Cisco TechDay

Application Centric Infrastructure

Design pro řešení na zelené louce i do stávajícího DC

DCA4

Miroslav Brzek, Systems Engineer

Agenda

- Modern DC infrastructure – Customer requirements
- What’s Application Centric Infrastructure (ACI)
- How ACI affects / enhances Data Center Infrastructure
    - Network fabric
    - Hypervisors
    - L4-L7 Services
- How ACI affects Applications
    - Security
    - Automation / Orchestration
- Migration to ACI

Modern DC infrastructure – Customer's Requests

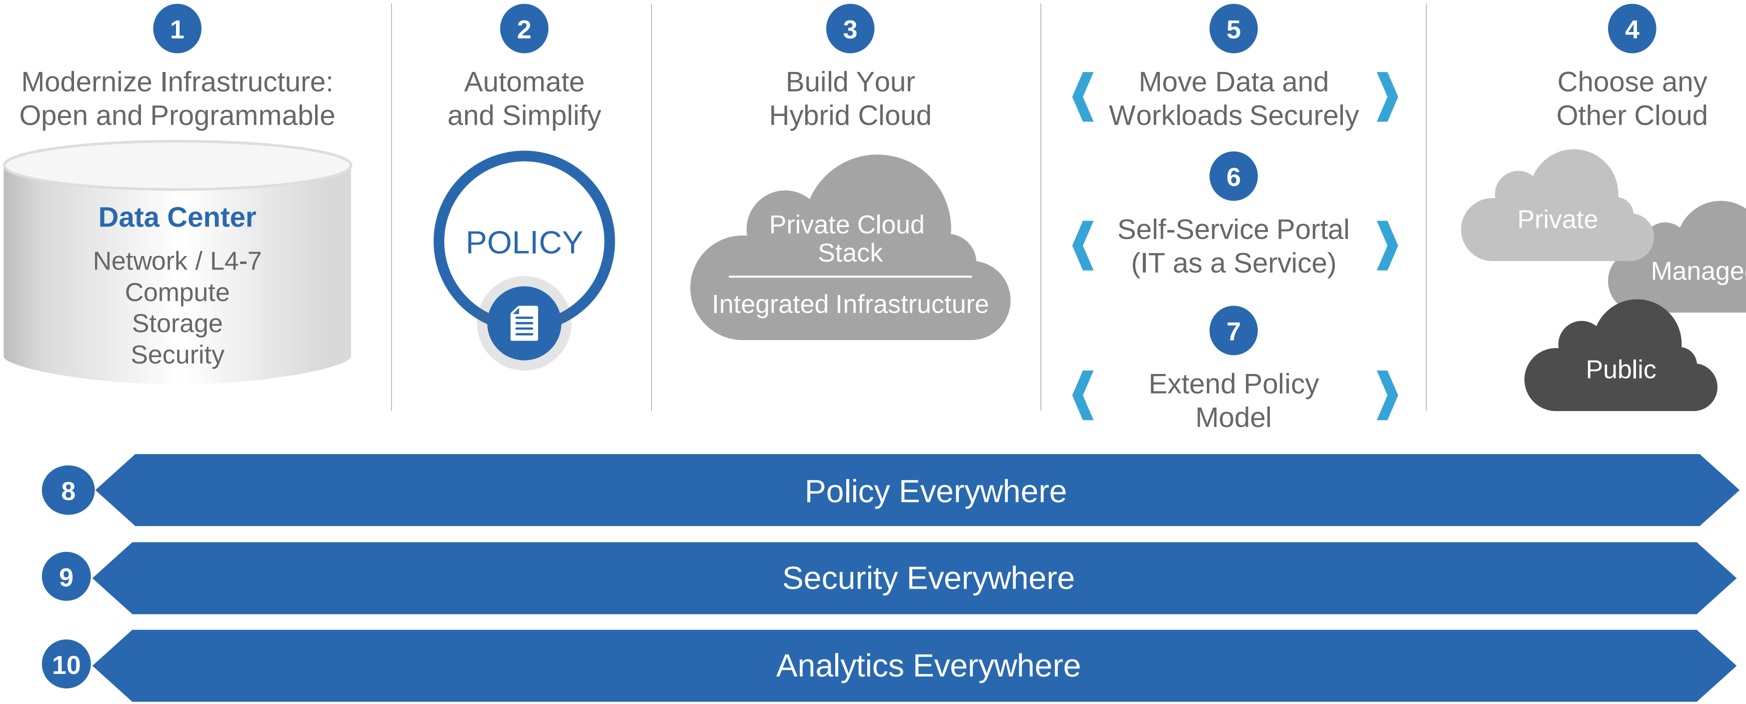

1

Modernize Infrastructure: Open and Programmable

Data Center

Network / L4-7
Compute
Storage
Security

2

Automate and Simplify

POLICY

3

Build Your Hybrid Cloud

Private Cloud Stack

Integrated Infrastructure

5

Move Data and Workloads Securely

6

Self-Service Portal (IT as a Service)

7

Extend Policy Model

4

Choose any Other Cloud

Private

Managed

Public

8 Policy Everywhere

9 Security Everywhere

10 Analytics Everywhere

Cisco Data Center Networking Strategy

Providing Choice in Automation and Programmability

Programmable Network

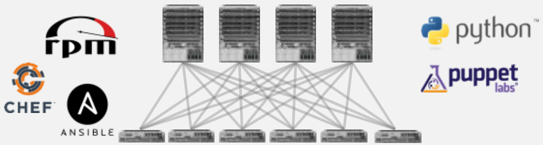

Modern NX-OS with enhanced NX-APIs

DevOps toolset used for Network Management
(Puppet, Chef, Ansible etc.)

Programmable Fabric

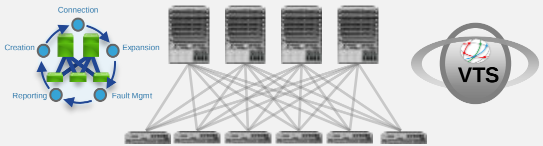

VxLAN-BGP EVPN standard-based

3rd party controller support

Cisco Controller for software overlay provisioning and management across N2K-N9K

Application Centric Infrastructure

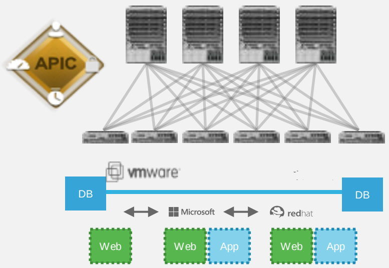

Turnkey integrated solution with security, centralized management, compliance and scale

Automated application centric-policy model with embedded security

Broad and deep ecosystem

Automation, API’s, Controllers and Tool-chain’s

Agenda

- Modern DC infrastructure – Customer requirements
- **What's Application Centric Infrastructure (ACI)**
- How ACI affects / enhances Data Center Infrastructure
  - Network fabric
  - Hypervisors
  - L4-L7 Services
- How ACI affects Applications
  - Security
  - Automation / Orchestration
- Migration to ACI

ACI: Policy-Driven DC Infrastructure

Answers customer’s requests

App Requirements Drive
Network Deployment/Operation

Agile

Policy Automation

Visibility

Scale and Performance

- Speed through Automation
- Physical and Virtual Endpoints with Consistent Policy
- Application Health Monitoring

Open

Open API’s

Partner Ecosystem

- Open APIs, Open Source and Open Standards
- Customer Choice And Interoperability
- Drives Innovation

Secure

Multi-Tenant Security

Compliance

- Whitelist Approach
- Multitenant Aware
- Simplified Compliance

Modern Data Center Network

It's All About the Application

An Application is more than just a VM

Interconnected components

How do we define the network for the application?

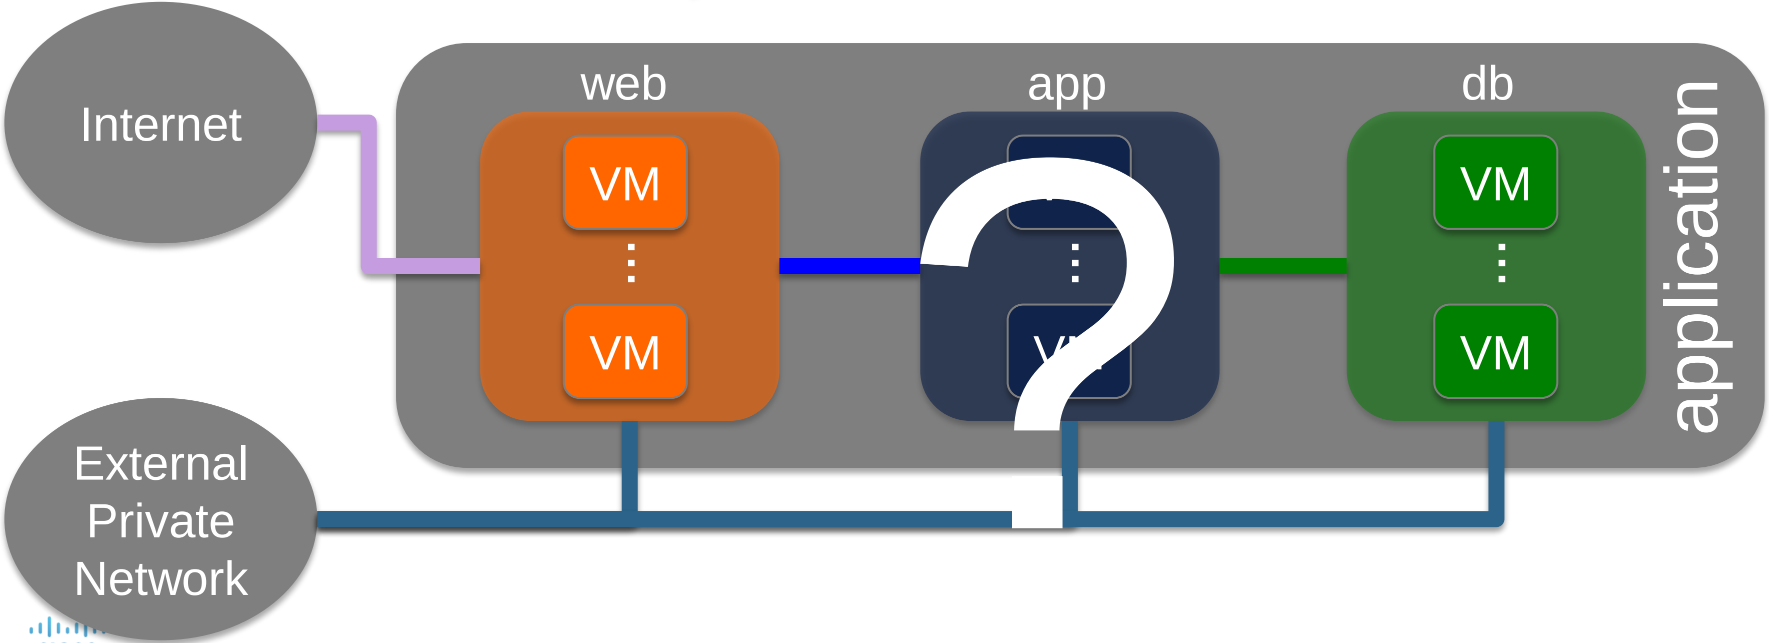

How do we define the network for the application today?

- Group applications by VLAN to segment them and to control the path between them
- Map IP subnets to those VLANs
  - Policy boundary
  - Security identifier
  - Application identifier
- Apply connectivity, policies (security, QoS) and network services based on those constructs

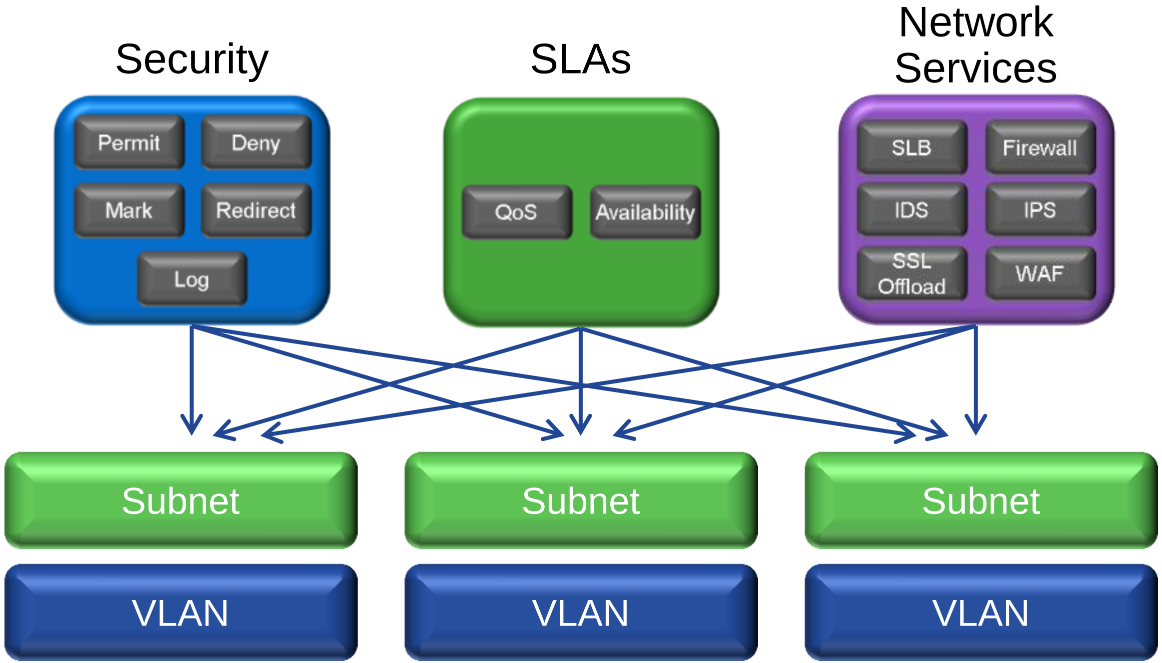

This leads to restrictions on how applications can be grouped and how policy can be applied

```
switch1(config)#
switch1(config)# int eth 1/1
switch1(config)# switch mode acc
```

```
switch2(config)#
switch2(config)# int eth 1/2 - 3
```

```
switch3(config)#
switch3(config)# int eth 1/4 - 5
switch3(config)# switch mode acc
```

```
switch4(config)#
switch4(config)# int eth 1/6
switch4(config)# switch mode acc
switch4(config)# switch acc vlan 333
switch4(config)# no shut
switch4(config)# int eth 1/7 - 9
switch4(config)# switch mode acc
switch4(config)# switch acc vlan 333
switch4(config)# no shut
```

```
switch5(config)#
switch5(config)# int eth 1/10 - 11
switch5(config)# switch mode acc
switch5(config)# switch acc vlan 444
switch5(config)# no shut
switch5(config)# int eth 1/11 - 15
switch5(config)# switch mode acc
switch5(config)# switch acc vlan 555
switch5(config)# no shut
switch5(config)# monitor session 1 source vlan 555
switch5(config)# monitor session 1 dest eth 1/16
```

```
switch6(config)#
switch6(config)# int eth 1/16 - 19
switch6(config)# switch mode acc
switch6(config)# switch acc vlan 777
switch6(config)# no shut
switch6(config)# monitor session 1 source vlan
777
switch6(config)# monitor session 1 dest eth 1/20
```

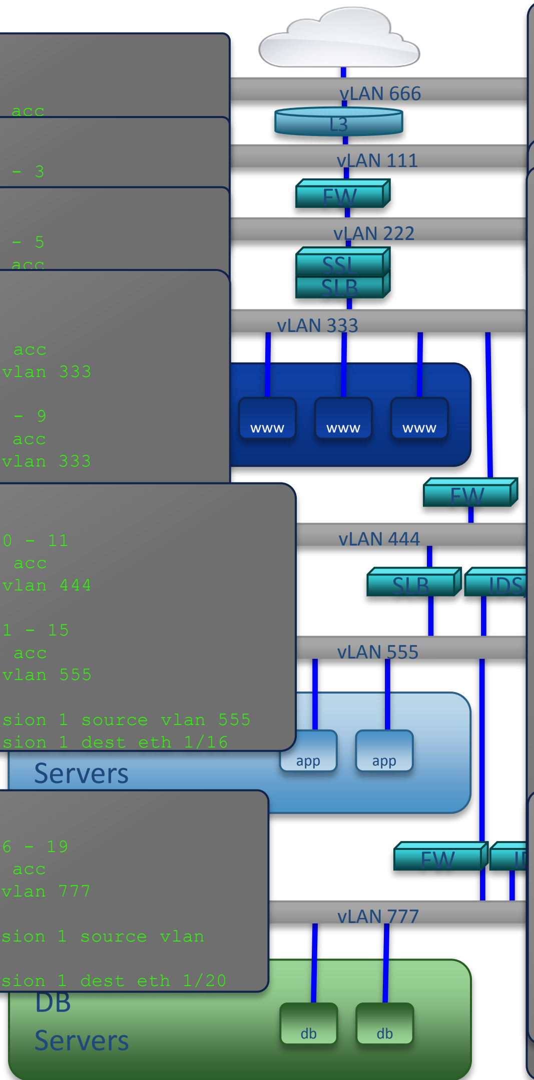

```
router(config)#
router(config)# int eth 1
router(config)# ip add 6.6.6.1 255.255.255.0
router(config)# not shut
router(config)# int eth 2
```

```
fw1(config)#
```

```
slb1 (CONFIG)
probe http http-probe
 interval 30
 expect status 200 200
rserver host websrvr1
  description foo web server
  ip address 3.3.3.1
  inservice
rserver host websrvr2
  description foo web server
  ip address 3.3.3.2
  inservice
```

```
fw2(config)#
fw2(config)# int eth 0/1
fw2(config)# nameif webfront 20
fw2(config)# int eth 0/2
fw2(config)# nameif appfront 50
```

```
slb2 (CONFIG)
rserver host appsrvr1
  description foo app server
  ip address 5.5.5.1
  inservice
rserver host appsrvr2
  description foo app server
  ip address 5.5.5.2
  inservice
rserver host appsrvr3
  description foo app server
```

```
fw3(config)#
fw3(config)# int eth 0/1
fw3(config)# nameif appfront 70
fw3(config)# int eth 0/2
fw3(config)# nameif dbfront 90
fw3(config)# object network db_cluster
fw3(config)# host 7.7.7.7
fw3(config)# nat (dbfront,appfront) static 5.5.5.50
fw3(config)# access-list web_to_app permit tcp any host 5.5.5.50 eq 1433
```

```
policy-map type loadbalance first-match FOO_APP-MATCH
      class FOO_APP
```

ACI policy model brings the concept of End-Point Group (EPG)

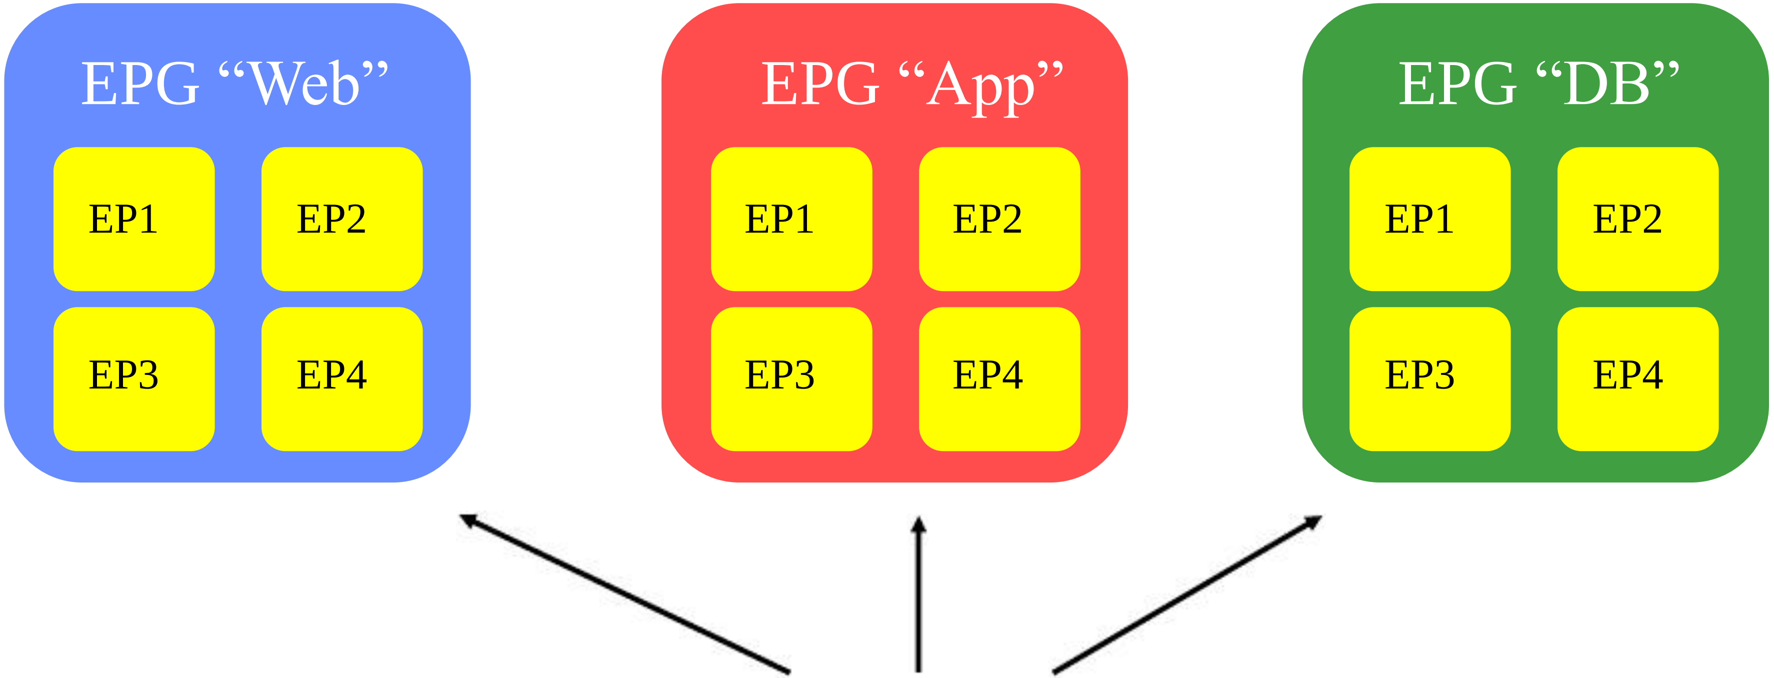

EPGs are a **grouping of end-points** representing **application or application components** **independent** of other network constructs.

End-Points end EPG membership

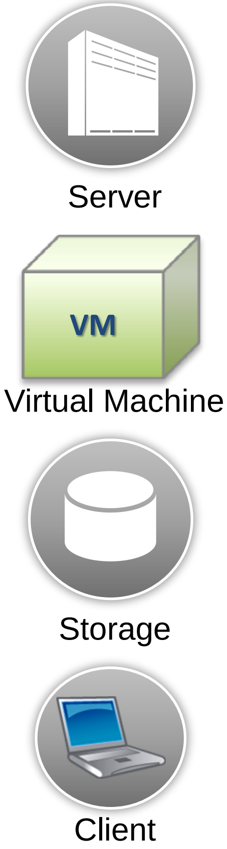

- Device connected to network directly or indirectly
- Has address (identity), location, attributes (version, patch level)
- Can be physical or virtual
- End Point Group (EPG) membership defined by:
  - Ingress physical port (leaf or FEX)
  - VLAN ID
  - VXLAN (VNID)
  - IP address
  - IP Prefix/Subnet (applicable to external/border leaf connectivity)
  - VM Attribute

Applying Policy between EPGs: ACI Contracts

Once we have our EPGs defined, we need to create policies to determine how they communicate with each other -> ACI Contract

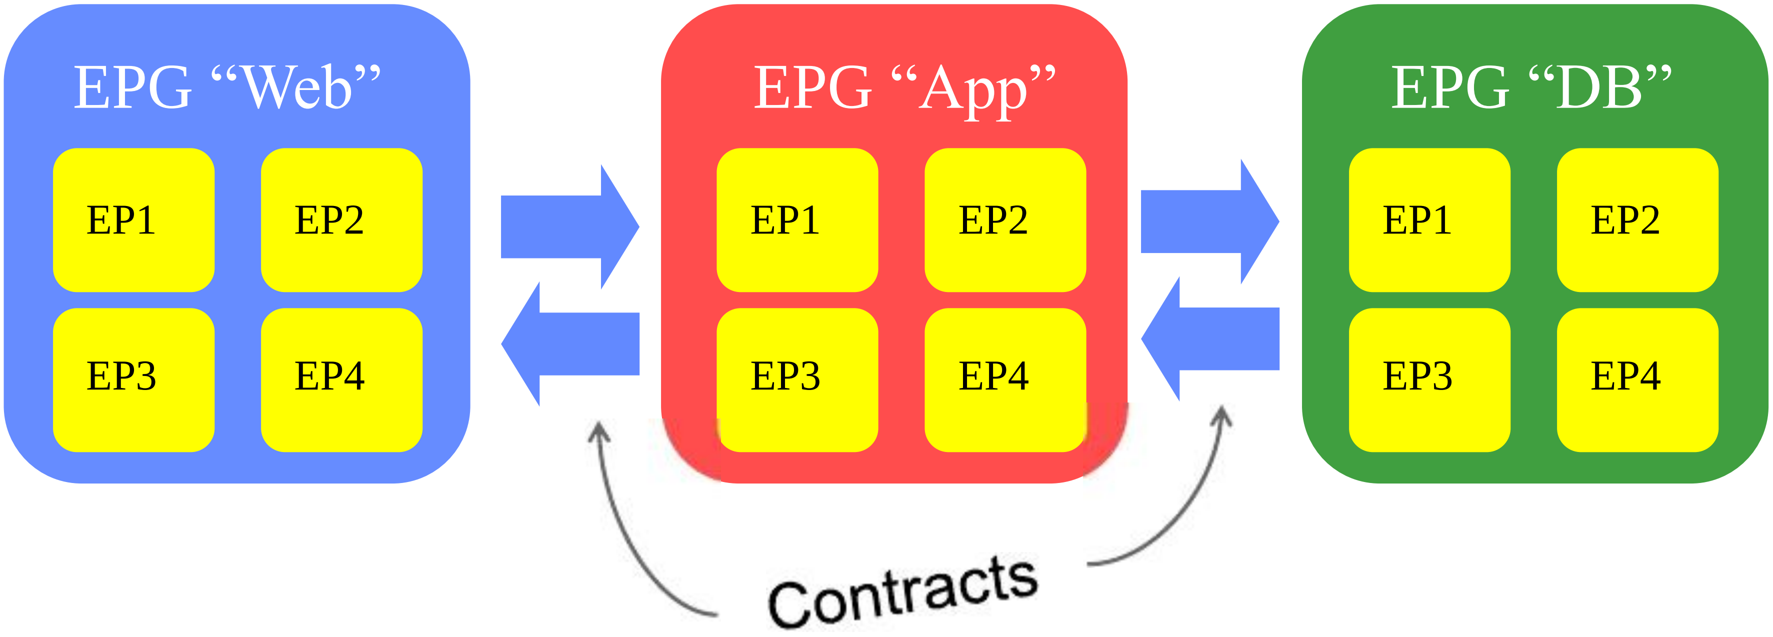

Applying Policy between EPGs: ACI Contracts

A contract typically refers to one or more ‘filters’ to define specific protocols & ports allowed between EPGs

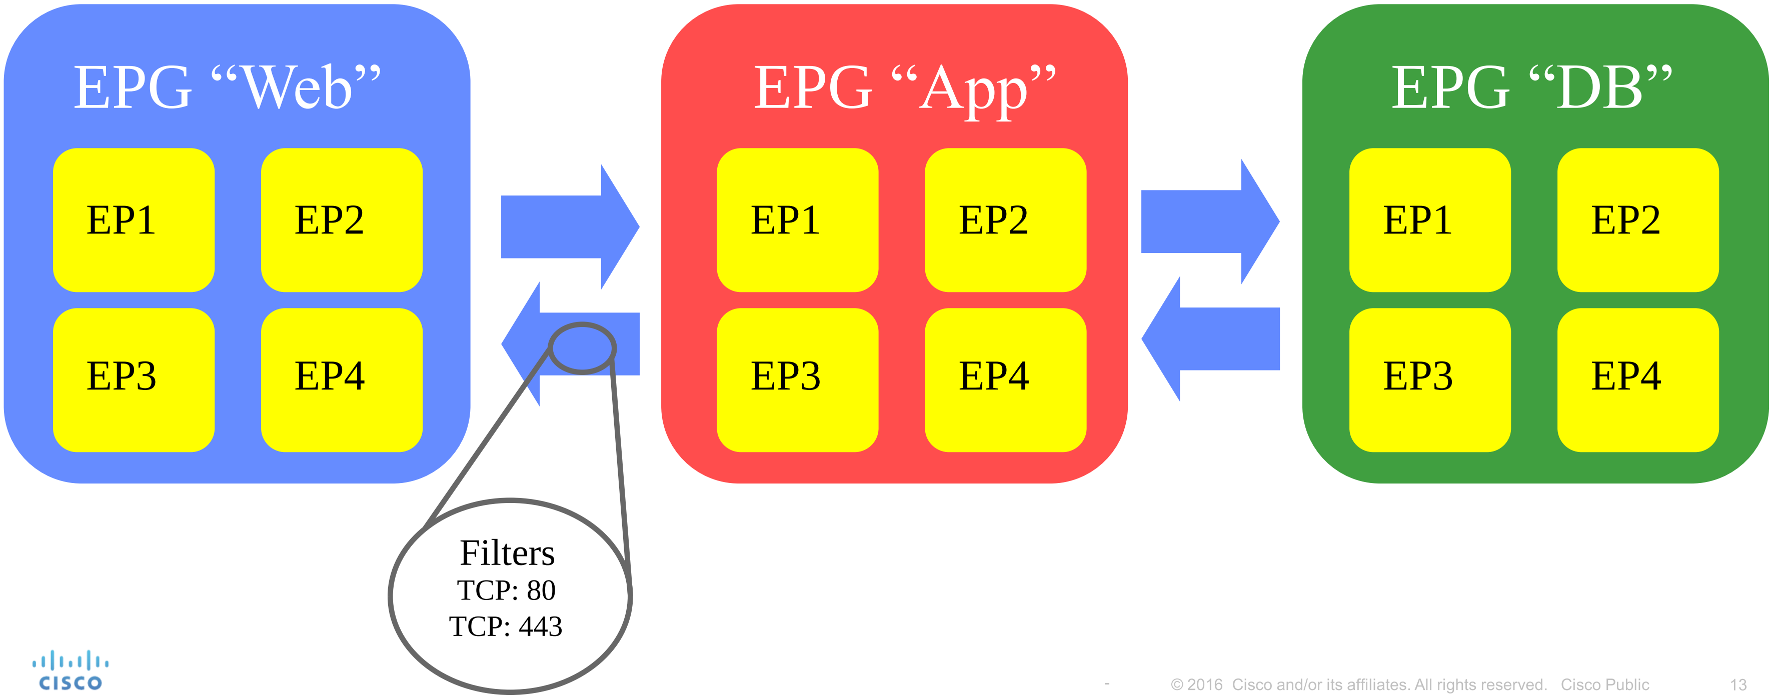

ACI Service Graph

Insertion of Layer 4 - 7 services with contracts

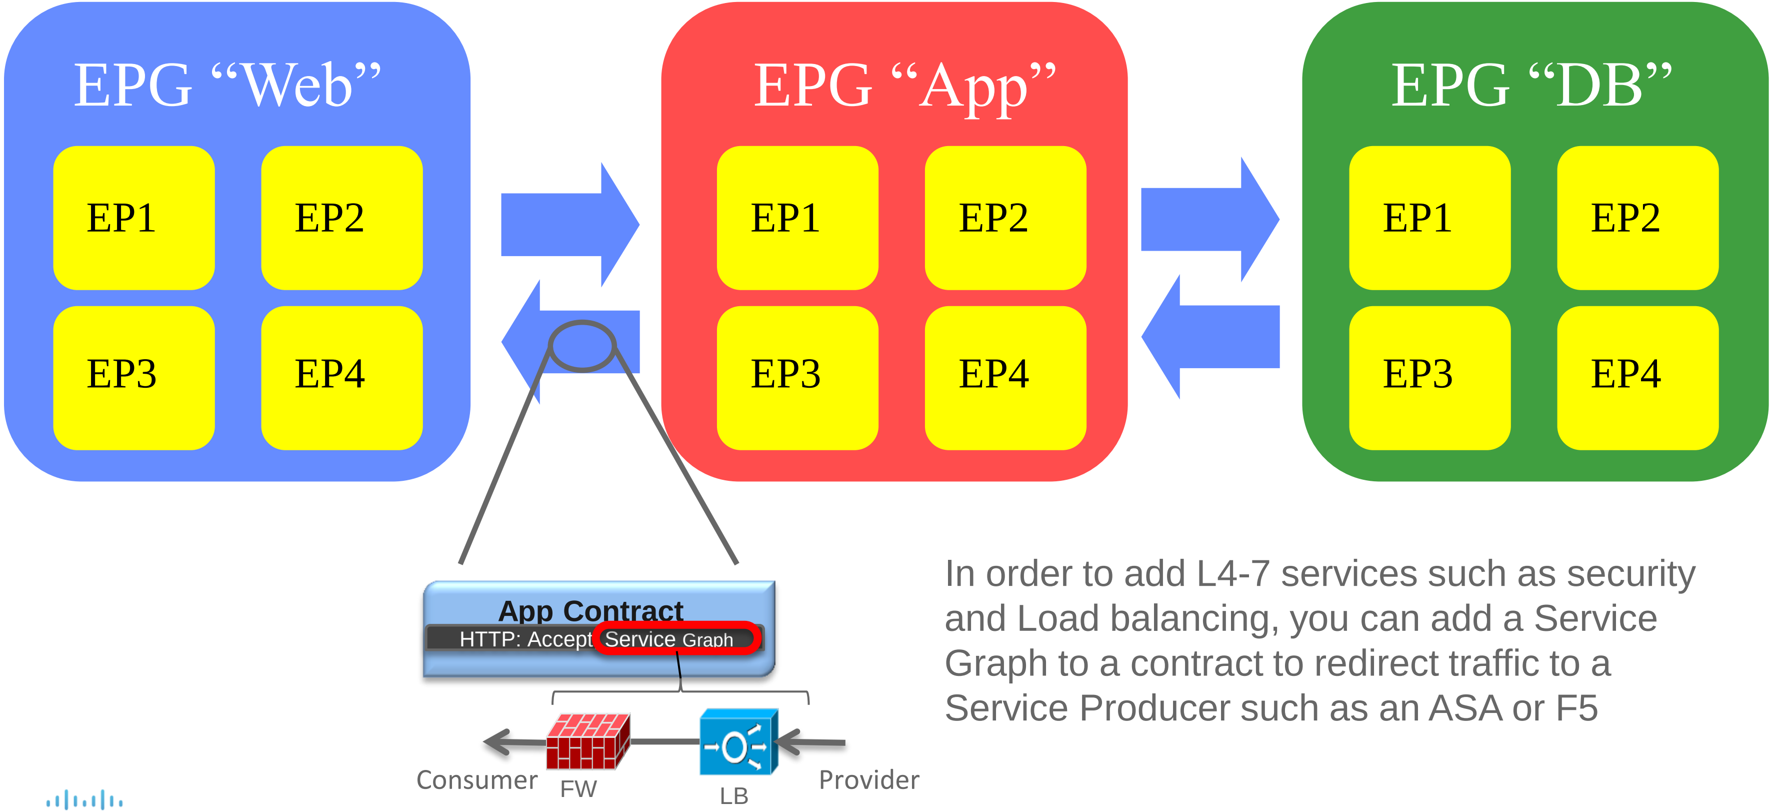

In order to add L4-7 services such as security and Load balancing, you can add a Service Graph to a contract to redirect traffic to a Service Producer such as an ASA or F5

EPGs @ ACI bring true network abstraction

Traditional Network Model

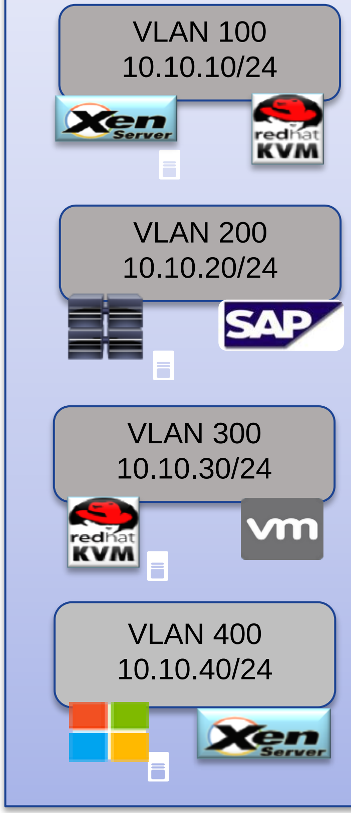

VLAN 100
10.10.10/24

VLAN 200
10.10.20/24

VLAN 300
10.10.30/24

VLAN 400
10.10.40/24

Apps Coupled to Location

Visibility At Network or VLAN Level

ACL-based Policy Per Interface

No Address Independence or Policy Mobility

Application Centric Infrastructure

Apps Decoupled from Location

Visibility At App or Group Level

Policy Between Groups

Complete Address Independence & Policy Mobility

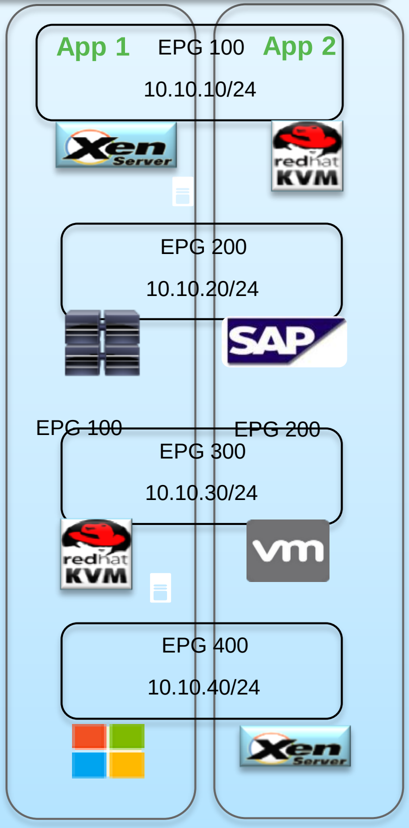

App 1 EPG 100 App 2
10.10.10/24

EPG 200
10.10.20/24

EPG 100 EPG 200
EPG 300
10.10.30/24

EPG 400
10.10.40/24

“Do I need to have a complete knowledge of my current application environment to fully use, benefit or leverage Cisco ACI ?”

ABSOLUTELY NOT !!!

ACI Design Flexibility

Network Centric or Application Centric Deployment

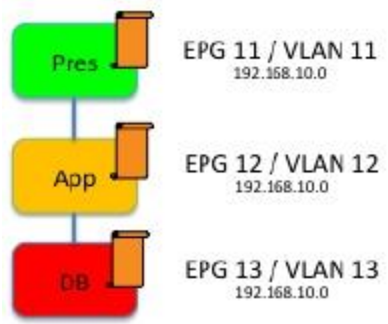

EPG = VLAN

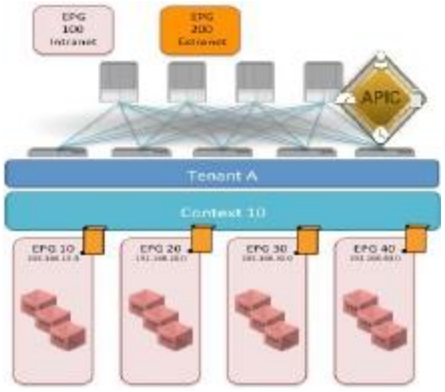

EPG = Application Group

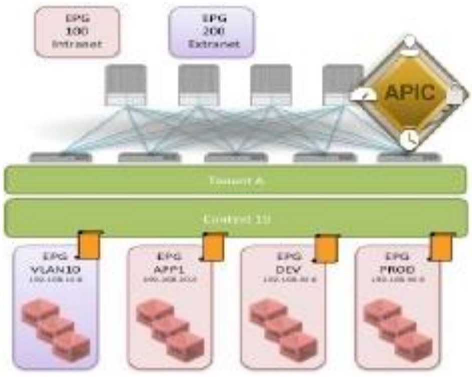

Hybrid Case

Align EPGs to a Traditional VLAN/Subnets

Align EPGs to an Individual Components of one or more Applications

How can we define the network for the application?

Defining Application Logic Through Policy

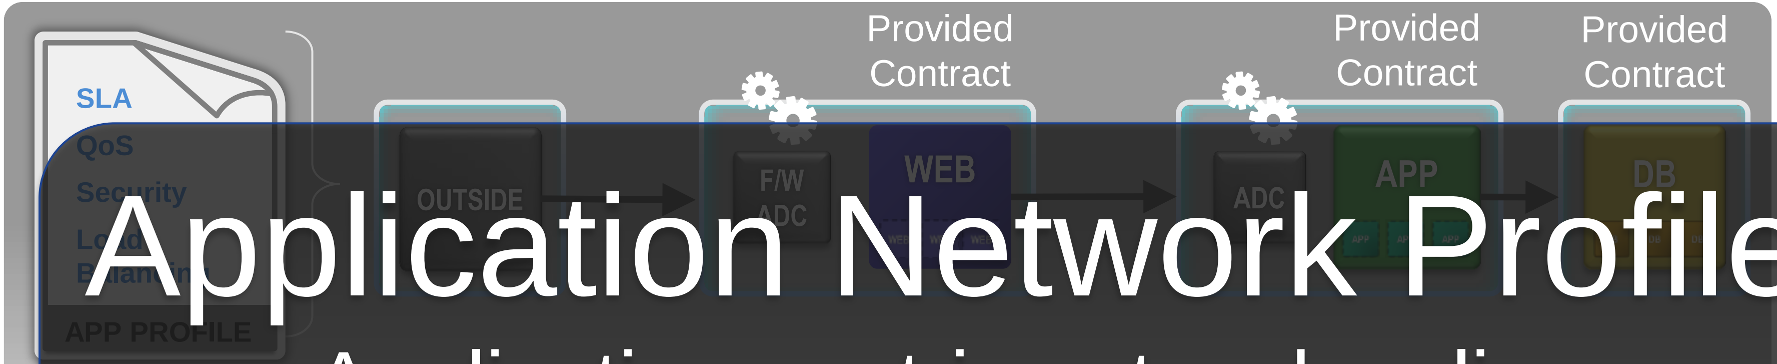

Application Network Profile

Application-centric network policy

What is an application policy?

1. Group: A set of virtual or physical workloads with the same policy
2. Contracts: A set of rules governing communication between groups
3. Service Chains: A set of network services between groups

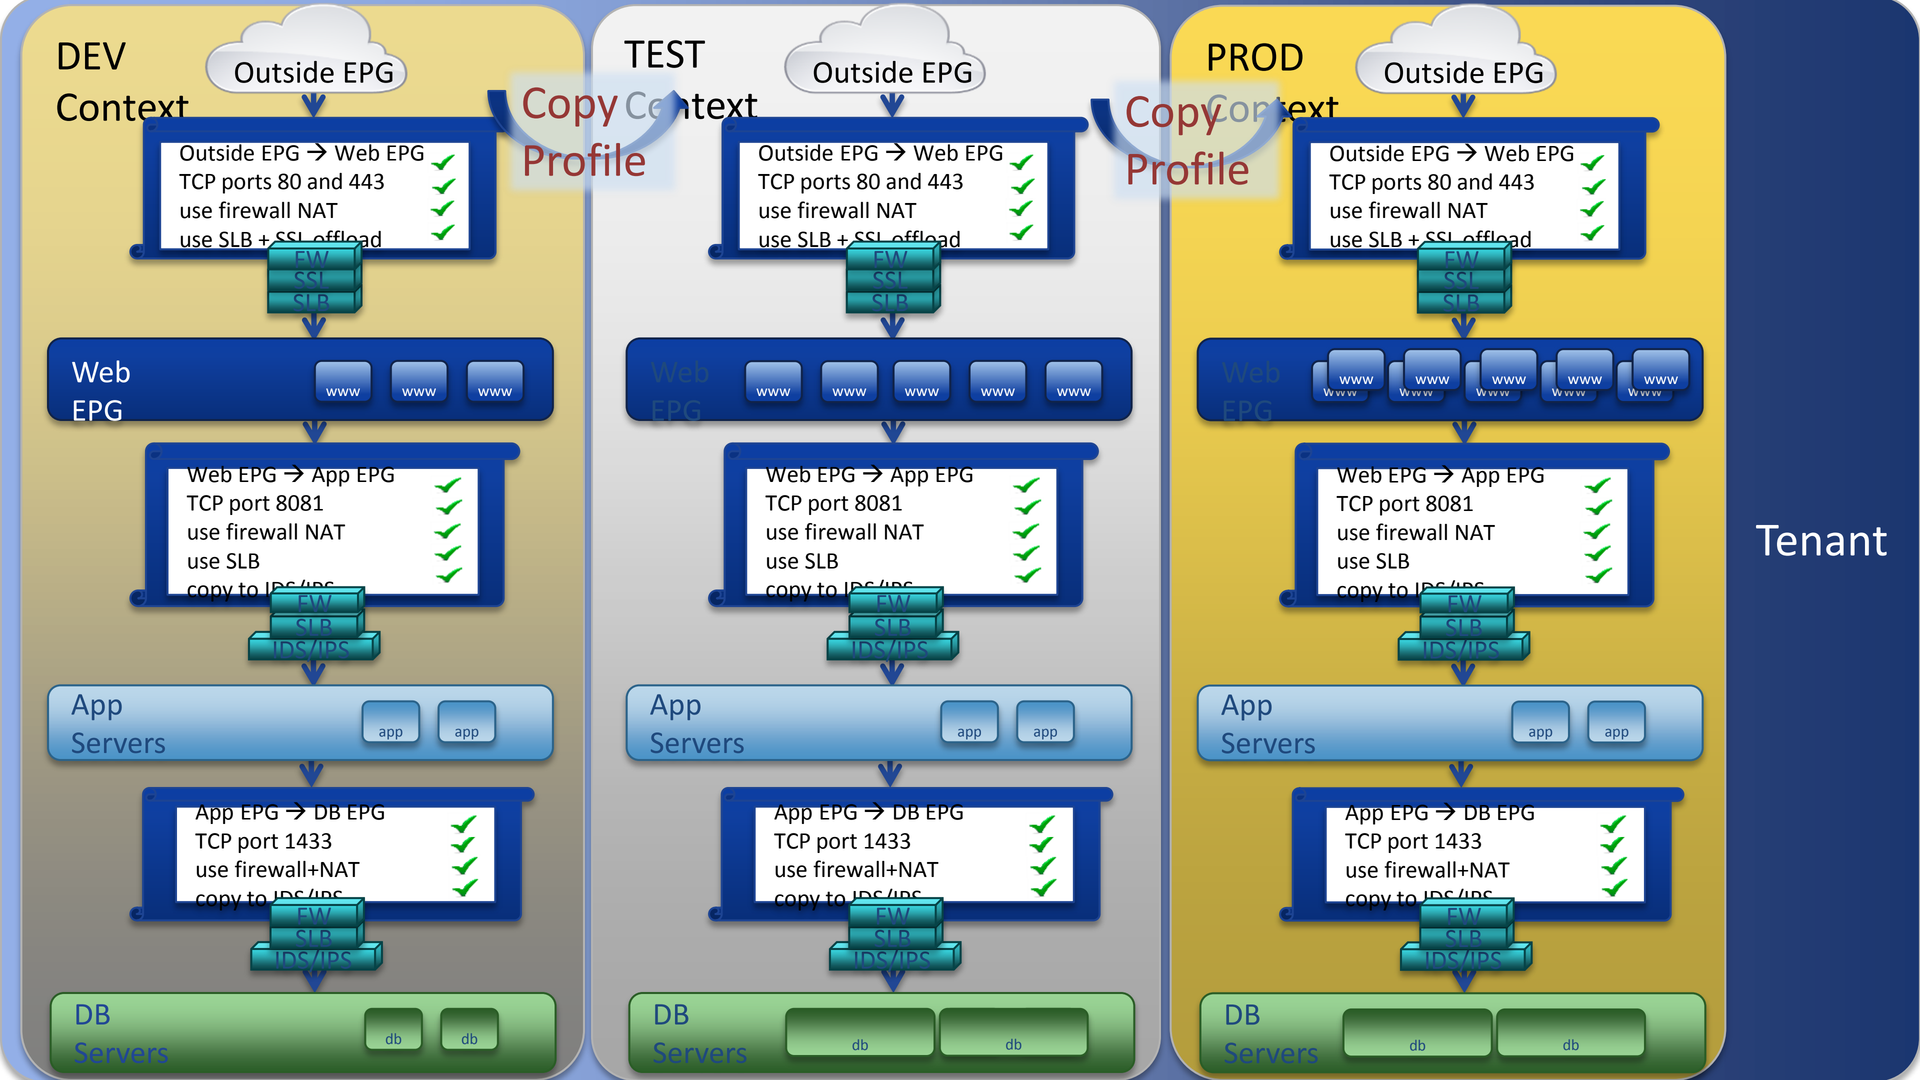

DEV Context

Outside EPG

Outside EPG → Web EPG ✓
TCP ports 80 and 443 ✓
use firewall NAT ✓
use SLB + SSL offload ✓

Web EPG

Web EPG → App EPG ✓
TCP port 8081 ✓
use firewall NAT ✓
use SLB ✓
copy to IDS/IPS ✓

App Servers

App EPG → DB EPG ✓
TCP port 1433 ✓
use firewall+NAT ✓
copy to IDS/IPS ✓

DB Servers

Copy Profile

TEST Context

Outside EPG

Outside EPG → Web EPG ✓
TCP ports 80 and 443 ✓
use firewall NAT ✓
use SLB + SSL offload ✓

Web EPG

Web EPG → App EPG ✓
TCP port 8081 ✓
use firewall NAT ✓
use SLB ✓
copy to IDS/IPS ✓

App Servers

App EPG → DB EPG ✓
TCP port 1433 ✓
use firewall+NAT ✓
copy to IDS/IPS ✓

DB Servers

Copy Profile

PROD Context

Outside EPG

Outside EPG → Web EPG ✓
TCP ports 80 and 443 ✓
use firewall NAT ✓
use SLB + SSL offload ✓

Web EPG

Web EPG → App EPG ✓
TCP port 8081 ✓
use firewall NAT ✓
use SLB ✓
copy to IDS/IPS ✓

App Servers

App EPG → DB EPG ✓
TCP port 1433 ✓
use firewall+NAT ✓
copy to IDS/IPS ✓

DB Servers

Tenant

Cisco ACI

Logical Network Provisioning of Stateless Hardware with ANP

Admin creates policies and profiles

Admin does **NOT** program the hardware

APIC pushes policies and profiles to HW

HW programs itself!

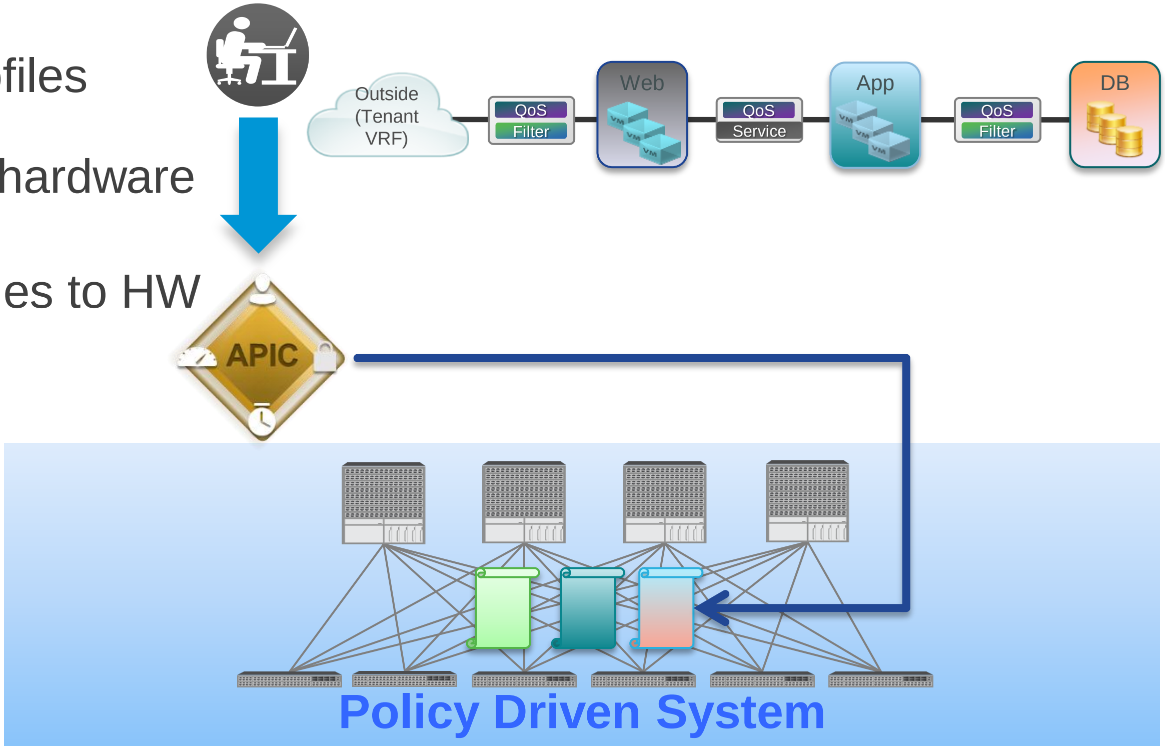

Policy Driven System

Agenda

- Modern DC infrastructure – Customer requirements
- What's Application Centric Infrastructure (ACI)
- **How ACI affects / enhances Data Center Infrastructure**
  - Network fabric
  - Hypervisors
  - L4-L7 Services
- How ACI affects Applications
  - Security
  - Automation / Orchestration
- Migration to ACI

Cisco ACI Fabric – IP network with an integrated GBP VXLAN overlay

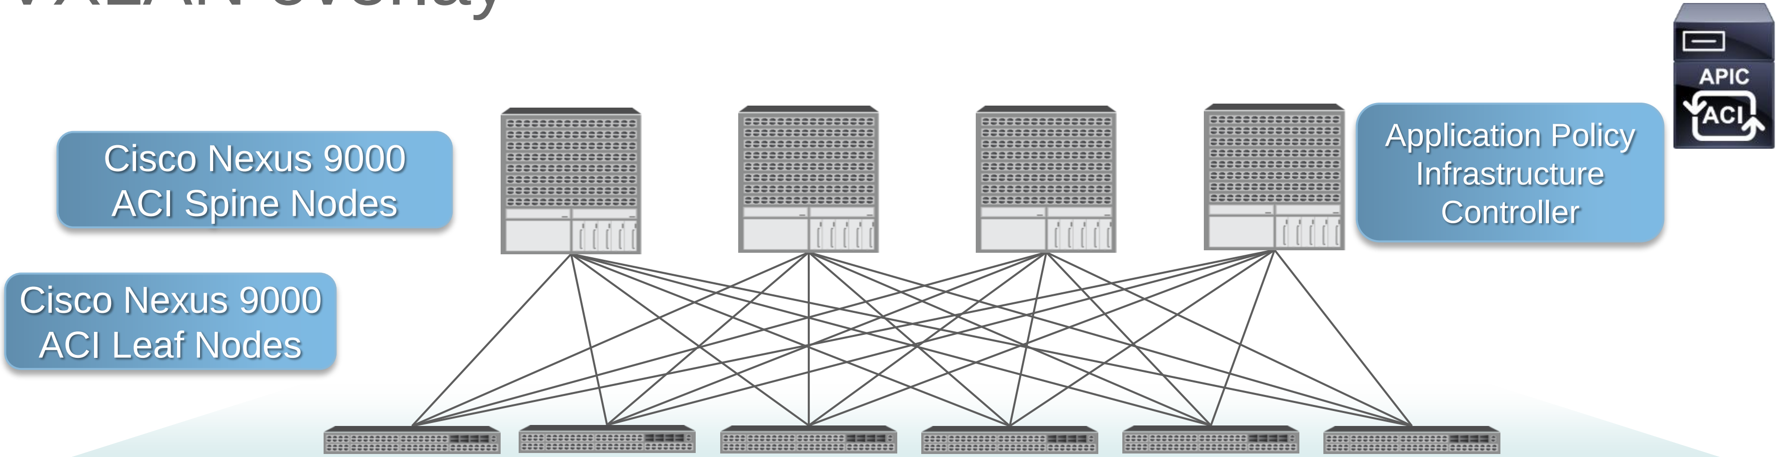

Cisco Nexus 9000 ACI Spine Nodes

Application Policy Infrastructure Controller

Cisco Nexus 9000 ACI Leaf Nodes

Cisco ACI Fabric provides:

- Full normalization of the ingress encapsulation mechanism used: 802.1Q VLAN, IETF VXLAN, and IETF NVGRE
- Distributed Layer 3 gateway to help ensure optimal forwarding for Layers 3 and 2 (No HSRP/VRRP required)
- Support for standard bridging and routing semantics without standard location constraints (any IP address anywhere)
- Service insertion and redirection
- Removal of flooding requirements for IP control plane (IP ARP and GARP packets are forwarded directly to the target endpoint address contained within ARP or GARP header)

Cisco ACI Fabric Multi-Tenancy Construct

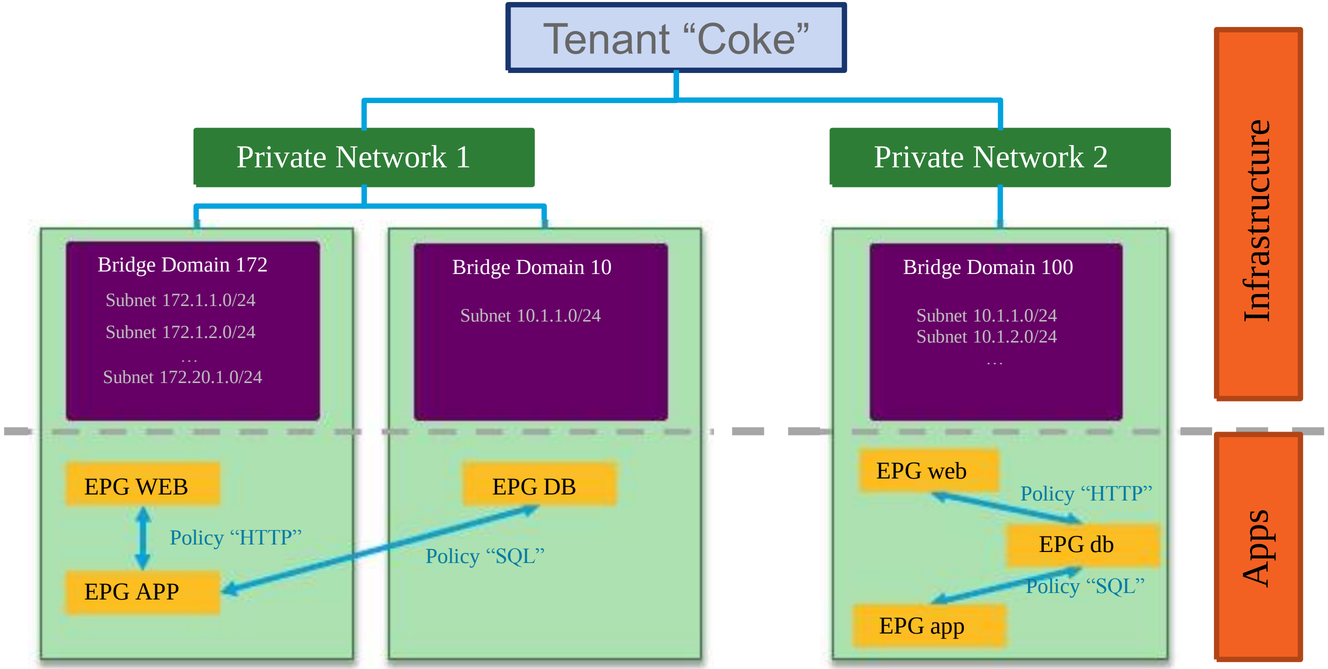

Tenant is a container for all network, security and L4–7 service policies
Tenant resources are isolated from each other

Private network (VRF or context) is used to allow isolated and potentially overlapping IP address space

Bridge domain is a L2 forwarding construct within the fabric, used to constrain broadcast and multicast traffic

EPGs exist within a single bridge domain only

Policy (Contract) is used to determine how EPGs communicate with each other

Cisco ACI Fabric Multi-Tenancy Construct

Mapping the Configuration to the Packet

- ACI Fabric leverages VXLAN Encapsulation to build network overlay
- The ACI VXLAN header is not associated with a specific L2 segment or L3 domain but provides a multi-function tagging mechanism used in ACI fabric.
- VXLAN Source Group/Source Class ID is used as a tag/label to identify the specific end point for each application function (EPG)
- Policy is enforced between an ingress or source application tier (EPG) and an egress or destination application tier (EPG)
- Policy can be enforced at source or destination

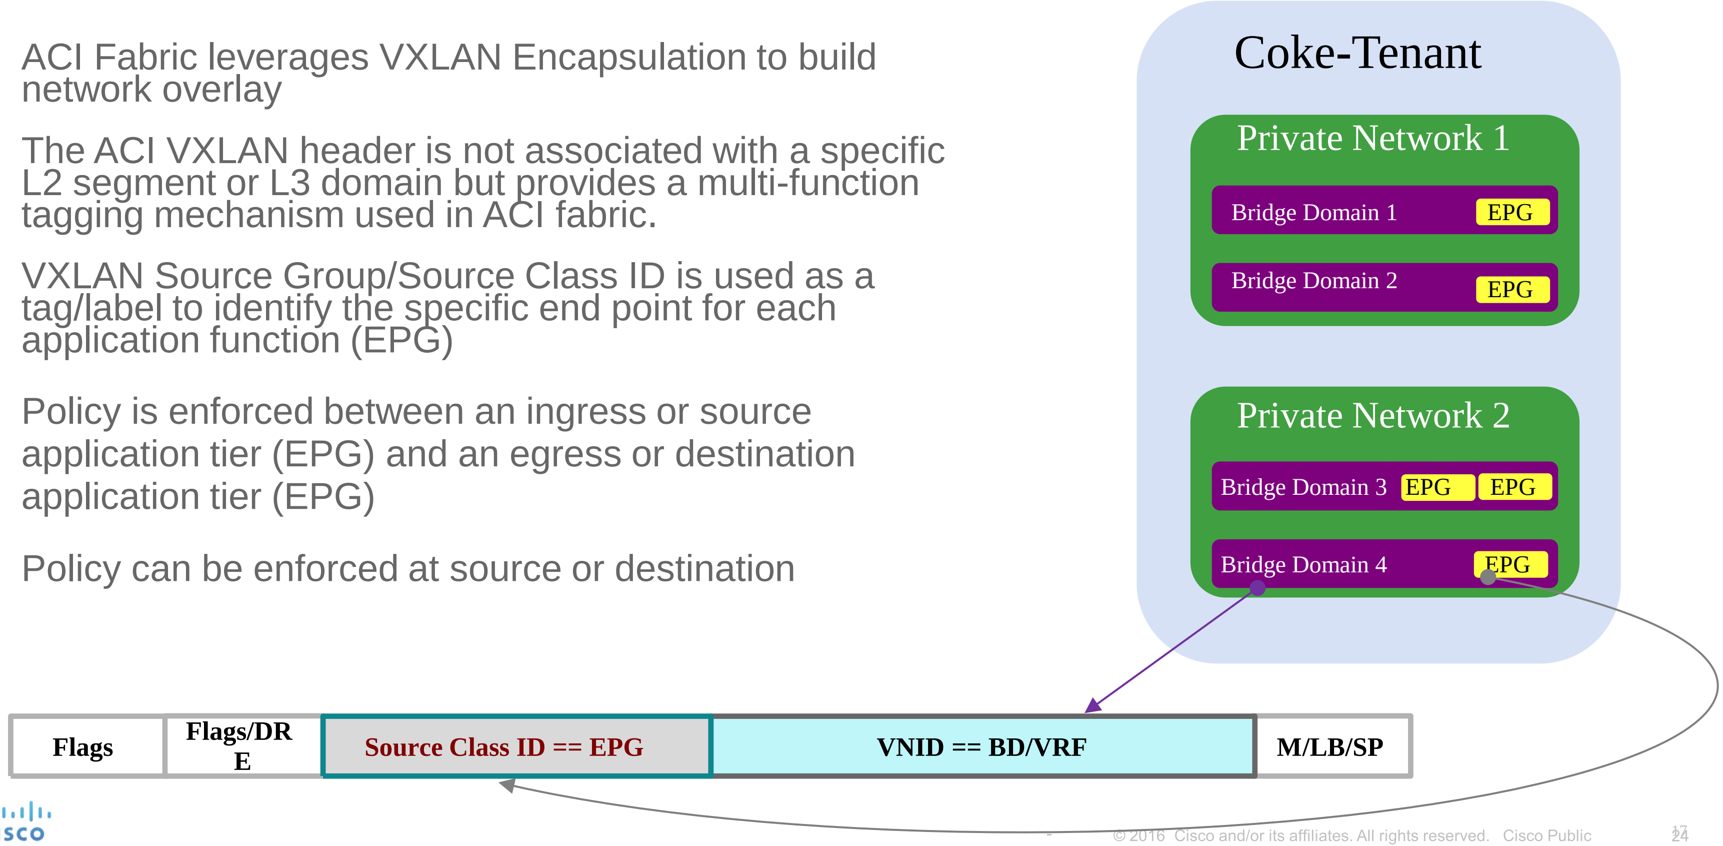

Cisco ACI Fabric Load Balancing

Focus on the Application Response Time

- Cisco® ACI fabric tracks the congestion along the full path between the ingress leaf and the egress leaf through the data plane (real-time measurements)
    - Congestion on switch-to-switch ports (external wires)
    - Congestion on internal ASIC-to-ASIC connections (internal wires)
- Fabric load-balances traffic on a "flowlet" basis
    - Dynamic shedding of active flows from congested to less congested paths
- Fabric prioritizes small (and early) flowlets
    - Provides DC-TCP behavior without having to modify host stacks
    - Ramps up large TCP flows faster

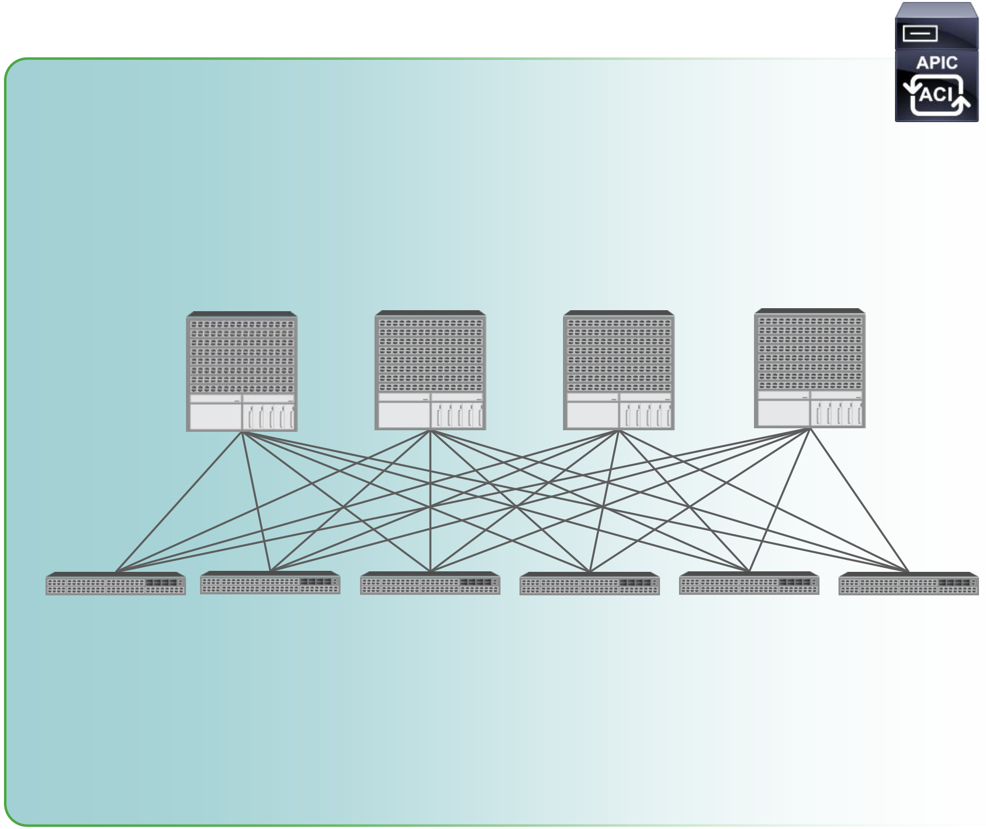

Application Awareness

Application-Level Visibility

Cisco® ACI Fabric provides the next generation of analytic capabilities

Per application, tenants, and infrastructure:

- Health scores
- Latency
- Atomic counters
- Resource consumption

Integrate with workload placement or migration

Tenant

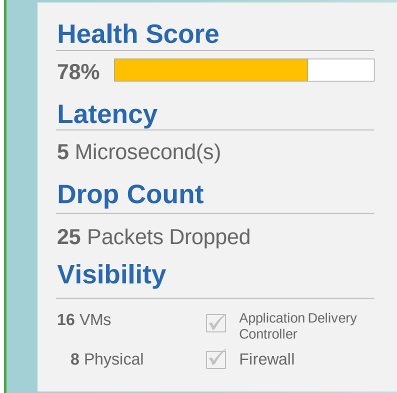

Health Score

78%

Latency

5 Microsecond(s)

Drop Count

25 Packets Dropped

Visibility

16 VMs ☑ Application Delivery Controller

8 Physical ☑ Firewall

Application

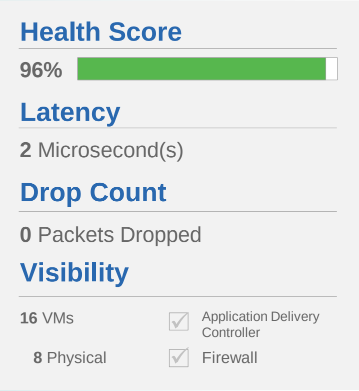

Health Score

96%

Latency

2 Microsecond(s)

Drop Count

0 Packets Dropped

Visibility

16 VMs ☑ Application Delivery Controller

8 Physical ☑ Firewall

VXLAN
Per-Hop Visibility

Physical and
Virtual as One

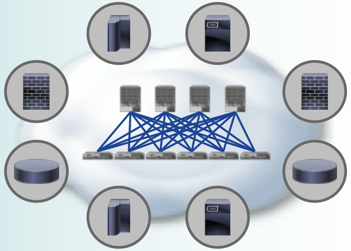

Cisco ACI - Multi-Hypervisor Fabric

Virtual Integration

APIC

- Integrated gateway for VLAN, VxLAN, and NVGRE networks from virtual to physical
- Normalization for NVGRE, VXLAN, and VLAN networks
- Customer not restricted by a choice of hypervisor
- Fabric is ready for multi-hypervisor

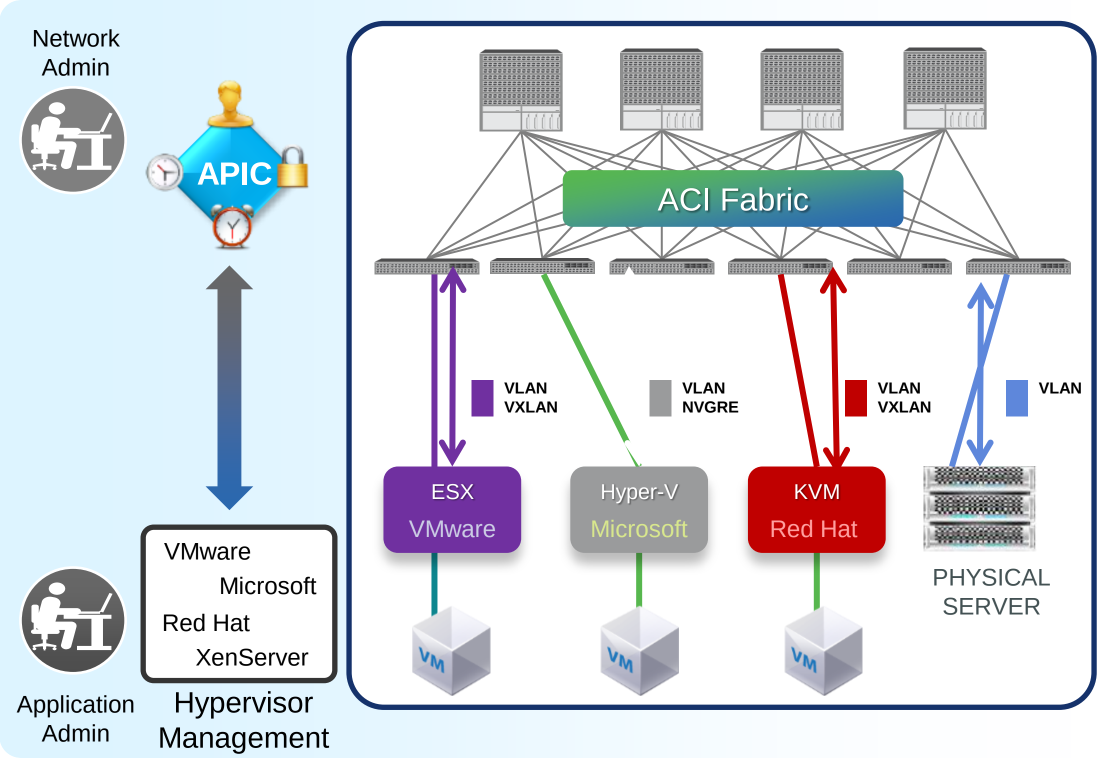

Network Admin

APIC

VMware
Microsoft
Red Hat
XenServer

Application Admin

Hypervisor Management

ACI Fabric

🟪 VLAN VXLAN

▨ VLAN NVGRE

🟥 VLAN VXLAN

🟦 VLAN

ESX VMware

Hyper-V Microsoft

KVM Red Hat

PHYSICAL SERVER

Hypervisor Integration with ACI

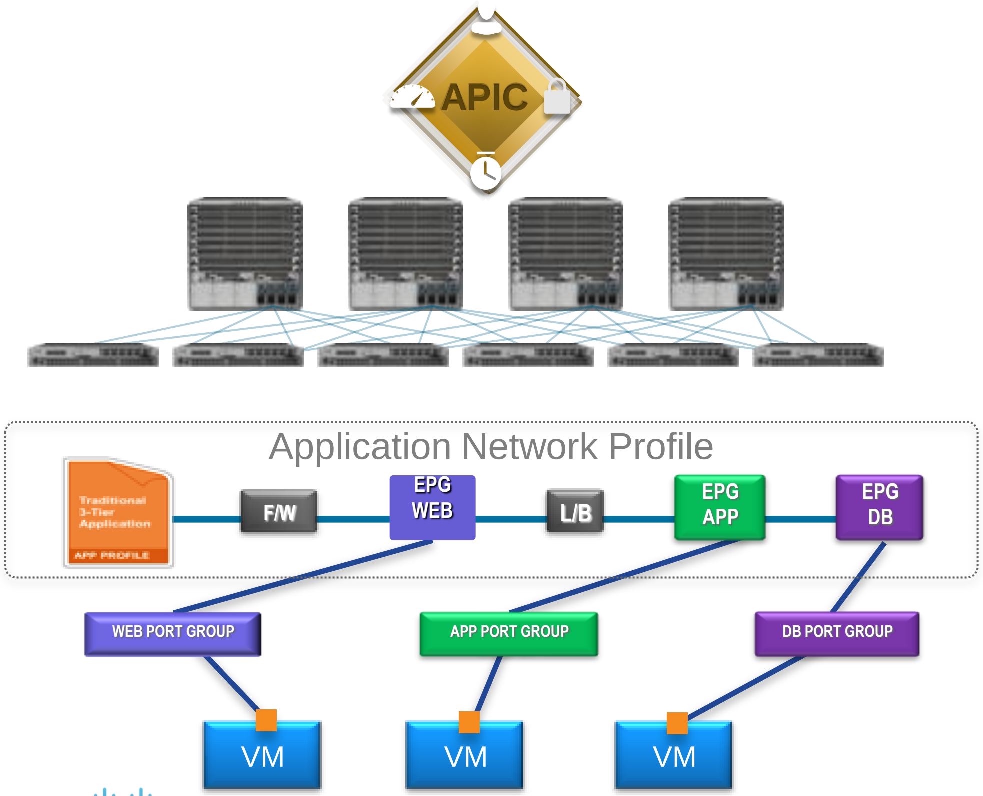

- ACI Fabric implements policy on Virtual Networks by mapping Endpoints to EPGs
- Endpoints in a Virtualized environment are represented as the vNICs
- VMM applies network configuration by placement of vNICs into:
  - Port Groups (VMWare),
  - VM Networks (Hyper-V)
  - Networks or Policy groups (OpenStack)
- EPGs are exposed to the VMM as a 1:1 mapping to Port Groups, VM Networks or OpenStack Networking.

ACI Layer 4 - 7 Service Integration

Centralized, Automated, and Supports Existing Model

- Automated and scalable L4-L7 service insertion
- Elastic service insertion architecture for physical and virtual services
- Packet match on a redirection rule sends the packet into a services graph
- Service Graph can be one or more service nodes pre-defined in a series
- APIC as central point of network control with policy coordination
- Supports existing operational model when integrated with existing services

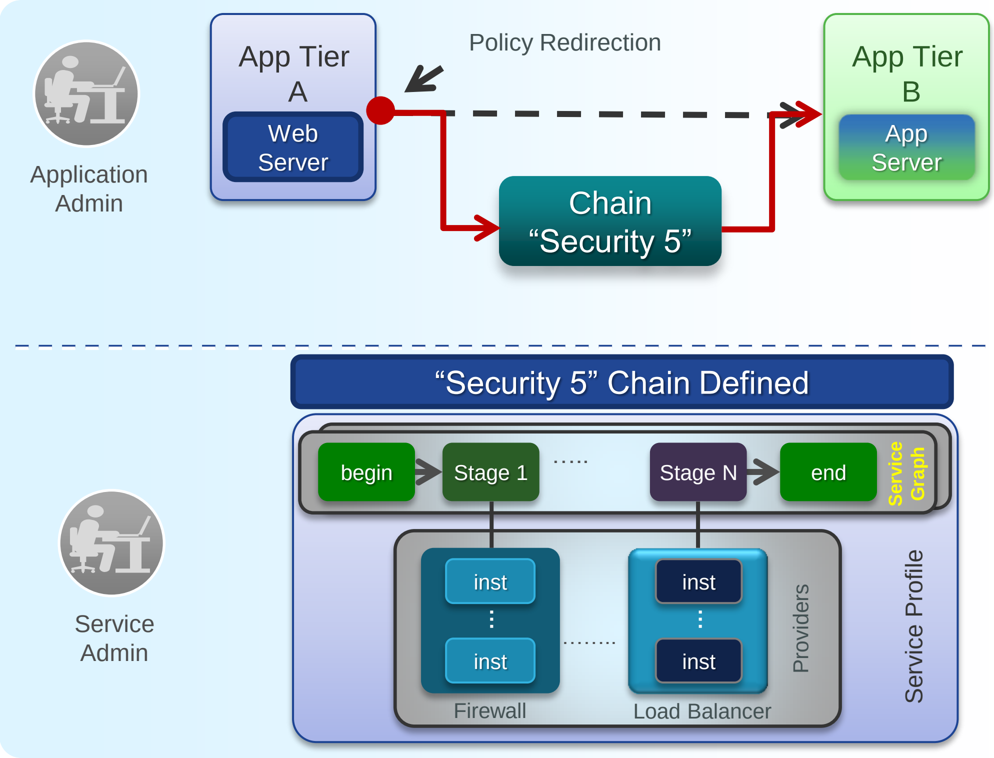

L4-L7 Service Automation: Support for All Devices

Any Device and Cluster Manager Support

L4-L7 Service Automation

L4-L7 Services

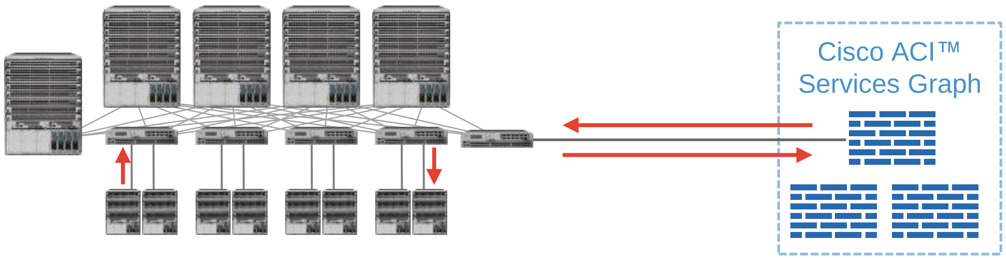

L4- L7 Device Package

- Centralized L4-L7 service configuration and management
- Full L4-L7 service automation (with device package)
- Large ecosystem and investment protection

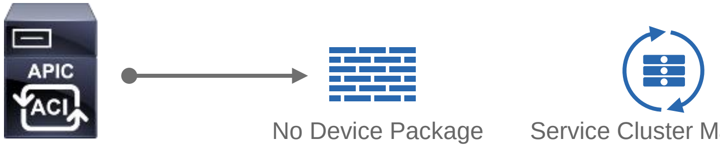
No Device Package

Service Cluster Manager

- Automated service insertion and chaining
- Support for any L4-L7 device
- New support for L4-L7 cluster managers

Cisco Application Policy Infrastructure Controller

Algorithmically Sharded Cluster

- Applications fully use clustered and replicated controller (N+1, N+2, etc.)
- Any node is able to service any user for any operation
- Seamless APIC node adds and deletes
- Fully automated APIC software cluster upgrade with redundancy during upgrade
- Cluster size driven by transaction rate requirements
- **APIC is not in the data path**

Single Point of Management Without a Single Point of Failure

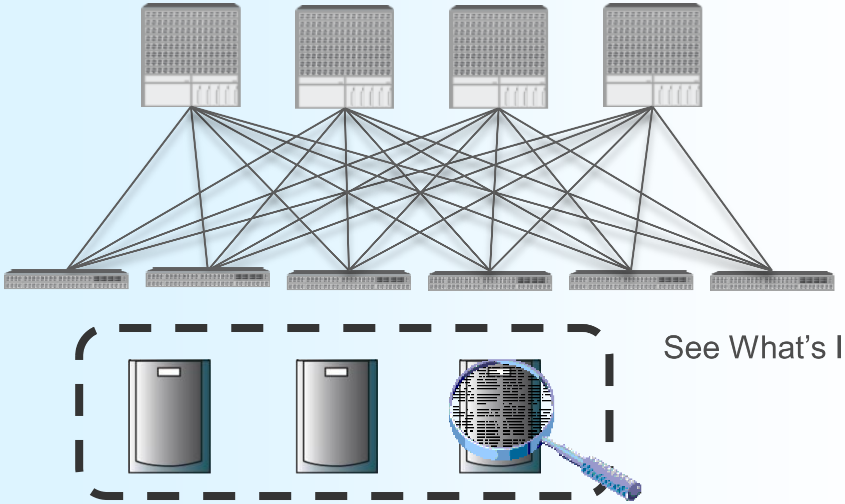

See What's Inside

APIC Cluster

Distributed, Synchronised, Replicated

Agenda

- Modern DC infrastructure – Customer requirements
- What's Application Centric Infrastructure (ACI)
- How ACI affects / enhances Data Center Infrastructure
  - Network fabric
  - Hypervisors
  - L4-L7 Services
- **How ACI affects Applications**
  - Security
  - Automation / Orchestration
- Migration to ACI

ACI Security

Whitelist, Multi-Tenant Isolation, Service Automation

L4 Distributed Stateless Firewall

L4-7 Security Via ACI Service Graph

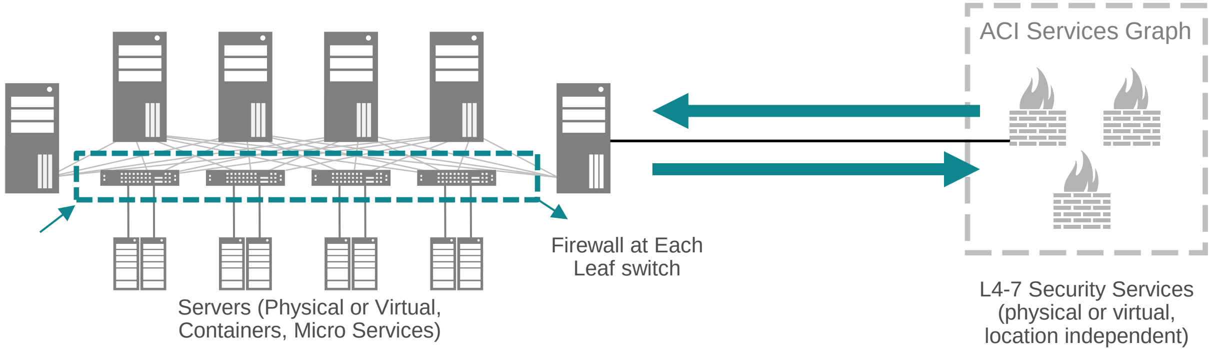
ACI Services Graph
Firewall at Each Leaf switch
Servers (Physical or Virtual, Containers, Micro Services)
L4-7 Security Services (physical or virtual, location independent)

Symantec. radware SOURCEfire intel Security FORTINET Check Point SOFTWARE TECHNOLOGIES LTD. CISCO paloalto NETWORKS

Embedded Security

- White-list Firewall Policy Model (line rate)
- Authenticated Northbound API (X.509)
- Encrypted Management Plane (TLS 1.2)
- PCI, FIPS

Micro-Segmentation

- Intra End-Point-Group (EPG) isolation
- Attribute Based isolation and quarantine

Security Automation

- Dynamic Service Insertion and Chaining
- Security Policy follows workloads
- Centralized Security provisioning and visibility

ACI Enables Segmentation Based on Business Needs

Policy Driven Micro-Segmentation and Intra-EPG Isolation

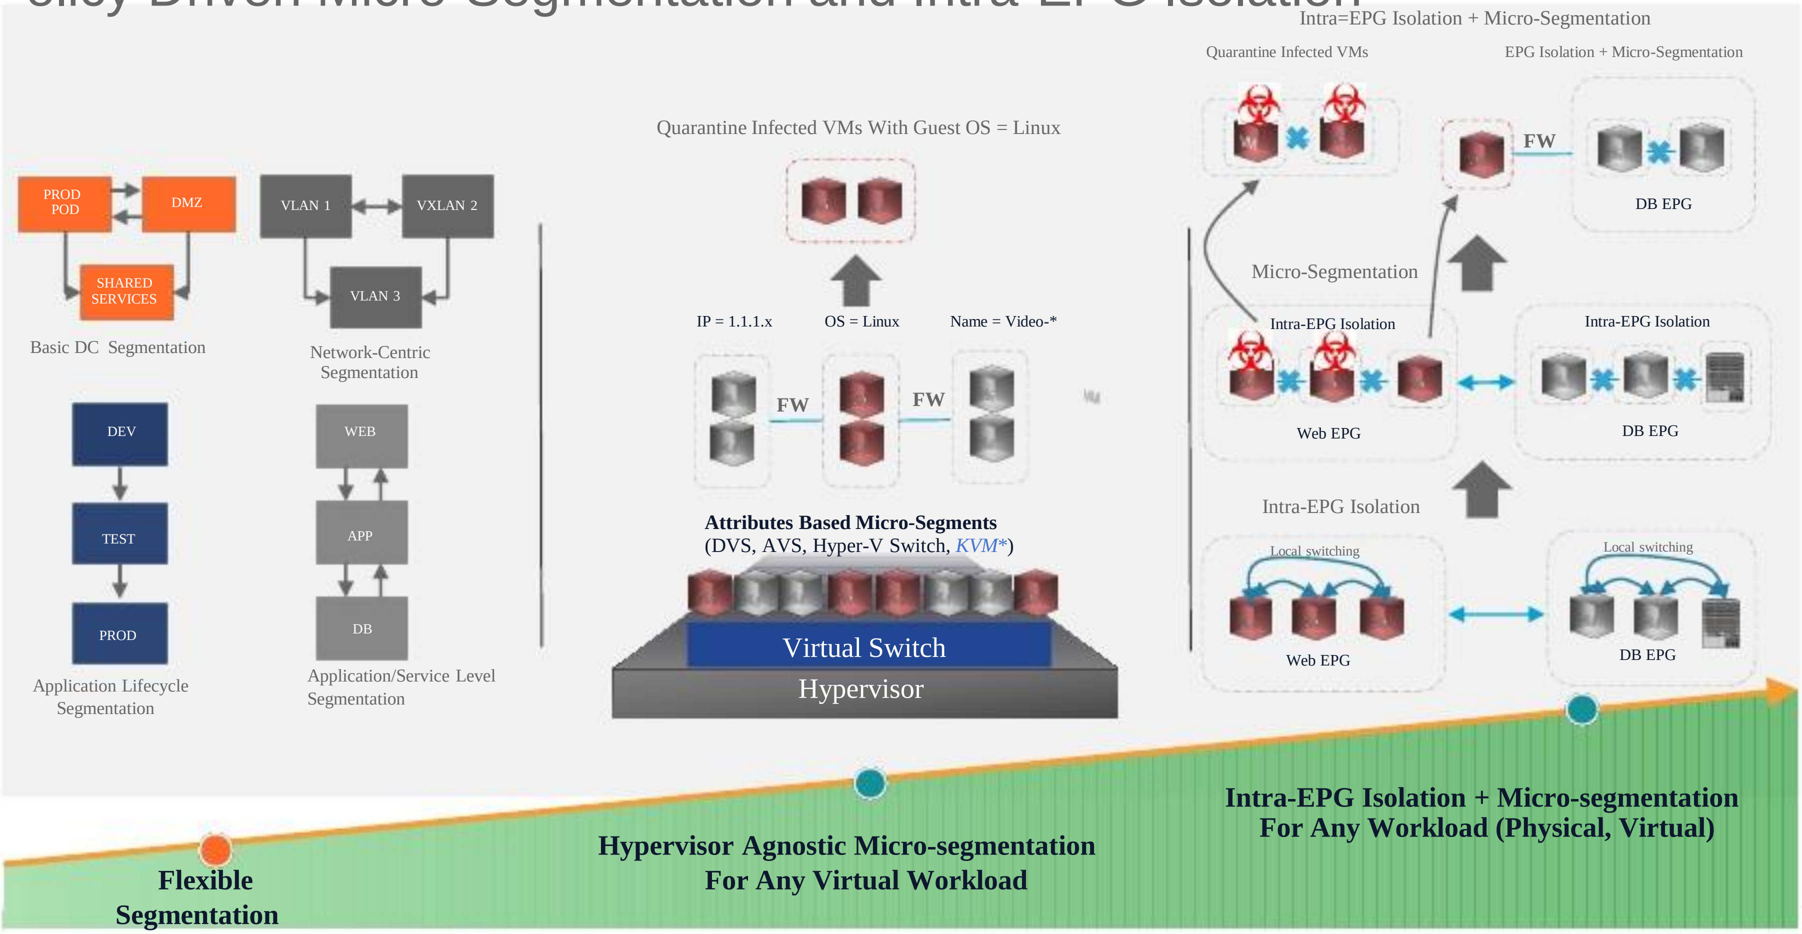

PROD POD

DMZ

SHARED SERVICES

Basic DC Segmentation

VLAN 1

VXLAN 2

VLAN 3

Network-Centric Segmentation

DEV

TEST

PROD

Application Lifecycle Segmentation

WEB

APP

DB

Application/Service Level Segmentation

Quarantine Infected VMs With Guest OS = Linux

IP = 1.1.1.x

OS = Linux

Name = Video-*

FW

FW

Attributes Based Micro-Segments
(DVS, AVS, Hyper-V Switch, *KVM**)

Virtual Switch

Hypervisor

Intra=EPG Isolation + Micro-Segmentation

Quarantine Infected VMs

EPG Isolation + Micro-Segmentation

FW

DB EPG

Micro-Segmentation

Intra-EPG Isolation

Web EPG

Intra-EPG Isolation

DB EPG

Intra-EPG Isolation

Local switching

Web EPG

Local switching

DB EPG

Flexible Segmentation

Hypervisor Agnostic Micro-segmentation For Any Virtual Workload

Intra-EPG Isolation + Micro-segmentation For Any Workload (Physical, Virtual)

Attribute-Based EPG/uSeg EPG

Isolate Malicious Virtual Machines

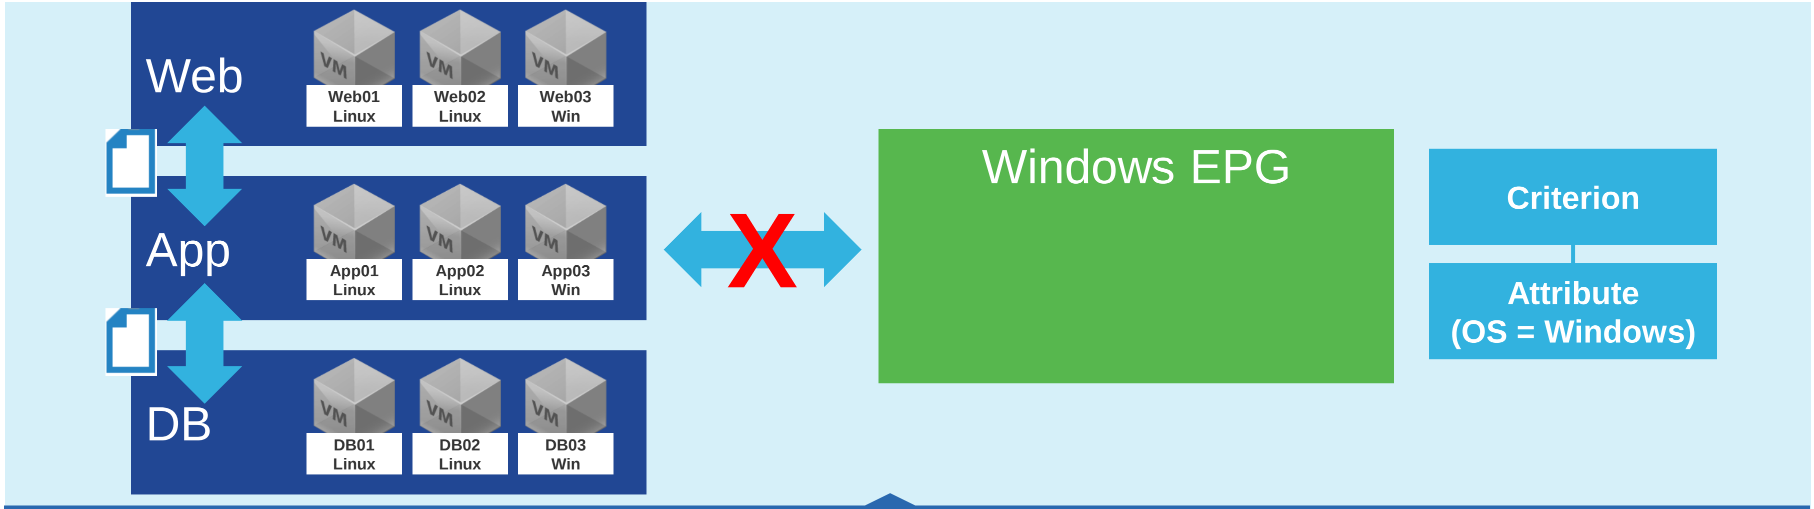

- Problem: A vulnerability is detected in a particular type of operating system (for example, Microsoft Windows). The network security administrator wants to isolate all Windows virtual machines.
- Solution: Define a security EPG with a criterion such as Operating System = Windows. No contracts are provided or consumed by this EPG. It will stop all inter-EPG communication for the matching virtual machines.
- No virtual machine attachment or detachment or placement in a different port group is needed.

Intra-EPG Segmentation

Problem

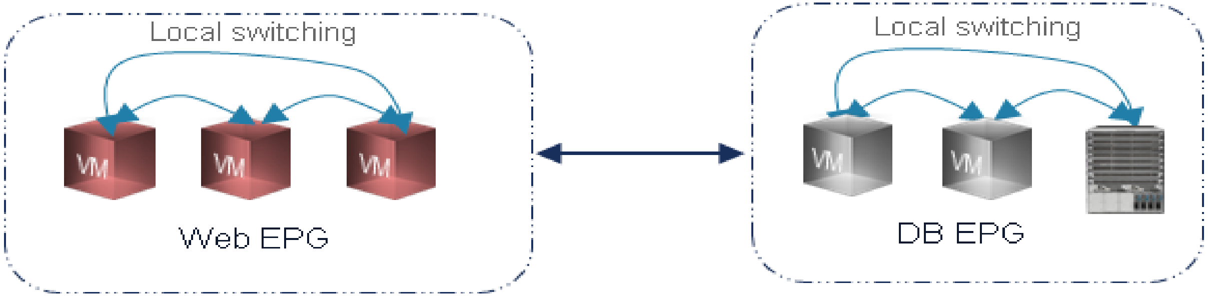

Intra-EPG Isolation Denies All Communication within an EPG

Solution

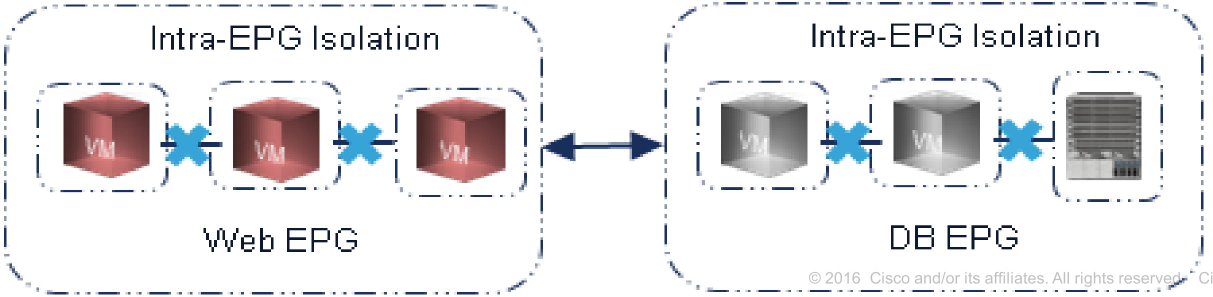

Why ACI is best for Micro Segmentation

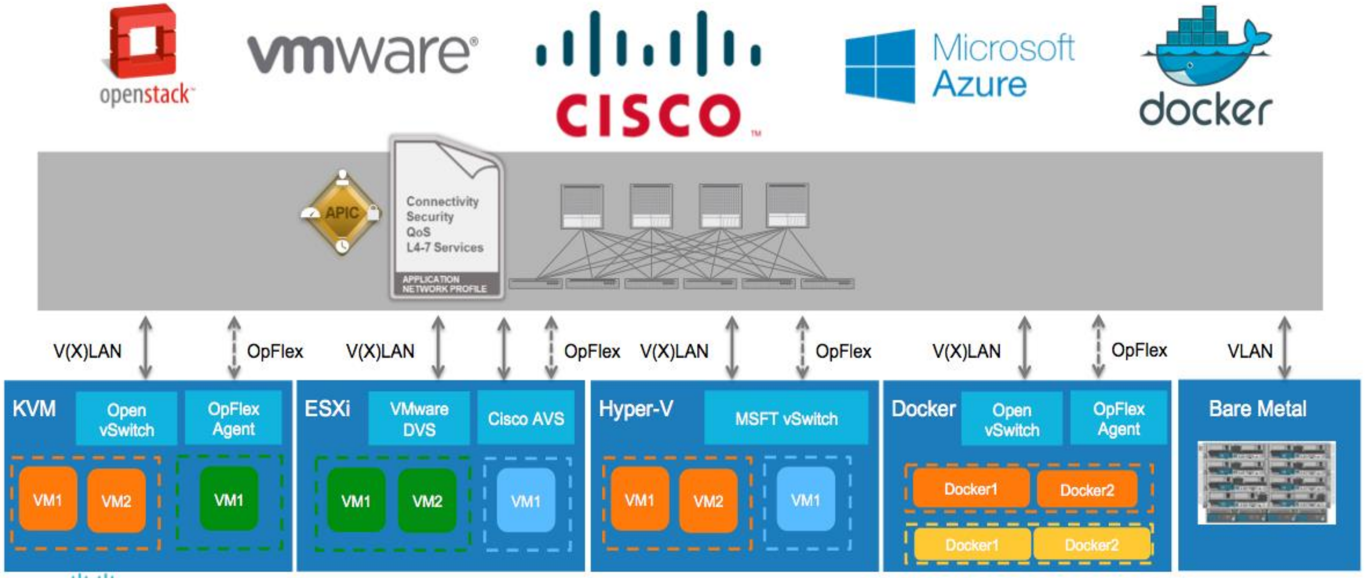

- Micro Segmentation works for all workloads (bare metal, virtual, containers, management, backup …)
- Same policy-model for vSphere, Hyper-V, OpenStack, Containers and Bare Metal.
- Works with standard virtual switch offerings, including VMware VDS, OVS, MSFT vSwitch (AVS is optional for vSphere)
- Stateful firewall when using Cisco AVS on vSphere at no extra cost with better performance at the VMware environment

Application decommission & the compliance / audit demand

“Due to compliance regulations, when an application gets decommissioned, every IT resource associated with that must be removed and/or wiped out”

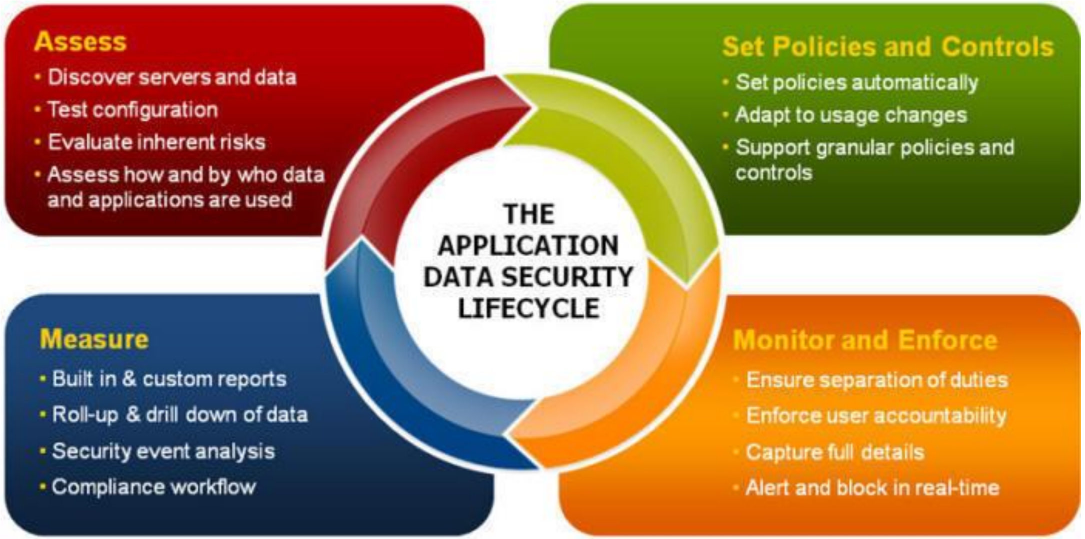

THE APPLICATION DATA SECURITY LIFECYCLE

Assess
- Discover servers and data
- Test configuration
- Evaluate inherent risks
- Assess how and by who data and applications are used

Set Policies and Controls
- Set policies automatically
- Adapt to usage changes
- Support granular policies and controls

Monitor and Enforce
- Ensure separation of duties
- Enforce user accountability
- Capture full details
- Alert and block in real-time

Measure
- Built in & custom reports
- Roll-up & drill down of data
- Security event analysis
- Compliance workflow

UCS allows one do dissociate service profile(s) associated with this application.
Audit OK !

Storage arrays can wipe-out the data or associated disks can be trashed.
Audit OK !

Current network approach and solutions don’t have a way to map application workflow and “remove” it.
Audit Fail ☹

APIC

ACI is the only one that can, inclusive programmatically and automated
Audit OK !

Orchestration with ACI

UCS Director

openstack

Windows Azure

vmware vRealize

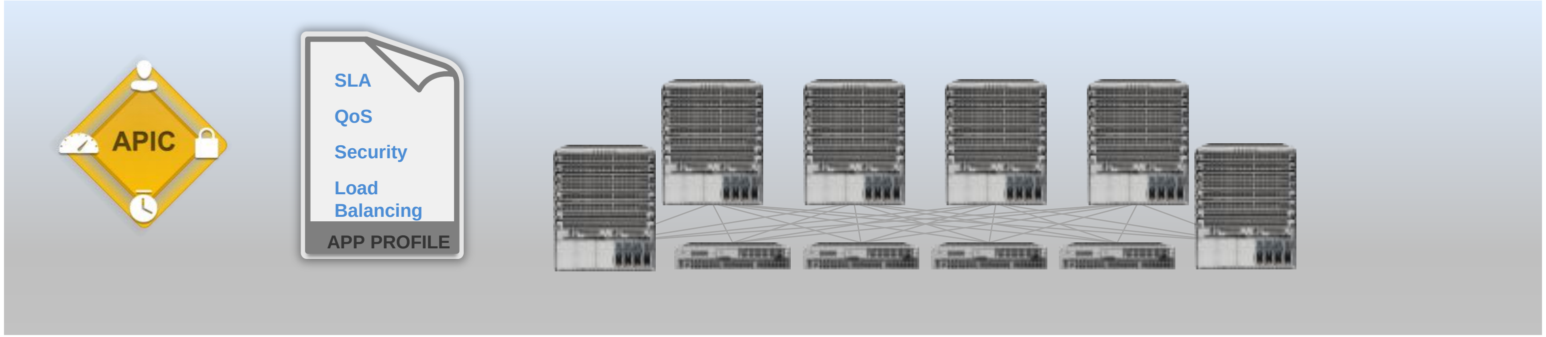

ACI Open APIs and Ecosystem

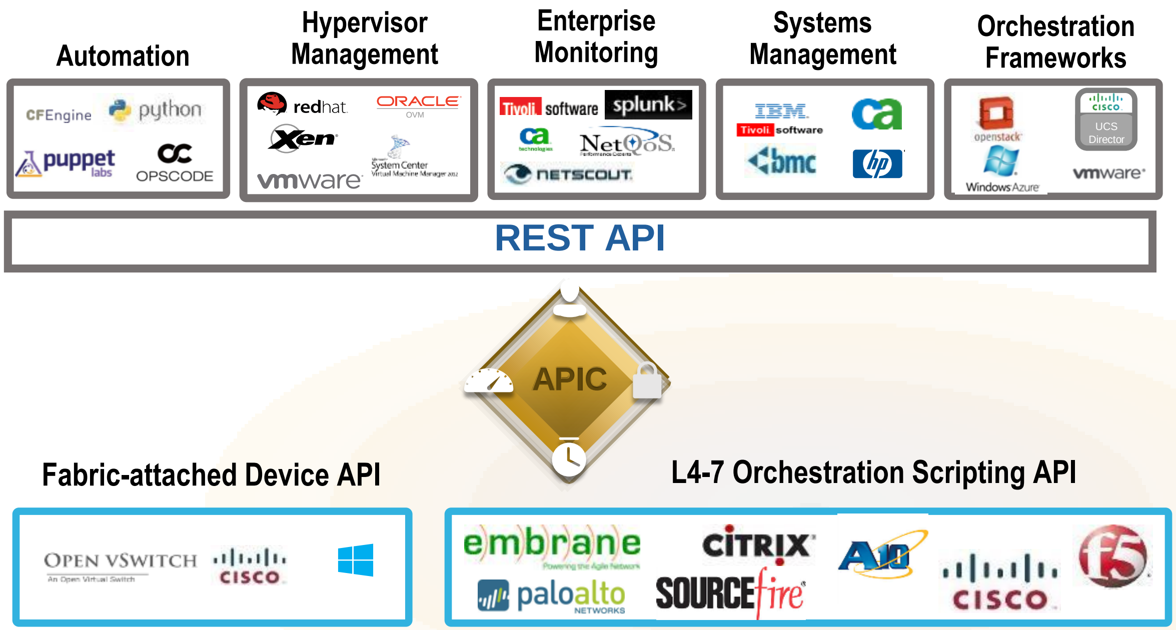

Automation

Hypervisor Management

Enterprise Monitoring

Systems Management

Orchestration Frameworks

REST API

Fabric-attached Device API

L4-7 Orchestration Scripting API

NORTHBOUND PROGRAMMABILITY LAYER

SOUTHBOUND PROGRAMMABILITY LAYER

APIC SUPPORTS A RICH ECOSYSTEM BUILT AROUND OPEN NORTHBOUND AND SOUTHBOUND APIS

Agenda

- Modern DC infrastructure – Customer requirements
- What’s Application Centric Infrastructure (ACI)
- How ACI affects / enhances Data Center Infrastructure
  - Network fabric
  - Hypervisors
  - L4-L7 Services
- How ACI affects Applications
  - Security
  - Automation / Orchestration
- **Migration to ACI**

Migration to ACI

Connecting Brownfield to new ACI network

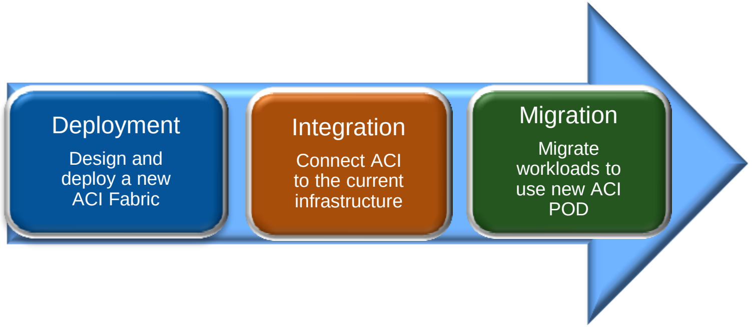

Deployment
Design and deploy a new ACI Fabric

Integration
Connect ACI to the current infrastructure

Migration
Migrate workloads to use new ACI POD

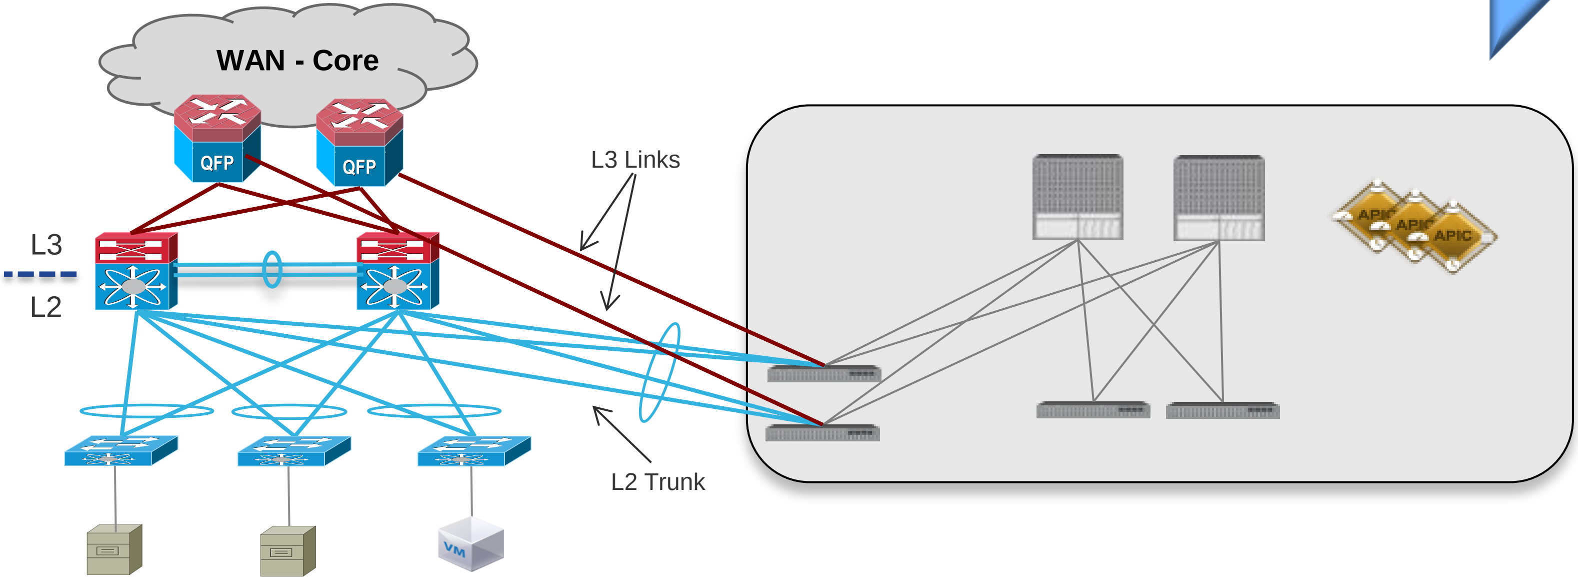

Cisco ACI - Delivering the Next-Generation DC Infrastructure for an Application-Centric World

6,000+
Nexus 9K and ACI Customers Globally

1400+
ACI Customers

50
Ecosystem Partners

NetQoS. ca technologies AVI Networks CFEngine PANDUIT splunk> radware Check Point SOFTWARE TECHNOLOGIES LTD cloudstack

Microsoft IBM NUTANIX MAPR NIKSUN SAP FORTINET puppet labs One Convergence

VCE f5 Symantec bmc Infoblox EMC² EMULEX SOURCEfire CANONICAL DATATORRENT

KillerIT Vnomic Zenoss vmware openstack VERITAS n3n CATBIRD python

citrix cloudera redhat NetApp A10 Networks NETSCOUT OPSCODE apprenda intel Security

ScienceLogic CliQr

APP App Based Automation

Automated L4-7 Stitching

Turn-key Network Automation

Děkujeme za pozornost.