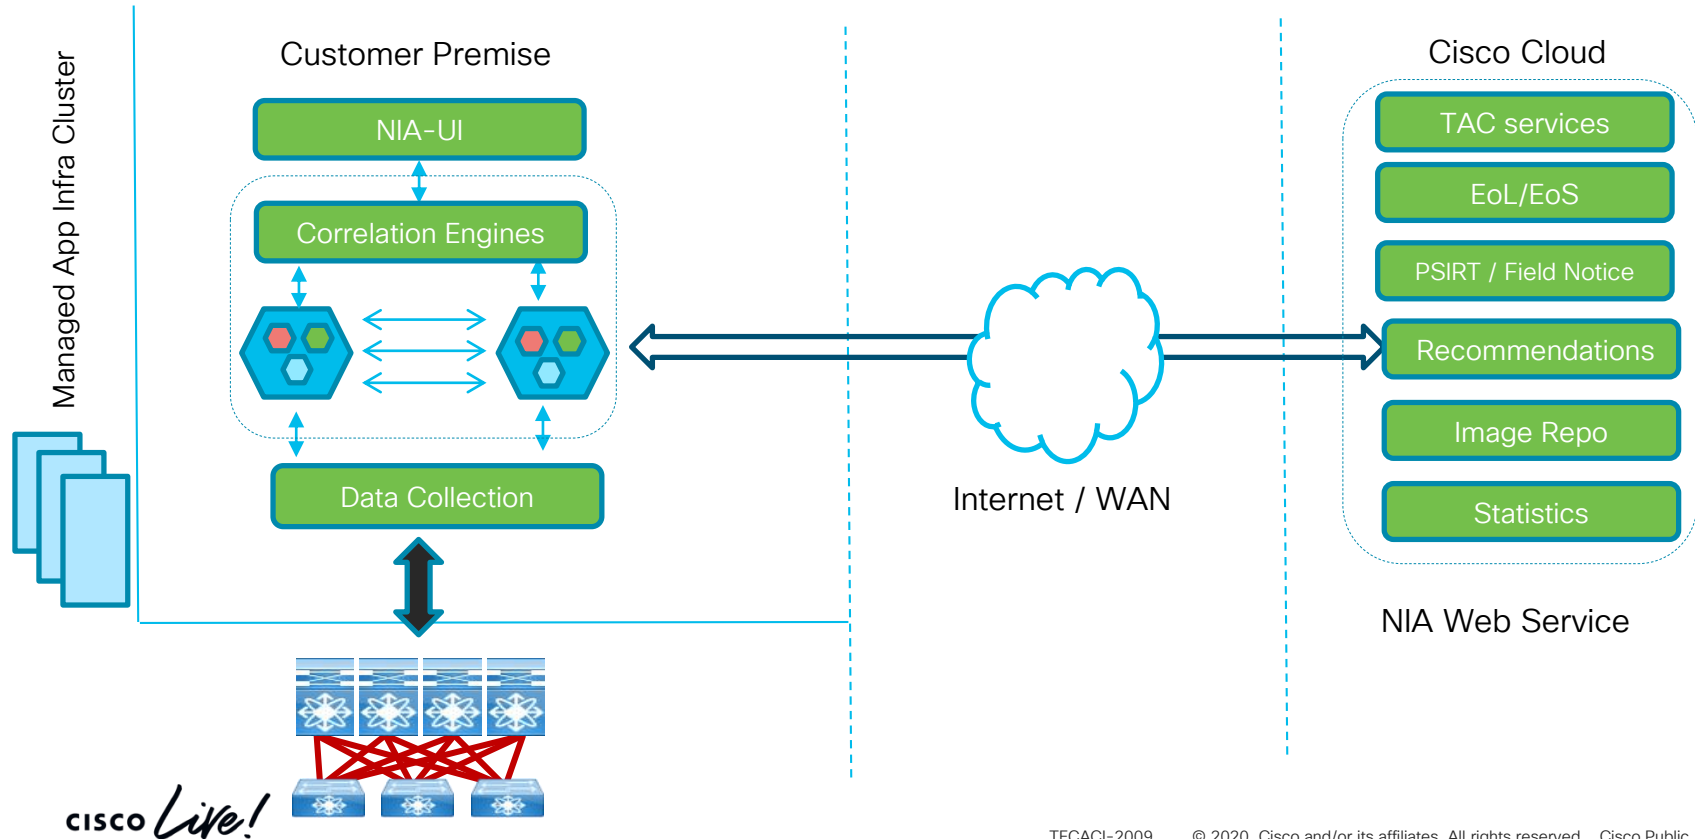
NIA helps prevent downtimes/outages

Significant OPEX, CAPEX and time savings

Network Insights Advisor – Customer Benefits



NIA High Level Architecture

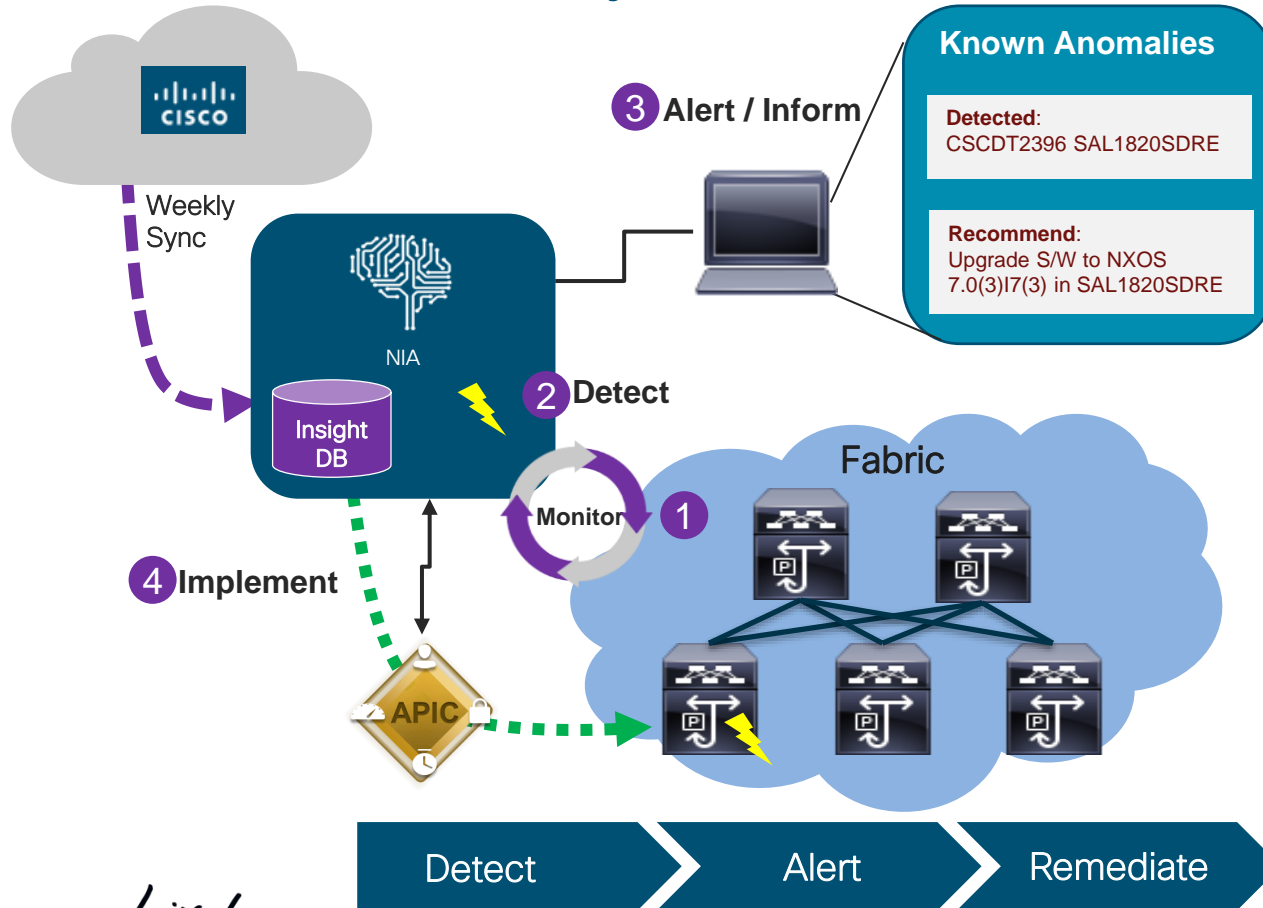


Network Insights Advisor Architecture

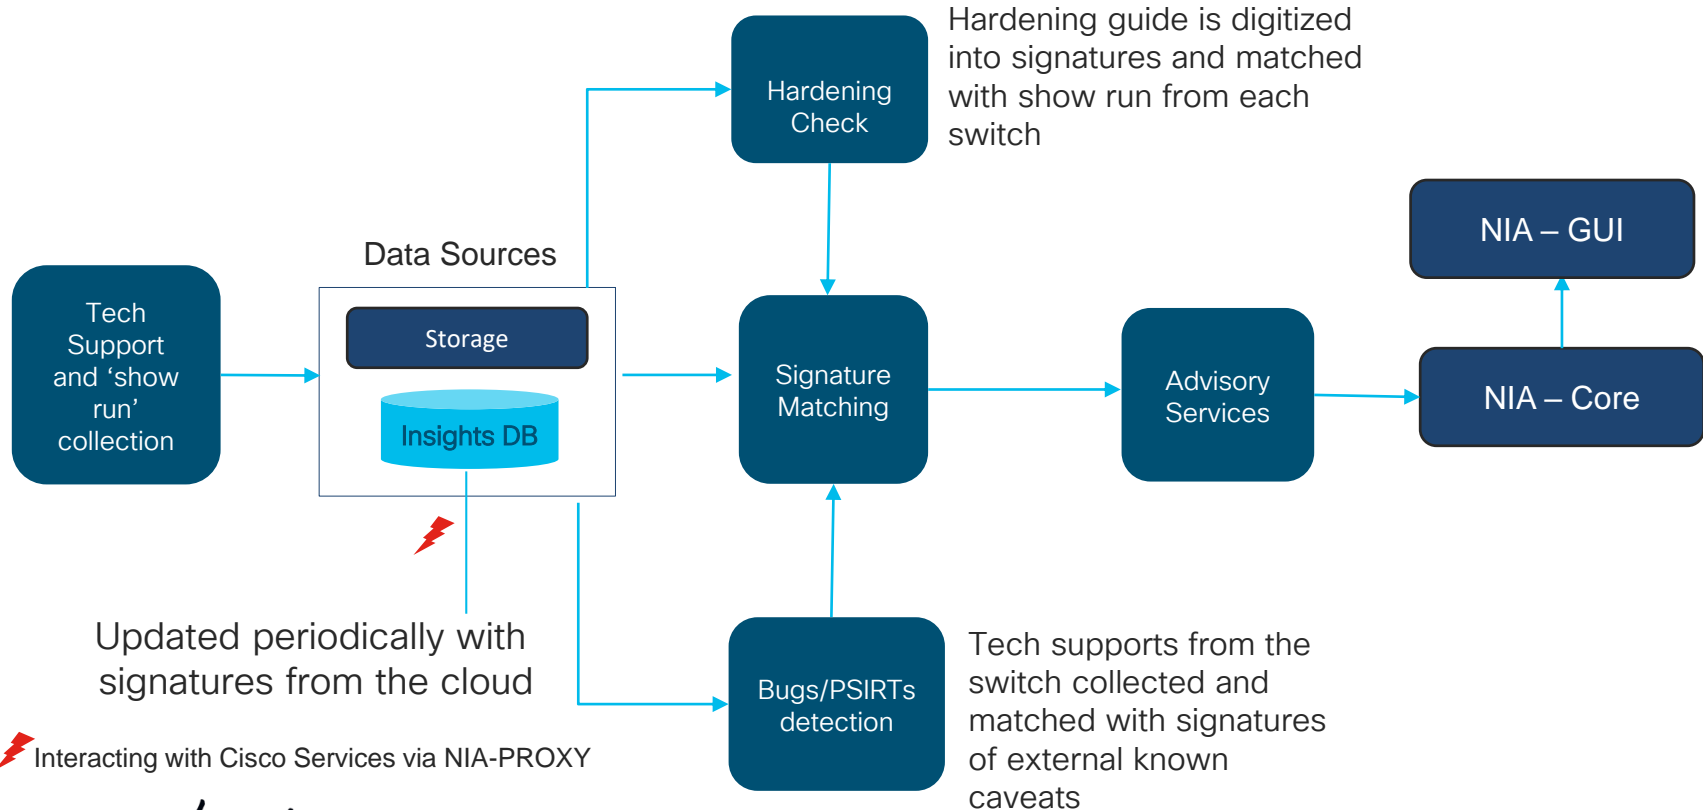
- NIA application resides on-premise within the controller and is based on a micro-services architecture. It collects information relevant information from the switches (tech-support/running config) and processes it for further analysis
- NIA also communicates with Cisco Cloud services periodically to get the latest signatures related to known bugs, field notices, EOL/EOS information on SW and HW
- NIA requires secure internet access from APIC to be able to interface with the cloud services
- From NIA TAC assist, users can collect logs in a timely manner which can be attached to an SR
- All the updates to the app will be available from within the app store and don't require controller or switch upgrade

Note: No user data will be sent to Cloud

Use Case – Notify About Anomalies



Network Insights issue detection



TAC ASSIST

Helps with collecting logs when a user notices abnormal behavior in the fabric. These logs can then be attached to an SR

Network Insights - Advisor

- Dashboard
- Advisories
- Notices
- Issues
- Anomalies
- Bugs
- PSIRTs
- Network
 - TAC Assist**

TAC Assist

Begin the Log Collection

You will be asked to select up to 5 devices to collect logs to assist TAC.

Log Collection

Activity Name

- TAC Assist
- TAC Assist

Collect Logs

Select up to 5 devices to collect logs to assist TAC.

Device	Fabric	Version	IP Address	Platform
<input type="checkbox"/> sj4_4nd_1-n9kv-2	mutate	9.2(1)	192.168.0.125	N9K-9000v
<input type="checkbox"/> sj4_4nd_1-n9kv-1	mutate	7.0(3)I7(1)	192.168.0.124	N9K-9000v
<input type="checkbox"/> sj4_4nd_1-n9kv-3	mutate	7.0(3)I7(1)	192.168.0.120	N9K-9000v
<input type="checkbox"/> sj4_4nd_1-n9kv-4	mutate	9.2(1)	192.168.0.126	N9K-9000v
<input checked="" type="checkbox"/> ACC11_OSLO	mutate	7.0(3)I7(2)	192.168.254.9	N9K-C9372PX

Page 1 of 1 | Objects Per Page: 10 rows | Displaying Objects 1 - 5 of 5

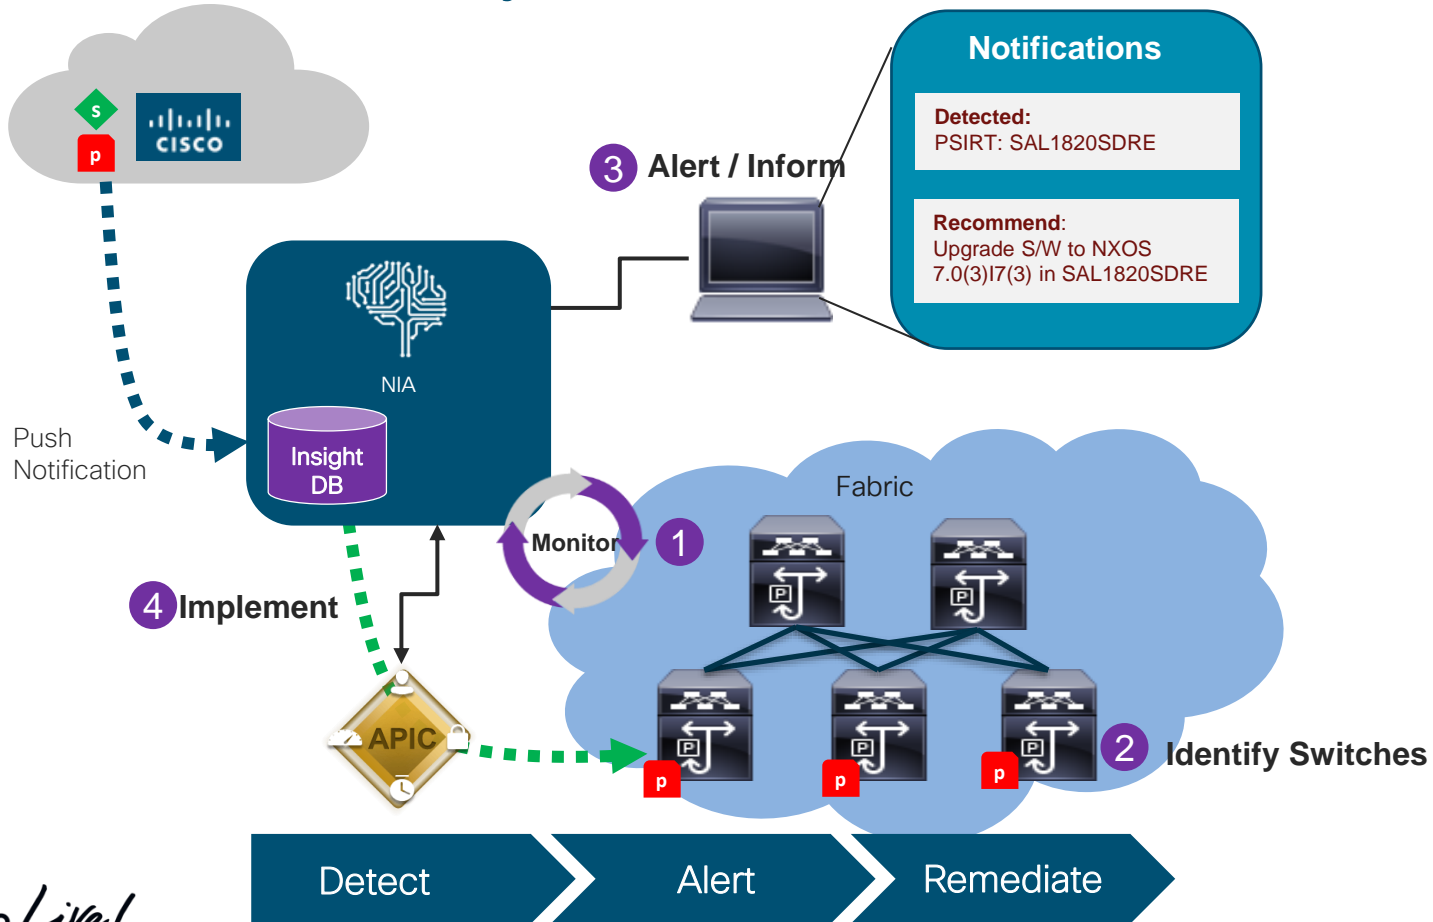
Cancel | **Collect Logs**

Action

- View details
- Stop

Begin

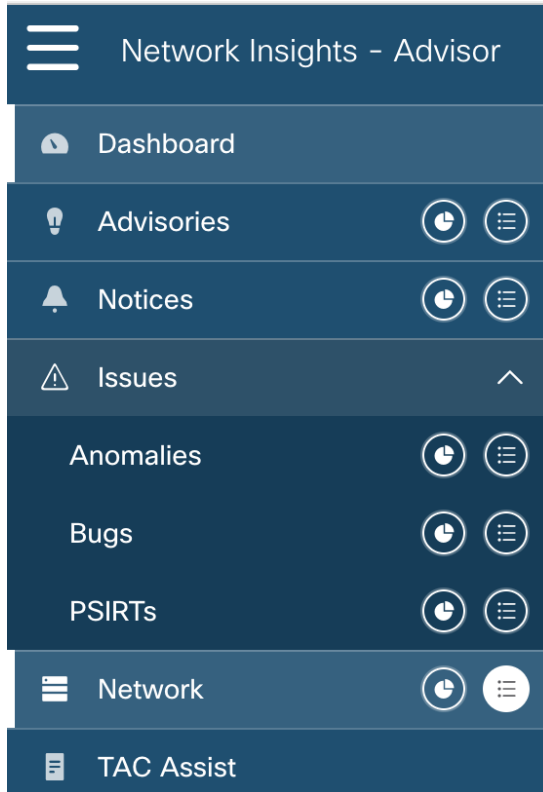
Use Case – Notify Me About New Releases



p PSIRT
s S/W Notify

Network Insights Advisor Targeted Use Cases

Proactive supportability insights



Dashboard "Give me a summary of issues"

Advisories

Provides advisories based on anomalies, bugs, PSIRTs and field notices. Measure upgrade impact

Anomalies

hardening checks, scale checks

Bugs and PSIRTs

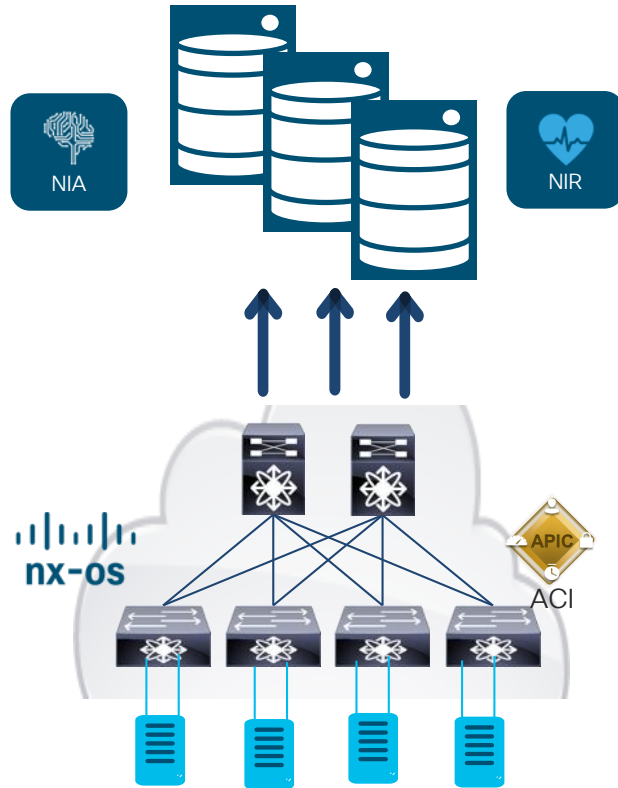
Known bugs and vulnerabilities in the system

Fabric wide analysis



Network Insights Solution Dependencies

Targeted Scale, ACI / DCNM / NX-OS Versions



Compute Footprint*

- Two Options - Node Server Appliance or OVA Cluster
- 3 Compute nodes each with **Memory**: 192GB, 2.4 TB Hard Drive, 400GB SSD, 32 vCPUs

Fabric Scale & Platforms*

- Fabric Scale up to 300 switches target at FCS
- Multi fabric support, up to 300 switches
- Directed Flow Monitoring up to 10k Flows/s
- Streaming intervals: 30sec S/S, 1sec H/W

APIC / DCNM & NX-OS Target Versions*

- APIC / ACI Minimum Release 4.1 onwards
- NX-OS 7.0(3)i7(6) onwards
- DCNM 11.3 (Arگون MR2) onwards

*under evaluation / pre engineering commit

Demo 4: ACI Operations




ACI Operations - Agenda

- Before getting started – setting the concepts stage

Visibility	Insights	Actions
Faults, events, stats, health, logs, trails	Application dependency	Incident Troubleshooting
Configuration	Containers Integration	Change Management
Capacity, Fabric metrics (utilization, flows, states, environmental, etc.), Telemetry	Anomalies detection (via SW & HW correlation) Trends	Increase Performance, Availability & Reliability Prevent Outages
Security Intent, Policy	Network Isolation Network Modeling	Segmentation Assurance

ACI Security

Automated Security With Built In Multi-Tenancy



APIC Hardening - Cent OS 7.2

Line Rate Security Enforcement

Open: Integrate Any Security Device

PCI, FIPS, CC, UC-APL, USG-v6



Embedded
Security



Micro-
Segmentation

Security
Automation

Encryption

Analytics

ACI Security Certifications – by Jun/2018

Certification	ACI
	Certified
	Certified
	Certified
	Certified
Vulnerability Scanners <ul style="list-style-type: none">• Nessus, Fuzzing, etc ...• Port Scan, AppScan	Certified (Ran every release)

ACI Hardening- Every Software Release

Flooding Attacks

SYN-FLOOD: Remain stable during SYN flooding attack

EST-FLOOD: Remain stable during ESTABLISHED flooding attack

LASTACK-FLOOD: Remain stable during LASTACK flooding attack

FINWAIT-FLOOD: Remain stable during FINWAIT flooding attack

CLOSING-FLOOD: Remain stable during CLOSING flooding attack

Port and Service Scans

DEF-CRED: No default authentication credentials

RECON-PORT-TCP: Remain stable during TCP port scan

RECON-PORT-UDP: Remain stable during UDP port scan

RECON-OSID: Remain stable during OS Fingerprinting

RECON-IP-PROT: Remain stable during IP protocol scan

NESSUS-SCAN: Known vulnerability scanner- Nessus

WEB-DEFECT: Known webserver and application defects

WEB-ID: Remain stable during web fingerprinting

Fuzzing

ESIC: UUT must endure malformed Ethernet packets

ICMPSIC: UUT must endure malformed ICMP packets

ISIC: UUT must endure malformed IPv4 packets

TCPSIC: UUT must endure malformed TCP packets

UDPSIC: UUT must endure malformed UDP packets

ICMPSIC6: UUT must endure malformed ICMPv6 packets

ISIC6: UUT must endure malformed IPv6 packets

TCPSIC6: UUT must endure malformed TCP over IPv6 packets

UDPSIC6: UUT must endure malformed UDP over IPv6 packets

Web Scan

Nexpose

IBM AppScan

OpenVas

Platform

Hardening

APIC+N9k

✓

ACI Multisite

✓

AVE

✓

Telemetry

✓

ACI 1st Hop Security Enhancements

Since ACI 3.2

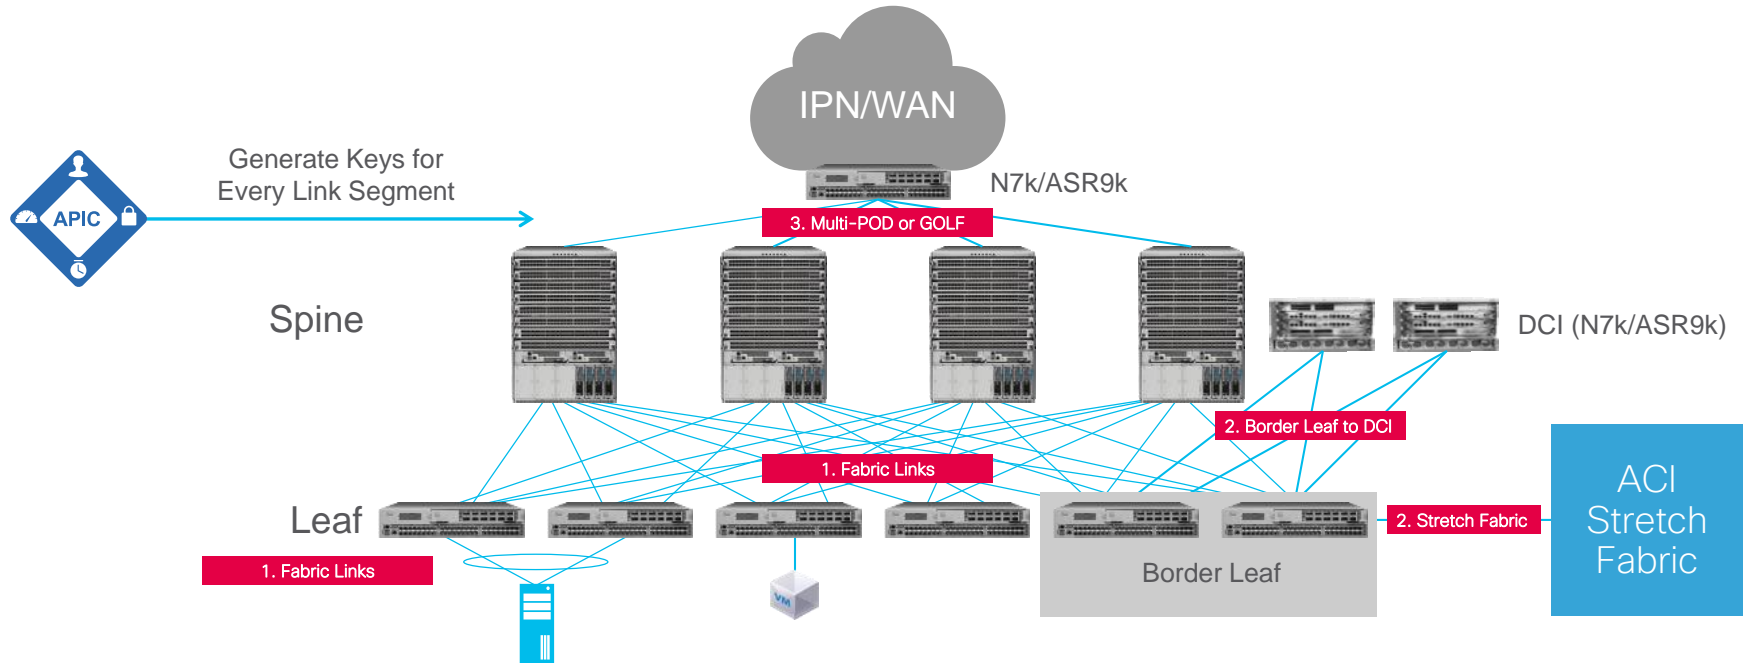
- Theft detection for VXLAN based endpoints
- Coexistence with multiple VPC interfaces
- IP Inspect feature support for IPv4 Local Endpoints
- IP Source Guard feature support for IPv4 Local Endpoints
- IPv4 Static endpoint configuration push to Cisco AVS

Caveats

- All FHS are always enabled for v4 and v6 together, HW cannot support each address family v4 or v6 separately
- FHS disables HW learning of endpoints in all ASIC, hence both v4 and v6 must be enabled together
- IP-Source Guard is enabled for v4 and v6 together (EX/FX leaf switches)
- In future ASIC, we will support FHS enablement for v4 or v6 separately

ACI MACSec

Since ACI 3.1

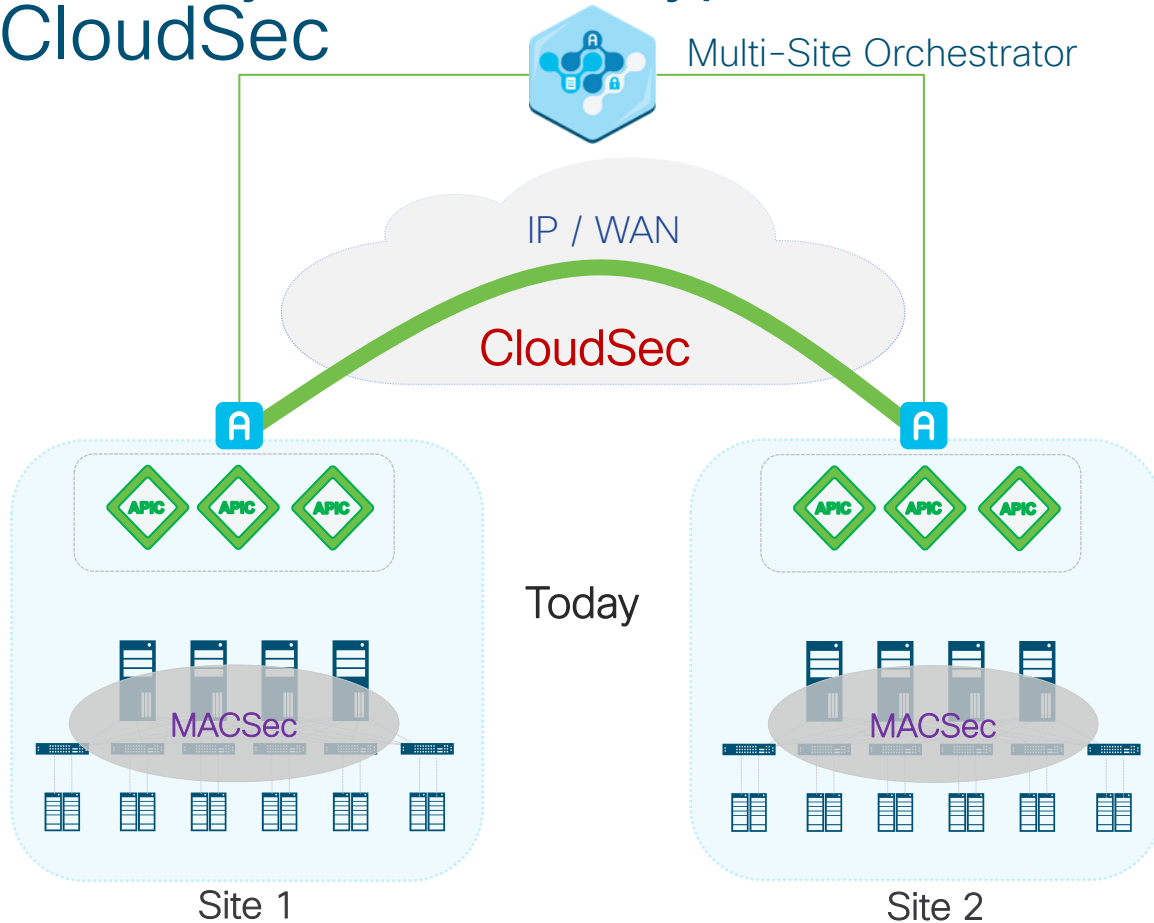


MACSEC Link Encryption
MKA Key Exchange

APIC Centralized Key
Management

CISCO *Live!*

ACI Anywhere Encrypted DCI Connectivity – CloudSec



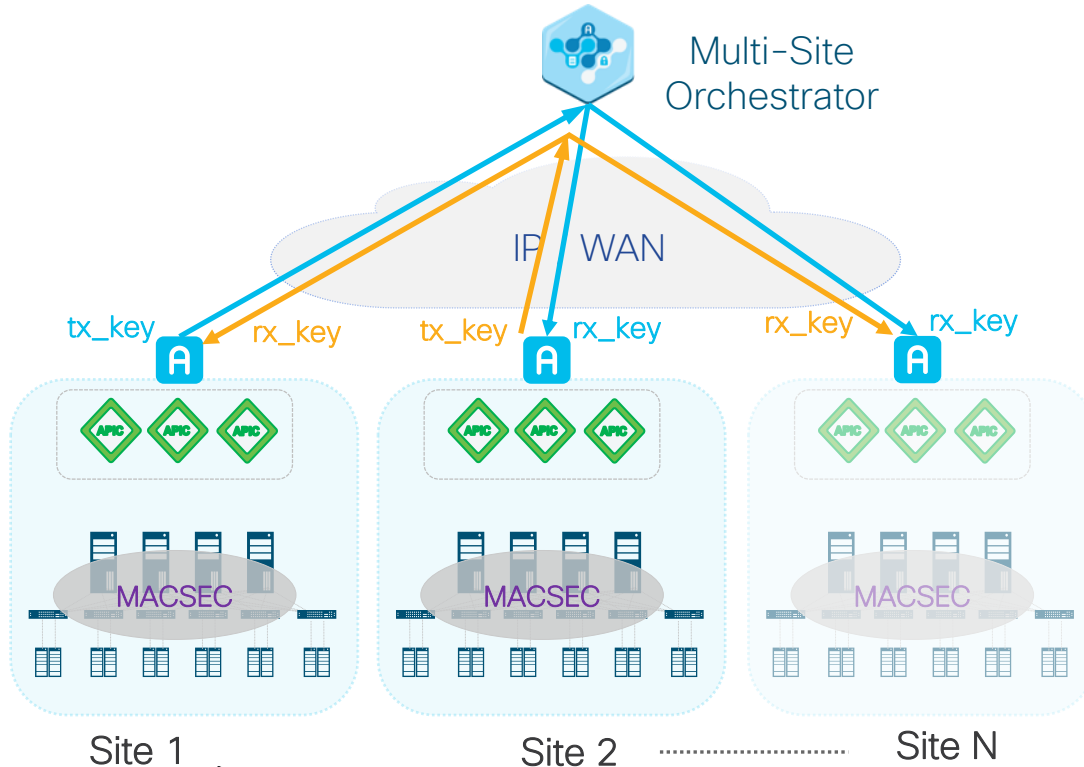
Encrypted VXLAN Overlay
for Inter-Site Traffic

MKA Key Exchange over
BGP-EVPN Protocol

Supported Hardware:
X9736C-FX LC
Nexus 9364C & 9332C
Nexus 9332C

ACI CloudSec Keys operations

Automated Key Distribution & Re-Key



- Multi-site Orchestrator driven
 - No protocol dependency
- Reliable and secure key transport
 - rx_key installed before tx_key
- Non disruptive re-key
 - Hitless make before break
- Always encrypted
 - In case of programming errors, stay encrypted with previous key


Site 1

CISCO *Live!*

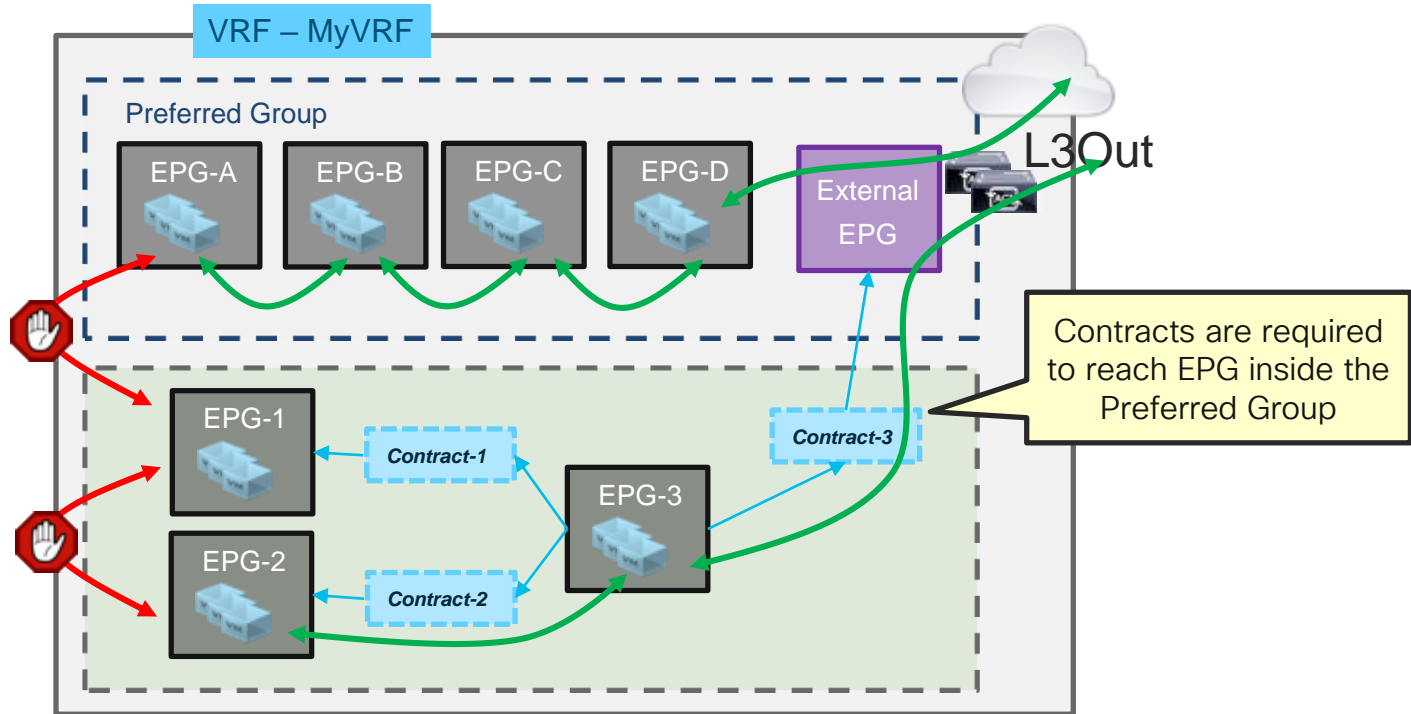
Site 2

Site N

The ACI Segmentation and Isolation Toolbox



Isolation: EPG Contract Preferred Groups



Intra EPG Isolation

Intra EPG Isolation

- Intra EPG Isolation **blocks communication between all endpoints** inside the group
- Supports mixing of Physical and Virtual endpoints in same EPG
- Can be configured on **all type of EPG**

Properties

Name: **baseEPG**

Description: optional

Tags:

Alias:

uSeg EPG: **false**

pcTag(sclass): **32772**

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

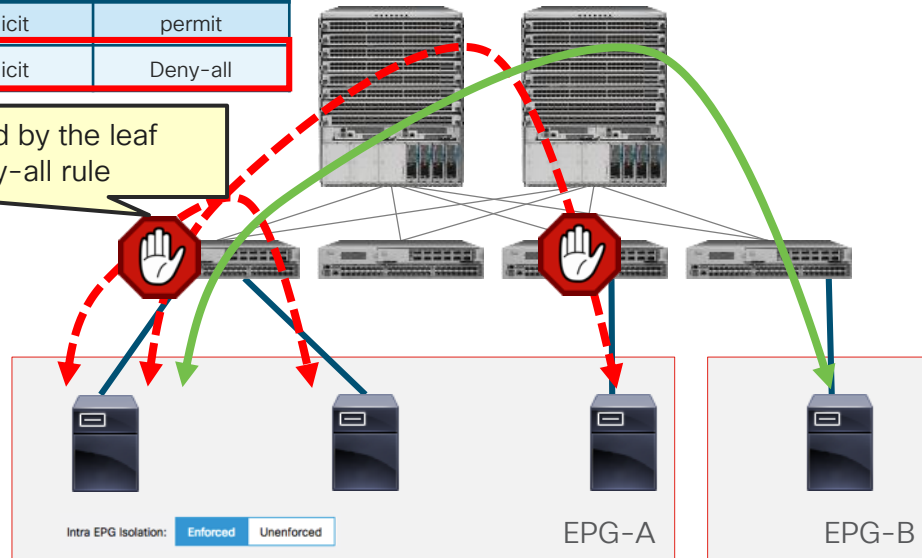
Intra EPG Isolation

```
<fvTenant name="Tenant1">
  <fvAp name="ap1">
    <fvAEPg isAttrBasedEPg="no" matchT="AtleastOne" name="baseEPG" pcEnfPref="enforced" prefGrMemb="exclude" prio="unspecified">
      <fvRsBd tnFvBDName="bd"/>
    </fvAEPg>
  </fvAp>
</fvTenant>
```

Intra EPG Isolation - Zoning Rules

Source	Destination	Filter	Action
EPG-A	EPG-B	implicit	permit
EPG-A	EPG-A	implicit	Deny-all

Intra EPG traffic will be dropped by the leaf because of the implicit deny-all rule



ACI Leaf Uses Zoning Rules to Forward or Drop the Traffic

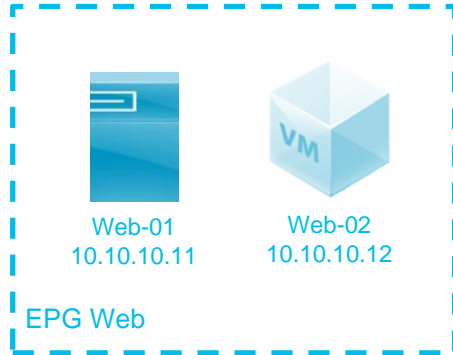


VRF - vrf1

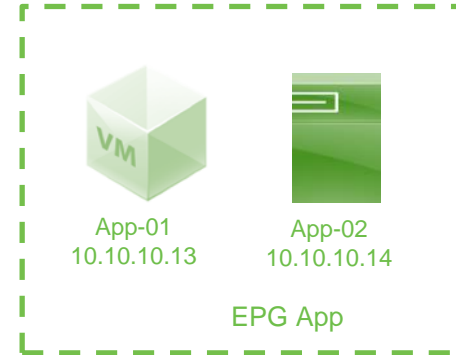
Policy Operational Stats Health Faults History

Associated EPGs Associated External Routed Networks

Name	Description	State	Issues	QoS	Encap	PC Tag
ap1/Blue		applied		Unspecified		16387
ap1/Green		applied		Unspecified		32771



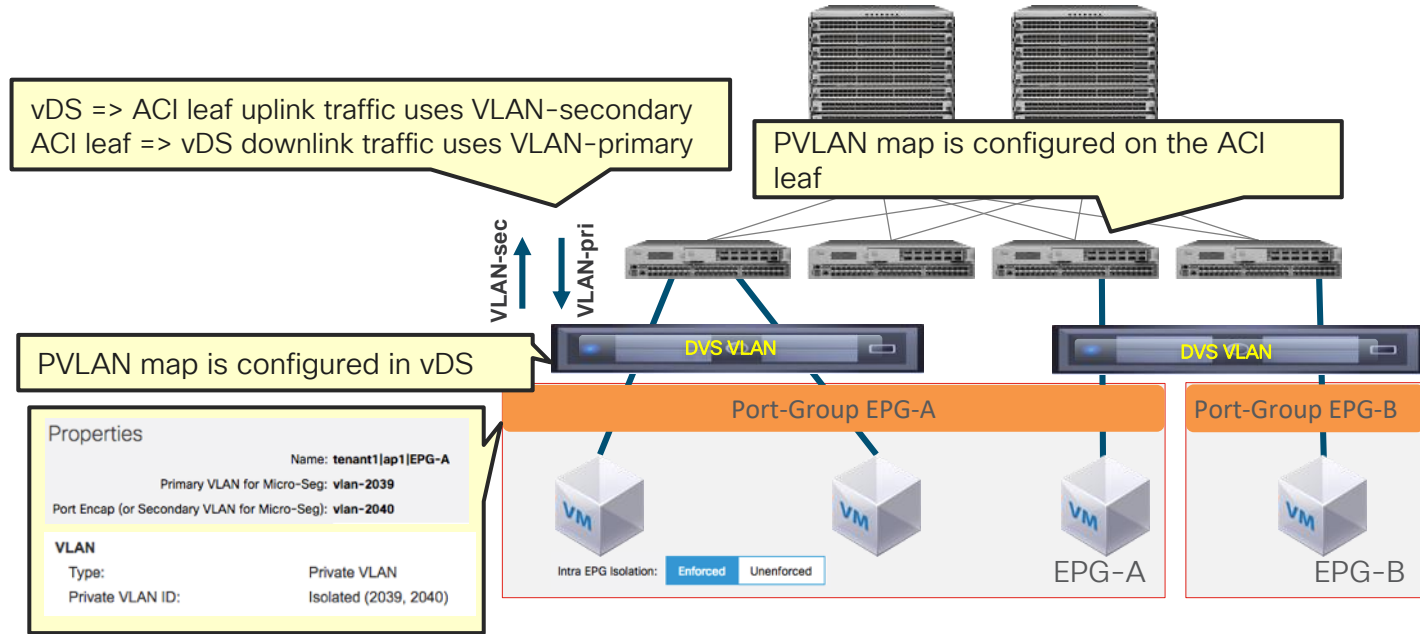
Once contract is created, it will get programmed on the ACI leaf as Zoning Rules. Leaf forwards/drops the packets based on those rules



```

leaf1# show zoning-rule scope 2162697 | egrep -E "Scope|32771|16387"
Rule ID  SrcEPG  DstEPG  FilterID operSt  Scope  Action  Priority
4616     16387    32771   5         enabled 2162697 permit  src_dst_any(8)
4617     32771    16387   5         enabled 2162697 permit  src_dst_any(8)
  
```

VMWare DVS Intra EPG Isolation



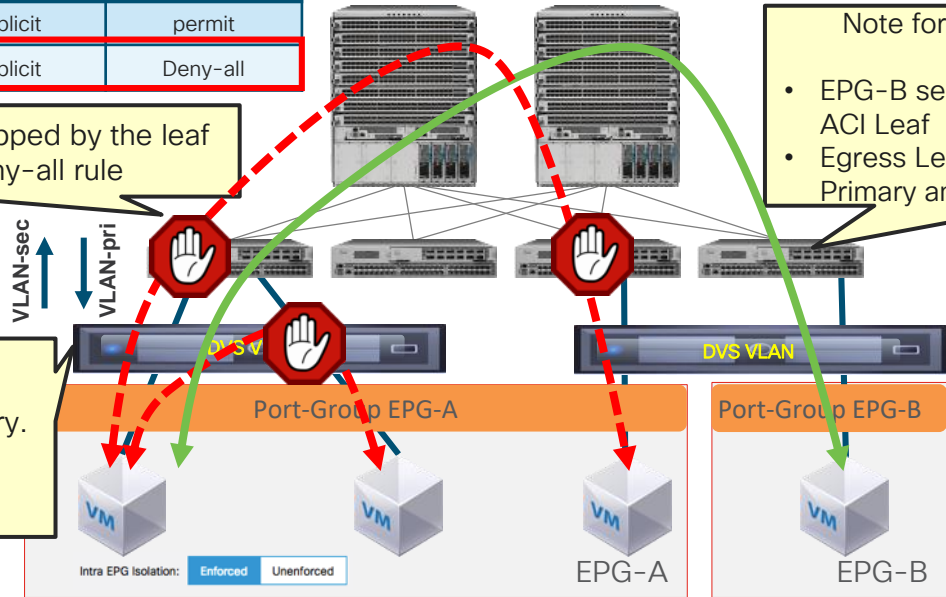
VMWare DVS Intra EPG Isolation

Source	Destination	Filter	Action
EPG-A	EPG-B	implicit	permit
EPG-A	EPG-A	implicit	Deny-all

Inter-ESXi host traffic will be dropped by the leaf because of the implicit deny-all rule

Note for Inter-EPG Traffic with Isolation Enabled:

- EPG-B sends traffic over regular VLAN to ACI Leaf
- Egress Leaf will encapsulate traffic in VLAN-Primary and send towards EPG-A VMs




- Intra-ESXi host traffic is encapsulated in VLAN-secondary.
- vDS denies local intra-EPG VM traffic via PVLAN

Intra-EPG Contracts


- Since release 3.0, ACI supports “Intra-EPG Contracts”
 - Allows whitelist policy enforcement of Intra-EPG traffic
 - Can co-exist with Inter-EPG contracts
 - Eliminates the need to create uSeg EPGs or deploy external FW for Intra-EPG segmentation
 - Enforcement is on Leaf switch (ie. Nexus 9000-EX models or above)
 - Same as regular contract scale

Intra-EPG Contract:

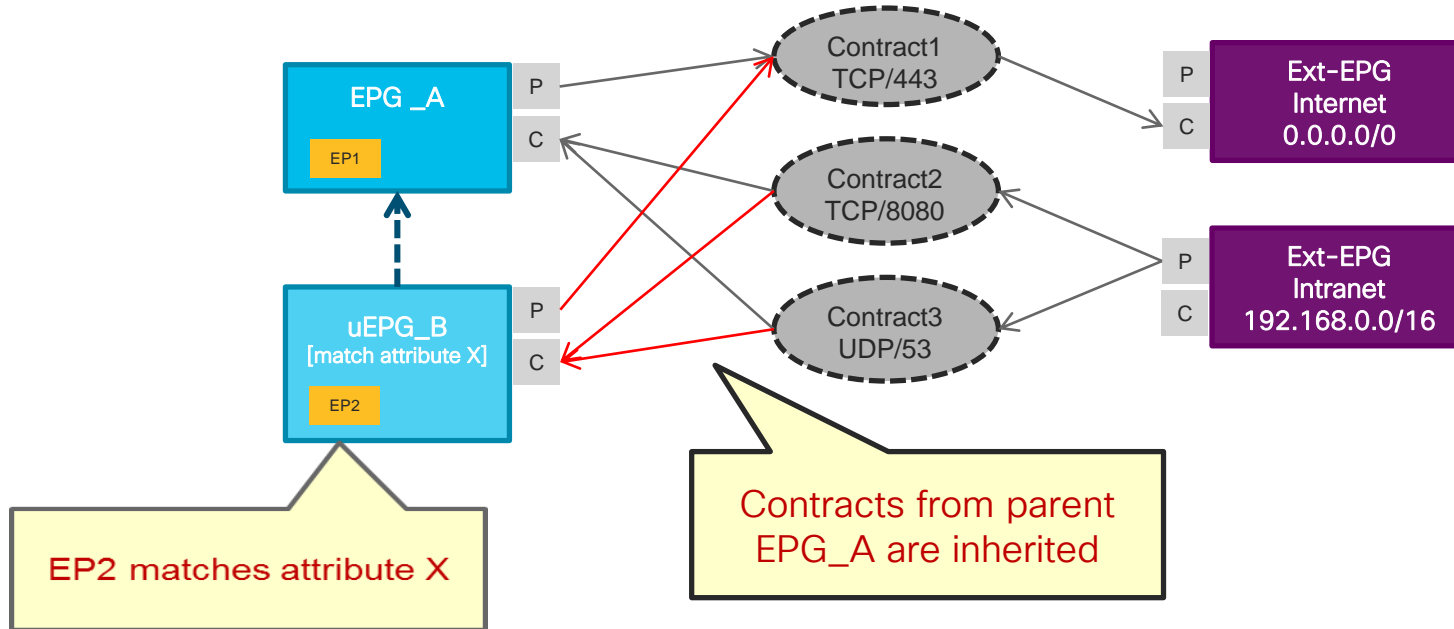
- Src-Class = Dest-Class
- Src-Class, Src-Class, Contract



Contract Inheritance



Contract Inheritance



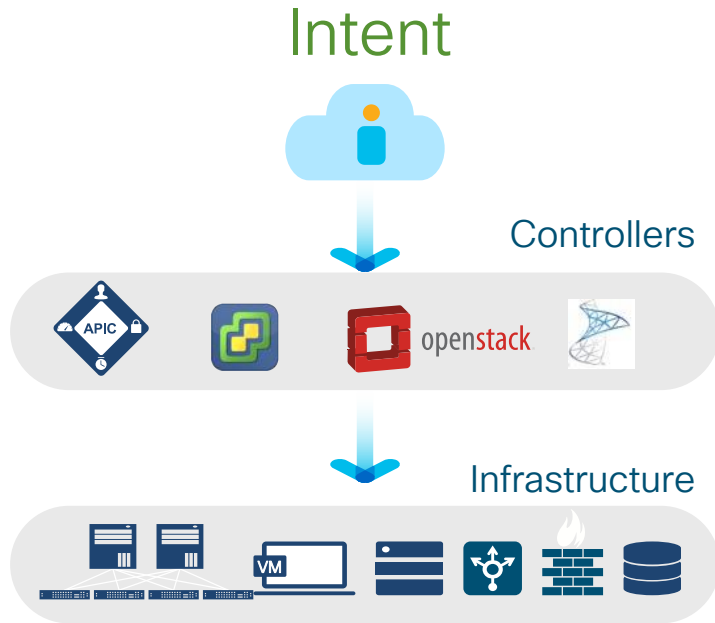
What is Network Assurance?

“**Assurance** is the **Guarantee** that
the Network is doing what you
Intended it to do”

Intent Encompasses Data Center Operations

Bridging, Routing, Security, QOS, Service Chaining, VM placement ... Compliance, Audits

Assurance Gap in Today's Networks



1

How do I have confidence that I don't have errors due to my changes?

2

How do I easily understand the state of my entire infrastructure?

3

How do I rapidly analyze the network to identify issues?

Key Insight: Networks are Deterministic

If we **understand** the entire state of the network, we can accurately **predict** it's behavior

... without a single packet flowing through it

Cisco Day 2 Operations Stack (aka “Opstack”)

Network Insights & Assurance



Insights: Health and
Availability

Assurance: Moving from
Reactive to Proactive

ACI & NX-OS Fabrics


ACI Fabrics

Opstack benefits

Architecture and Planning teams



Network Operations



Network Administration and Maintenance



NAE
Policy Explorer

Verification of design/connectivity mandates
Ad-hoc policy exploration
Ad-hoc Connectivity and Segmentation
Analysis

Accelerate ACI on-ramp

Light weight book-keeping procedures

NAE

Capacity Planning
Design and verify compliance mandate and posture
Design and verify security mandate and posture

Proactive Assurance and Compliance
Faster incident and problem management
Shrink change management windows
Accelerate ACI on-ramp

Execute high confidence production
maintenance and upgrades

NIR/NIA

Trend based capacity planning
Design trend-based environmental site
operating procedures

Fabric wide anomaly detection based on
- Resource monitoring
- Flow monitoring
Faster low-level troubleshooting and diagnostics

Proactively reduce vulnerability exposure
Improve Site reliability

cisco *Live!*

Opstack- Available via Premier License

Network Insight & Network Assurance

NAE Policy Explorer *

- Network Policy exploration
- Ad-hoc connectivity and segmentation discovery

Network Assurance Engine

- Policy/ Control/Data plane Assurance
- Incident and Problem Management
- Compliance and Audit

Network Insights Resources *

- Fabric wide resource utilization & trends
- Anomaly detection – environmental, config & operational resources, interface errors
- End-to-end flow path, latency and drop reason

Network Insights Advisor **

- Notifications of EOS/EOL of H/W & S/W
- Security Advisory Notification Updates (PSIRTs)
- Recommended S/W Release Updates and upgrade impact analysis
- TAC assist


Architecture and Planning



Network Administration and Maintenance



Network Operations



* Available as App on APIC


Cisco Network Assurance Engine



Based on mathematical models of the network

Continuously verifies and validates the entire ACI network


Proactively delivers the confidence that the ACI network is operating correctly



Always-On
Network Assurance

Comprehensive, Intelligent, Continuous

Cisco Network Assurance Engine: How It Works



Data Collection

Captures all non-packet data:
intent, policy, state across
data center network

Comprehensive Network Modeling

Mathematically accurate models
spanning underlay, overlay and
virtualization layers

Intelligent Analysis

5000+ domain knowledge-based
error scenarios built-in, codified
remediation steps

Use Cases & Benefits

Making Operations Fundamentally Proactive



PREDICT THE IMPACT OF CHANGES

- Drive change agility
- Minimize human errors and eliminate configuration drift
- Accelerate migrations



PROACTIVELY VERIFY NETWORK-WIDE BEHAVIOR

- Ensure connectivity
- Proactively eliminate potential network outages or vulnerabilities
- Enhance SLAs



ASSURE NETWORK SECURITY POLICY AND COMPLIANCE

- Reduce security risk
- Achieve provable compliance by design, continuously

User Interface: Centered Around “Smart Events”

Change Management

Search and Visualization

Incidence and Problem Management

Real-time Change Management Alerts

Alert	App EPG	Count
APP_EPG_IS_NOT_DEPLOYED		1
OVERLAPPING_SUBNETS_ACROSS_VRFs_D	Communication	1
PERMIT_POLICY_SHADOWED_BY_DENY_PO	Deny policy is overriding Permit policy	1

Description: EPGs cannot communicate due to a deny rule overriding a permit rule.

Affected Objects:

Provider EPG	Consumer EPG
EPG_11	EPG_6

Permit Contract	Subject	Filter	Filter Entry
Cont_100	Sub_1	Flt_100	6027_0

Failure Condition: Deny policy is overriding Permit policy

View Control

View Filters: Tenants & Apps, VRFs & SVs

Severities: Critical, High, Medium, Low

Events Count: 11, 1-10, 1-15, 1

Tenant Endpoints Events

Severity	Event Date	App	Tenant
Critical	2023-08-01	IP_ADDRESS_WRONG	Employees
High	2023-08-01	IP_ADDRESS_WRONG	Employees
Medium	2023-08-01	IP_ADDRESS_WRONG	Employees

Suggested Next Steps:

- Login into the APIC UI and verify the presence of the endpoints in the operational tab of the EPGs.
- Identify the EPGs that actually own the IPs and change the IP address on the host/device that has the incorrect IP address.
- Clear the endpoint on the leaf by opening a SSH session to each leaf and entering the following command:
leafIf clear system internal epg endpoint command

Smart Events: What, Where, Why, and How

Ex.: NAE UI for Compliance Analysis

Cisco Network Assurance Engine is currently running with a TRIAL license. [View License Details](#)

Assurance Group: Can5_BM28_mod28.tar.gz-20181108-1739

Navigation: Dashboard, Change Management, Verify & Diagnose, Epoch Analysis, Optimize, **Compliance**, Smart Events

Time Range: 01:01 PM 07/20/2018 to 03 PM

Zoom Level: All, 1 m, 1 w, 1 d, 12 h, 6 h, 1 h, Custom

Dashboard

Epoch | Trend

Smart Events

Severity	Count
Critical	325
Major	965
Minor	204
Warning	294
Info	3,240
Total	5,028

Real-time Change Analysis

Tenant	Critical	Major	Minor
abc02	0	48	0

VRF	Issues
abc02_ctx02	24

Tenant	Forwar...	Security	Fabric	Others
abc02	48	0	0	0

Annotations:

- view compliance analysis results (points to Compliance Analysis in the dropdown)
- Configure compliance requirements (points to Manage Compliance Requirements in the dropdown)

Agenda



Data Center / Virtualization Cisco education offerings

Course	Description	Cisco Certification
Introducing Cisco Data Center Networking (DCICN) Introducing Cisco Data Center Technologies (DCICT)	Get job-ready foundational-level certification and skills in installing, configuring, and maintaining next generation data centers.	CCNA® Data Center
Implementing Cisco Data Center Unified Computing (DCUCI) Implementing Cisco Data Center Infrastructure (DCII) Implementing Cisco Data Center Virtualization and Automation (DCVAI) Designing Cisco Data Center Infrastructure (DCID) Troubleshooting Cisco Data Center Infrastructure (DCIT)	Obtain professional level skills to design, configure, implement, troubleshoot next generation data center infrastructure.	CCNP® Data Center
Product Training Portfolio:DCAC9K, DCINX9K, DCMDS, DCUCS, DCNX1K, DCNX5K, DCNX7K, CACND, DSACI, HFLEX UCSD, UCSDACI, DCUCEN	Gain hands-on skills using Cisco solutions to configure, deploy, manage and troubleshoot unified computing, policy-driven and virtualized data center infrastructure.	
Designing the FlexPod® Solution (FPDESIGN) Implementing and Administering the FlexPod® Solution (FPIMPADM)	Learn how to design, implement and administer FlexPod® solutions	Cisco and NetApp Certified FlexPod® Specialist
Designing the VersaStack Solution (VSDESIGN) Implementing and Administering the VersaStack Solution (VSIMP)	Learn how to design, implement and administer VersaStack solutions	

For more details, please visit: <http://learningnetwork.cisco.com>

Questions? Visit the Learning@Cisco Booth



Network Programmability Cisco education offerings

Course	Description	Cisco Certification
Developing with Cisco Network Programmability (NPDEV)	Provides Application Developers with comprehensive curriculum to develop infrastructure programming skills; Addresses needs of software engineers who automate network infrastructure and/or utilize APIs and toolkits to interface with SDN controllers and individual devices	Cisco Network Programmability Developer (NPDEV) Specialist Certification
Designing and Implementing Cisco Network Programmability (NPDES)	Provides network engineers with comprehensive soup-to-nuts curriculum to develop and validate automation and programming skills; Directly addresses the evolving role of network engineers towards more programmability, automation and orchestration	Cisco Network Programmability Design and Implementation (NPDES) Specialist Certification
Programming for Network Engineers (PRNE)	Learn the fundamentals of Python programming – within the context of performing functions relevant to network engineers. Use Network Programming to simplify or automate tasks	Recommended pre-requisite for NPDES and NPDEV Specialist Certifications
Cisco Digital Network Architecture Implementation Essentials (DNAIE)	This training provides students with the guiding principles and core elements of Cisco's Digital Network Architecture (DNA) architecture and its solution components including; APIC-EM, NFV, Analytics, Security and Fabric.	

For more details, please visit: <http://learningnetwork.cisco.com>

Questions? Visit the Learning@Cisco Booth



Complete your online session survey



- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

Continue your education



Demos in the
Cisco Showcase



Walk-In Labs



Meet the Engineer
1:1 meetings



Related sessions



Thank you





You make **possible**