Data Center Architectures evolution - SDN to the rescue

Fotis Baroumas

Technical Solution Architect – Data center

Agenda

- **Introduction to Datacenter.**
- **Datacenter Networking evolution.**
- **VXLAN BGP EVPN Fabric overview.**
- **Application Centric Infrastructure (ACI) overview.**
  - **Cisco ACI components.**
  - **Network services integrations.**
  - **Microsegmentation overview.**
  - **ACI operations.**

By 2023

…the number of applications running in the data center and edge locations increases by 300%

61%

…of organizations cite manual processes and siloed teams as top bottlenecks in app delivery pipelines

75%

…of Global 2000 IT organizations will adopt automated operations practices to transform their IT workforce to support unprecedented scale by 2023

Data Center Economics

Overall Spend Distribution

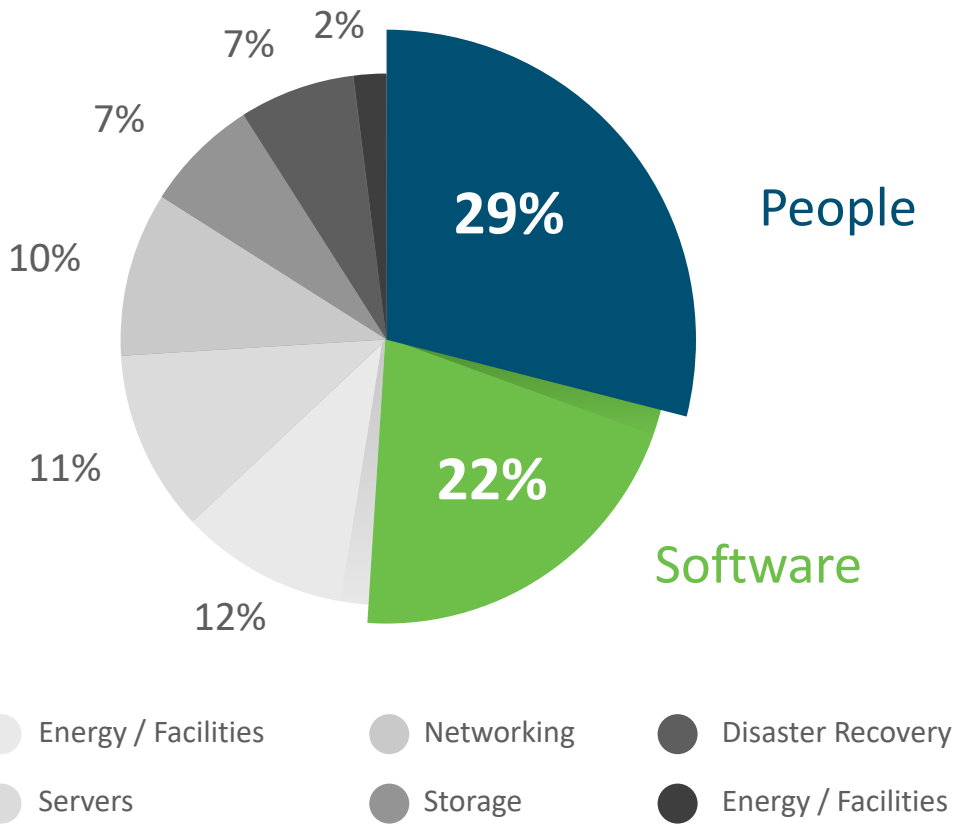

Source: Gartner, Cisco IT, “Data Center Cost Portfolio”

Server-Related Spend

WW New Server, Power, Cooling, Management, and Administration Spending Share

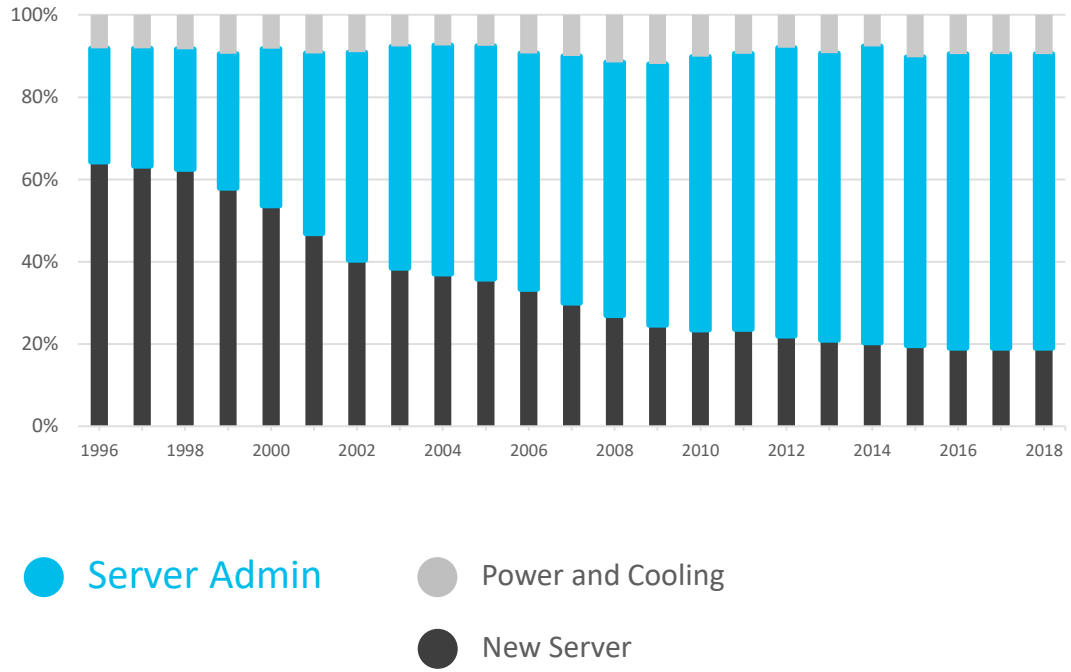

Source: IDC #250082, “Worldwide Server, Power and Cooling, and Management and Administration Spending 2014–2018 Forecast,” August 2014

Datacenter Networking Evolution

Application Centric Infrastructure (ACI)

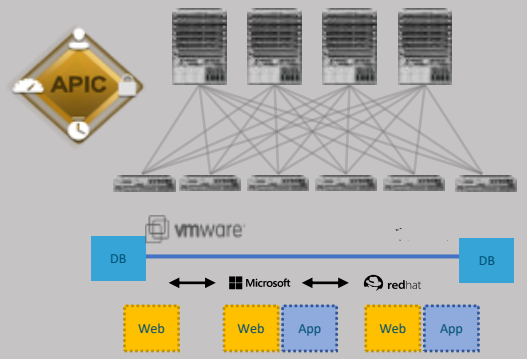

Turnkey integrated solution with security, centralized management, compliance and scale

Automated application centric-policy model with embedded security

Broad and deep ecosystem

VXLAN BGP EVPN Fabric

VxLAN-BGP EVPN standard-based

3rd party controller support

Cisco Controller for software overlay provisioning and management

Traditional Tiered

Modern NX-OS with enhanced NX-APIs

Command Line or Cisco Controller for management

VXLAN BGP EVPN Fabric Overview

Introducing VXLAN

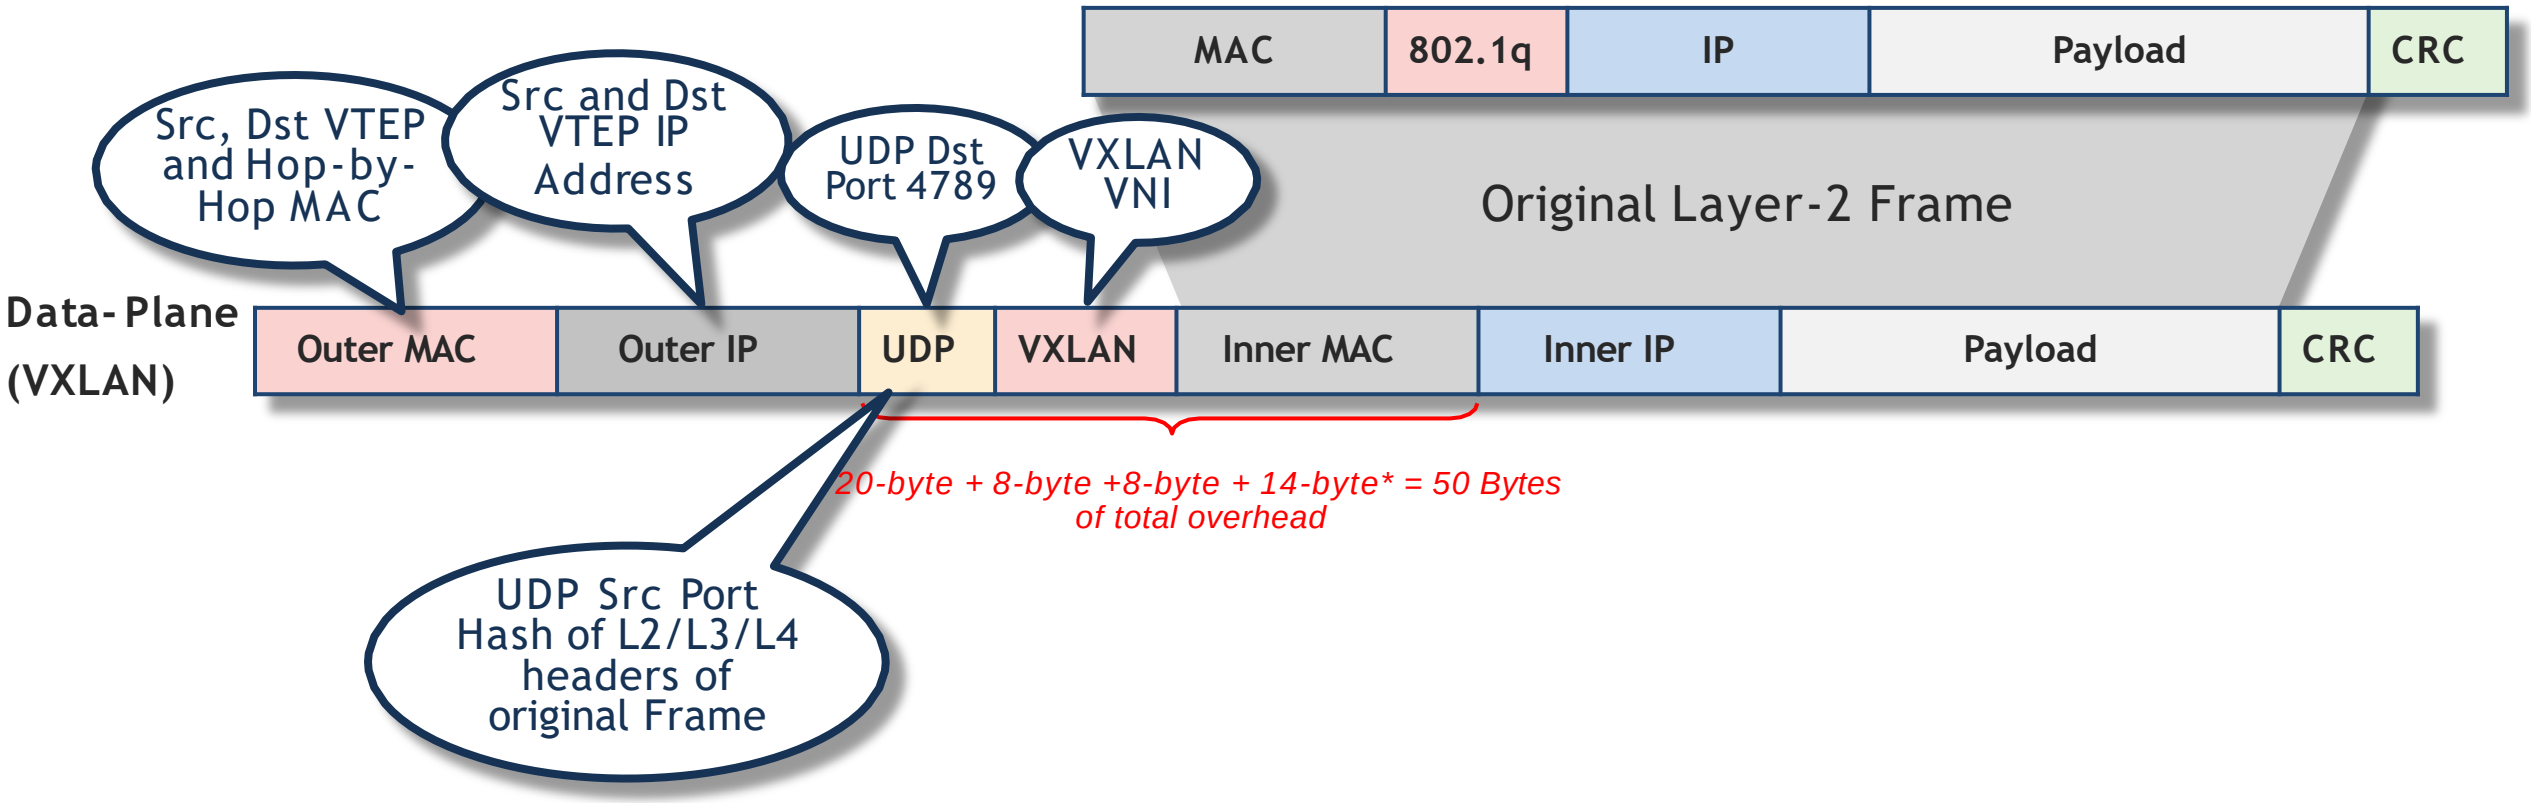

VXLAN Data Plane Packet

- VXLAN is point to multi-point tunneling mechanism to extend Layer 2 networks over an IP network

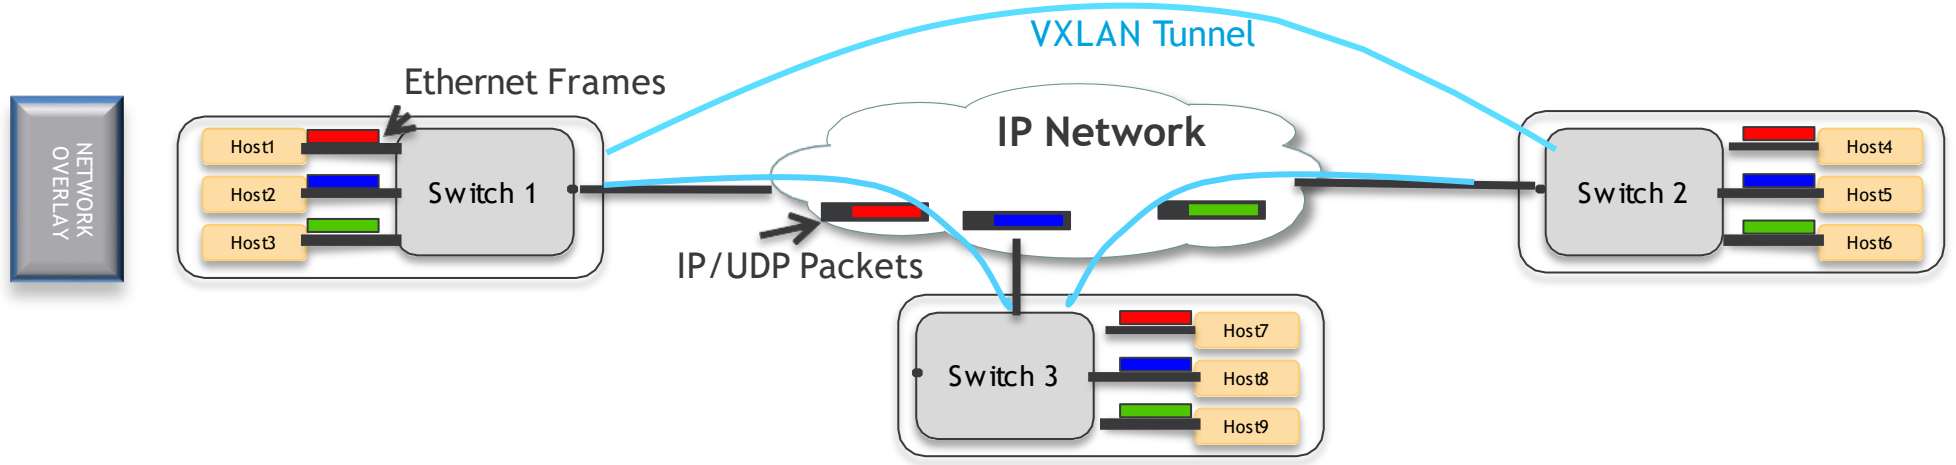

- VXLAN uses MAC in UDP encapsulation (UDP destination port 4789). VXLAN Adds 50 Bytes to the original frame.

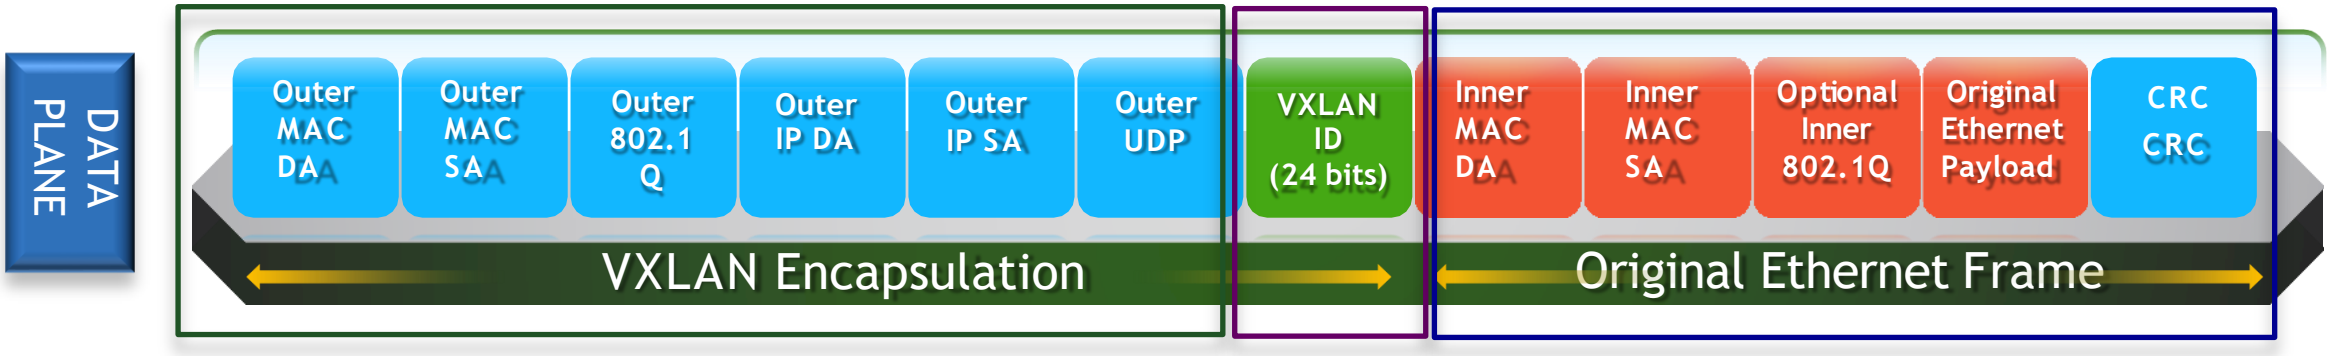

VXLAN: The Building Blocks

VTEP: VXLAN Tunnel End-Point
VNI/VNID: VXLAN Network Identifier
NVE: Network Virtualization Edge

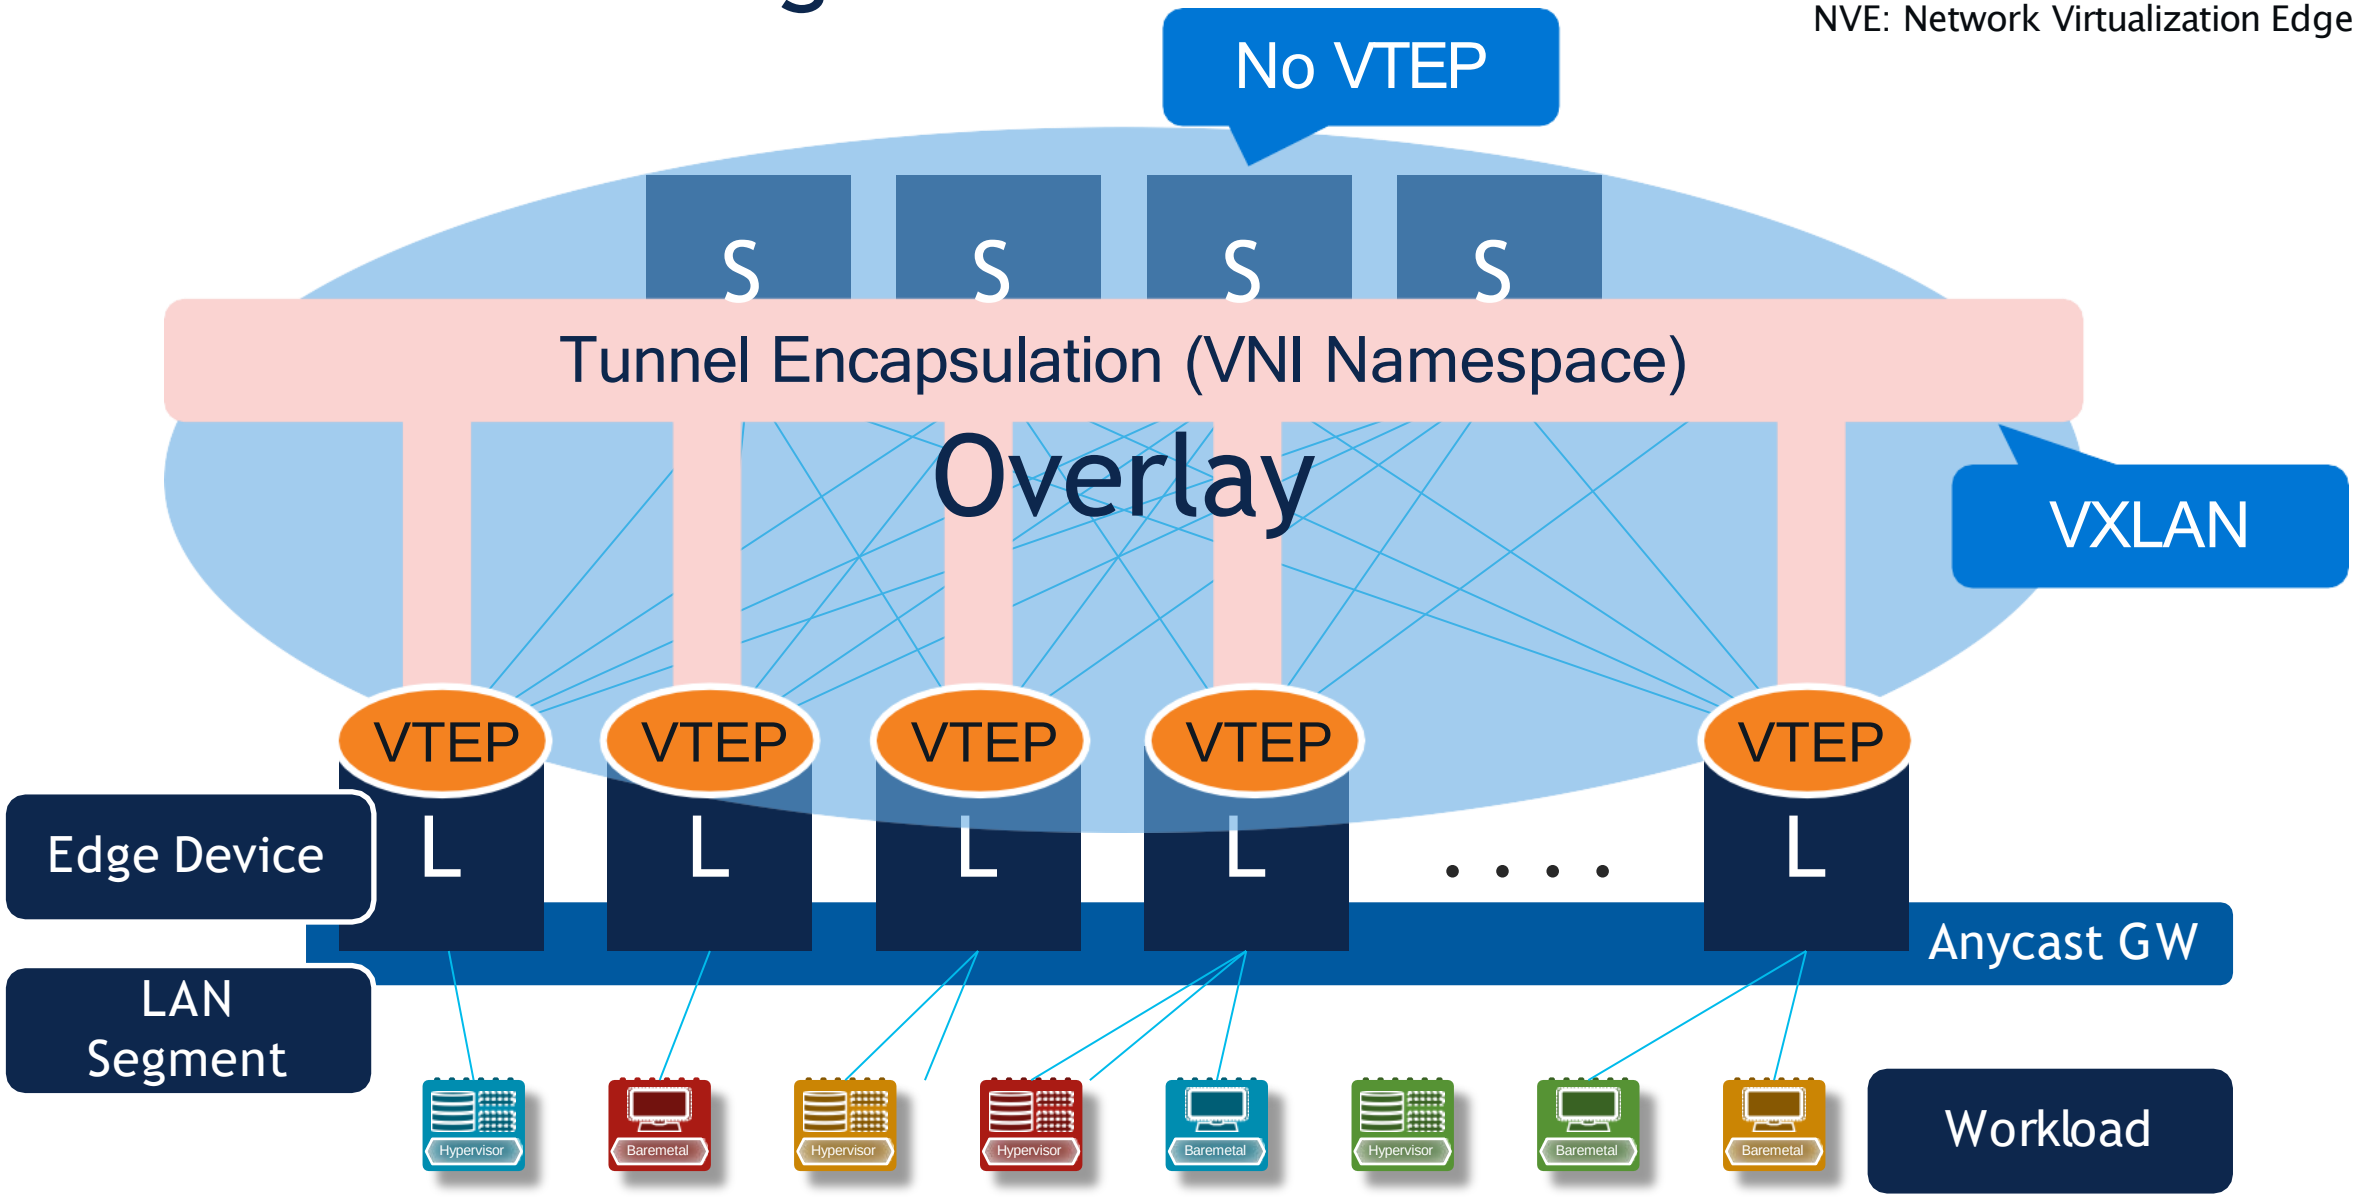

VXLAN BGP EVPN Configuration

VTEP Virtual Networks and Overlay Interface

```
Vlan 10
 vn-segment 5010
Vlan 20
 vn-segment 5020

Vlan 1000
 !Layer 3 VNI
 vn-segment 9999
Vlan 2000
 !Layer 3 VNI
 vn-segment 9998

interface Vlan10
  no shutdown
  vrf member VRF-RED
  ip address 192.168.10.254/24 tag 12345
  ipv6 address 2001::1/64 tag 12345
  fabric forwarding mode anycast-gateway

interface Vlan20
  no shutdown
  vrf member VRF-BLUE
  ip address 192.168.20.254/24 tag 12345
  ipv6 address 2002::1/64 tag 12345
  fabric forwarding mode anycast-gateway

interface nve1
  source-interface loopback0
  host-reachability protocol bgp
  member vni 5010
    mcast-group 239.1.1.1
  member vni 5020
    mcast-group  239.1.1.1
  member vni 9999 associate-vrf
  member vni 9998 associate-vrf
```

Layer 2 VNI

Layer 3 VNI

Map L2VNI to NVE

Associate L3VNI to NVE

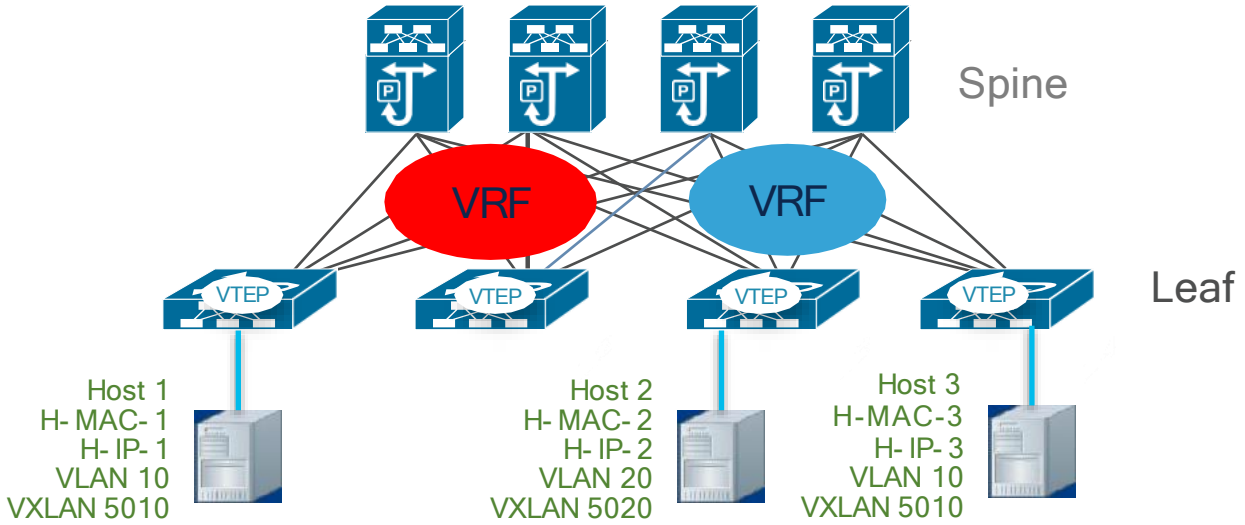

IP VRF

MAC VRF

```
vrf context VRF-RED
  vni 9999
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
evpn
  vni 5010 l2
    rd auto
    route-target both auto
```

```
vrf context VRF-BLUE
  vni 9998
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
evpn
  vni 5020 l2
    rd auto
    route-target both auto
```

VXLAN Benefits

Customer Needs	VXLAN Delivered
Any workload anywhere – VLANs limited by L3 boundaries	Any Workload anywhere- across Layer 3 boundaries
VM Mobility	Seamless VM Mobility
Scale above 4k Segments (VLAN limitation)	Scale up to 16M segments
Efficient use of bandwidth	Leverages ECMP for optimal path usage over the transport network
Secure Multi-tenancy	Traffic & Address Isolation

Application Centric Infrastructure (ACI) Overview

What is Cisco ACI

- Cisco ACI is the industry's most secure, open, and comprehensive Software-Defined Networking (SDN) solution.

- Cisco® Application Centric Infrastructure (Cisco ACI®) is part of intent-based networking framework to enable agility and resiliency in the data center.

- Cisco ACI enables automation that accelerates infrastructure deployment and governance, simplifies management to easily move workloads across a multifabric and multicloud framework.

Cisco ACI components

There are two main parts to an ACI fabric:

1) Physical fabric built with Nexus 9000 switches and based on a Leaf / Spine architecture.

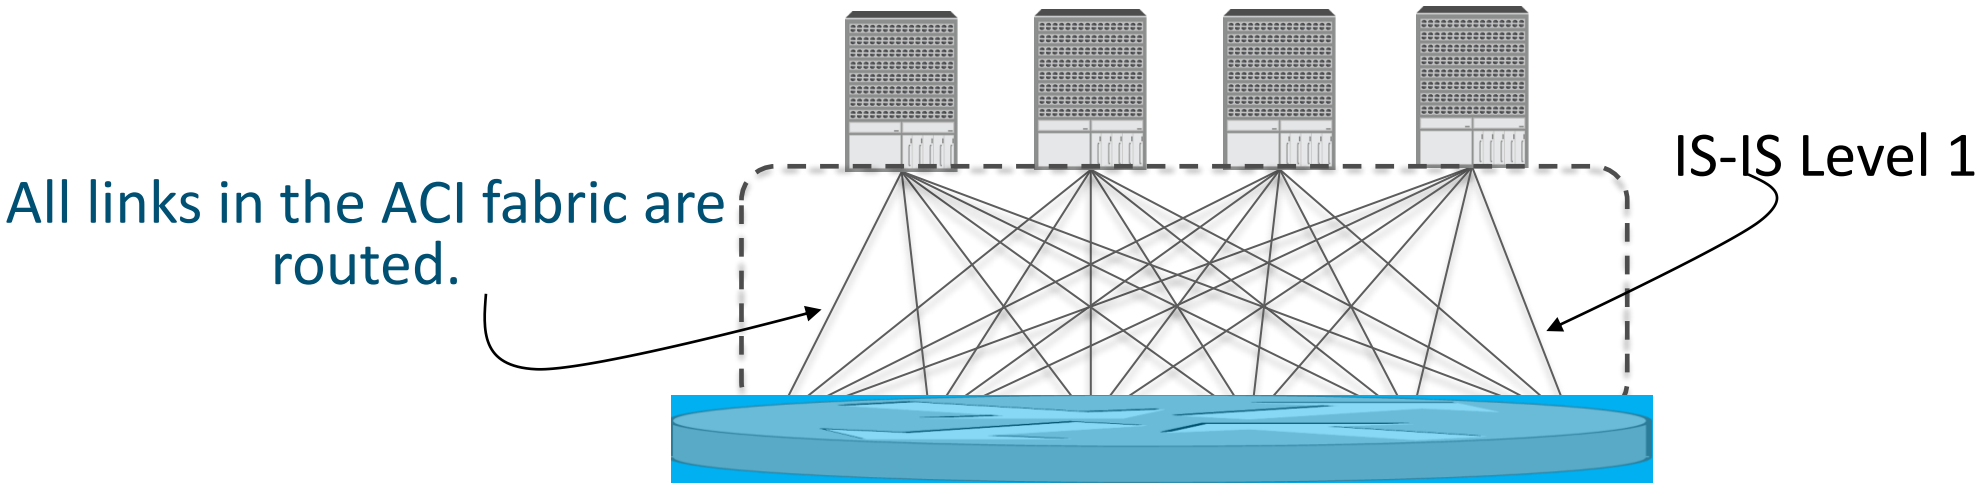

There are two main parts to an ACI fabric:

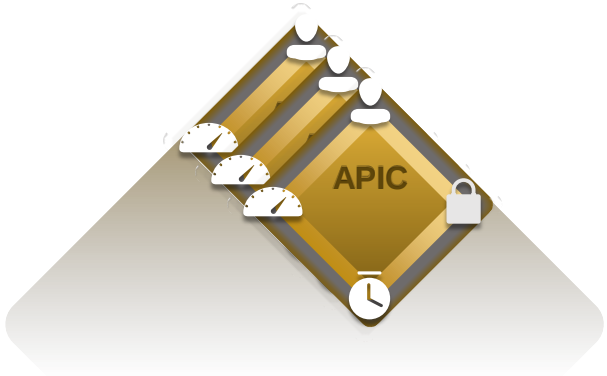

2) Application Policy Infrastructure Controller (APIC):
a unified DC network policy management system.

Smallest ACI Fabric in 2 DCs

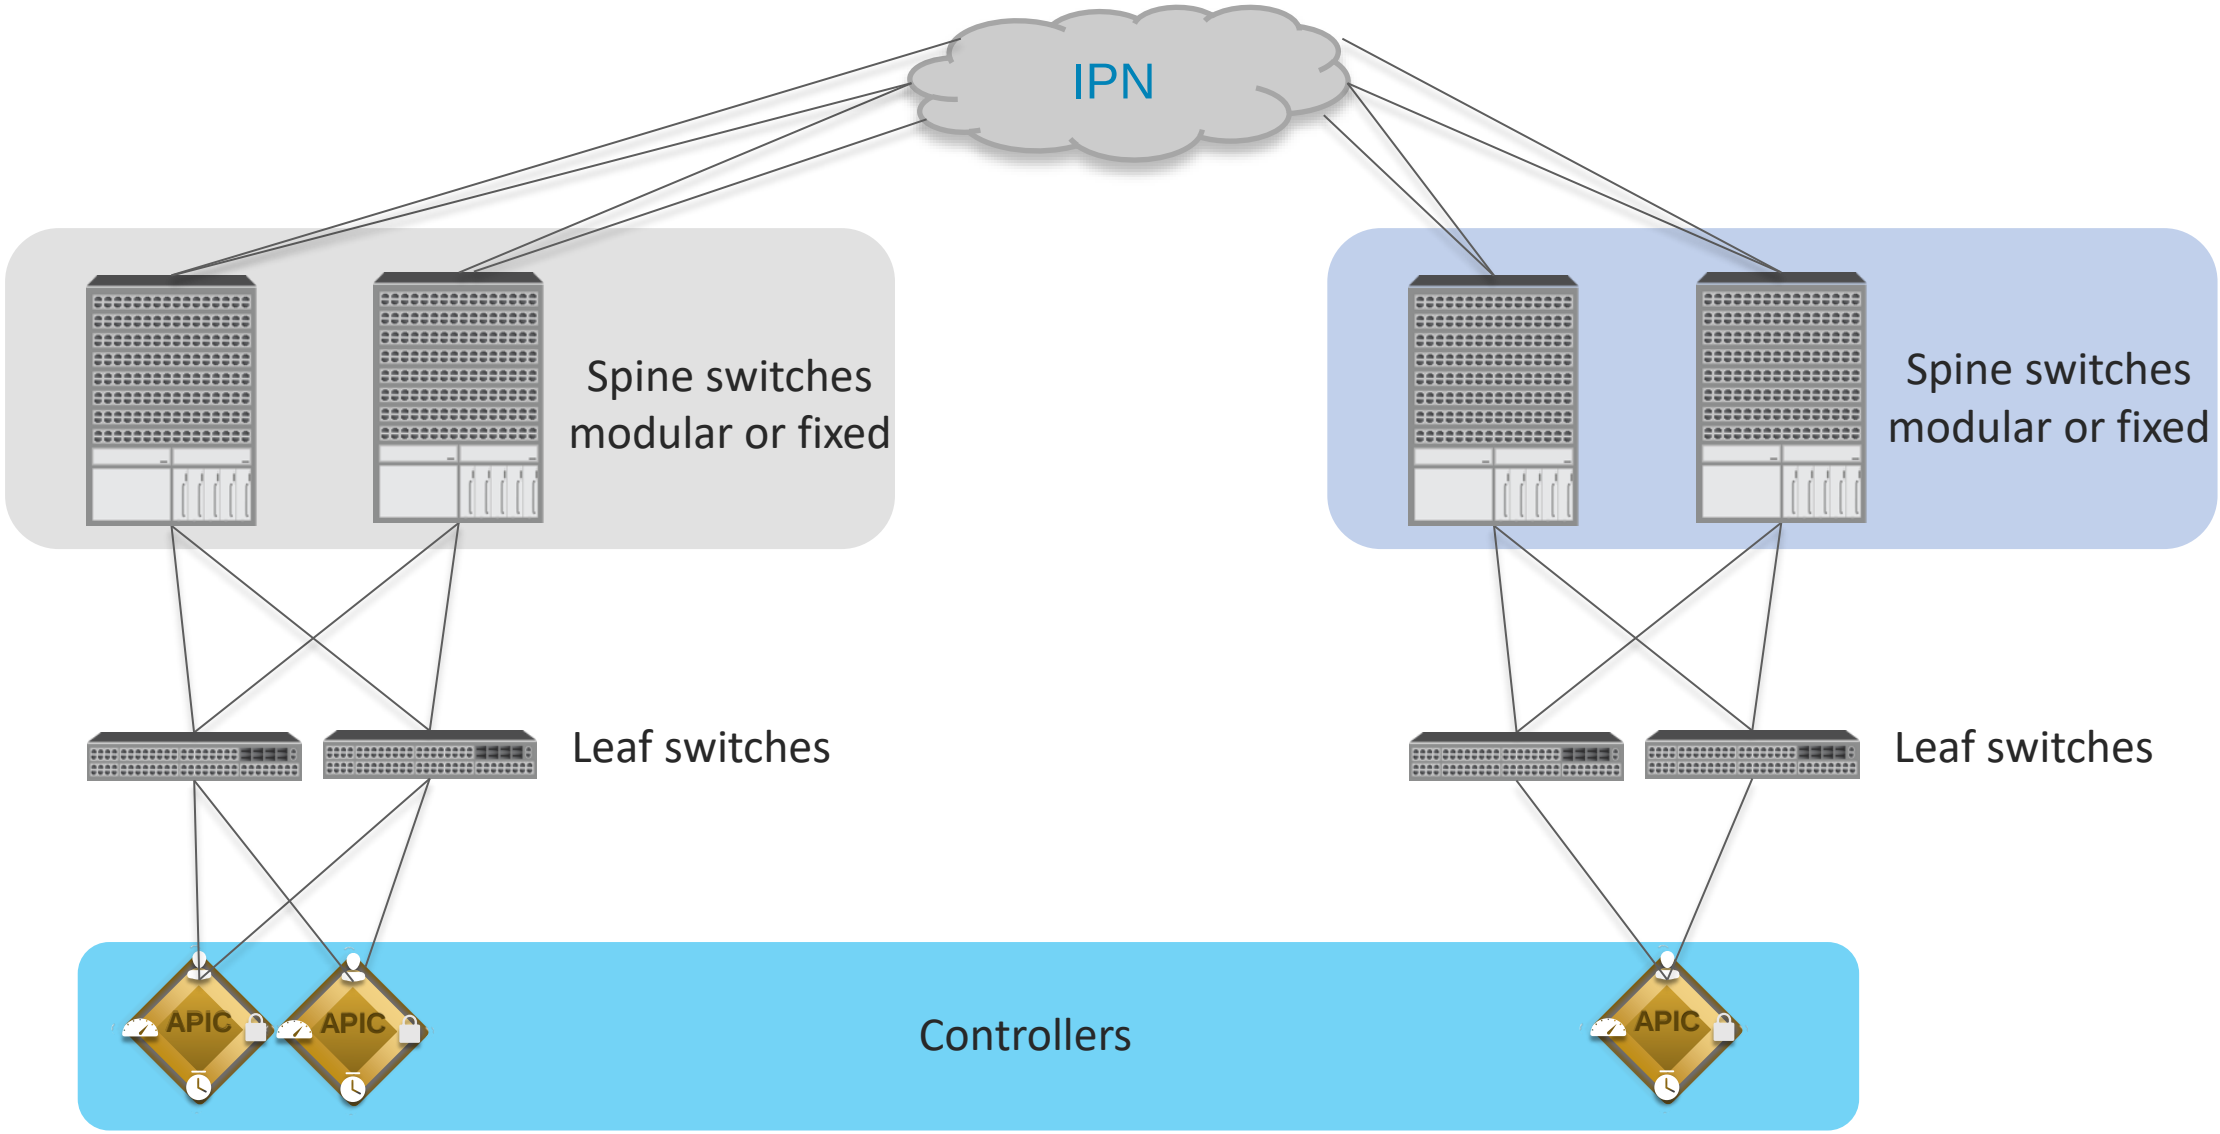

Central Management of the entire network.

- Single point for configuration and troubleshooting.
- Reduced possibility of human error.
- No extra effort as you scale.
- Easy deployment of new leaf switches.

Simplified topology Spine-leaf

VXLAN
data plane

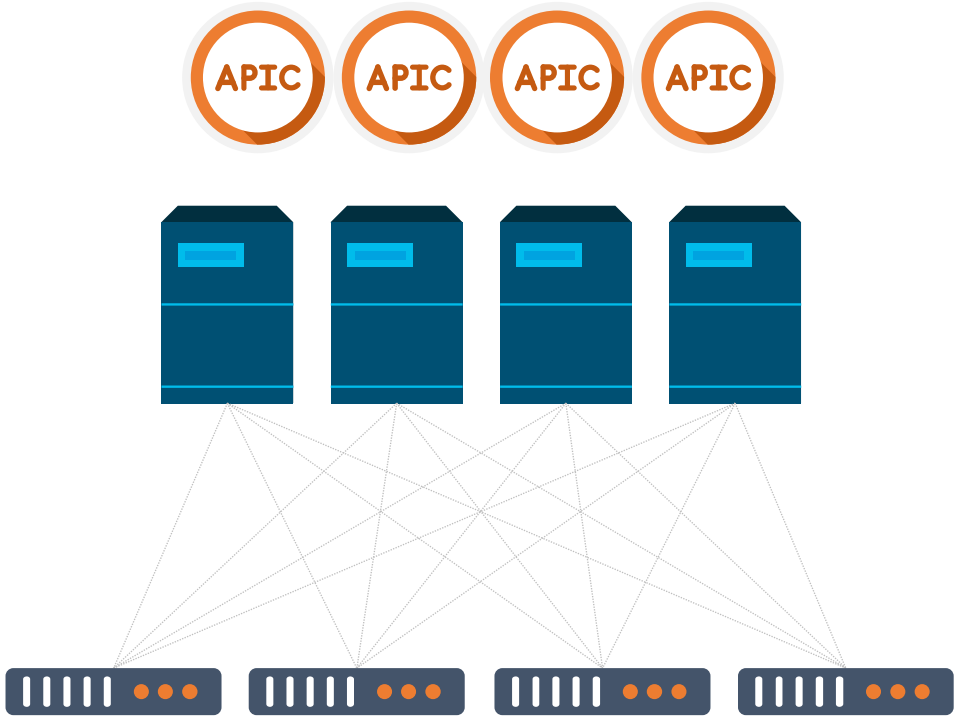

Spine and leaf
topology

- No spanning-tree
- More redundancy

- More throughput
- Faster convergence

ACI Anywhere

Any Workload, Any Location, Any Cloud

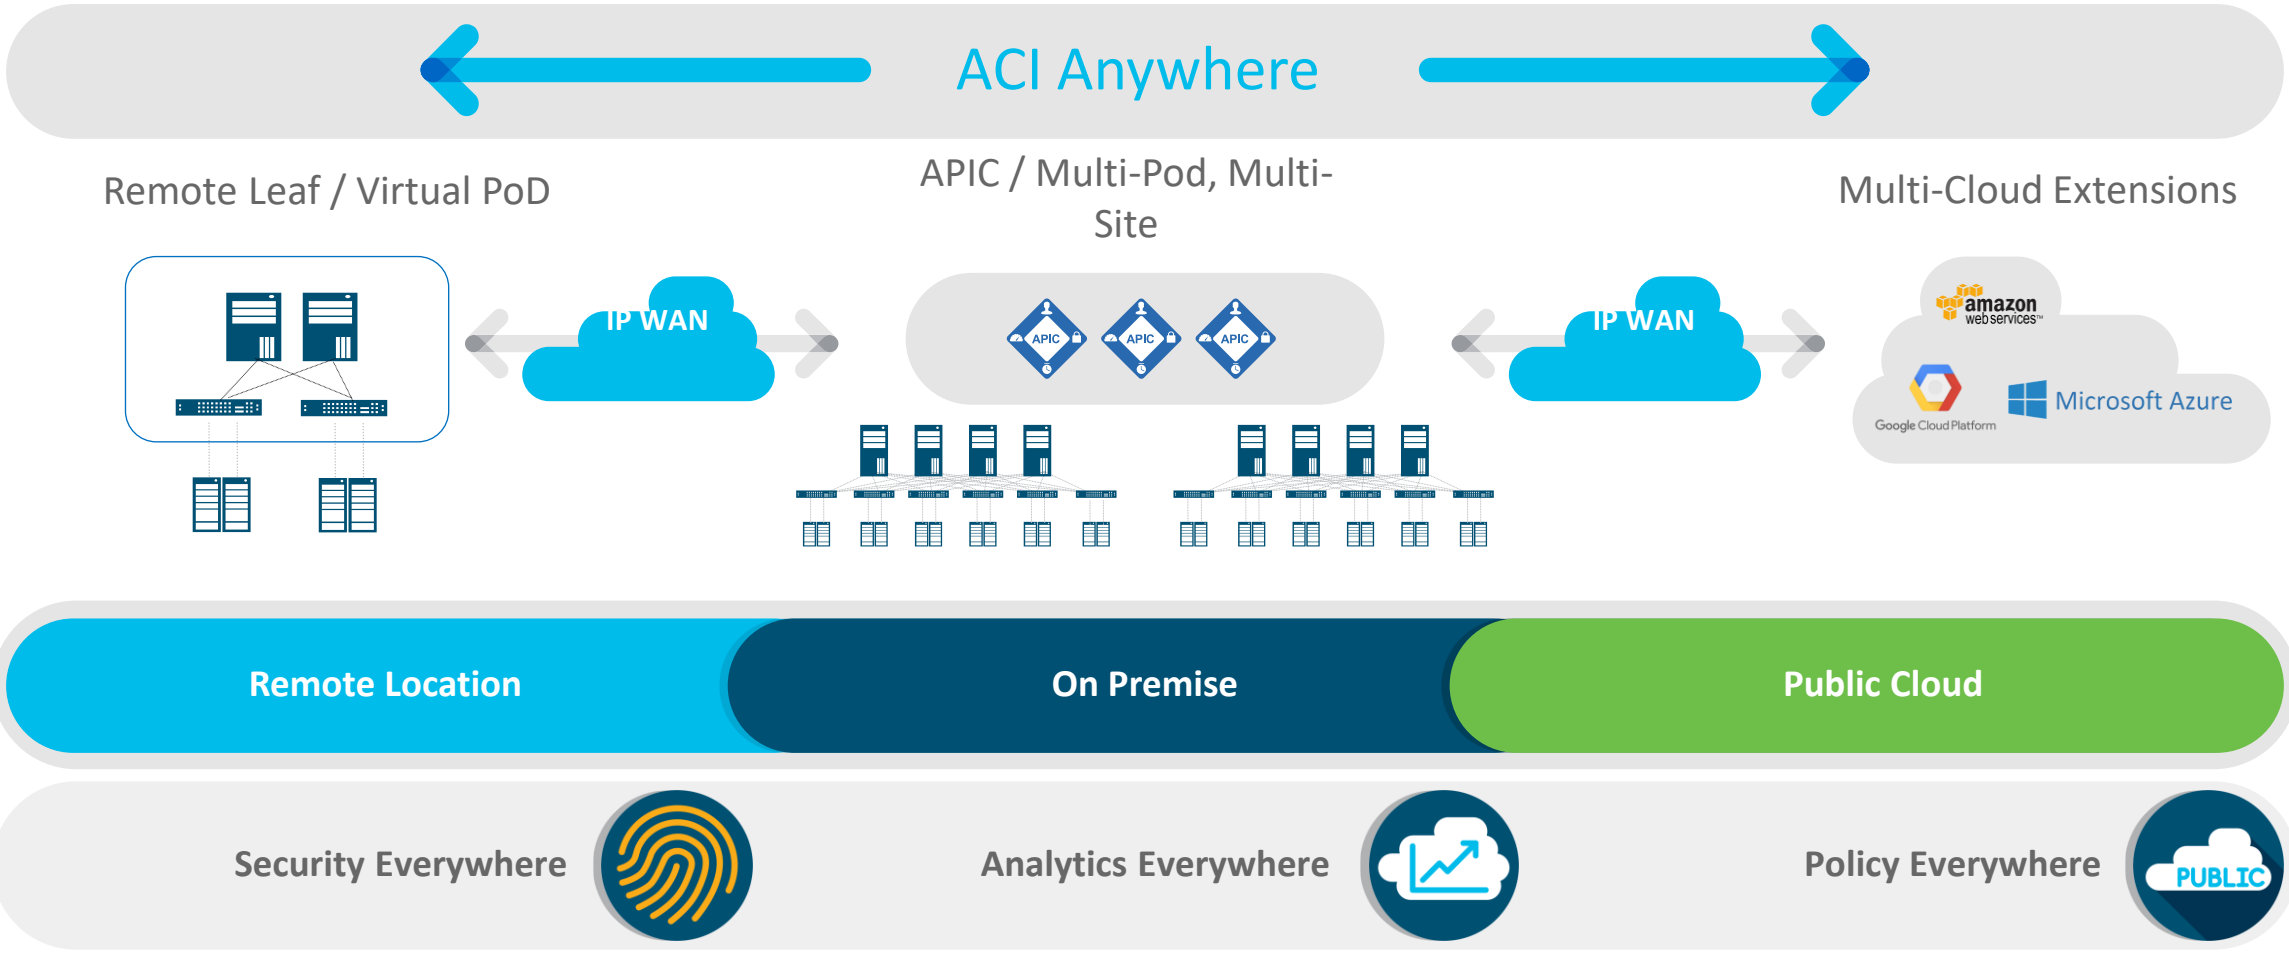

Network Services Integration

Furthering the reach with ACI integrations

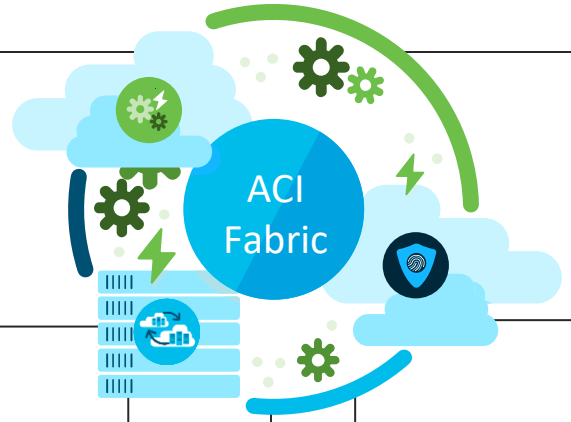

Observability and Analysis

splunk> ScienceLogic zenoss Own IT.
APPDYNAMICS SevOne

Public Cloud

Azure Stack Google Cloud Platform
Azure amazon web services

Automation

SALTSTACK ANSIBLE CHEF puppet labs Terraform

ITSM

bmc Cherwell
servicenow

ADC

HAPROXY A10
CITRIX radware
f5 AVI Networks

Workload Management

docker openstack
Hyper-V vmware

Security

FORTINET CISCO McAfee
Check Point SOFTWARE TECHNOLOGIES LTD. paloalto NETWORKS
radware tufin
Symantec
Infoblox algosec

Connecting ACI to Hypervisor Environment

- ACI can manage existing Hypervisor environments to automate provisioning of network resources for VMs
- Dynamically provisions and configures VLANs on Leafs and interfaces
- Provides ACI insight into VMMs and allows dynamic configuration of Virtual networks
- VMM Domain policy creates a DVS with the name of the VMM Domain policy

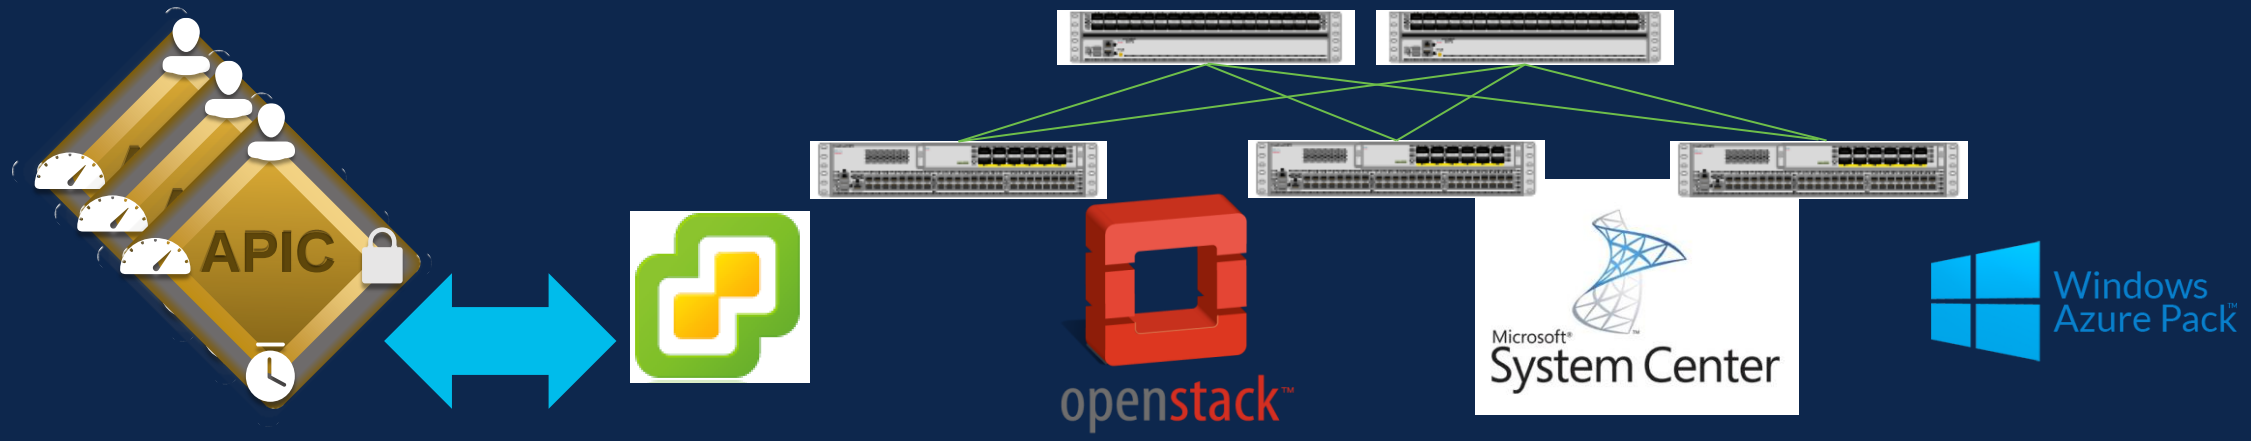

ACI Microsegmentation Overview

What is Micro Segmentation?

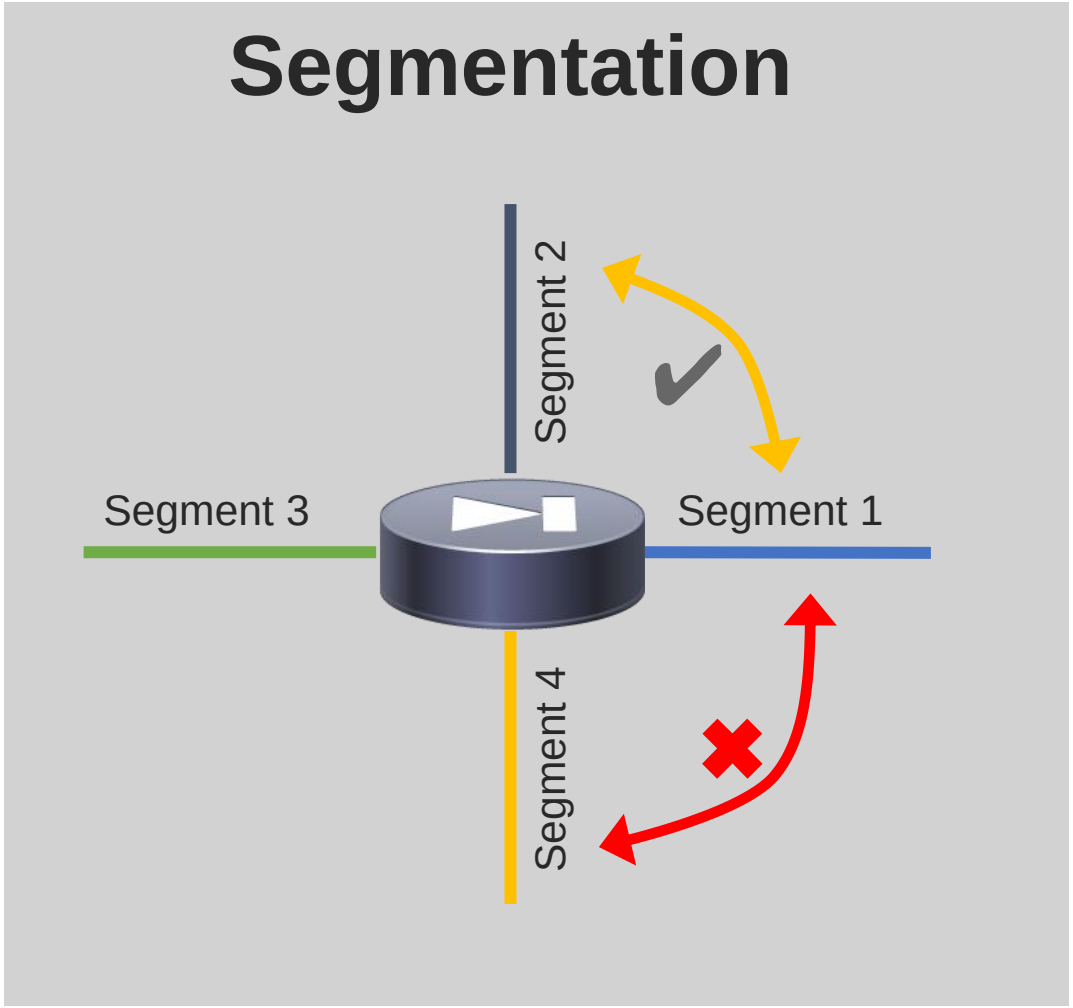

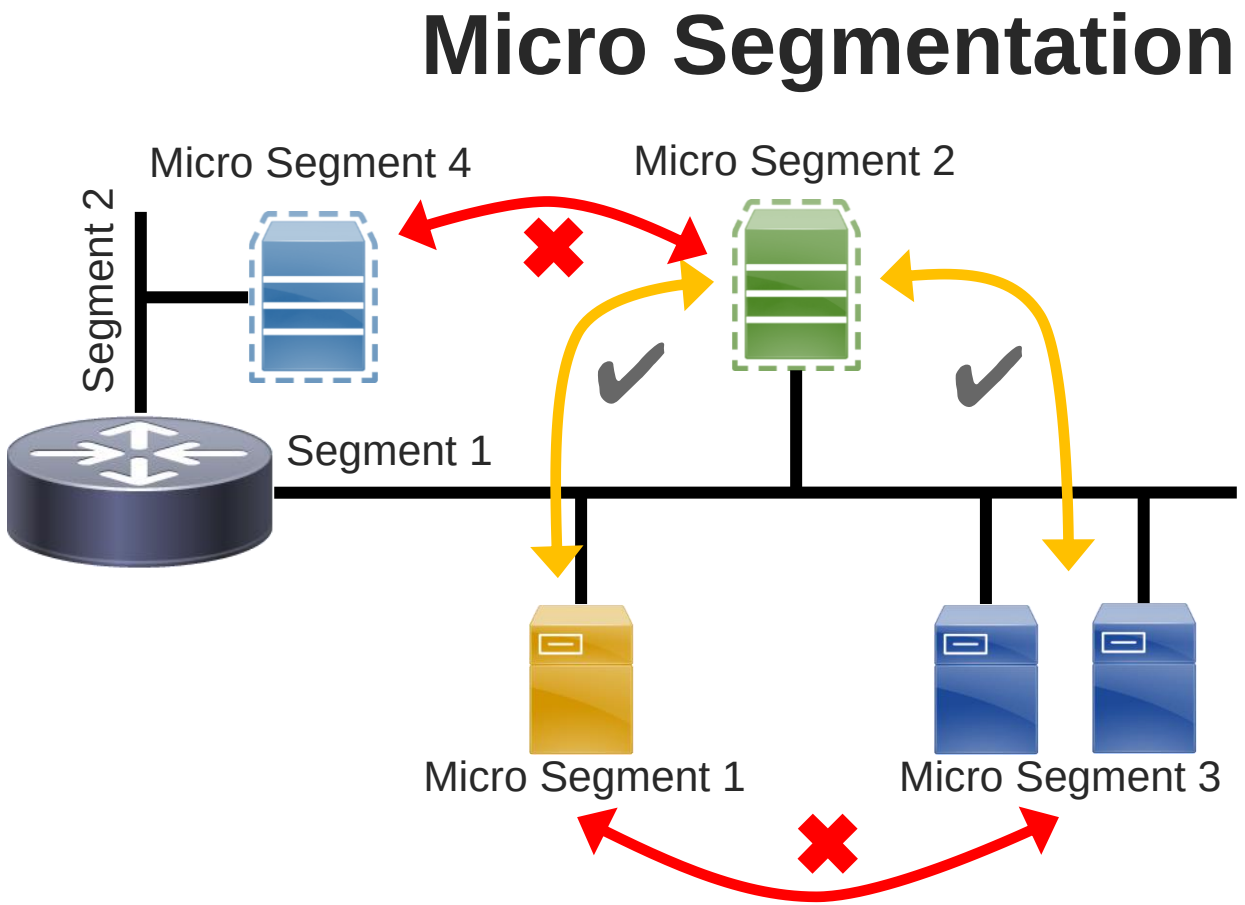

Micro Segment = Endpoint or Group of Endpoints

Segment = Broadcast domain / VLAN / Subnet

Why Micro Segmentation?

- Perimeter security is not enough: once breached, lateral movement can allow attackers to compromise much more
- Improve the security posture inside the Data center
- Minimize segment size and provide smallest exposure to lateral movement
- Changing data center landscape in the multi-cloud era

Inside our 4 Walls

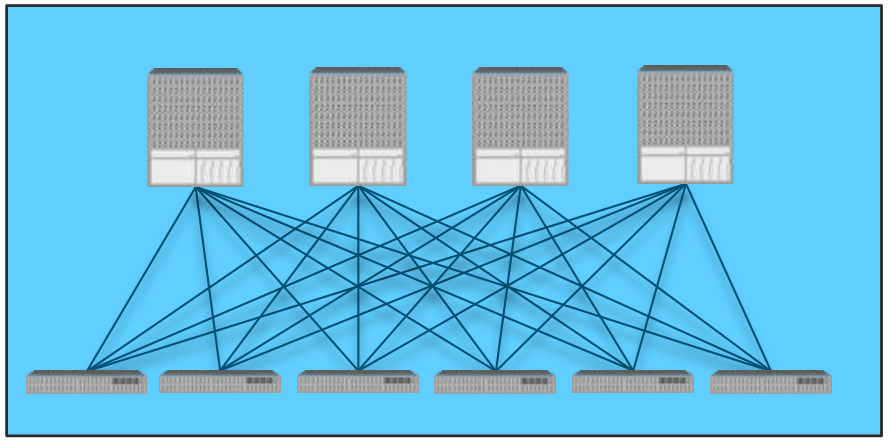

Outside our 4 Walls

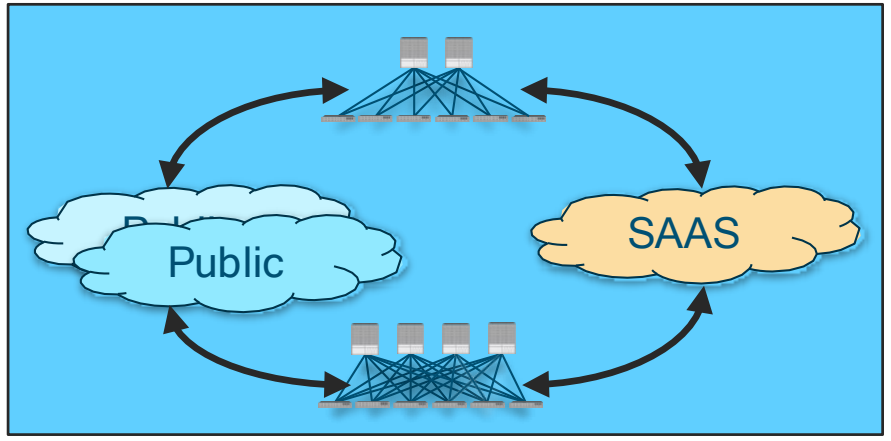

A Micro Segmentation Use Case

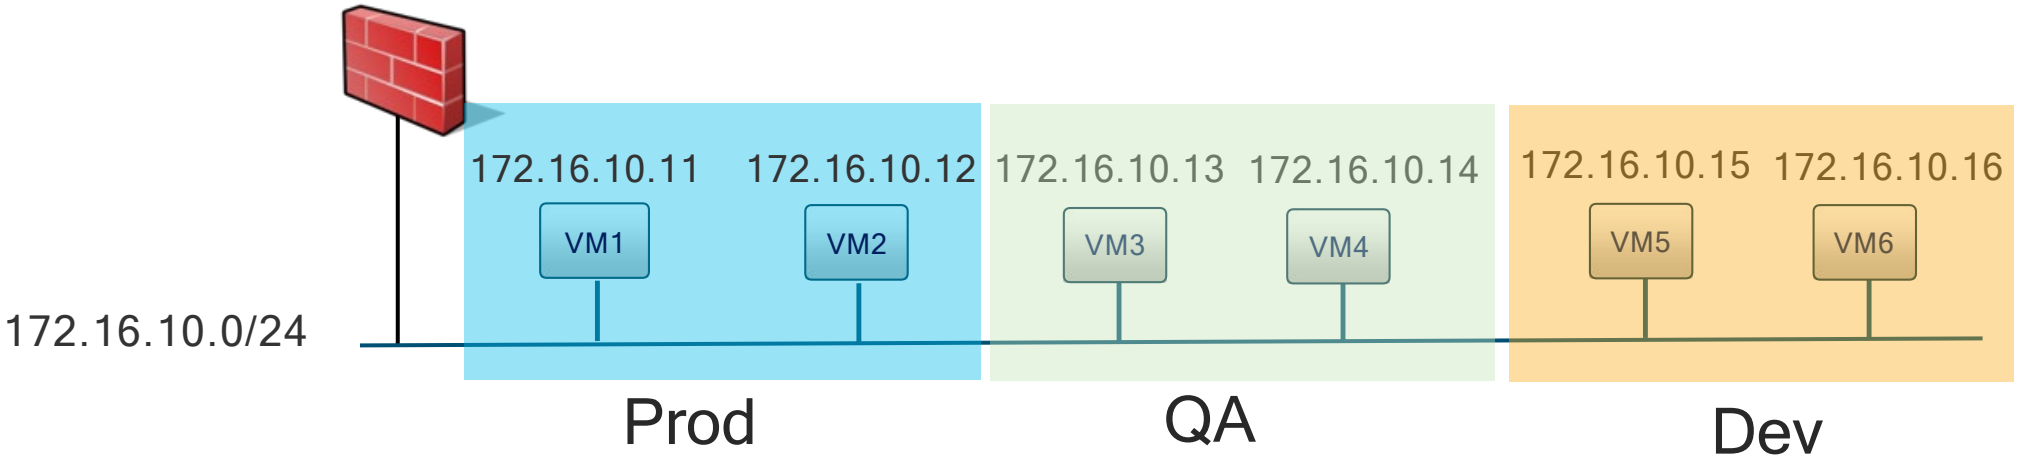

End Point Group

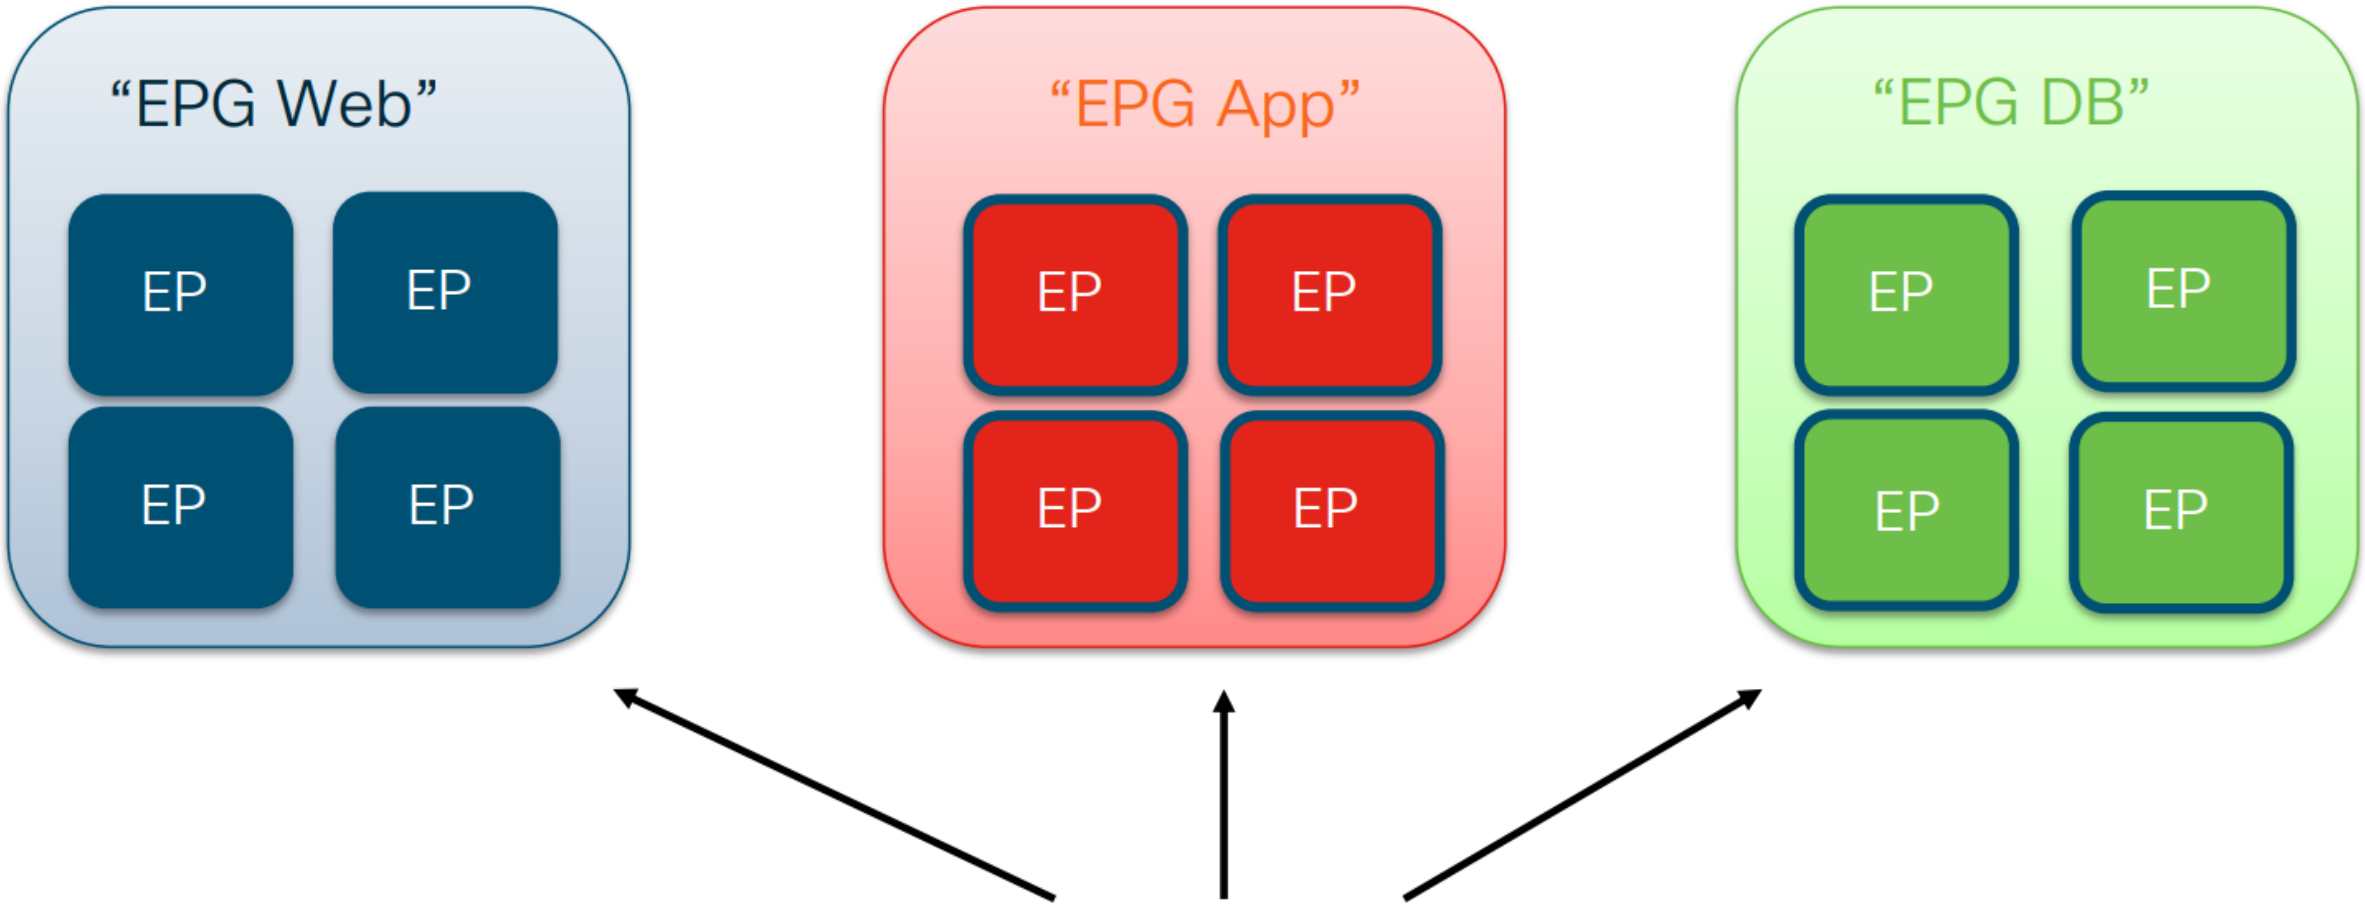

In the ACI model, we do this using the End Point Group (EPG).

Contracts – IP Agnostic Policy Definitions

- Contracts are semantics to specify EPG to EPG communication in ACI
- Communication policy includes filters (ACLs), QoS and Service Graphs
- Contract filters are similar to Access Control Lists
- Contracts can be defined between EPGs or between L3out External EPGs and regular EPGs

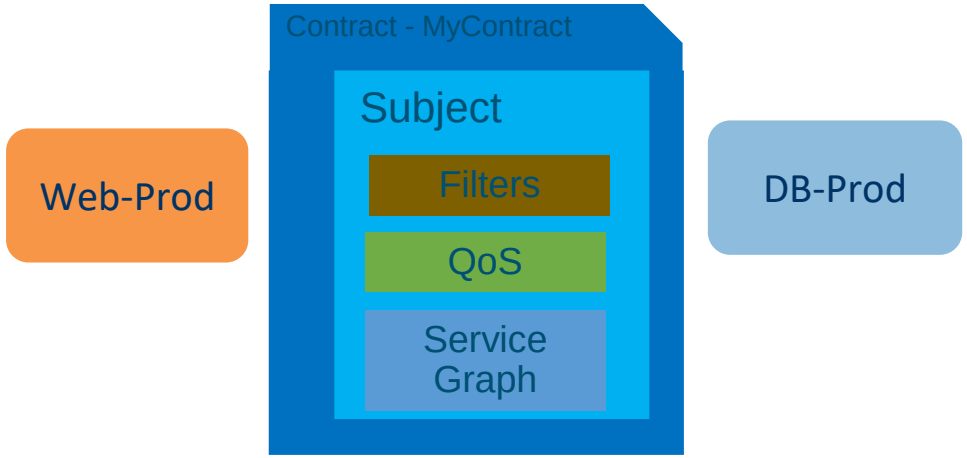

By default communication between EPGs is not allowed in absence of contracts

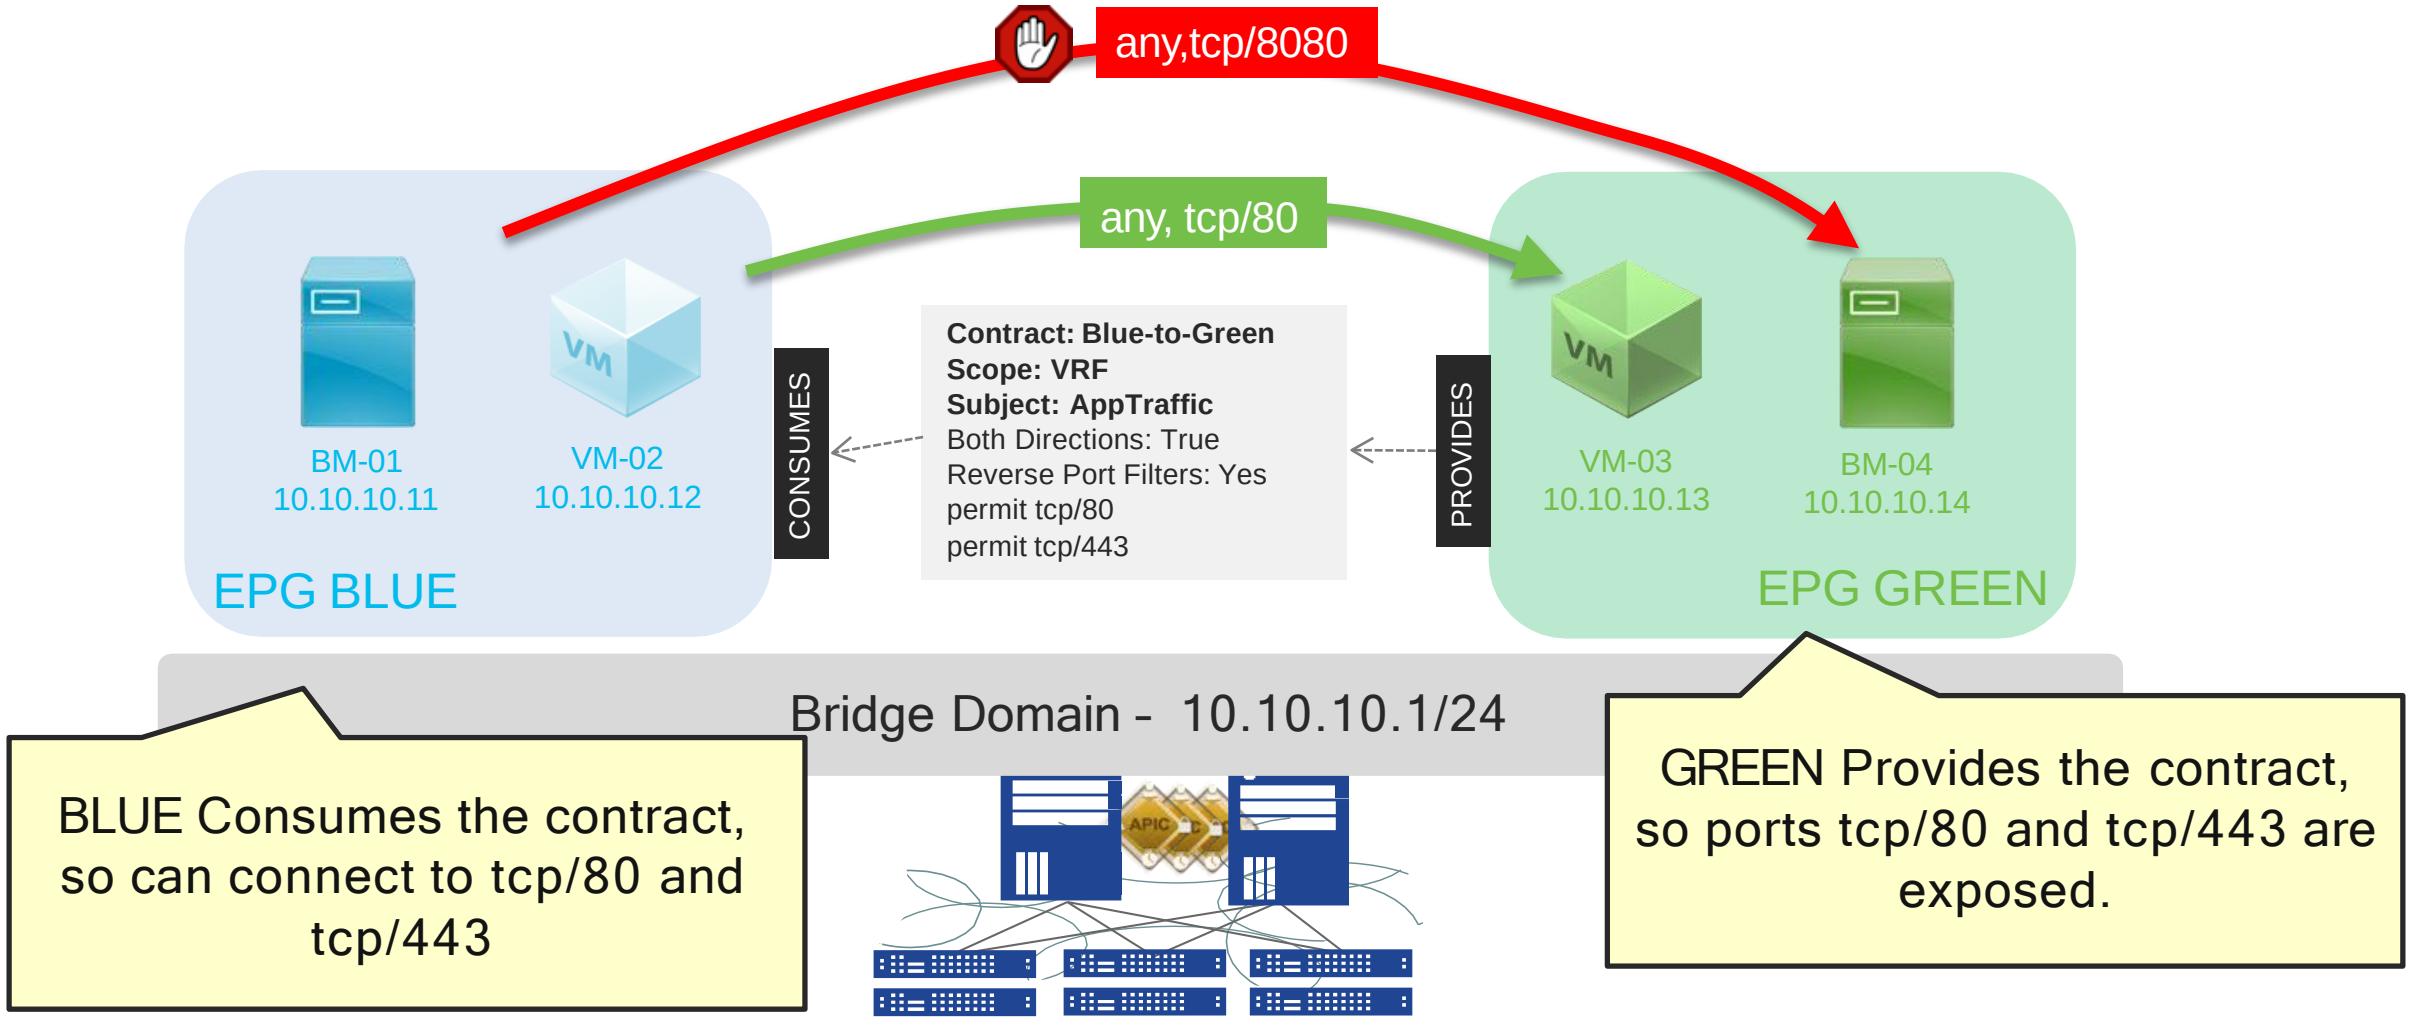

Cisco ACI Supports Flexible East-West Security Models

L4 Stateless Security

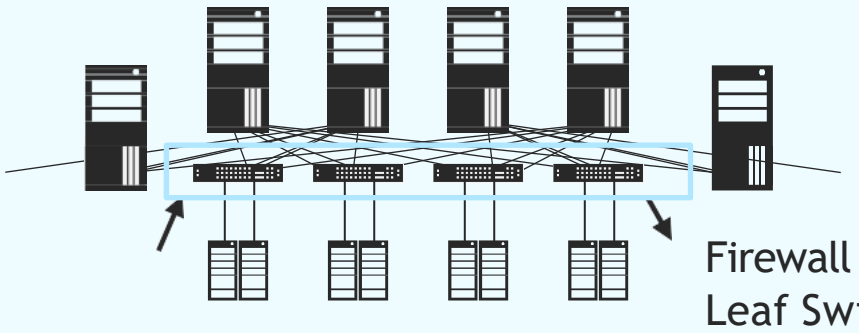

L4 Distributed Stateless Firewall

L4 Stateless Firewall Attached to Every Server Port

Line Rate Policy Enforcement

Policy Follows Workloads

L4-7 Visibility and Control

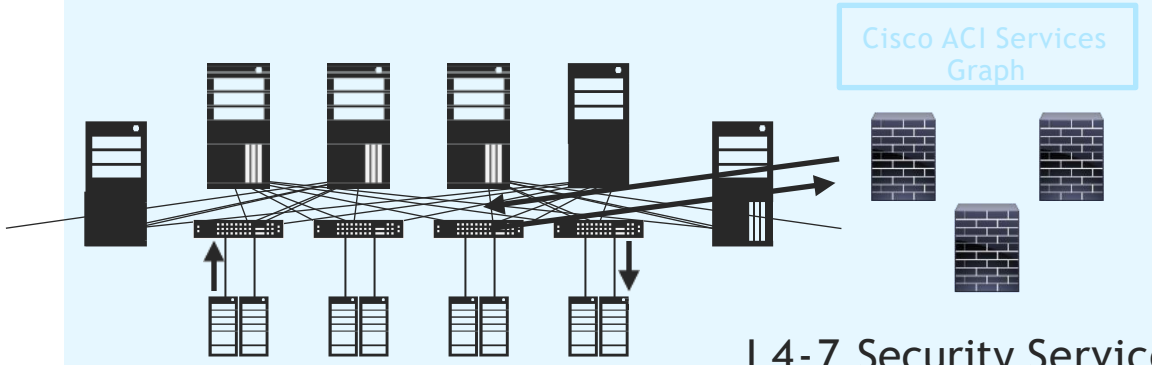

L4–7 Security via Cisco ACI™ Service Graph

Advanced Protection with NGFW, IPS/IDS, DDoS Services Insertion

Sizing at Scale: Can add ASA Cluster

L4-7 Security Policy Applied Consistently for Any Workload

ACI Operations

Day -2 Operations

Navigating to the System Health Dashboard will identify the switch that has a diminished health score

Double clicking on that leaf will allow navigation into the faults raised on that device. Here we click on rtp_leaf1

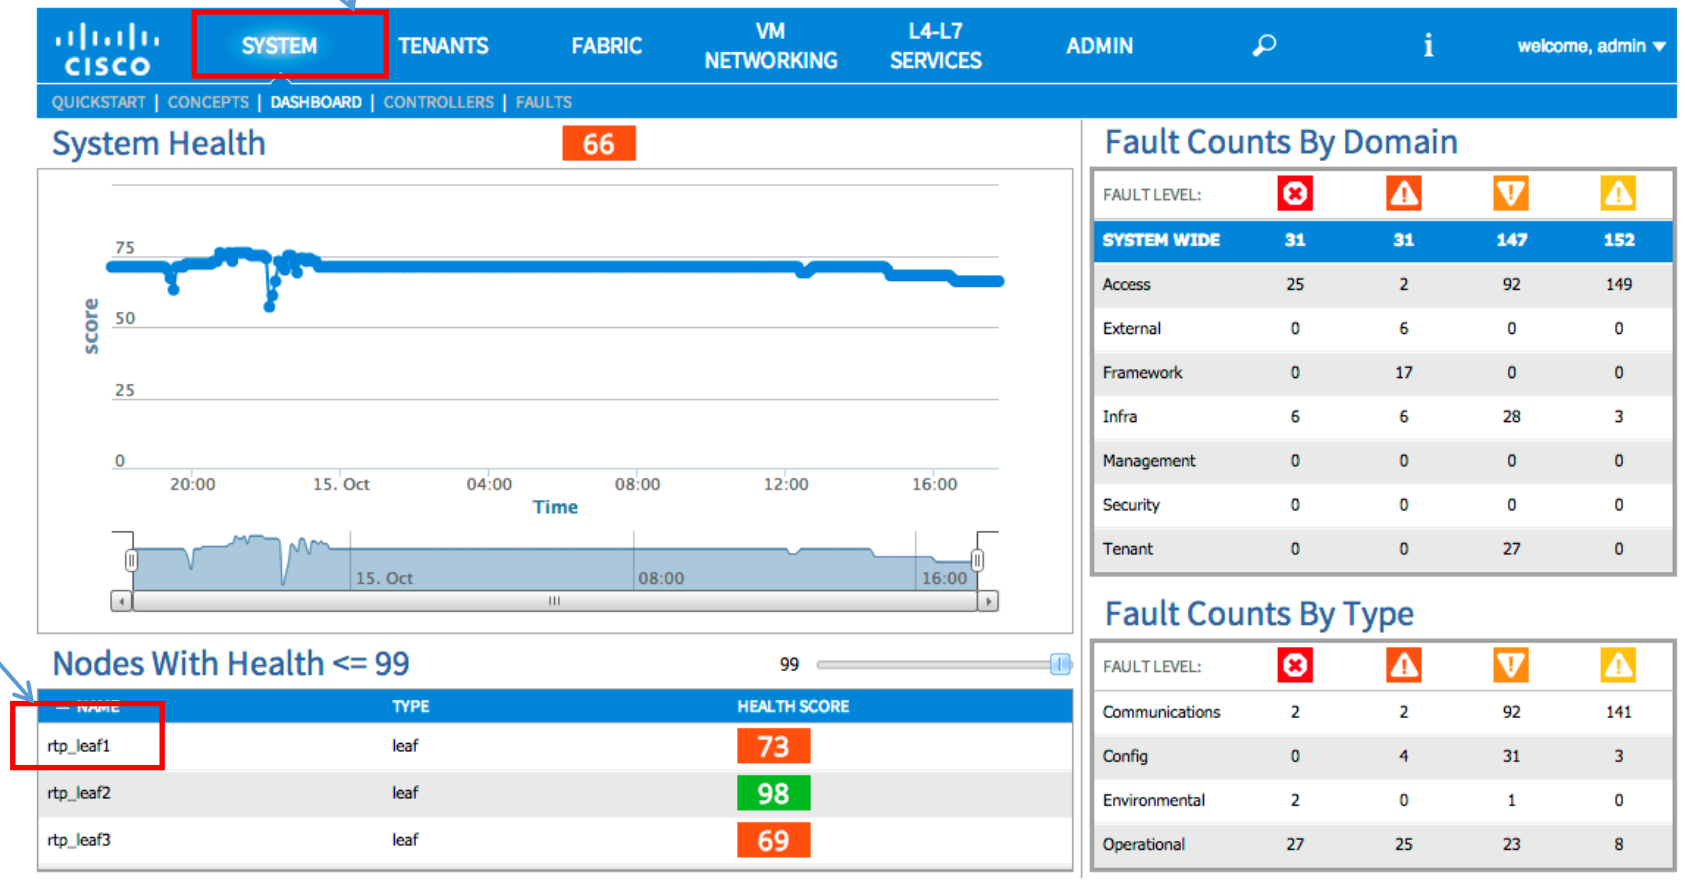

Infrastructure and Services Backups - Snapshots

- Rollback feature allows config rollback between 2 snapshots
- Can also compare differences between a previous S S

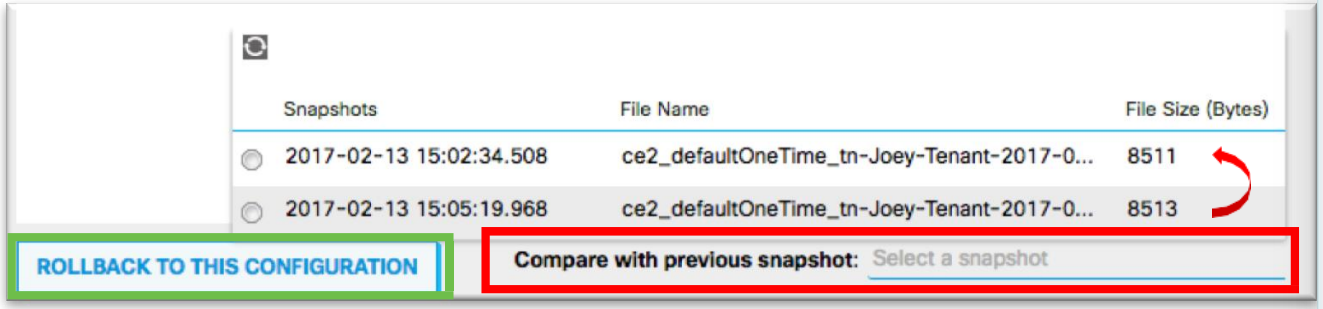

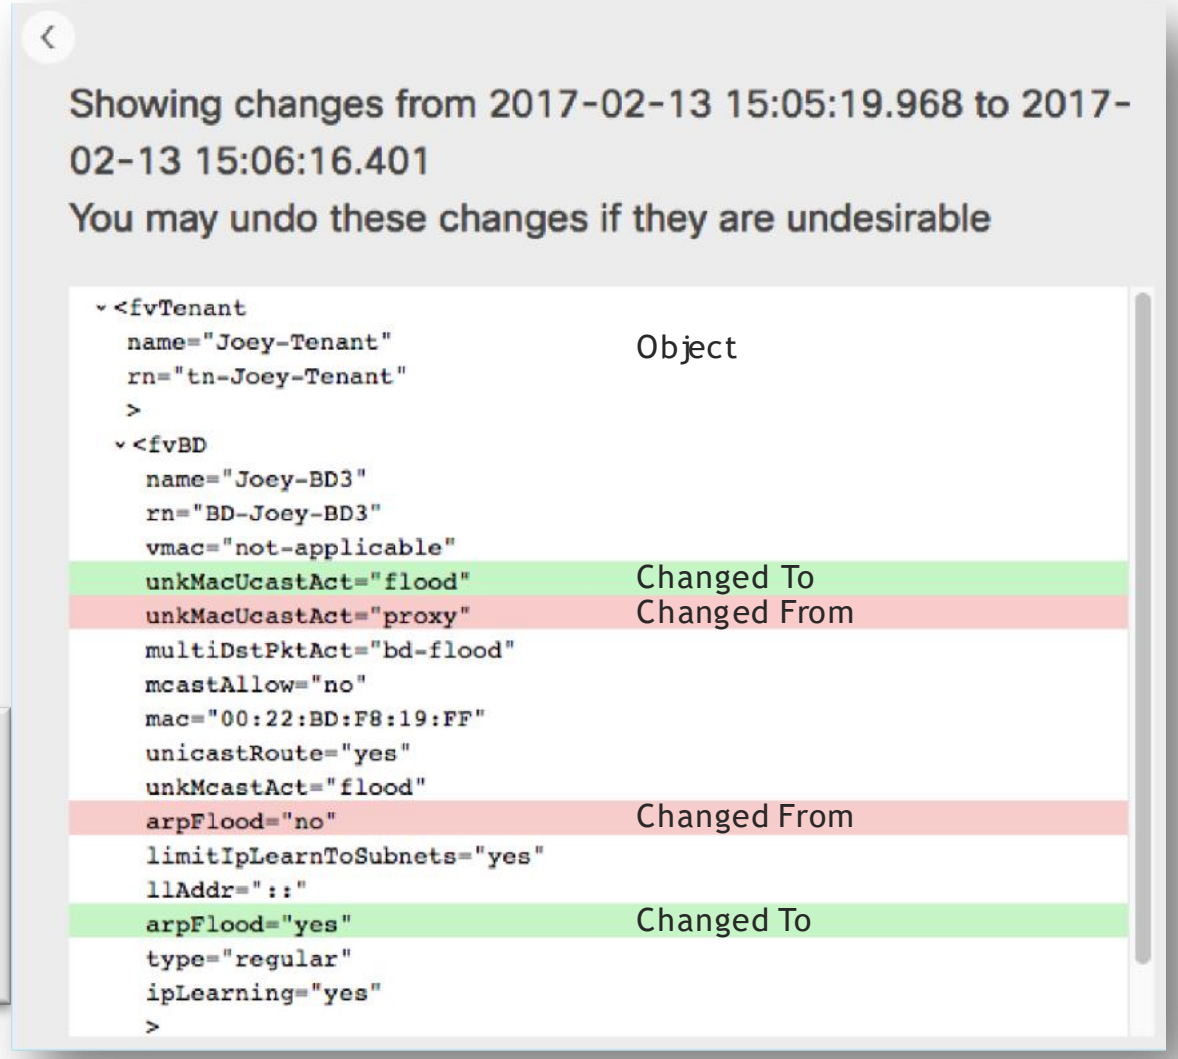

Troubleshoot specific flow

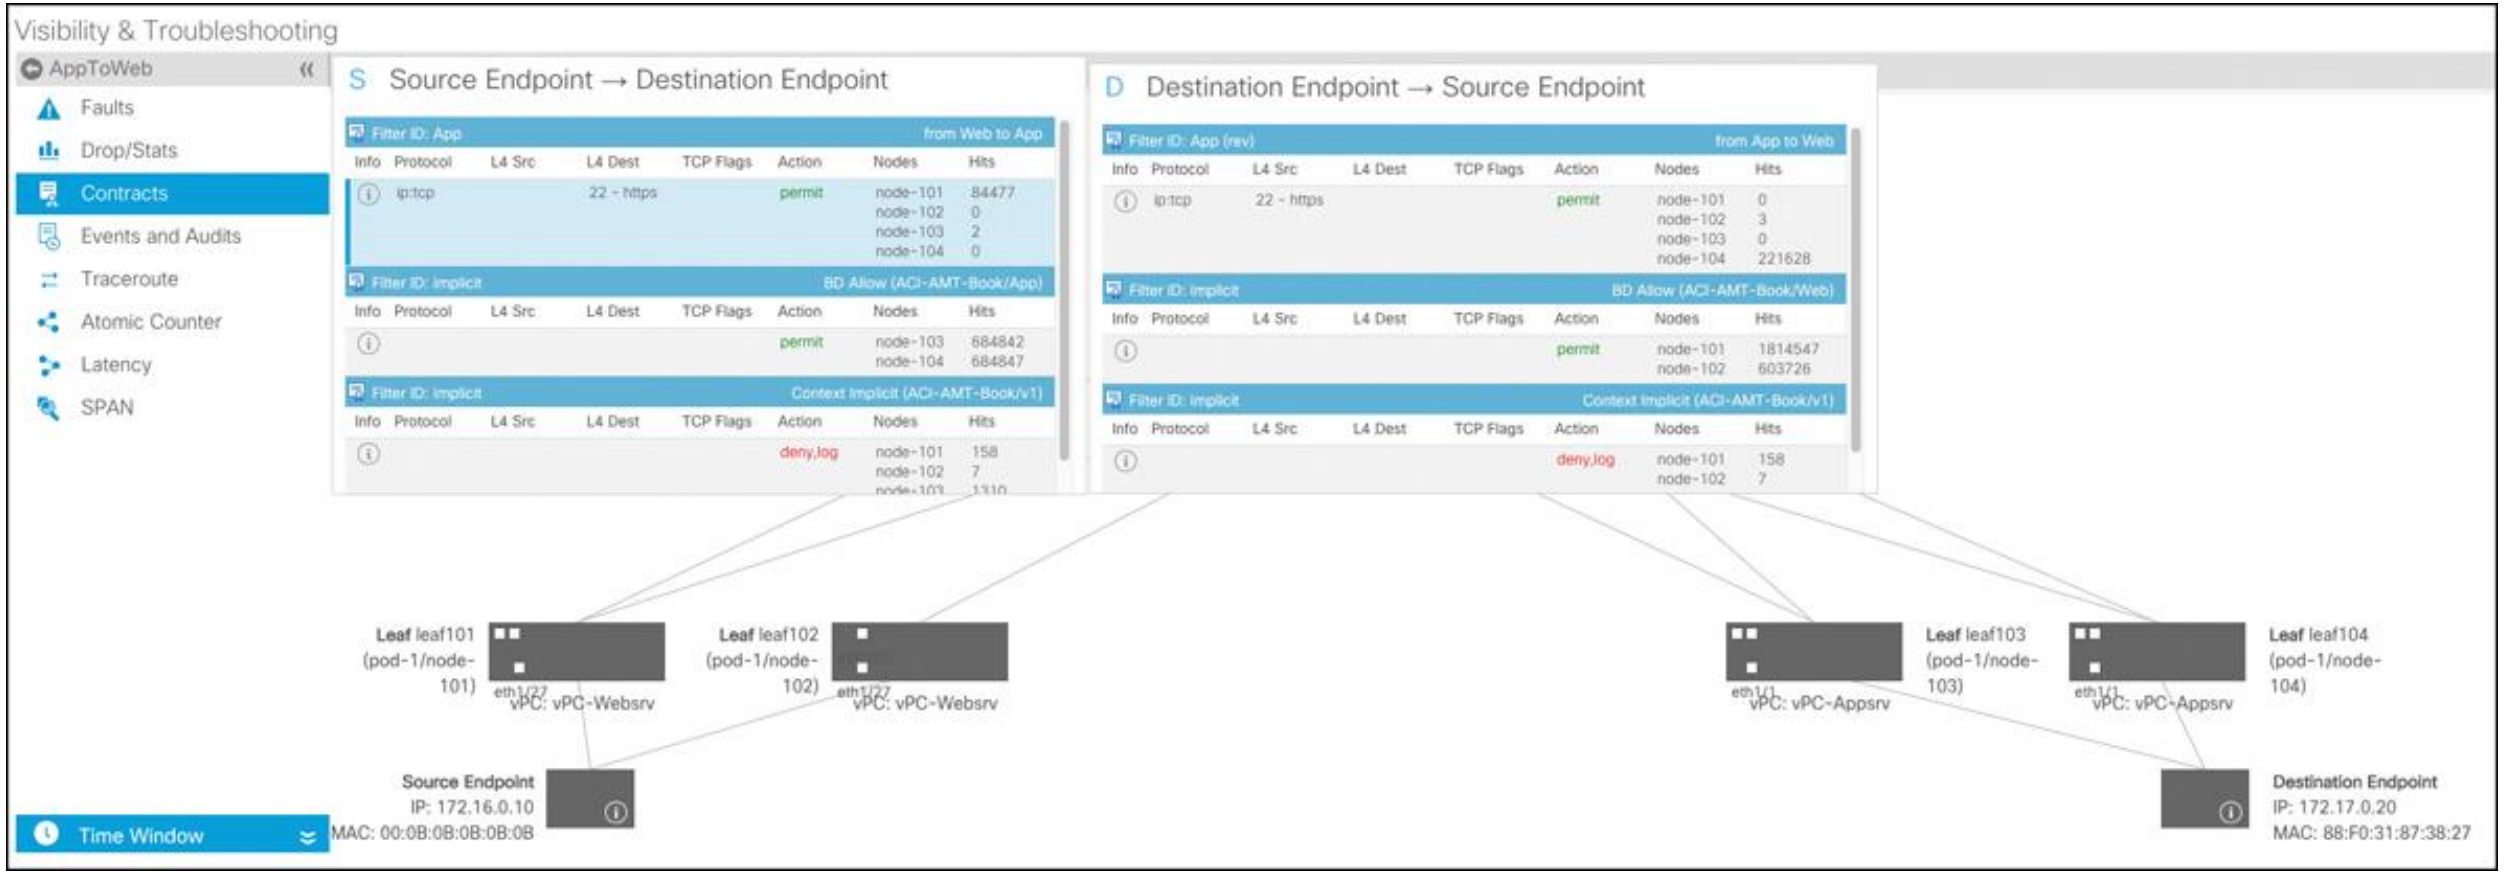

Troubleshoot specific endpoint

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

Visibility & Troubleshooting | Capacity Dashboard | ACI Optimizer | **EP Tracker** | Visualization

EP Tracker

End Point Search

00:00:DE:AD:00:01 [Search]

Learned At	Tenant	Application	EPG	IP
Pod:1, Leaf:101, Port:eth1/23	AutoLab	APP1	EPG1	192.168.101.1

State Transitions

Date	IP	MAC	EPG	Action	Node	Interface	Encap
2020/08/10 18:24:39	192.168.101.1	00:00:DE:AD:00:01	AutoLab/APP1/EPG1	attached	Pod-1/Node-101	eth1/23	vlan-8
2020/08/10 18:24:16	192.168.101.1	00:00:DE:AD:00:01	AutoLab/APP1/EPG1	detached	Pod-1/Node-101	eth1/23	vlan-8
2020/08/10 18:21:59	192.168.101.1	00:00:DE:AD:00:01	AutoLab/APP1/EPG1	attached	Pod-1/Node-101	eth1/23	vlan-8
2020/08/10 18:19:01	192.168.101.1	00:00:DE:AD:00:01	AutoLab/APP1/EPG1	detached	Pod-1/Node-101	eth1/23	vlan-3801
2020/08/07 17:10:21	192.168.101.1	00:00:DE:AD:00:01	AutoLab/APP1/EPG1	attached	Pod-1/Node-101	eth1/23	vlan-3801
2020/08/07 17:03:21	192.168.101.1	00:00:DE:AD:00:01	AutoLab/APP1/EPG1	detached	Pod-1/Node-101	eth1/23	vlan-3801

Summary

Challenges of Traditional Network

Complicated Topology

Core/Dist/Access layer separation

Spanning Tree for Loop Free Topology

No default security isolation

CLI to every Device

Harder as we scale

Lots of copy and paste

Static Configuration
No Automation

Coordination between Network and Server Team

ACI Benefits

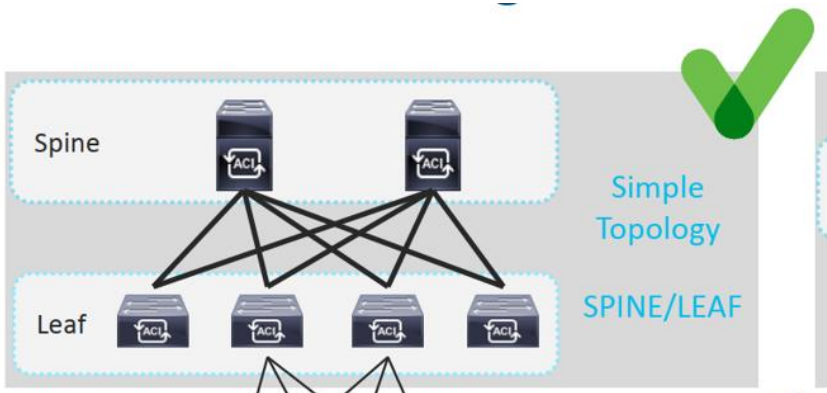

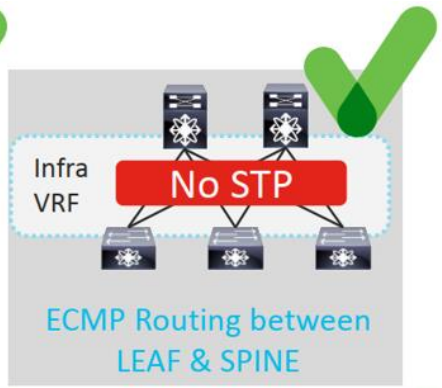

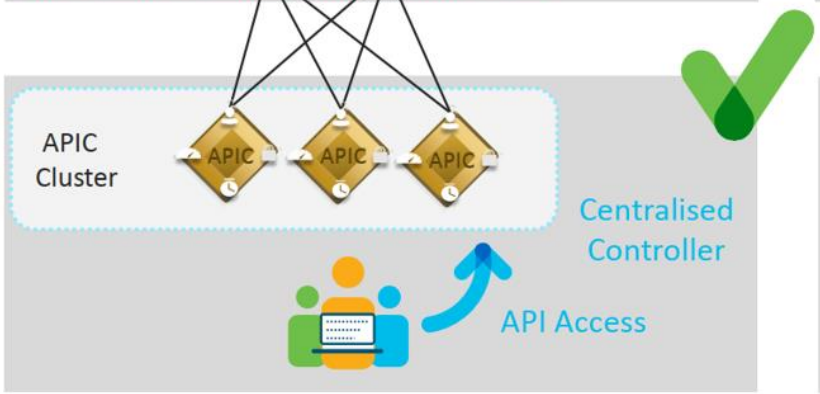

Simplified topology
- Easiest tshoot.
- Quick deployment of new switches.

Loop free network
- Easiest tshoot.
- Increased availability.

Advanced security – Micro segmentation
- Zero trust security model
- Communication that is explicitly defined is allowed.

Centralized management
- Single point of management for whole infrastructure
- Reduced possibility of human error.

Rest API support
- Automation for deployment.
- Config

Integration with 3rd parties
- Advanced visibility for network admins.
- Easiest tshoot.

Useful Links

- Cisco ACI Design guide

https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html

- Cisco ACI Best Practises

https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-aci-best-practices-quick-summary.html

- Cisco ACI white paper list

https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-listing.html

- Cisco ACI videos in Youtube

https://www.youtube.com/channel/UC-U0ud423cfgHls0bV2jxXw


Thank you