ACI Hypervisor Integration

Tiger Team Bootcamp

ACI Hypervisor Integration Overview

Virtualization Landscape

- VMware leading
- Microsoft Hyper-V is the biggest challenger in Enterprise space
- Red Hat gaining traction with adoption of OpenStack

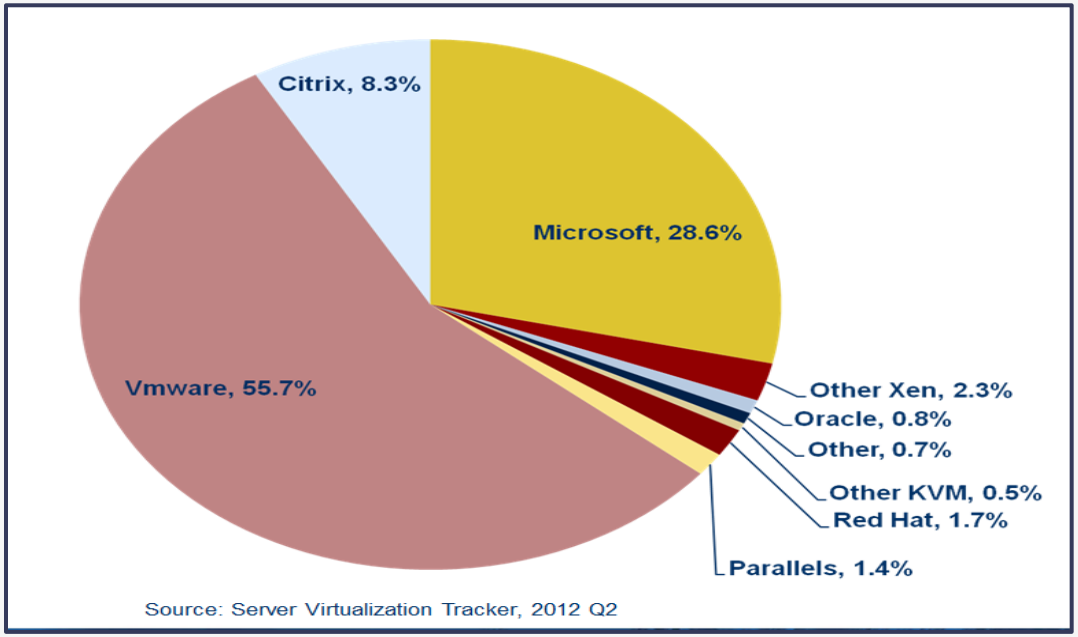

Key Virtualization Trends

18% have virtualized the majority of their compute infrastructure

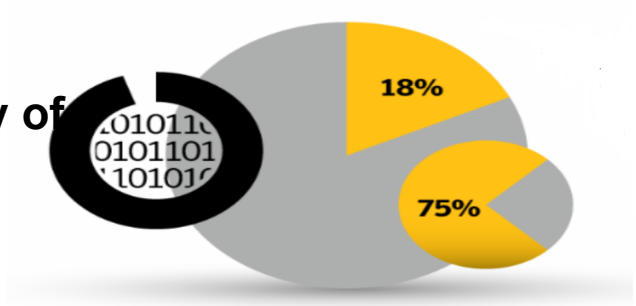

52% plan on full virtualization within two years

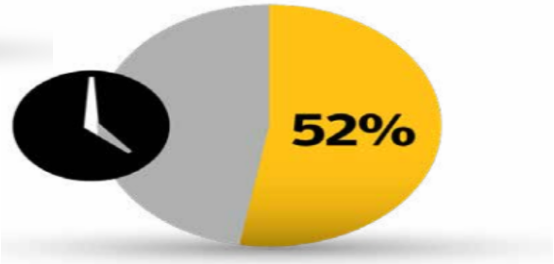

83% are changing hyper

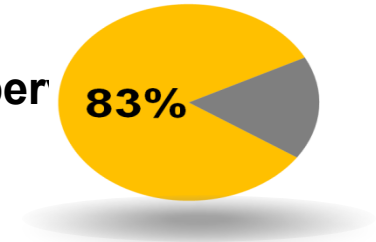

83% moving to a multi-hypervisor environment

Source: Symantec's Windows Server 2012 Migration/Virtualization Survey

Policy Coordination with VM Managers

Virtual Integration

- Network policy coordination with virtualization managers
- Automatic virtual end point detection and policy placement
- Policies consistently implemented in virtual and physical
- Network policy stays sticky with VM

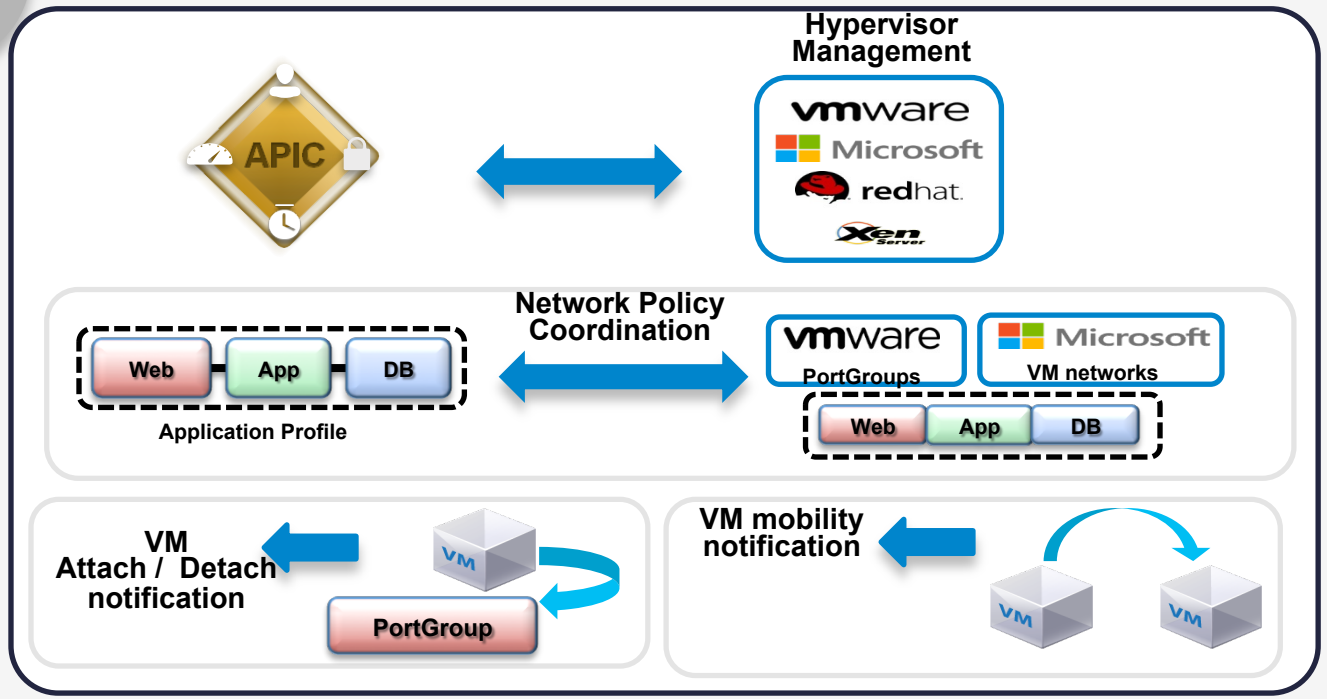

ACI Fabric – Integrated Overlay
Multi-Hypervisor Encapsulation Normalization

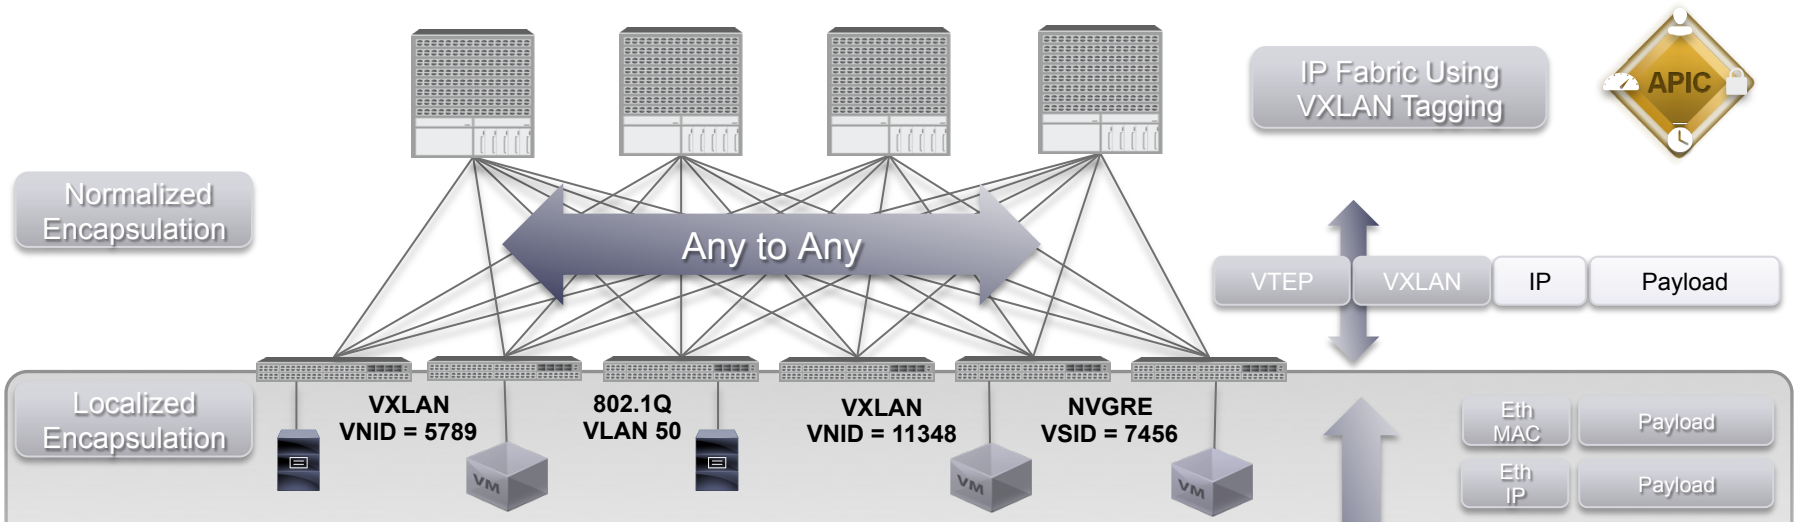

- All traffic within the ACI Fabric is encapsulated with a VXLAN eader
- External VLAN, VXLAN, NVGRE tags are mapped at ingress to an internal VXLAN tag
- Forwarding is not limited to, nor constrained within, the encapsulation type or encapsulation ‘overlay’ network
- External identifies are localized to the Leaf or Leaf port, allowing re-use and/or translation if required

Hypervisor Integration with ACI
VMM Domains

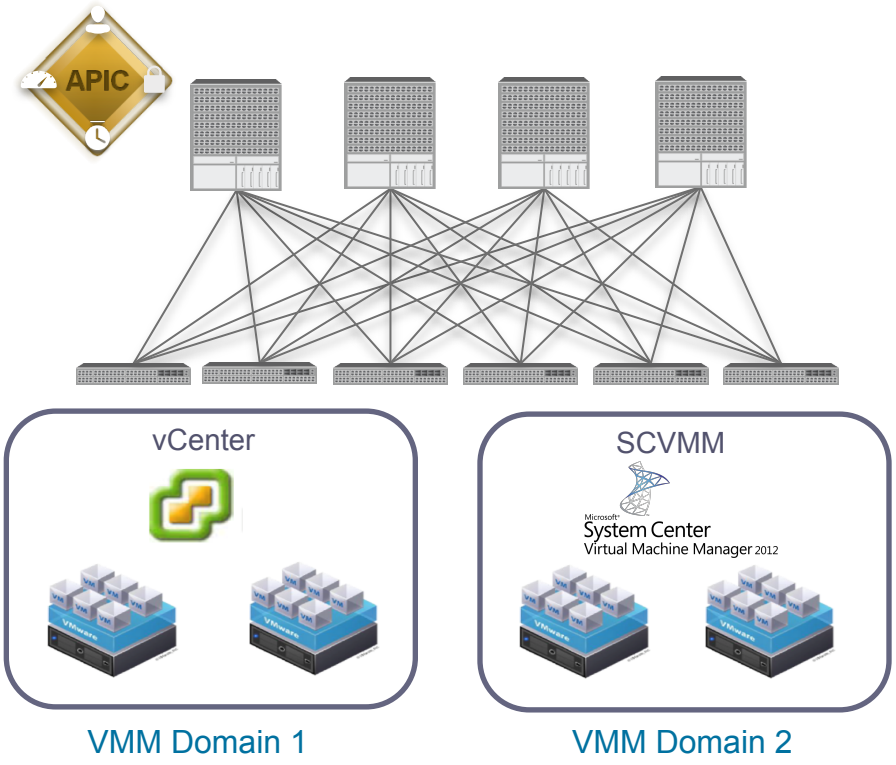

- Multiple Virtual Machine Managers (VMMs) likely on a single Fabric
- Each VMM and associated Virtual hosts are grouped within APIC
- Called VMM Domain

Hypervisor Integration with ACI
VMM Domains & VLANs

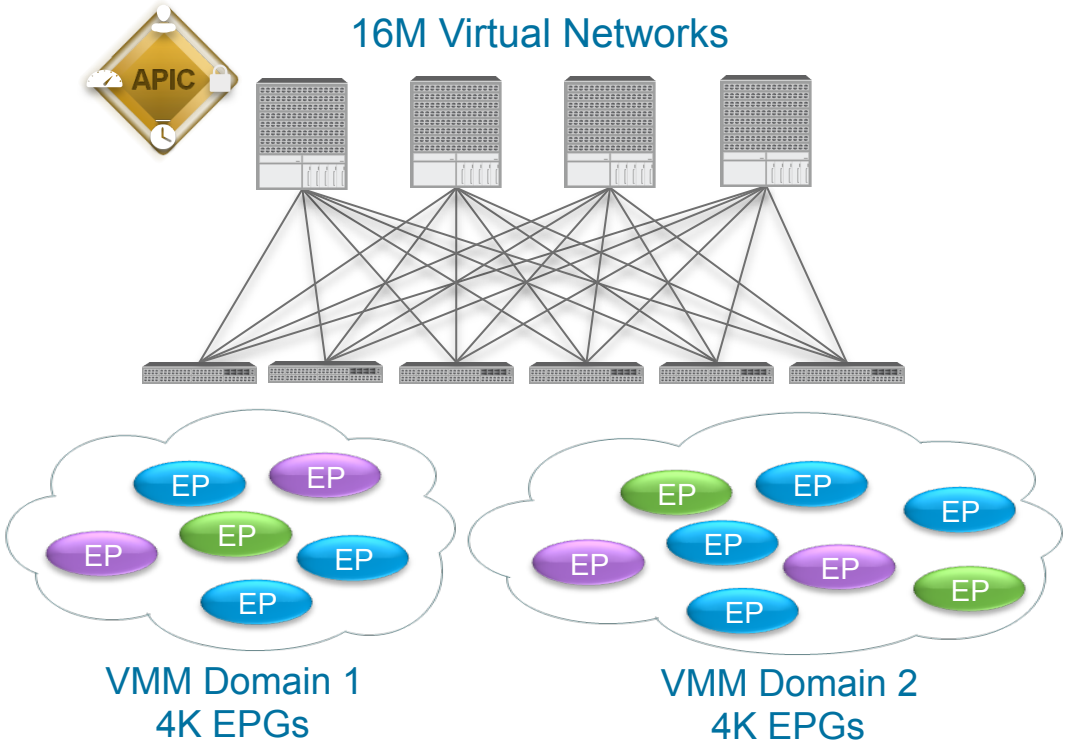

- VLAN ID only gives 4K EPGs (12 bits)
- Scale by creating **“pockets”** of 4K EPGs
- Map to scope of live migration
- Place VM anywhere
- Live migrate within VMM domain

EPG Spanning across VMM Domains

- EPGs can take different network identities across VMM Domain
- Applications can be deployed across VMM Domains
- VM Mobility is not allowed between VMM Domain due to vCenter/SCVMM limitation

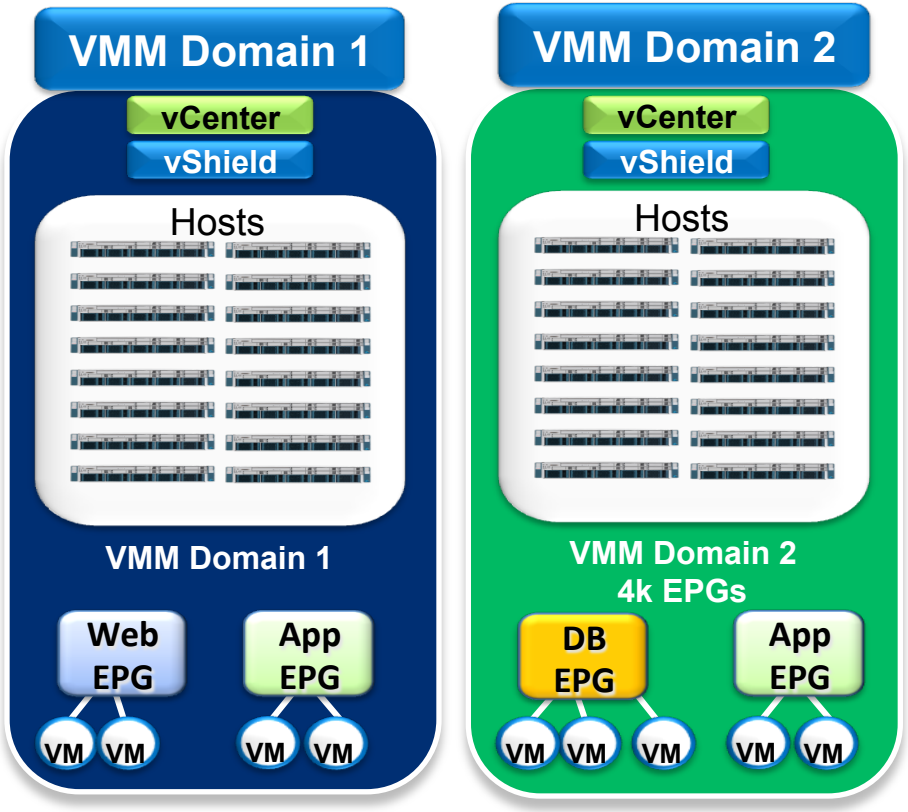

Recommended Practice for VLAN Networks

- Well separated VMM Domains
  - VLAN name space can be overlapping
- Separate VLAN name space when VMM domains share TOR

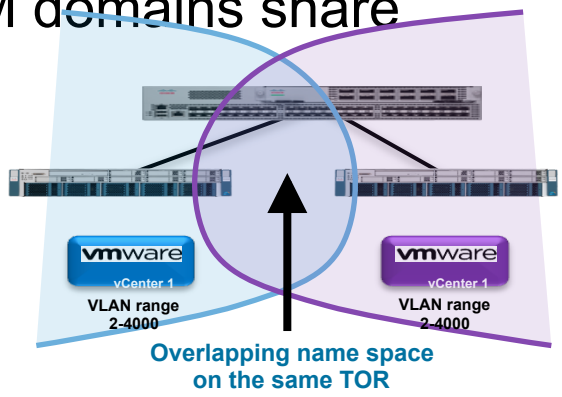

VMM Definition to avoid

Best Practice for VMM Domain definition

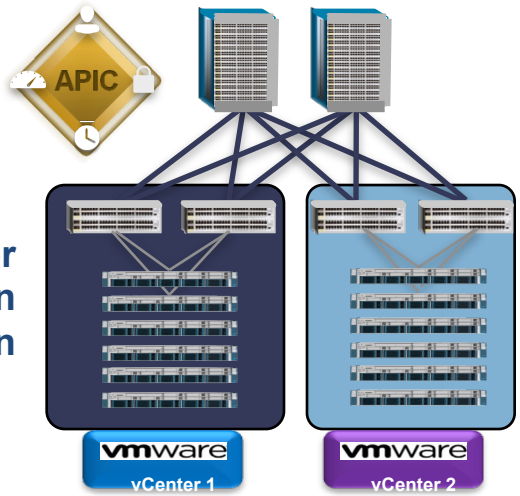

Integration with VMWare DVS

ACI Fabric and VMWare DVS Integration

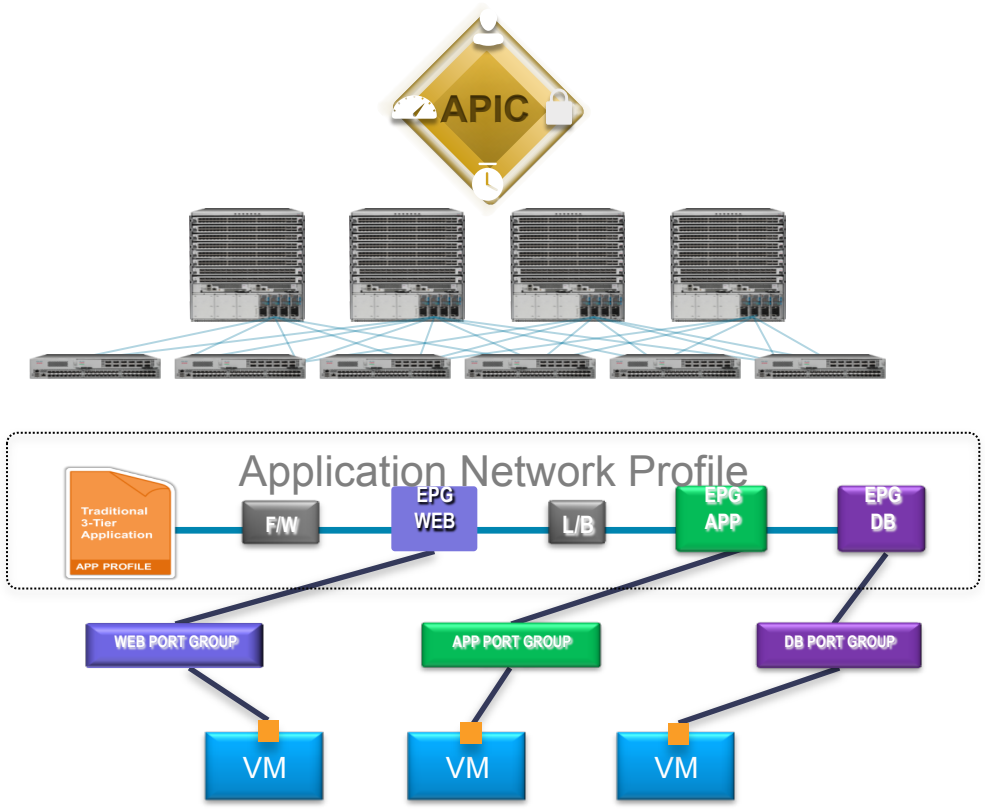

- How does ACI Fabric implement policy ?
  - Assigning EPs to EPGs
- What are EPs in virtual environment ?
  - VM vNICs
- How does VMware apply network configuration ?
  - Port Groups
- How are EPGs exposed to VMware ?
  - Map EPGs to Port Groups

End-Point identification

- Physical end-point

  Baremetal server

- Virtual end-point

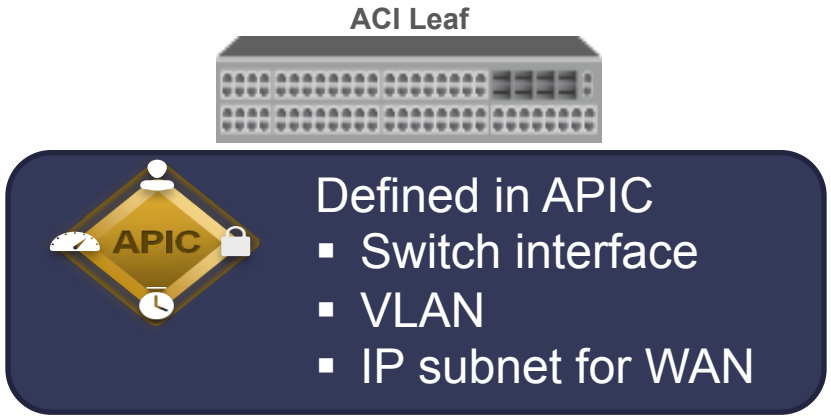

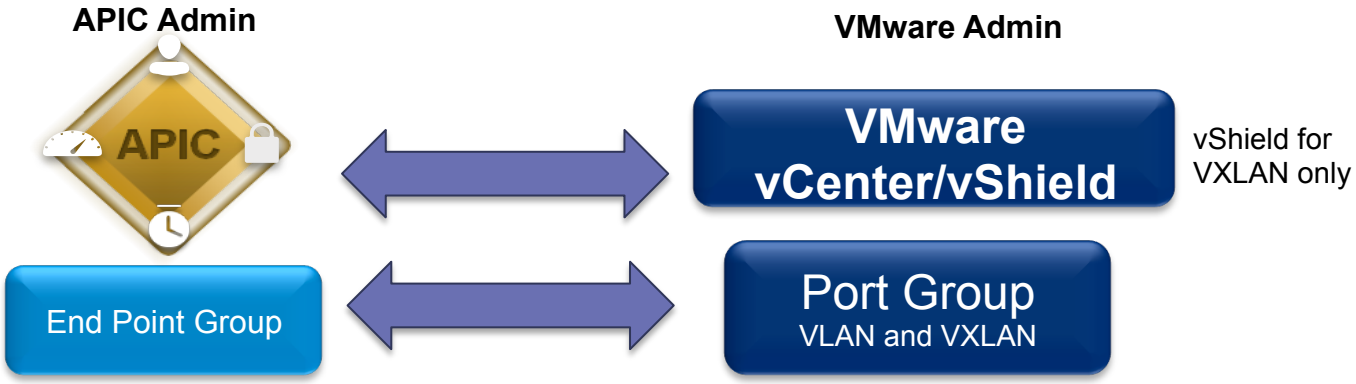

Cisco ACI Hypervisor Integration – VMWare DVS

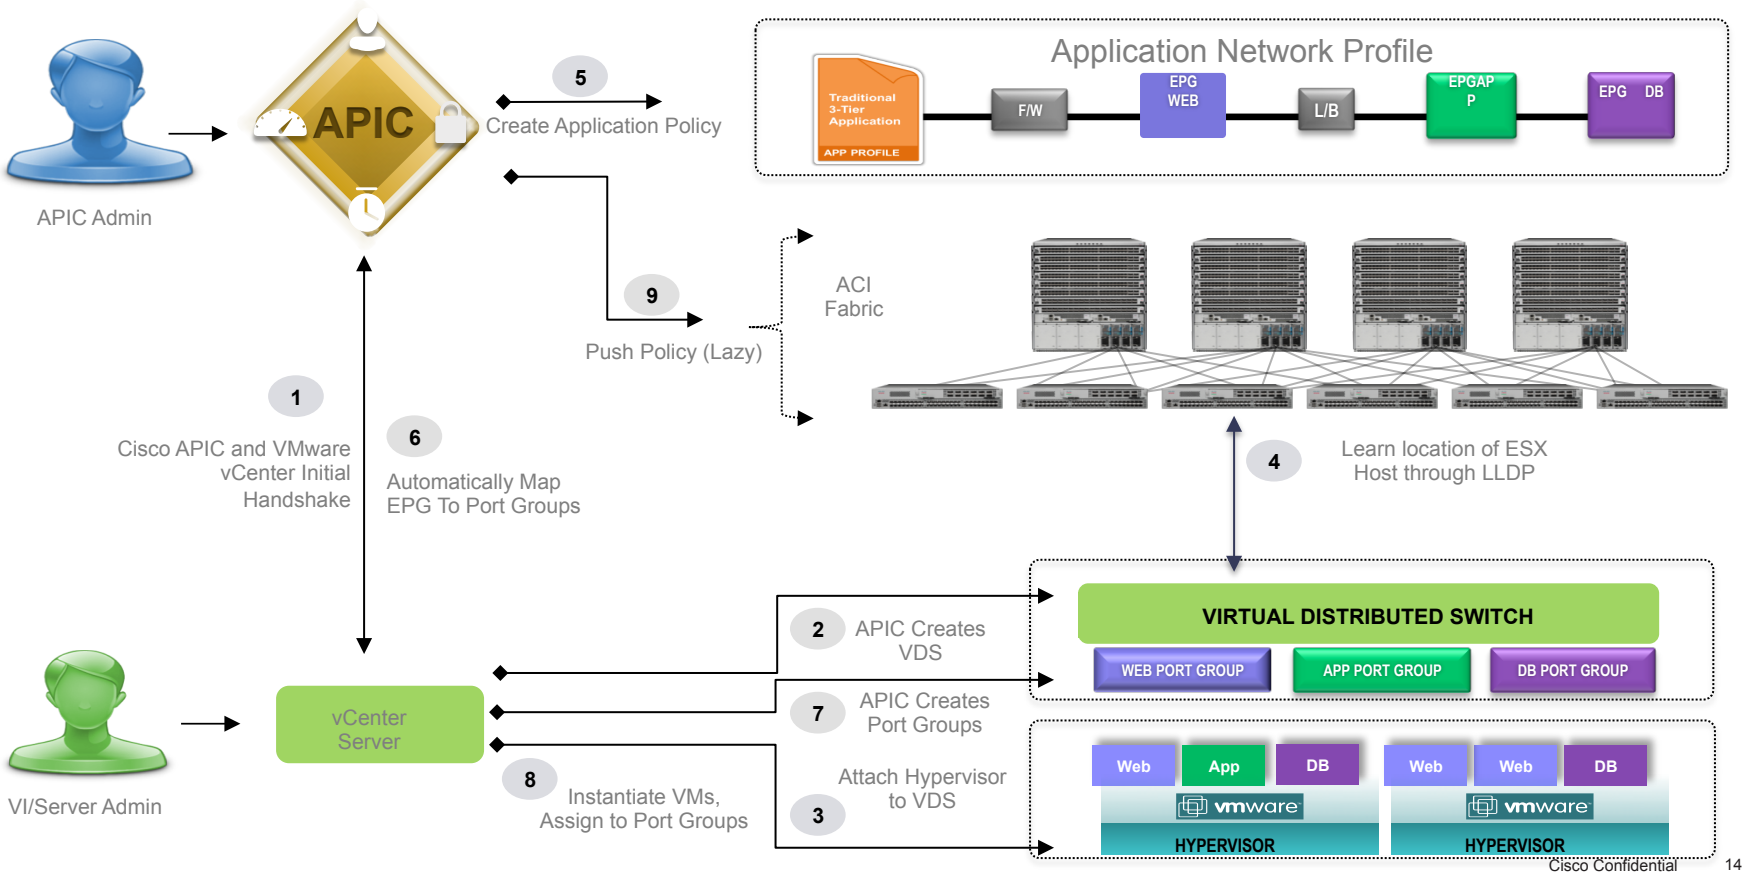

Pre-Staging VMware Integration

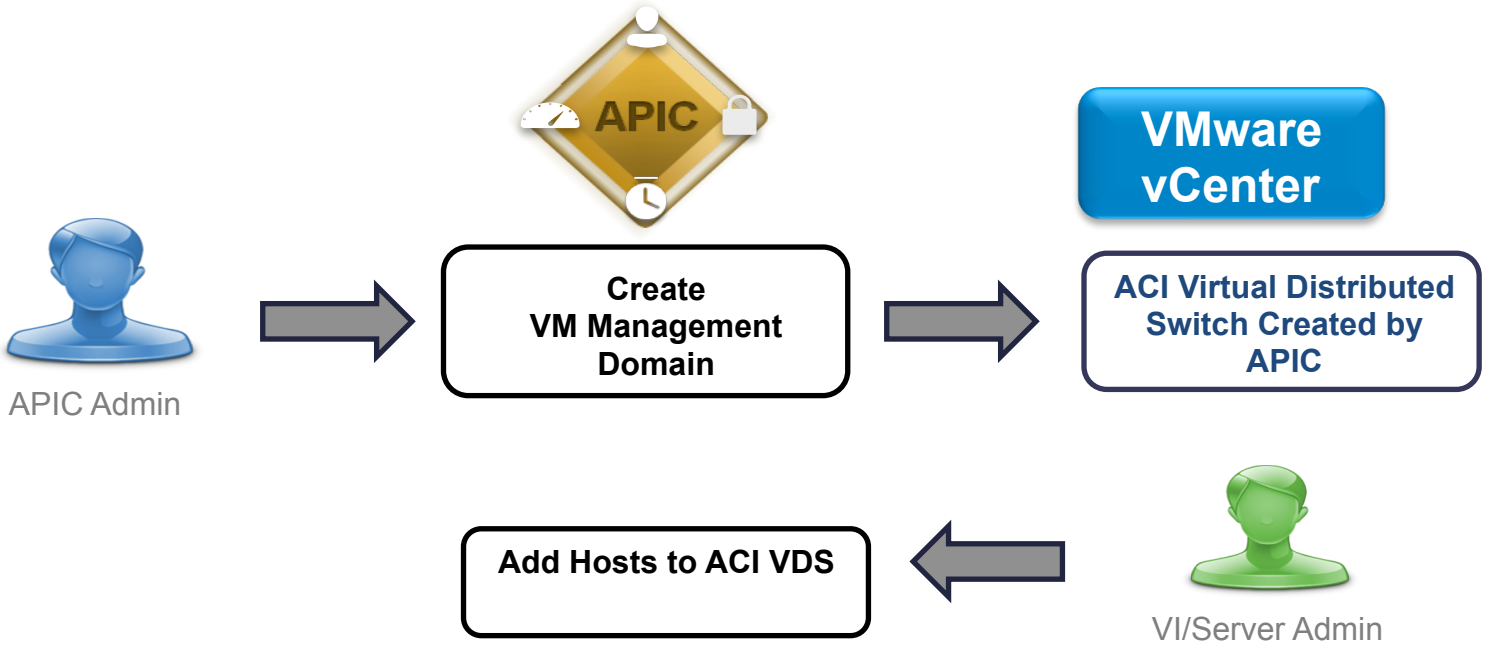

VMware Integration – App Instantiation

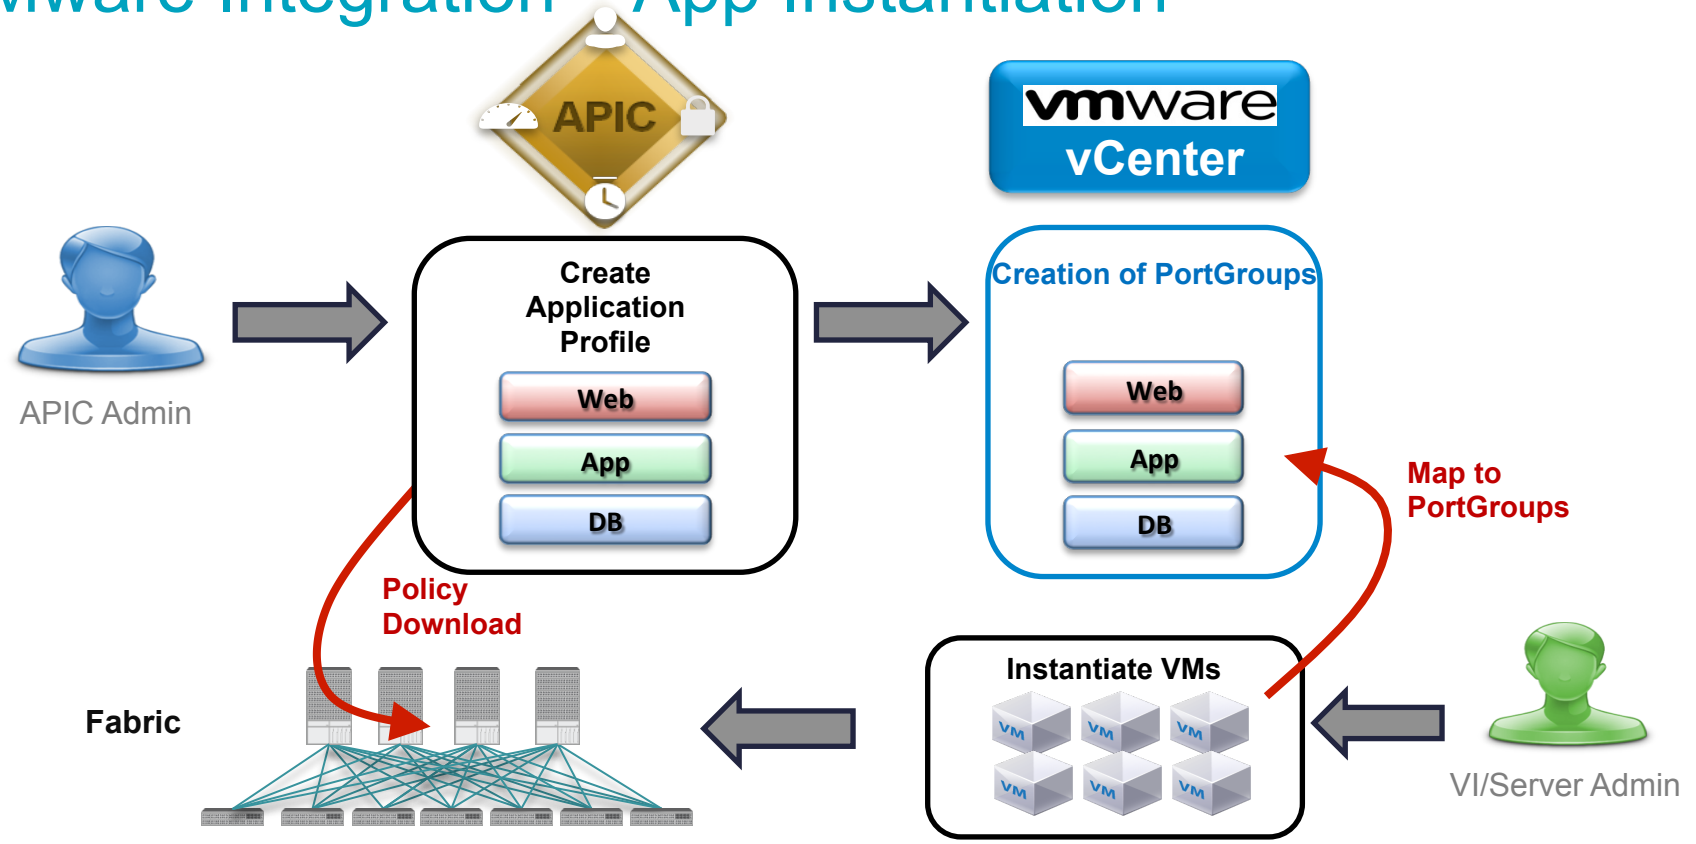

ACI VMWare Integration – Create VMM Domain

Create vCenter Domain

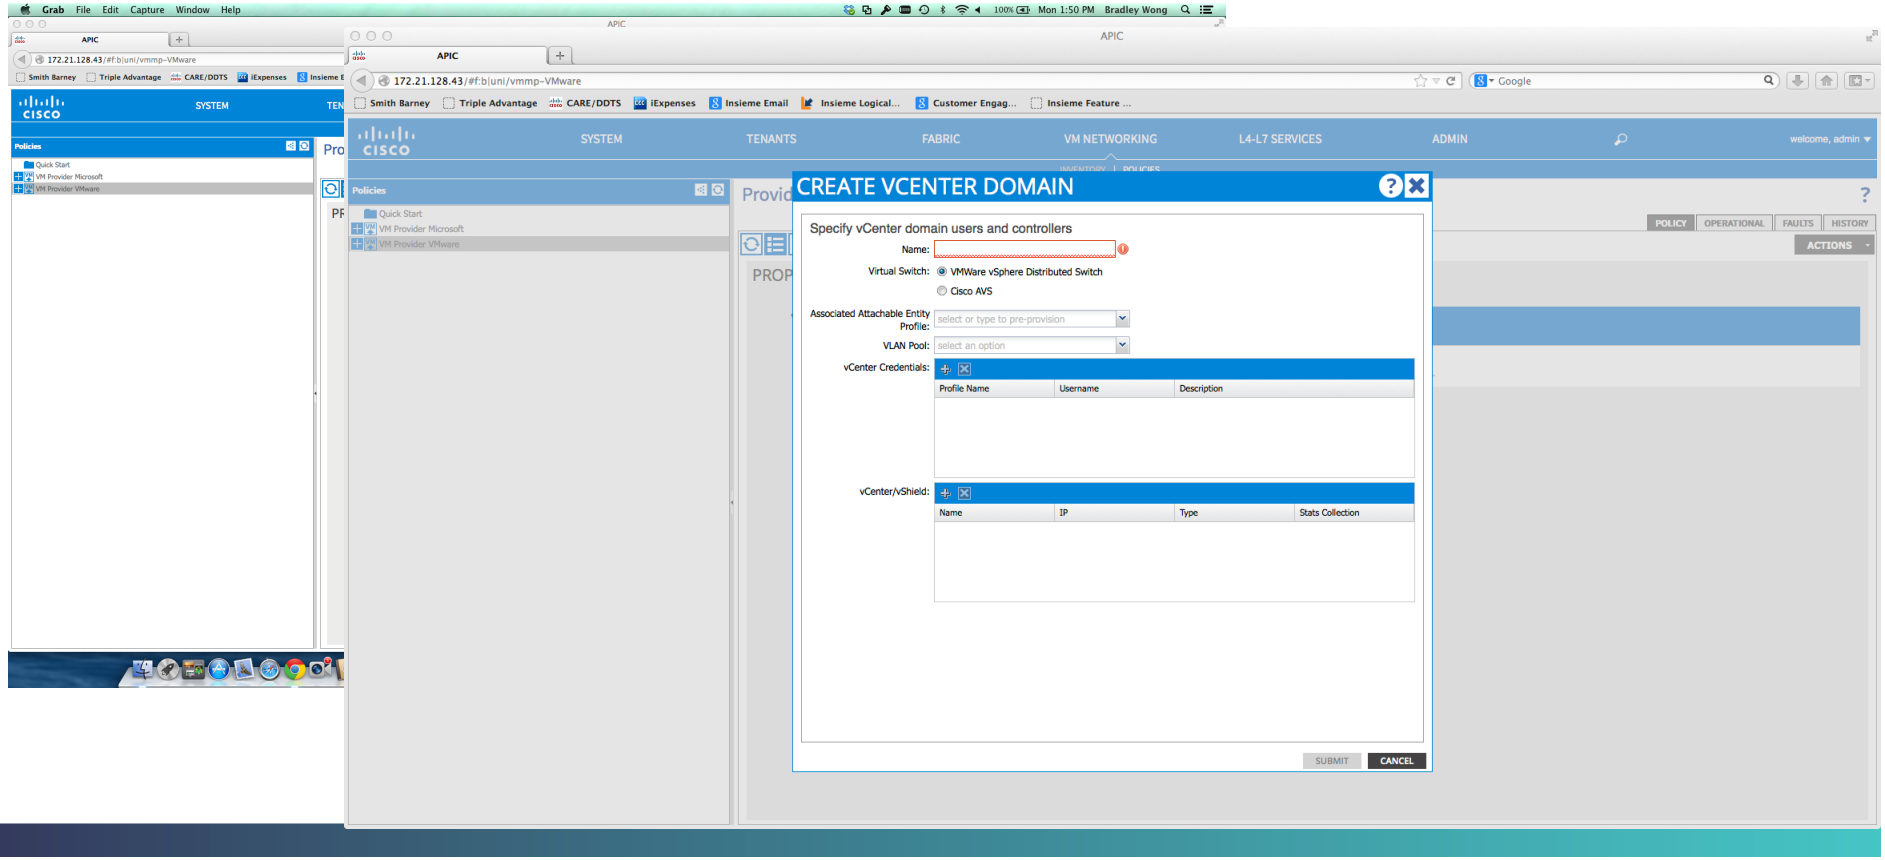

ACI VMWare Integration – Create VMM Domain

Create VLAN Namespace

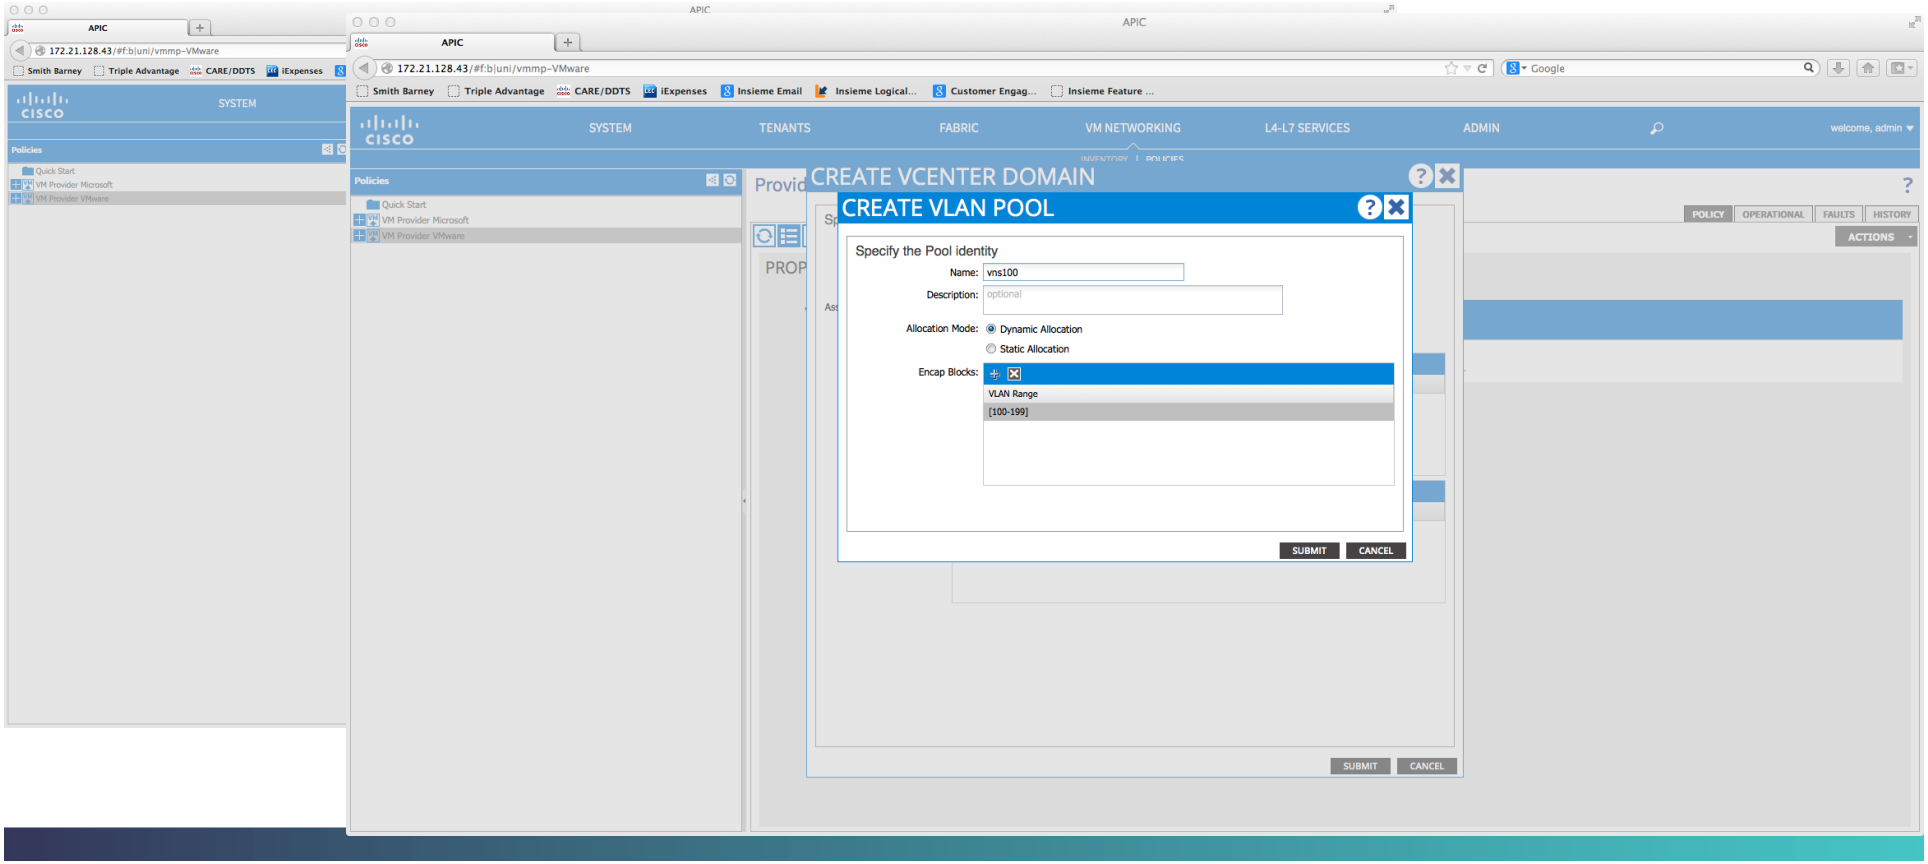

ACI VMWare Integration – Create VMM Domain

Create vCenter Credentials

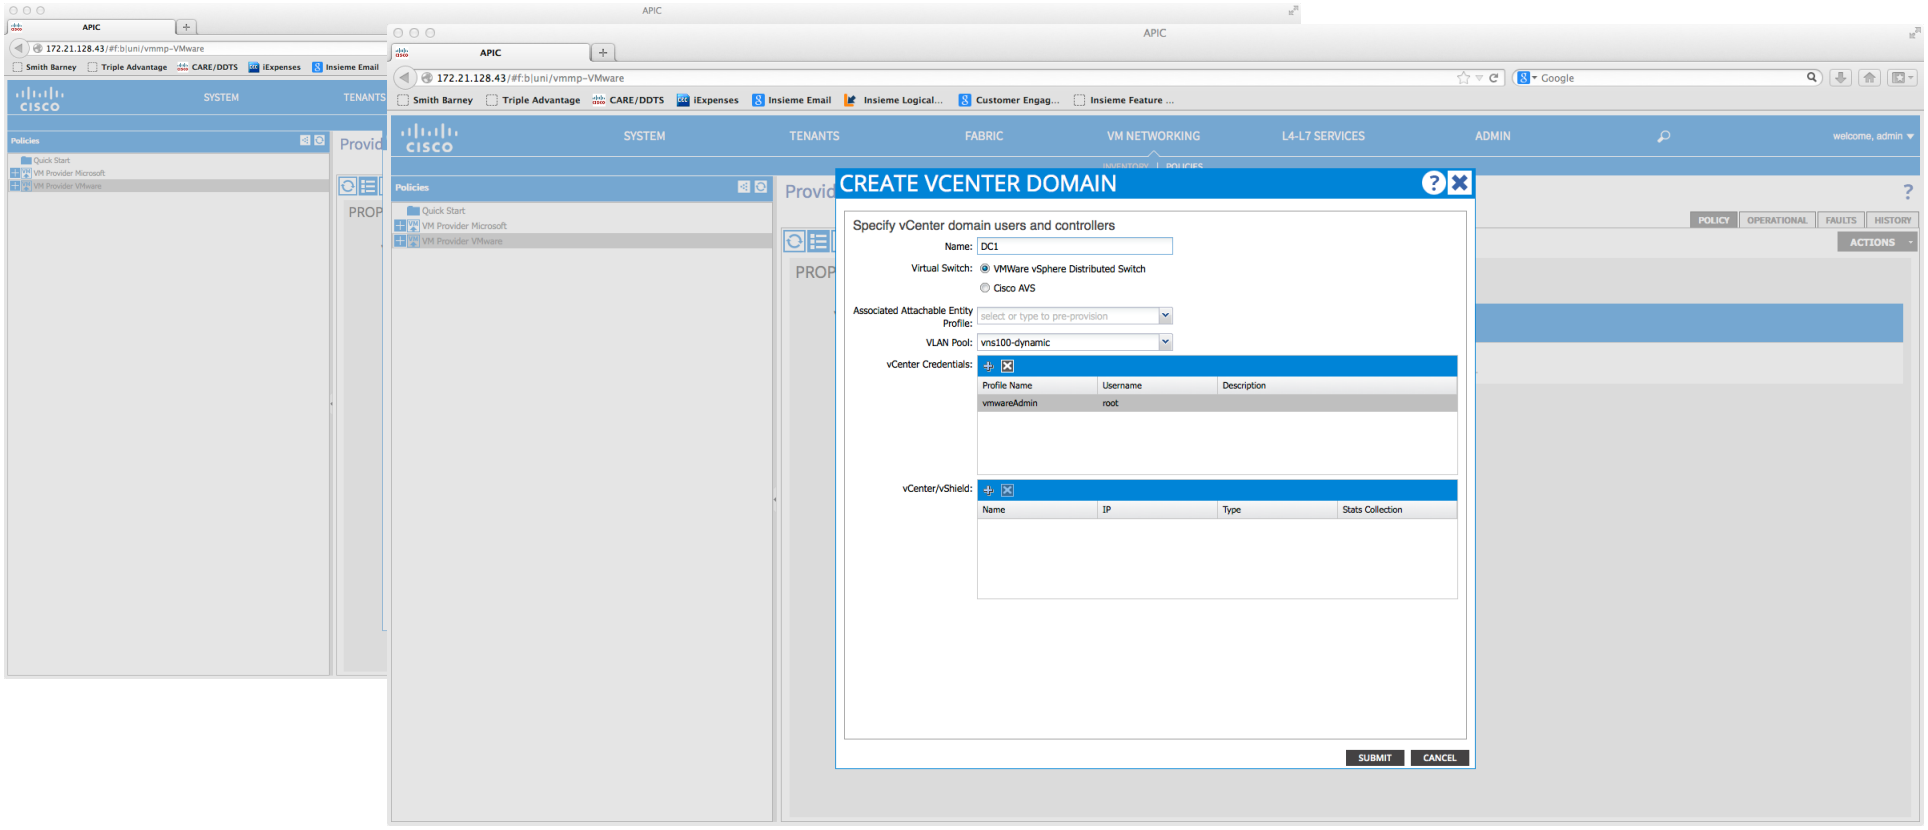

ACI VMWare Integration – Create VMM Domain

Create vCenter Controller Association

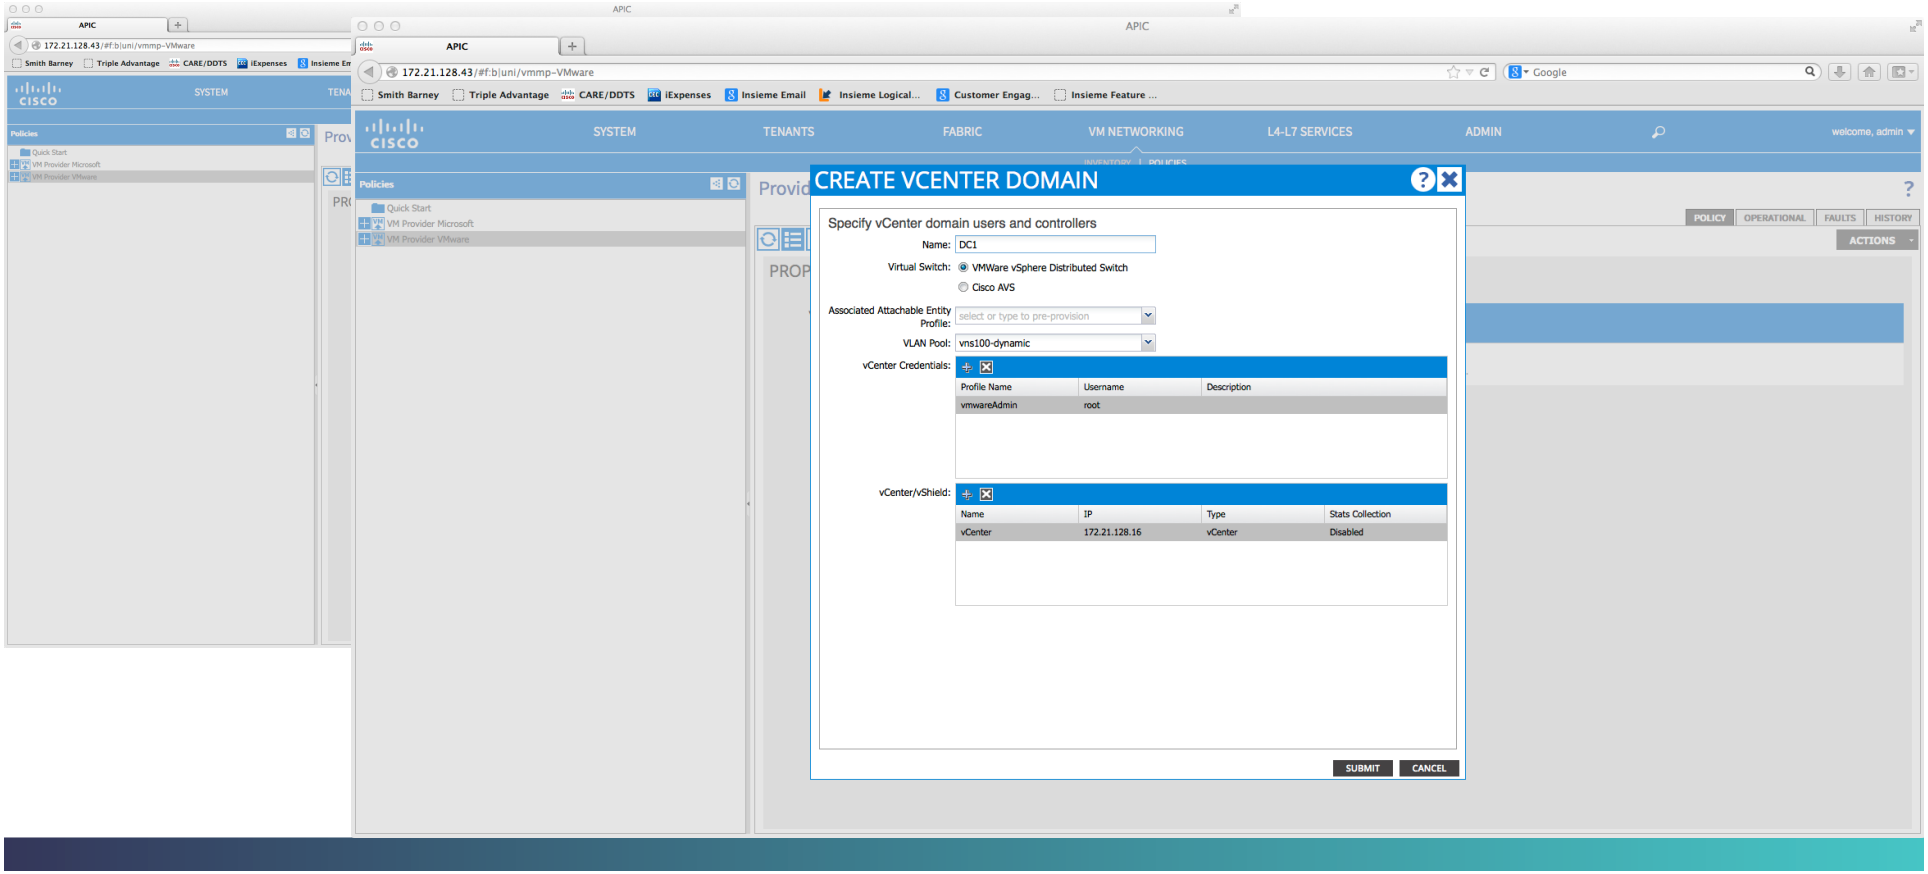

ACI VMWare Integration – Create VMM Domain

Display VMM Domain Association

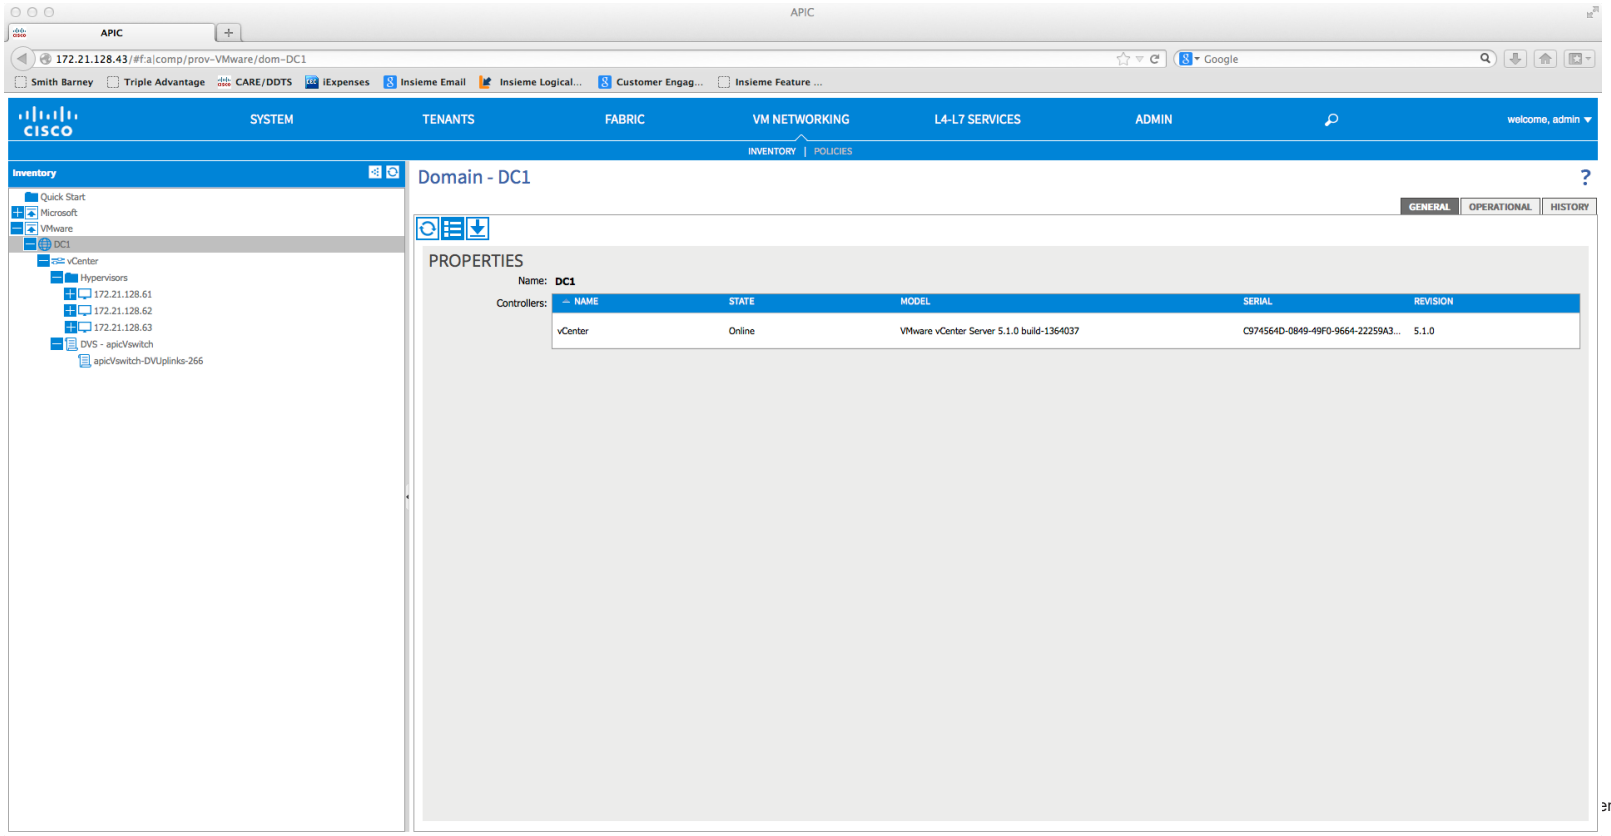

ACI VMWare Integration – Create VMM Domain

Display new DVS on vCenter

ACI VMWare Integration – Associate EPG to VMM

Add VMM Domain Association

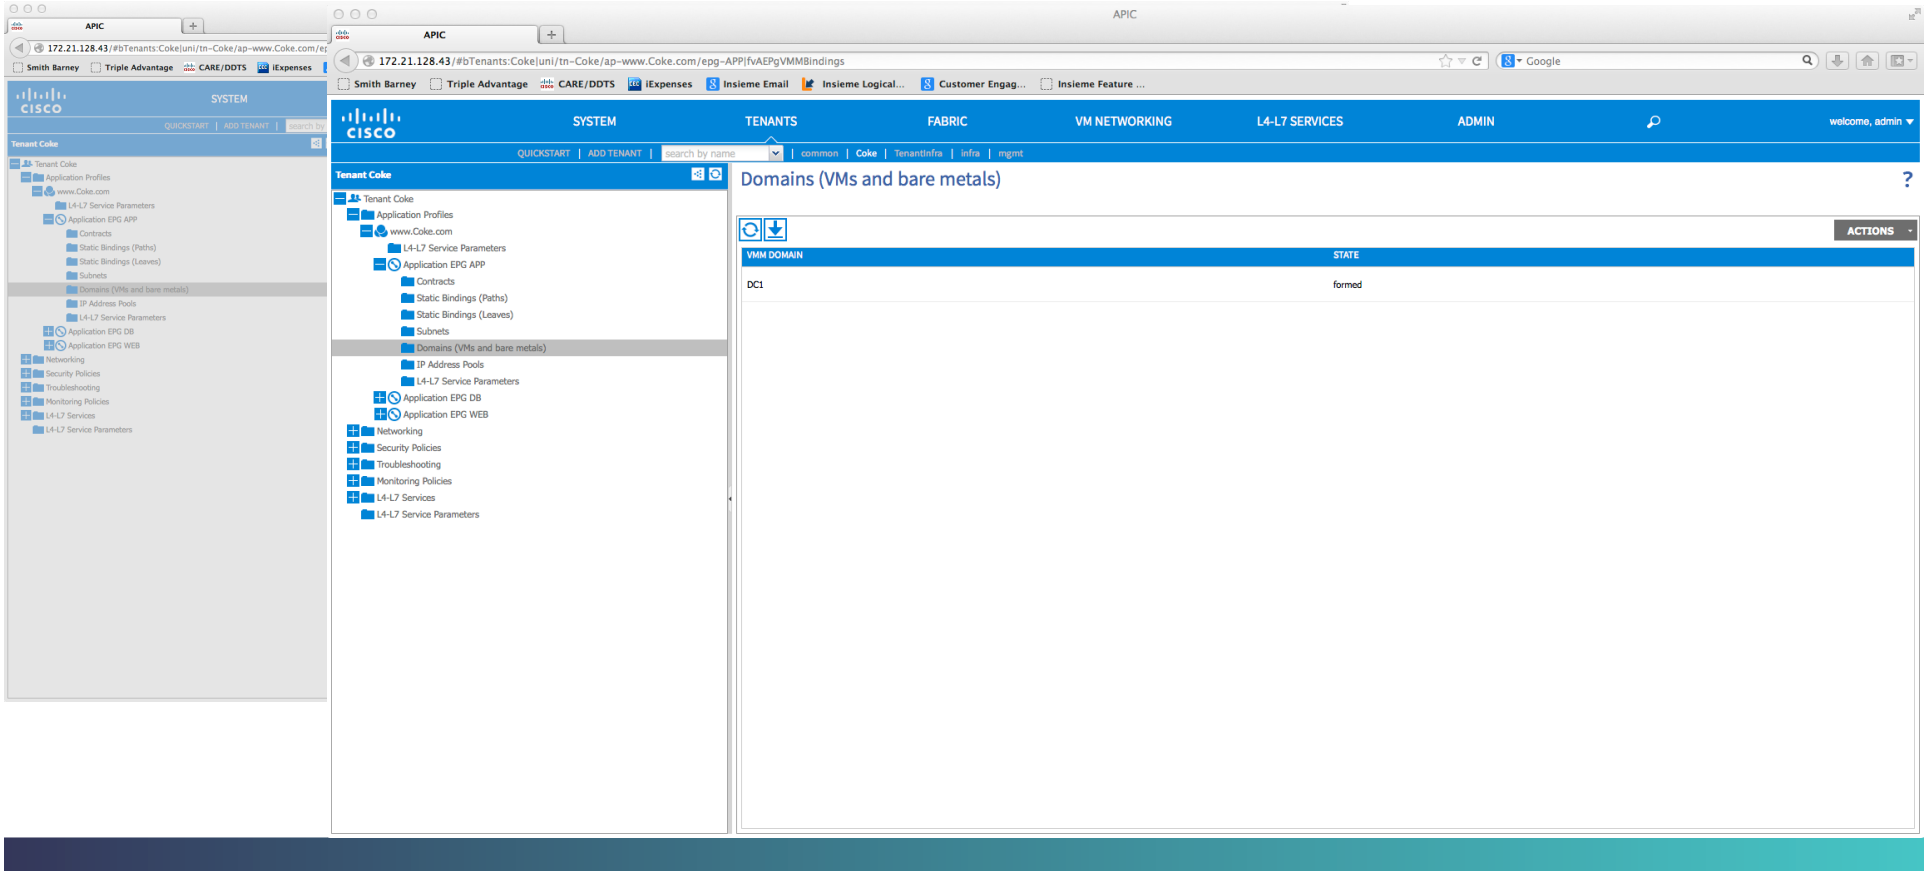

ACI VMWare Integration – Associate EPG to VMM

Add VMM Domain Association

Integration with Microsoft

Microsoft integration
DEMO ONLY

- Microsoft integration with fabric is a DEMO ONLY at FCS
- Integration using VLAN networking only, no NVGRE
- Full feature support is targeted for FCS+6 with SCVMM 2012

Microsoft Azure Pack Integration

- Integration with Microsoft requires:
  - Windows Server 2012
  - Systems Center 2012 R2 with SPF
  - Windows Azure Pack
- Azure Pack provides single pane of glass for Definition, creation, management of their cloud service
- Divided into Provider (Admin) portal and Consumer Self-Service (Tenant) portal

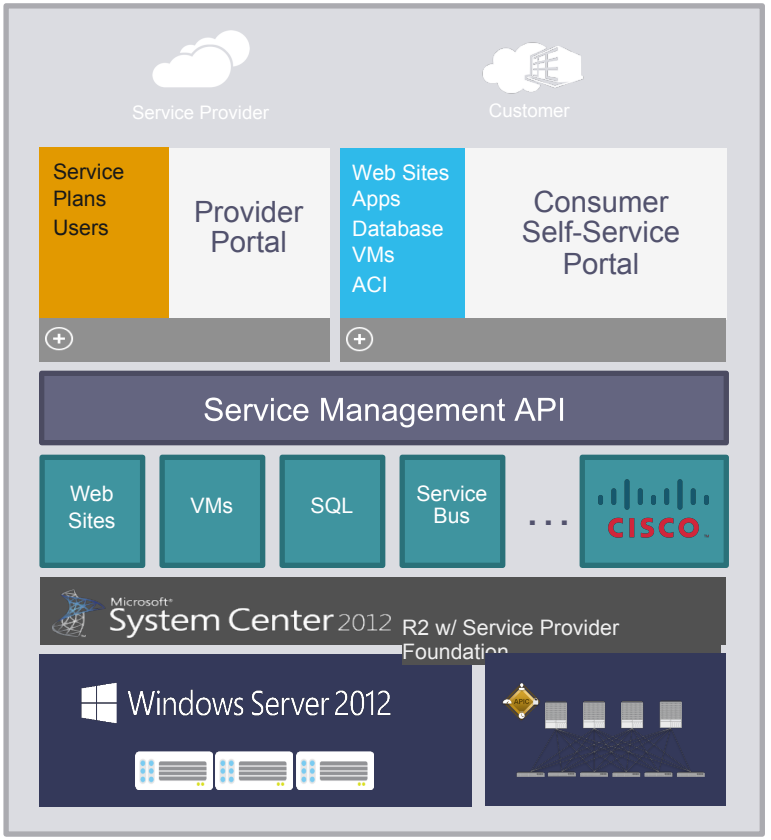

Microsoft Azure Pack Integration
Admin Experience

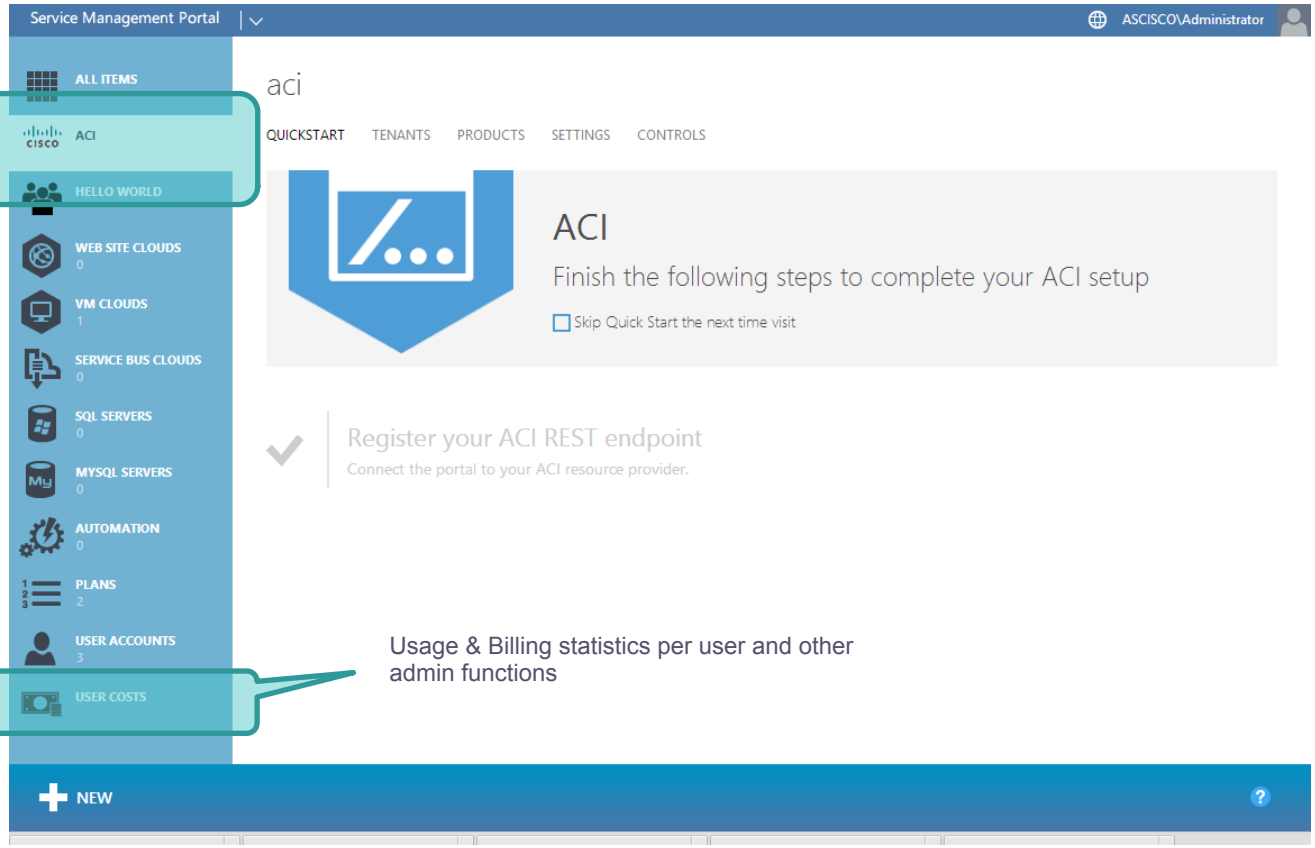

Add & Configure service providers for this deployment (APIC IP Address, Login Credentials, etc.)

Usage & Billing statistics per user and other admin functions

Microsoft Azure Pack Integration
Tenant Experience

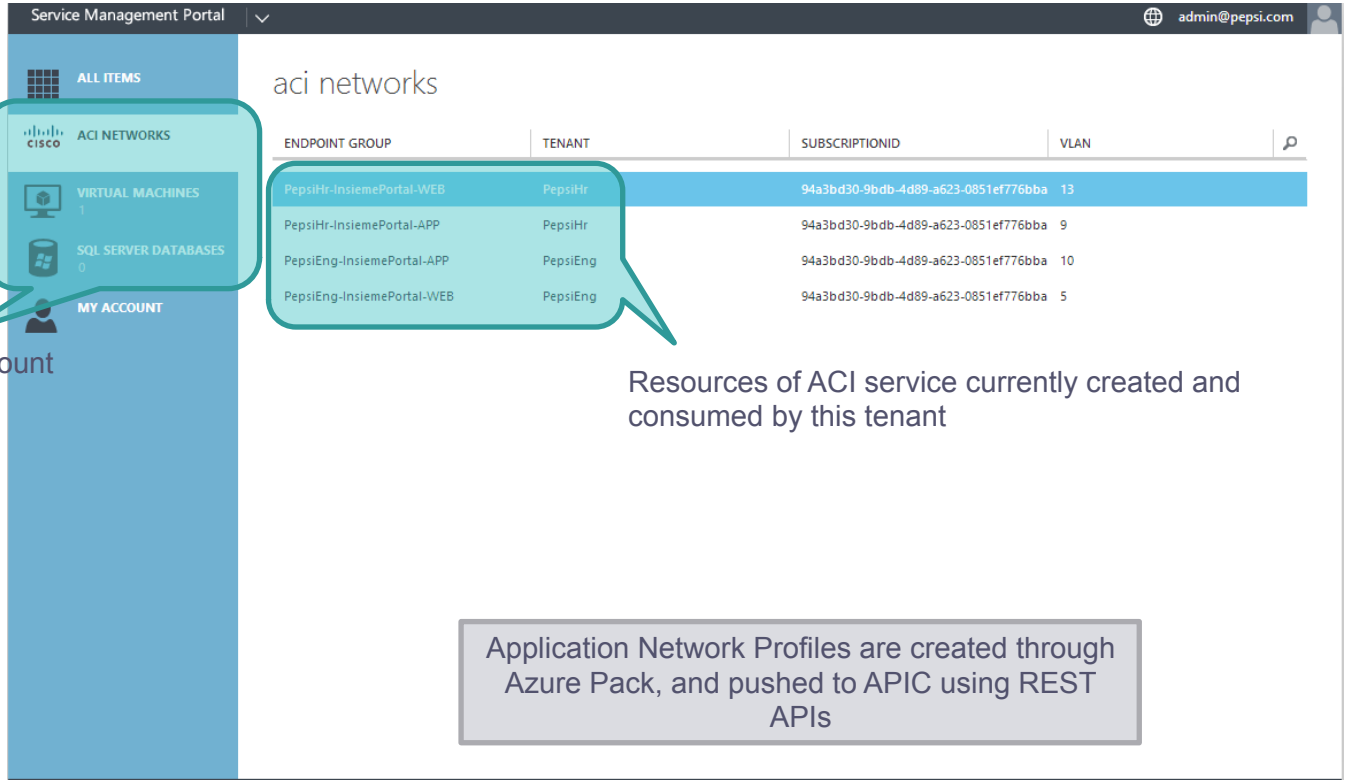

Services this account has access to

Resources of ACI service currently created and consumed by this tenant

Application Network Profiles are created through Azure Pack, and pushed to APIC using REST APIs

ACI Azure Pack Integration

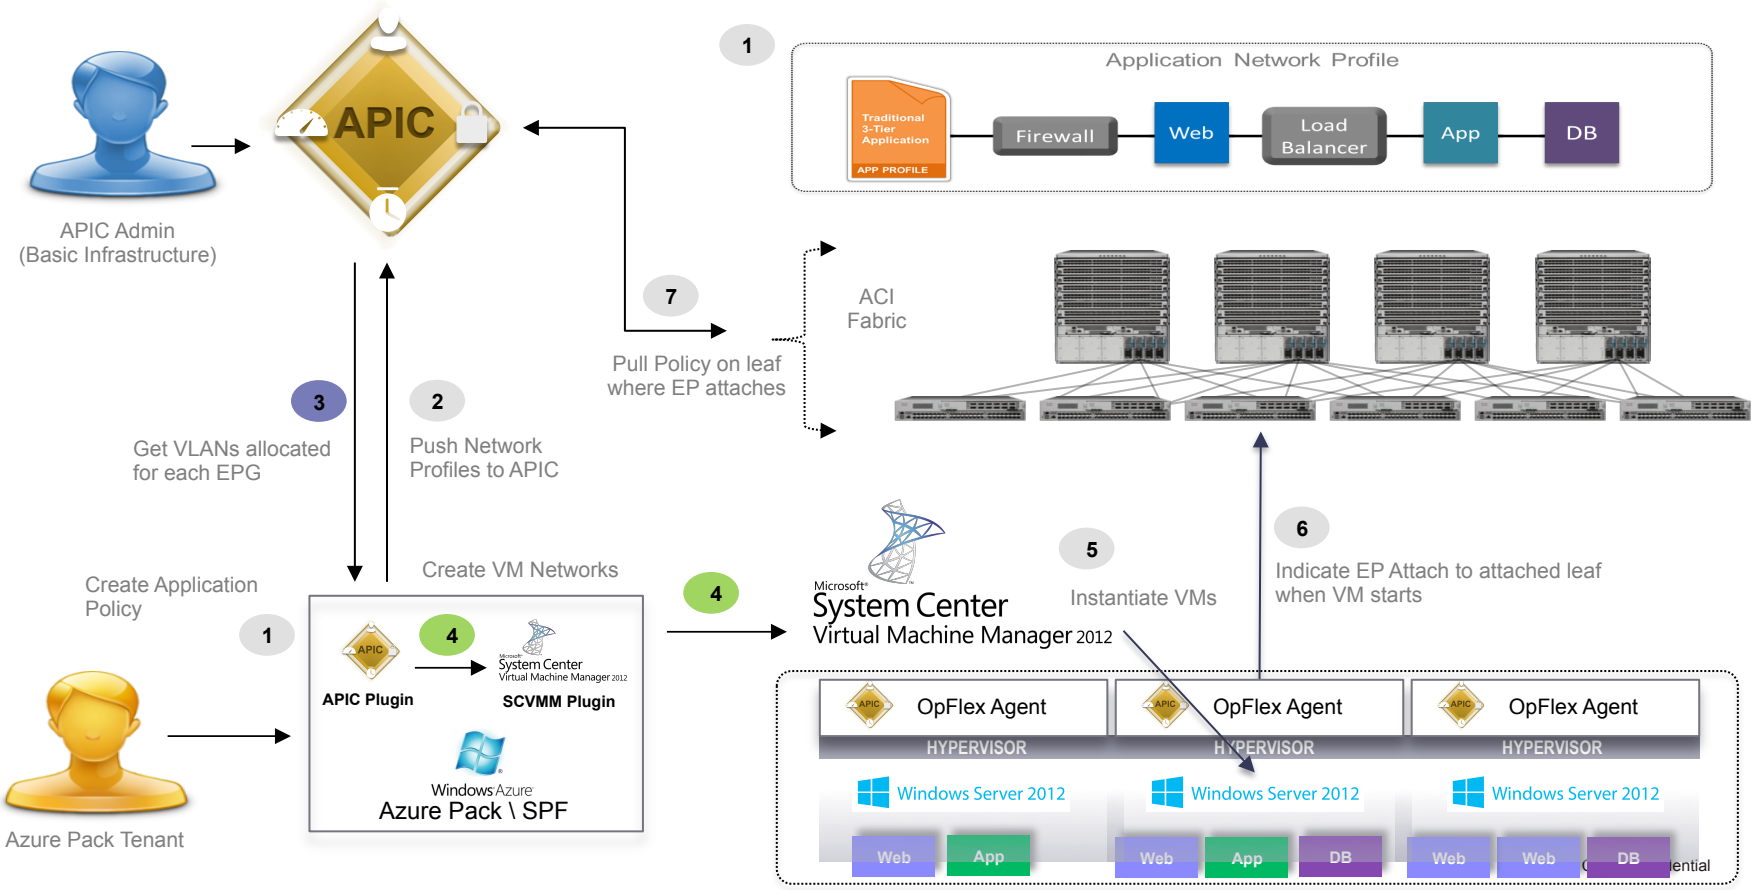

Integration with OpFlex

OpFlex – A Flexible, Extensible Policy Protocol

OPFLEX is a new extensible policy resolution protocol designed for declarative management of any datacenter infrastructure. Unlike legacy protocols such as OVSDB, OPFLEX was designed to offer:

- Declarative resolution – Push + Pull API support
- Abstract policies rather than device-specific configuration
- Flexible, extensible definition of using XML / JSON
- Support for any device – vswitch, physical switch, network services, servers, etc.

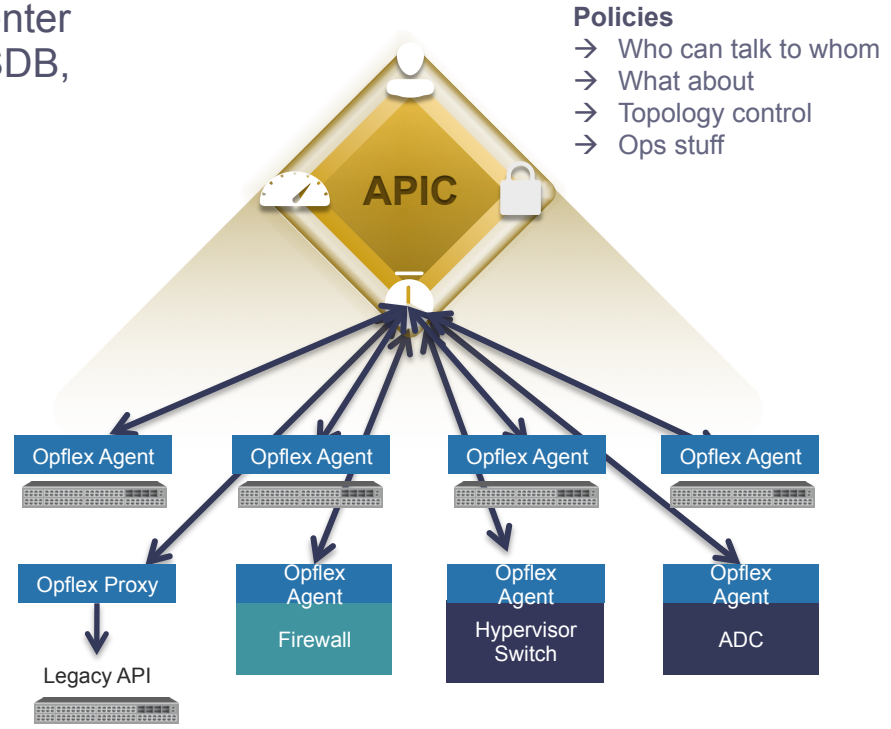

Hypervisor Integration with ACI
Endpoint Discovery

- Virtual Endpoints are discovered for reachability & policy purposes via 2 methods:
- **Control Plane Learning:**
  - Out-of-Band Handshake: vCenter APIs
  - Inband Handshake: OpFlex-enabled Host (N1KV, Windows Server 2012, etc.)
- **Data Path Learning:** Distributed switch learning

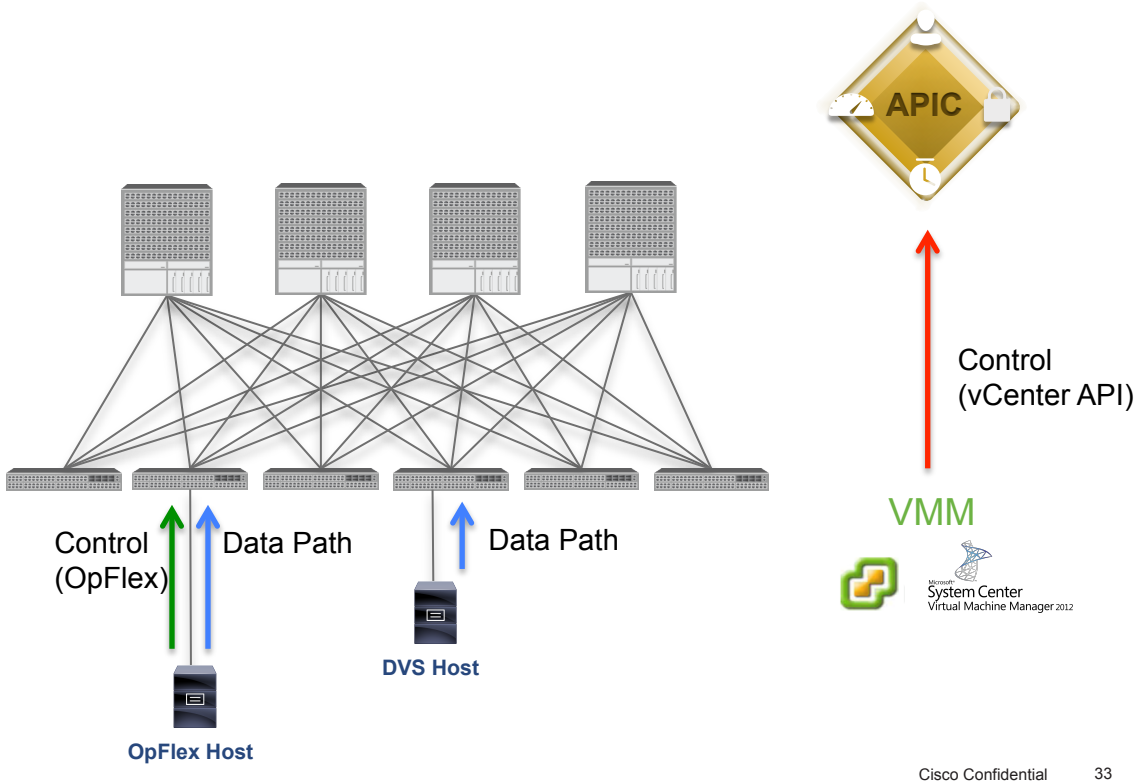

Hypervisor Integration with ACI
Policy Resolution Immediacy

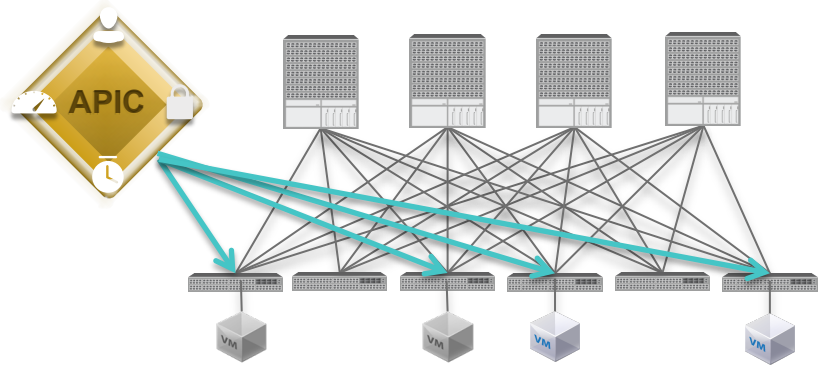

- Policies are pushed to Leaf nodes based on Resolution Immediacy defined upon association of EPG to VMM Domain
  - **Immediate**: All policies (VLAN / NVGRE / VXLAN bindings, Contracts, Filters) pushed to leaf node upon Hypervisor pNIC attachment. LLDP or OpFlex used to resolve Hypervisor to Leaf node attachment.
  - **Lazy:** Policies only pushed to leaf node upon pNIC attachment AND vNIC association with port-group (EPG)
- Policy programming in Leaf node hardware based on Instrumentation Immediacy
  - **Immediate:** Policies programmed in Policy CAM once received by APIC as defined by Resolution Immediacy Policy
  - Lazy: Polices programmed in hardware Policy CAM only when reachability is learnt through data path

Design Considerations
VLAN-Based Hypervisor Networks

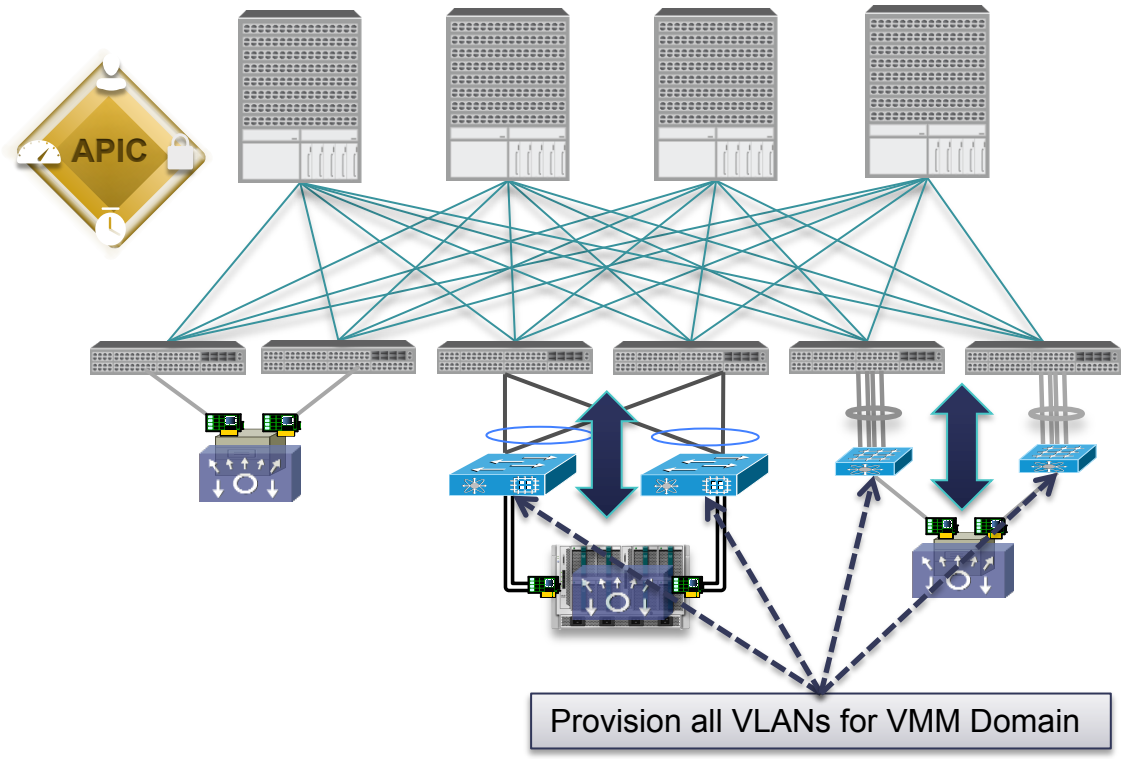

- Hosts are assigned VLAN ID to EPG binding through VMM & APIC Integration
- Intermediate L2 nodes not managed – need to manage VLANs on these for each VMM Domain
- Endpoint location discovered through “stitching” LLDP TLVs (non OpFlex-enabled Hosts)

Design Considerations
VXLAN & NVGRE-based Hypervisor Networks

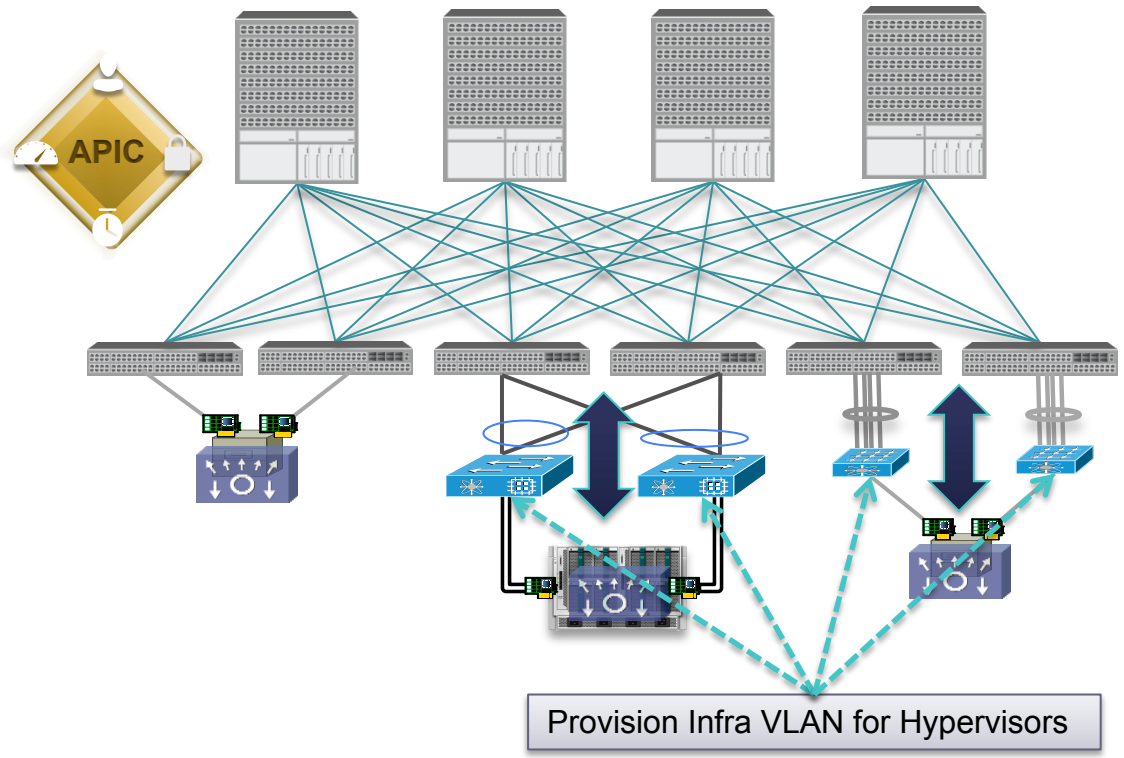

- Hosts are assigned VNID and VSID to EPG binding through VMM & APIC Integration
- Infra-VLAN is extended out to front-panel tenant ports - Infra-VLAN needs to be provisioned on intermediate L2 Nodes
- Endpoint location discovered though “stitching” LLDP TLVs (non OpFlex-enabled Hosts)

Roadmap Virtualization Integration

Fabric FCS

VMware

- VMware vCenter and vCNS 5.1/5.5
- Network Policy coordination
- VM Attach event detection and policy placement
- VMotion event detection and network policy mobility
- VM network stats collection

Microsoft - **Demo only, no production release**

- Network Policy coordination with SCVMM 2012
- VM attach/detach, VM Mobility event detection and policy placement

Nexus 1000V

- Non-switching and local switching modes
- OpFlex policy model implementation
- APIC Network policy extension into N1KV
- APIC centralized controller for N1KV VEMs

Fabric FCS+6

VMware

- VMware 2014

Microsoft

- SCVMM 2014 vNext+

KVM OVS with OpenStack Integration

Integration Features comparison between VMware, MS and OpenStack

	VMware	Microsoft	OpenStack
Products APIC integrated with	VMware vCenter, vShield, vCloud Director 5.1	System Center Virtual Machine Manager (SCVMM 2012)	Havana Release
Networks	VLAN and VxLAN	VLAN only Demo at Fabric FCS	VLAN & VXLAN using Quantum Provider network extension
Network Policy coordination	APIC creating a PortGroup in vCenter/ vShield for VLAN and VxLAN networks	VM Networks in SCVMM	OpenStack Quantum plug-in creating APIC Application profiles APIC EPG mapping to networks & Security Groups
VM Movement detection and network policy mobility	vCenter notification to APIC and detection on the wire	Via Insieme Host Agent installed in Hyper-V relaying VM mobility notification	
VM attach to networks detection	vCenter notification to APIC and detection on the wire	Via Insieme Host Agent installed in Hyper-V relaying VM attach/detach events	Via the Quantum Plug-in
VM Network Statistics collection	VM health, status, statistics Per EPG stats		Via OpenStack Ceilometer
Application Visibility	Per application, knowledge of all end points (virtual and physical) attached to EPGs and L4-7 Services	Per application, knowledge of all end points (virtual and physical) attached to EPGs and L4-7 Services	Per application, knowledge of all end points (virtual and physical) attached to EPGs and L4-7 Services

Thank you.