Cisco Next Generation Data Center Design

Maximizing Application Performance, Security And Economics To Accelerate Digital Transformation

Youssef Boukari

Solution Architect Director

3S TECHNOLOGY CITY

INNOVATION STARTS HERE.

3S Standard Sharing Software

3S TECHNOLOGY CITY

Agenda

- Challenges of IT
- ACI Architecture & Deployments
- ACI Use Cases
- Application Driven Data Center & Cisco VTS
- Scalable Fabric, Network Virtualization & EVPN/VXLAN
- Use Cases
- Q & A

3S TECHNOLOGY CITY

Applications Are Changing

Type
Consumption
Delivery

Big Data, Distributed Apps, Mobile

Cloud–public, Private, Hybrid

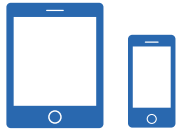

Anywhere, Anytime, Any Device

78% The network is even more critical to delivering applications than a year ago*

Deciding the Application Location(s)

Public, Private, Both?
Build, Buy, Rent?

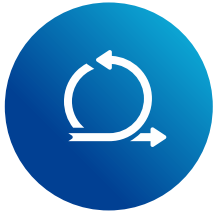

Empowering LoB & App. Developers

PublicCloud-like agility and simplicity
Self-Service Operations

Mitigating Risk

Securing Apps, Users, Data.
Compliance. Data sovereignty.

3S TECHNOLOGY CITY

Cisco Data Center Architecture Design

3S TECHNOLOGY CITY

Cisco's Differentiation:

Integration of DC/Cloud products to deliver the ASAP architecture

Consistent policy-defined model across entire hybrid cloud domain

Maximize Application Performance, economics and security Continuously

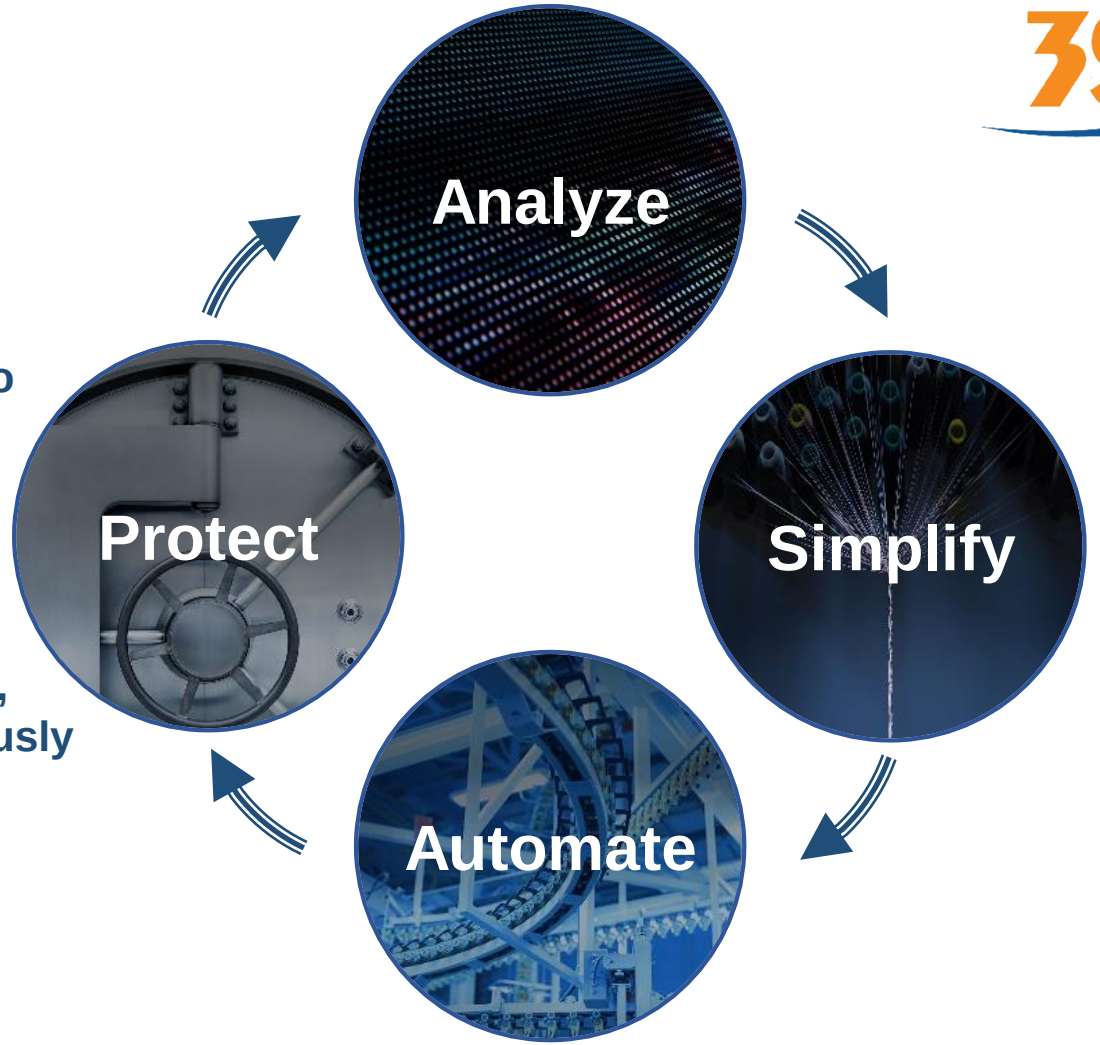


Analyze
Protect
Simplify
Automate
Automation
Analytics
Orchestration
SDN
APIC
ACI
Cisco Tetration Analytics
Cisco ONE
Enterprise
Cloud Suite
Service
Manager
(PSC)
Hardware Platforms
NEXUS
Networking
UCS
Compute
MDS 9000
Storage
HyperFlex
Best of Breed Solutions
Converged
Infrastructure
Hyper-Converged
Infrastructure

3S TECHNOLOGY CITY

Customer Journey with Cisco ASAP Data Center

Optimize Infrastructure	Simplify Operations	Build Cloud-native App Stack	Chose your Hybrid Cloud	Real-time Analytics
Nexus, MDS UCS Next-gen Firewalls Converged/ Hyperconverged			Cisco CloudCenter 	Cisco Tetration Analytics™
Performance	Automation	Containers	App. Benchmark	Monitor every flow
Scale	Unify Policy	Self-Service	Extend Policy	App dependency
Security	50+ eco-system	DevOps Tools	Securely move data & workloads	On-prem & Cloud

3S TECHNOLOGY CITY

ACI Real World Deployment

Requirements: Business Drivers & Solutions for Network Segmentation

- Multi-tenancy
- Security and Separation
- Traffic Engineering
- Scalable
- Flexible topology
- Minimise oversubscription
- Scale out and scale up
- Scalable L4-7 Service Layer
- No spanning tree
- Incremental scale
- Virtual FW/LB per tenant
- Flexible placement
- Incremental capacity

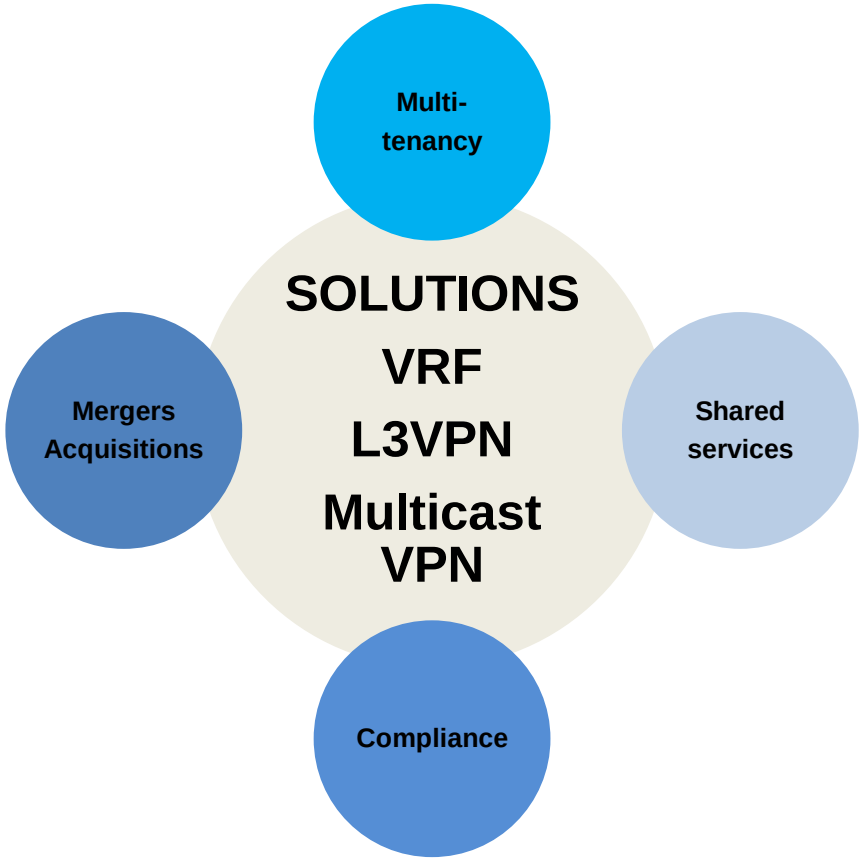

Customer Deployment: Application Centric Infrastructure (ACI)

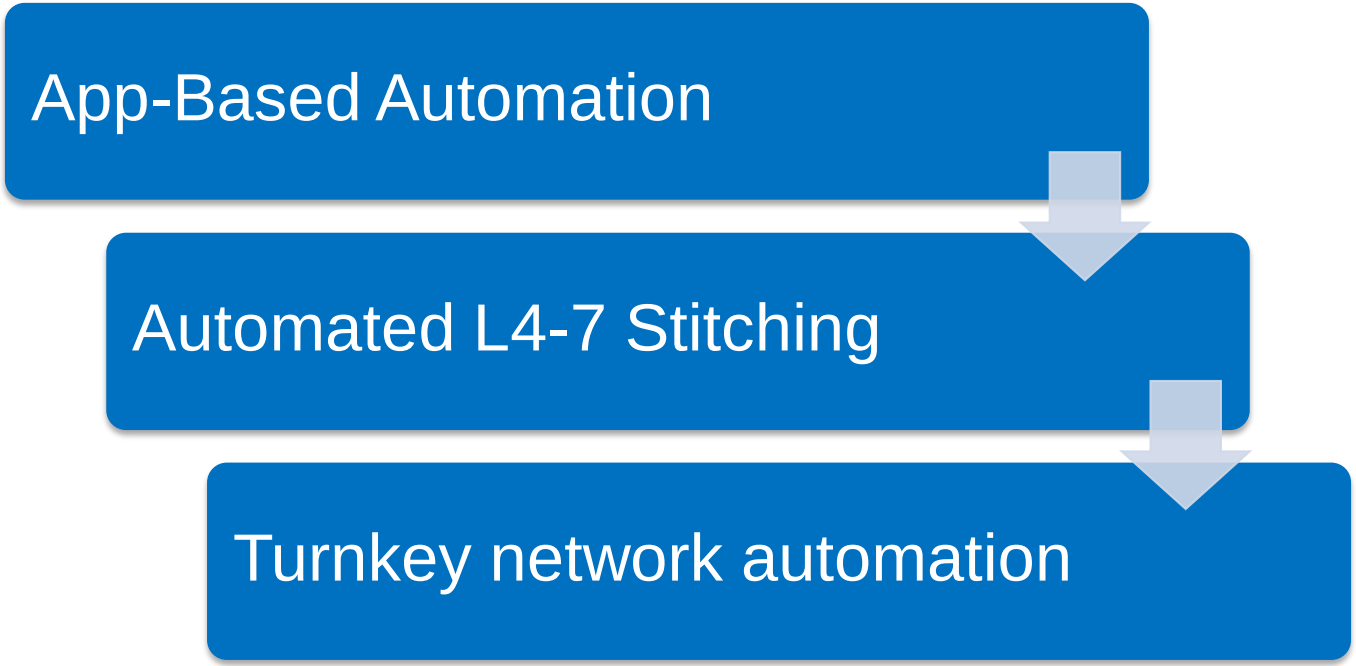

ACI Fabric Overview

Spine and Leaf Architecture / Design

Attaching the ACI APIC(s)

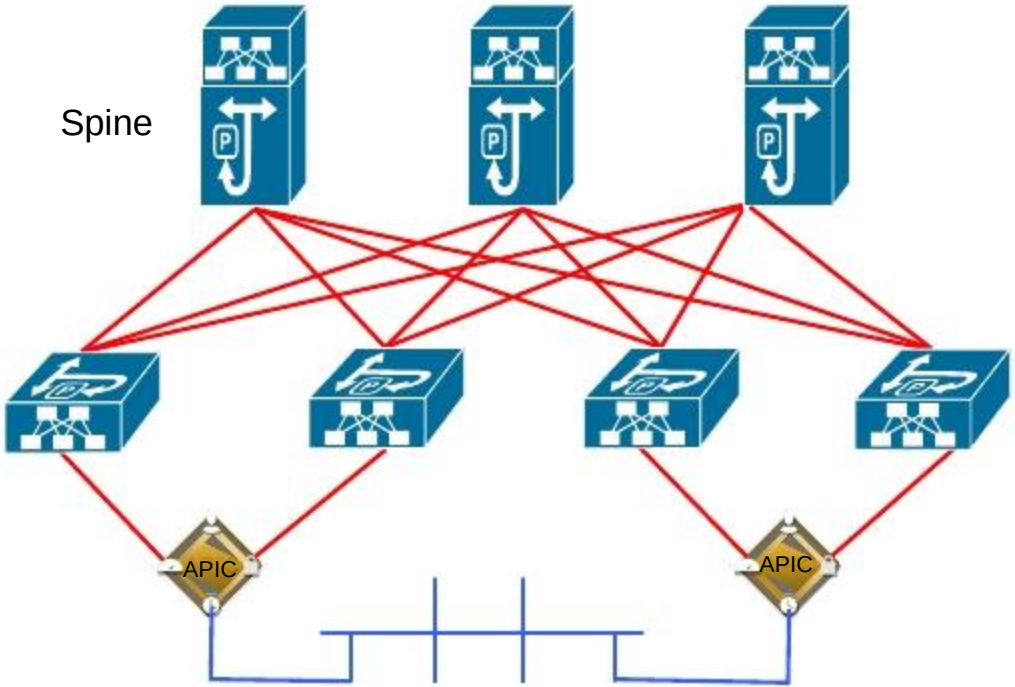

Defining Terms

Tenant: Logical separator for: Customer, BU, group etc. separates traffic, admin, visibility, etc.

Context: Equivalent to a VRF, separates routing instances, can be used as an admin separation

End-Point Group (EPG): Container for objects requiring the same policy treatment, i.e. app tiers, or services.

Bridge Domain: Not a VLAN, simply a container for subnets. It can be used to define a L2 boundary.

Logical Model Overview

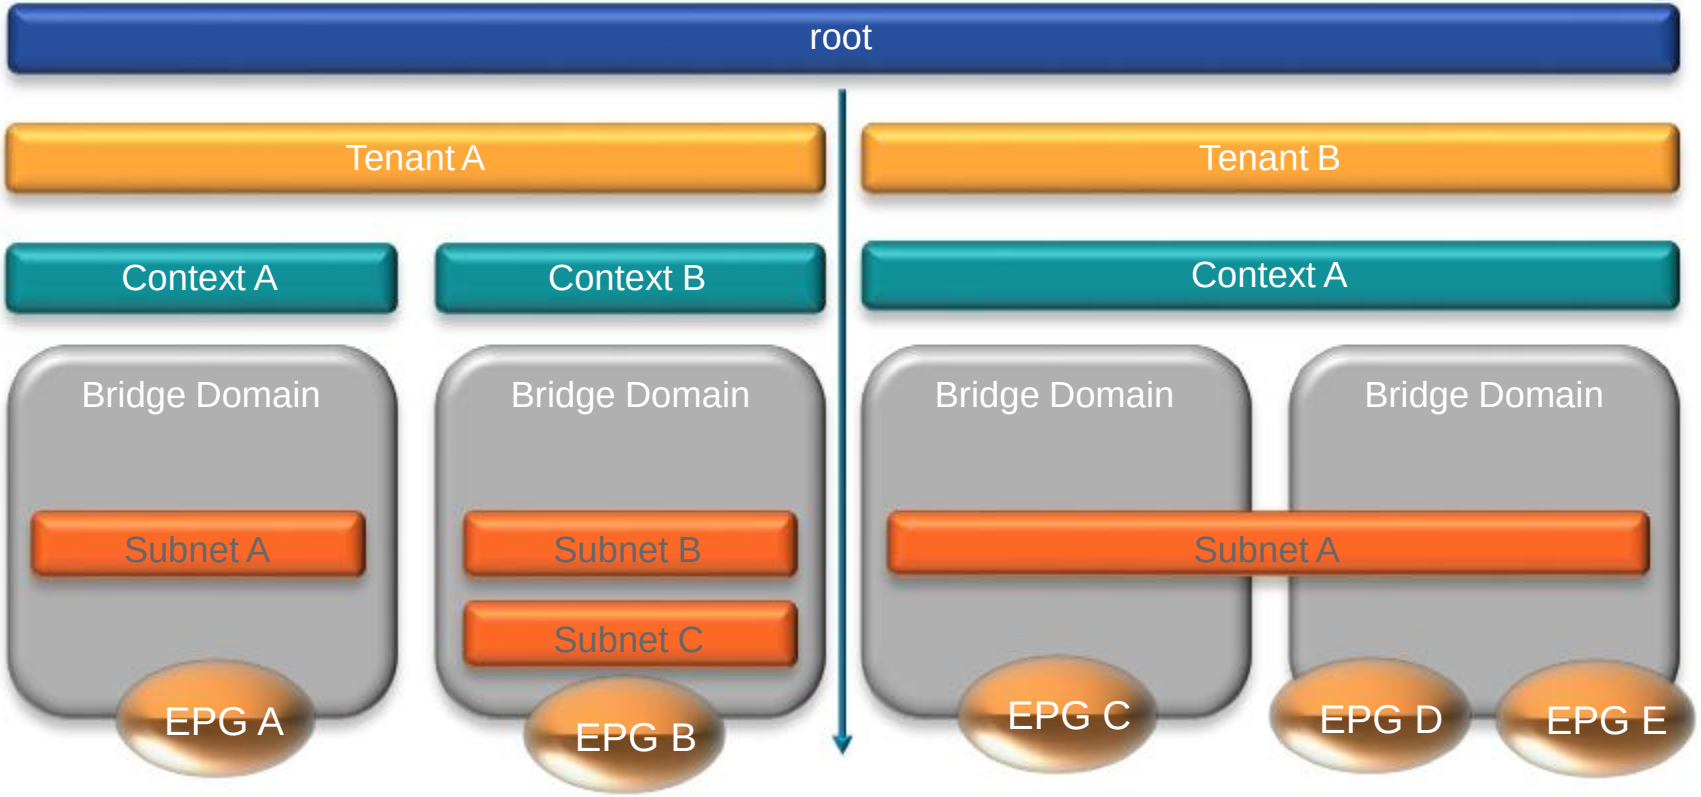

Context and subnets are independent between tenants

ACI Policy Model

High Level Overview

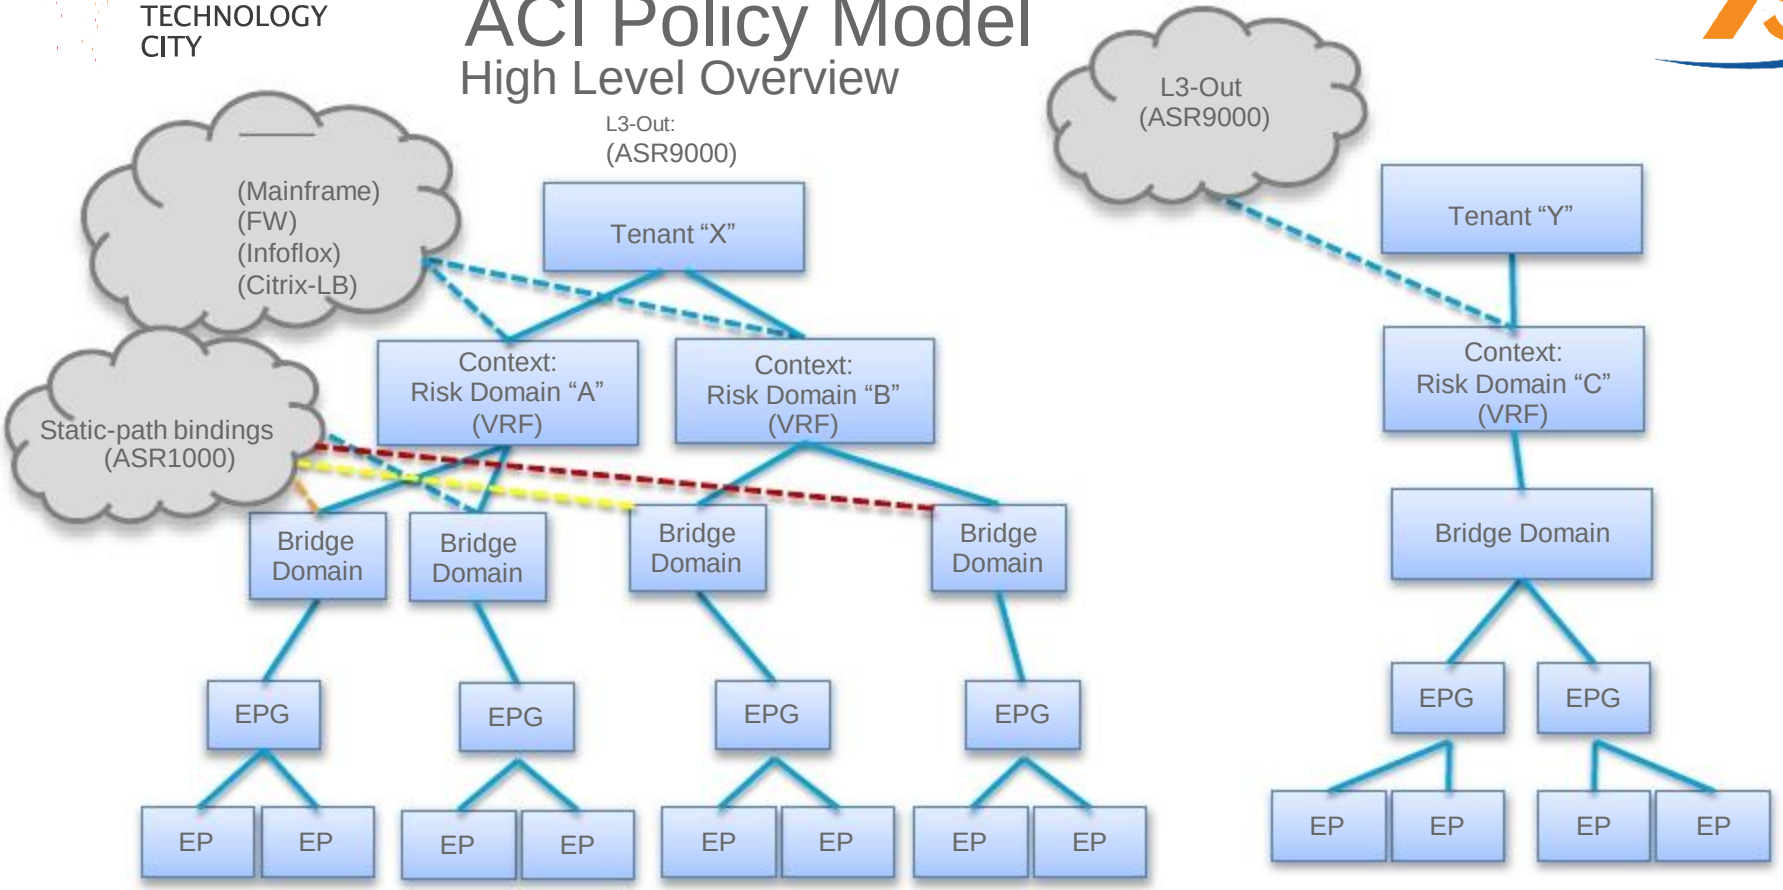

ACI Fabric

Attaching the Compute Resource to the Fabric

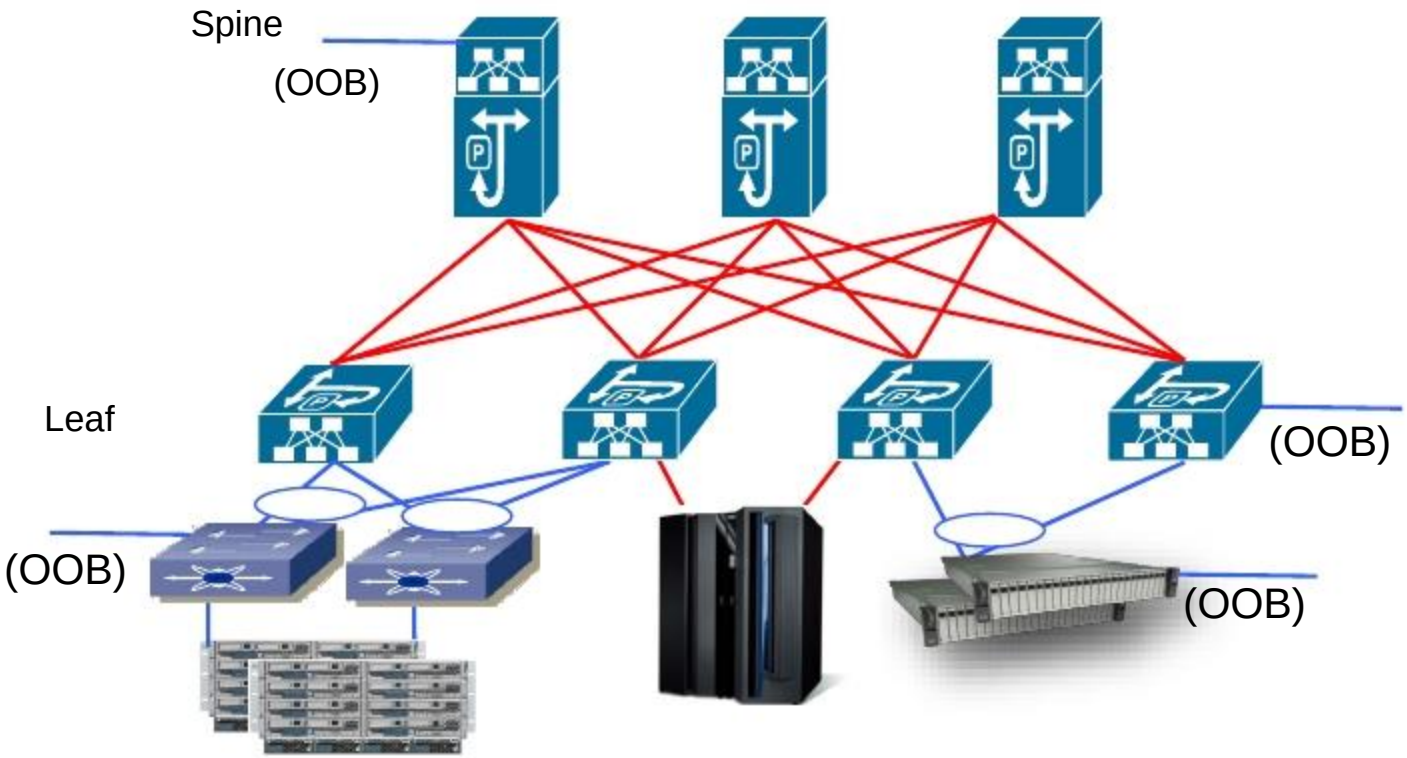

ACI Fabric

Attaching the Services to the Fabric

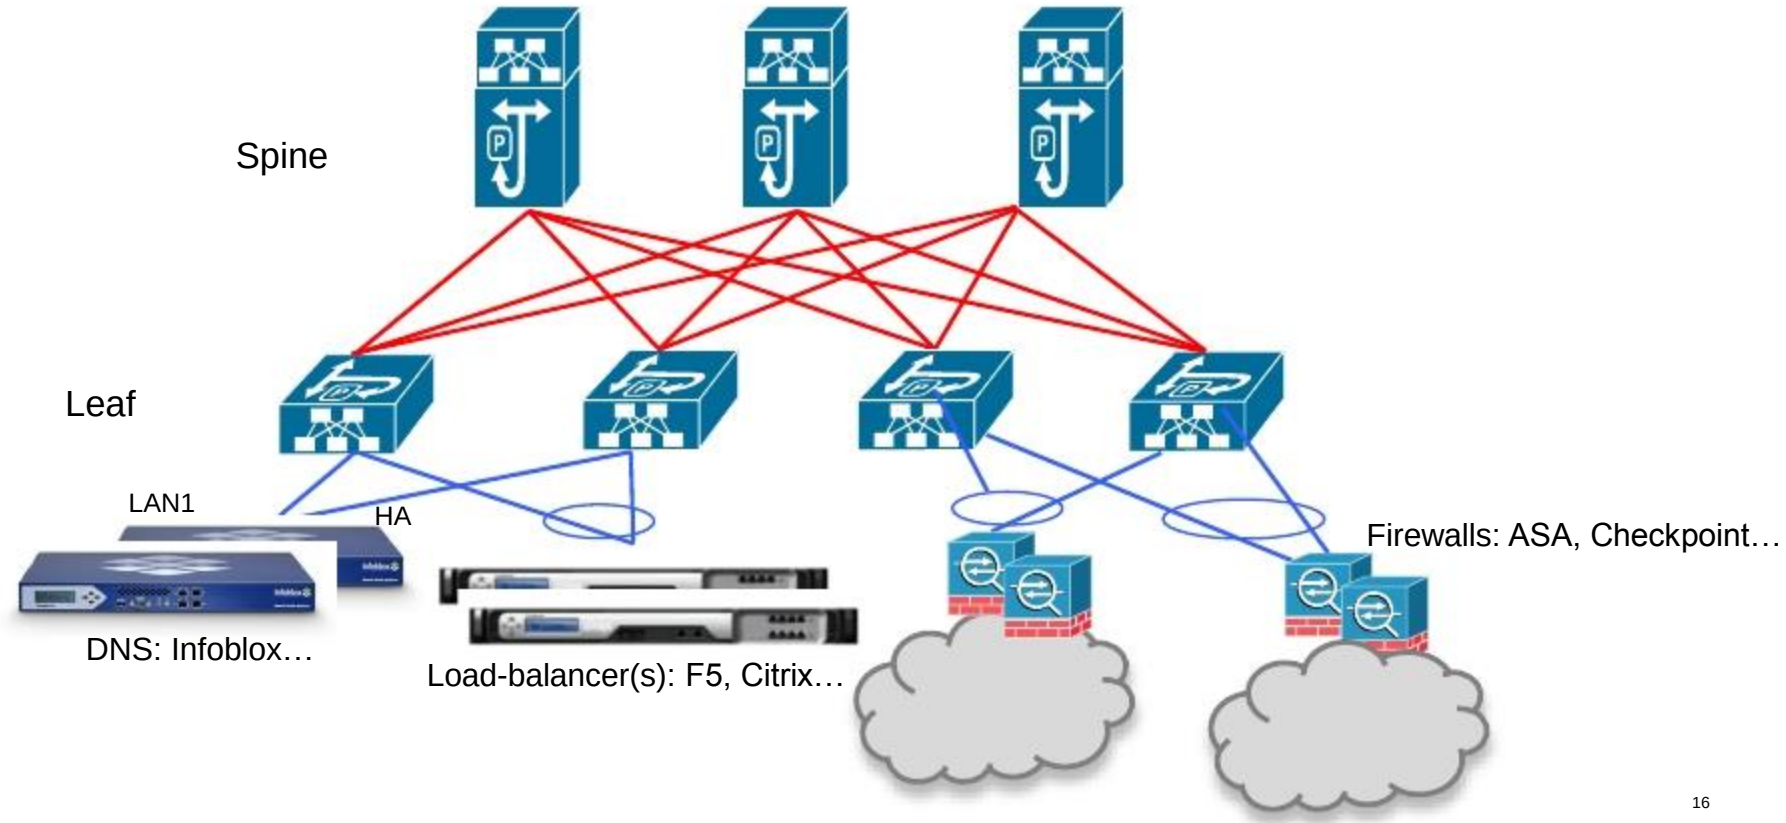

ACI Fabric

Attaching the VMM/Orchestration to the Fabric

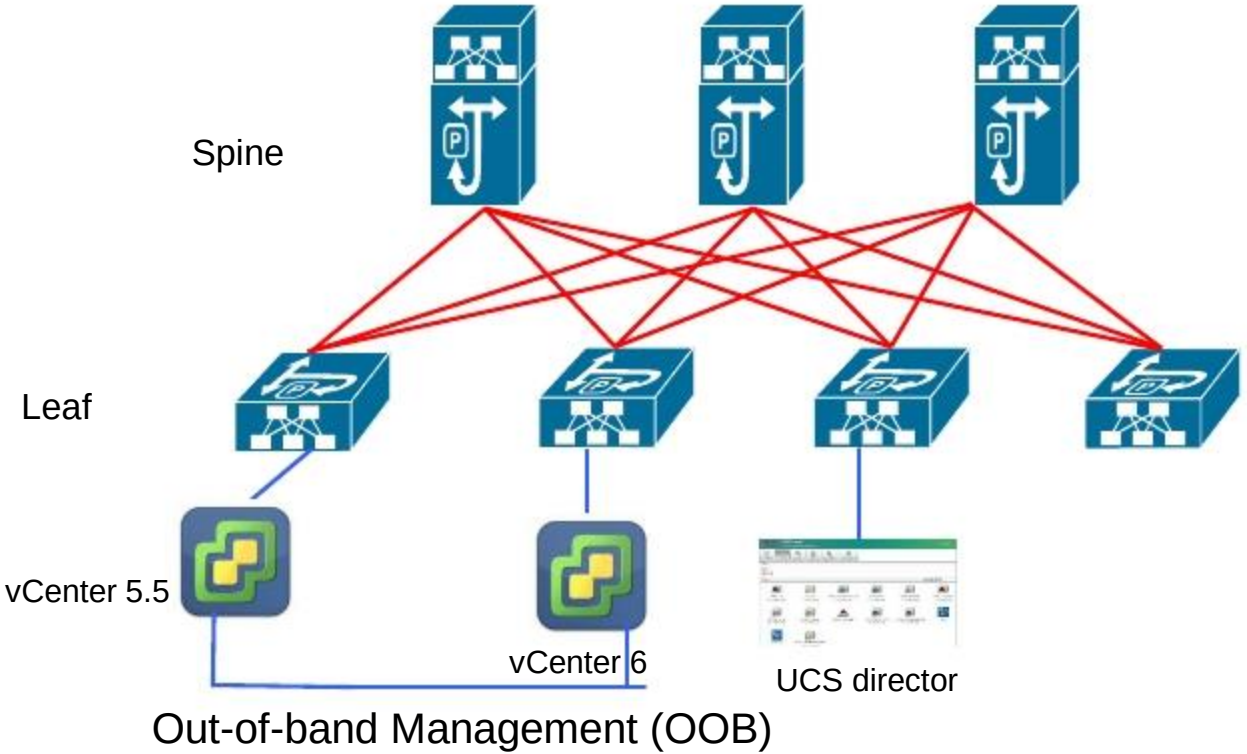

ACI Fabric

Attaching the External WAN/Enterprise to the Fabric

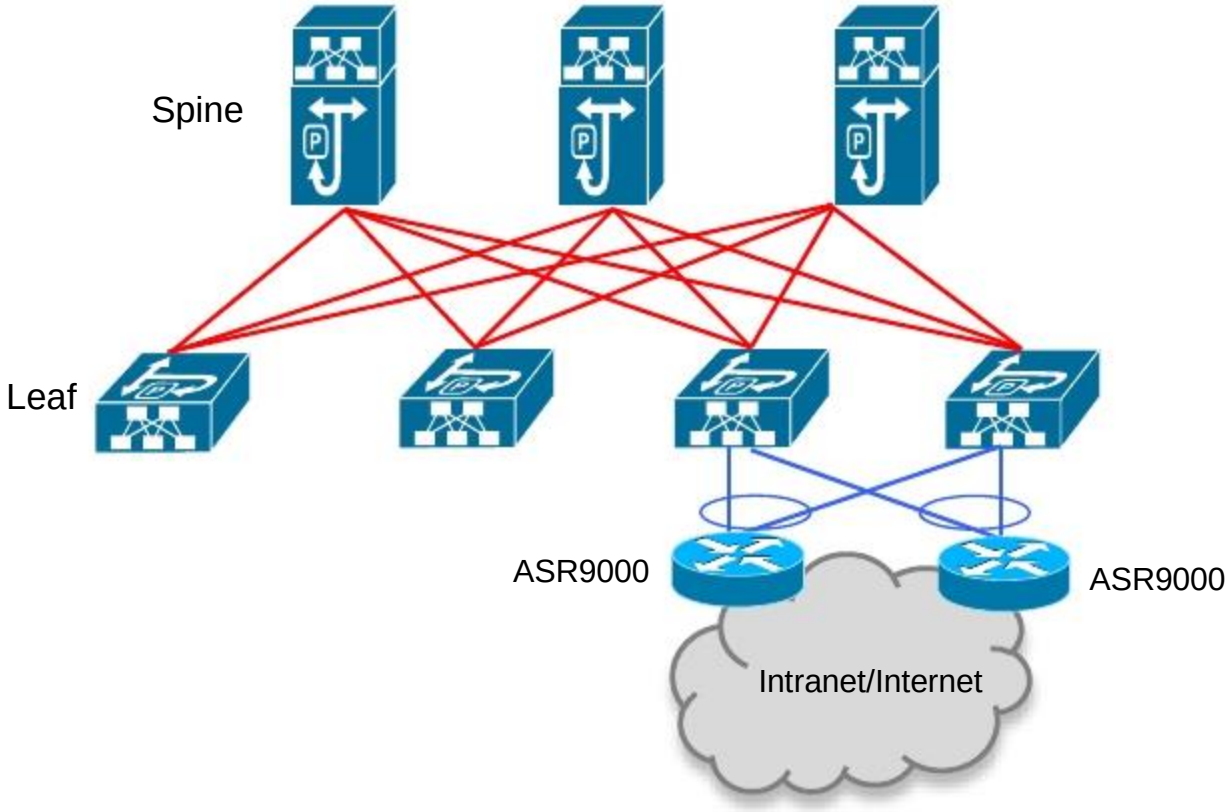

VLAN = EPG

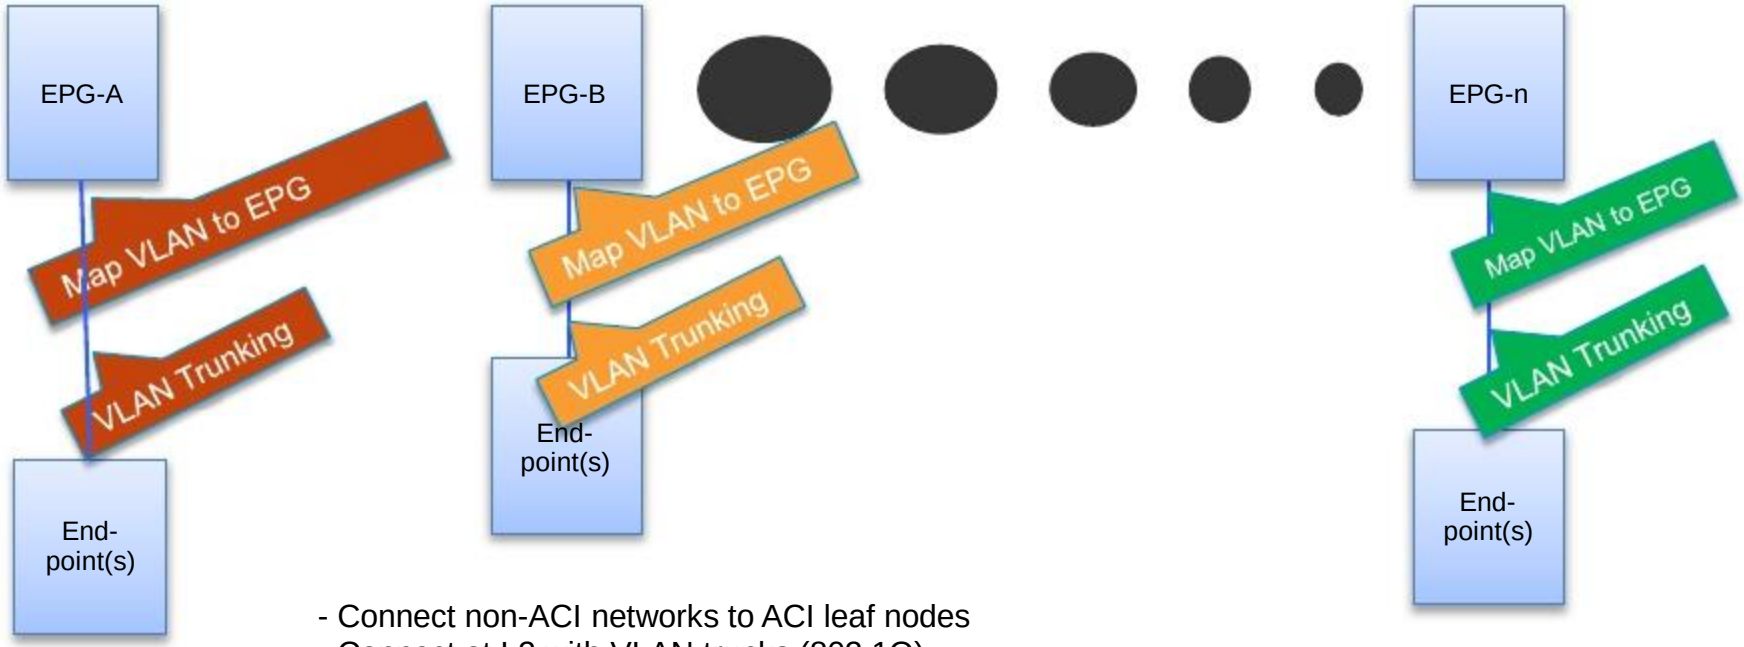

- Connect non-ACI networks to ACI leaf nodes
- Connect at L2 with VLAN trunks (802.1Q)
- Objective: Map VLANs to EPGs, extend policy model to non-ACI networks

ACI Policy Model: EPG To EPG Communication

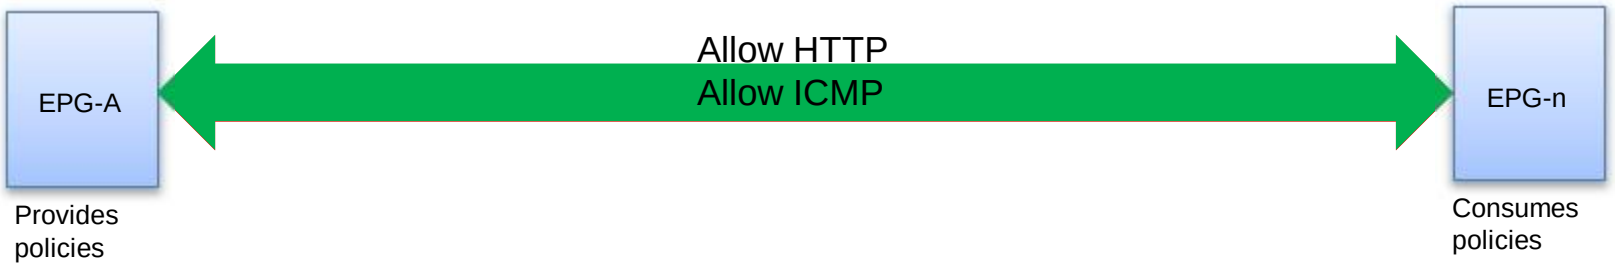

Zero Trust Security Model

- Need to define a Contract (Policy); - A contract is used to specify the interaction between two EPG(s), a provider/consumer pair.
- The goal is to provide a global policy view that focuses on improving automation and scalability.

3S TECHNOLOGY CITY

DNS/DHCP Integration: Infoblox

Infoblox Anycast (DNS/DHCP) L3-Out ACI Deployment

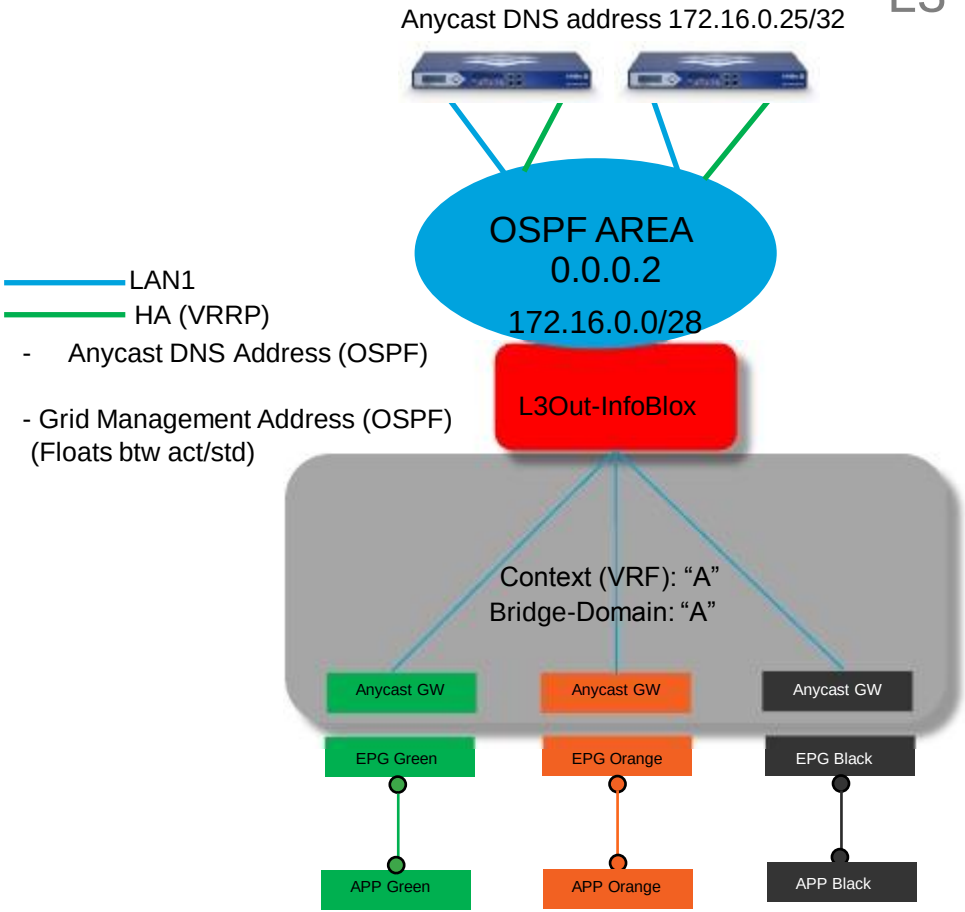

Access Interface (Untagged)

Leaf advertises default-route to the Infoblox. "External Network Instance Profile advertise 0.0.0.0/0 to Infoblox – like OSPF Stub no-summary.

Infoblox OSPF Priority = 0

OSPF Network Type: Broadcast

HA Active / Standby Anycast Management VIP

Physical: Infoblox1 LAN1/HA connects to Leaf1. Infoblox2 LAN1/HA connects to Leaf2. (2 OSPF peers)

LAN and HA interfaces all have to be in the same EPG/BD/Subnet.

Passive nodes listen to VRRP advertisements on the HA port while Active nodes listen on the LAN port.

Peering is on leaf interface, the SVI for the default gateway

Default route leak policy being used as an alternative to a pre-existing default-route. The VRF-Intra, it is being injected via the

ASR9000 (OSPF) or configure a static-route via the FW (security

Infoblox Grid Geographical Redundancy

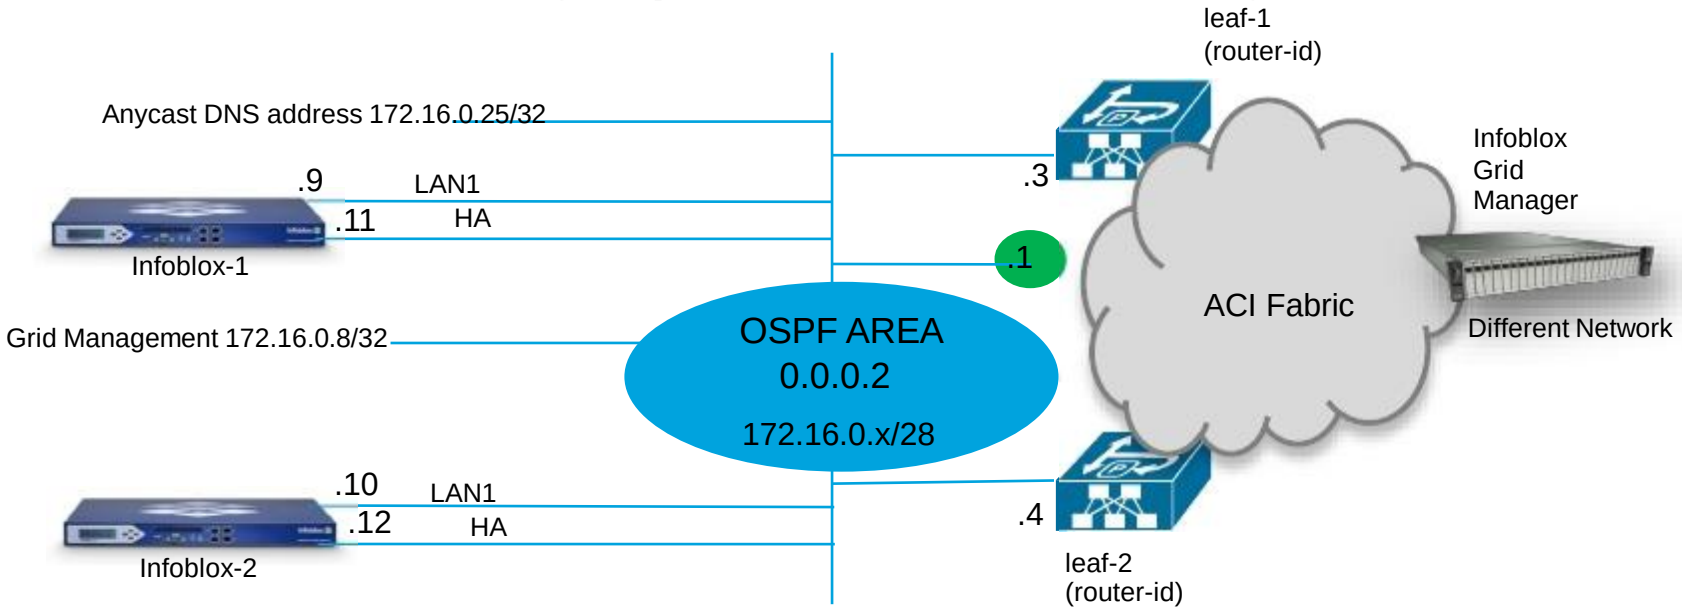

Floating IP .1 (SVI); this is the default gateway for the Infoblox Grid management.

L3-Outside Configuration: OSPF

1) Configure L3Out for OSPF

2) Select Context / VRF

3) Define OSPF Area, in this case OSPF Area 0.0.0.2

4) Define OSPF Area type, in this case regular OSPF Area

5) The external routed domain, policy for managing the physical infrastructure, such as ports/VLANS, that can be used by an L3 routed outside network.

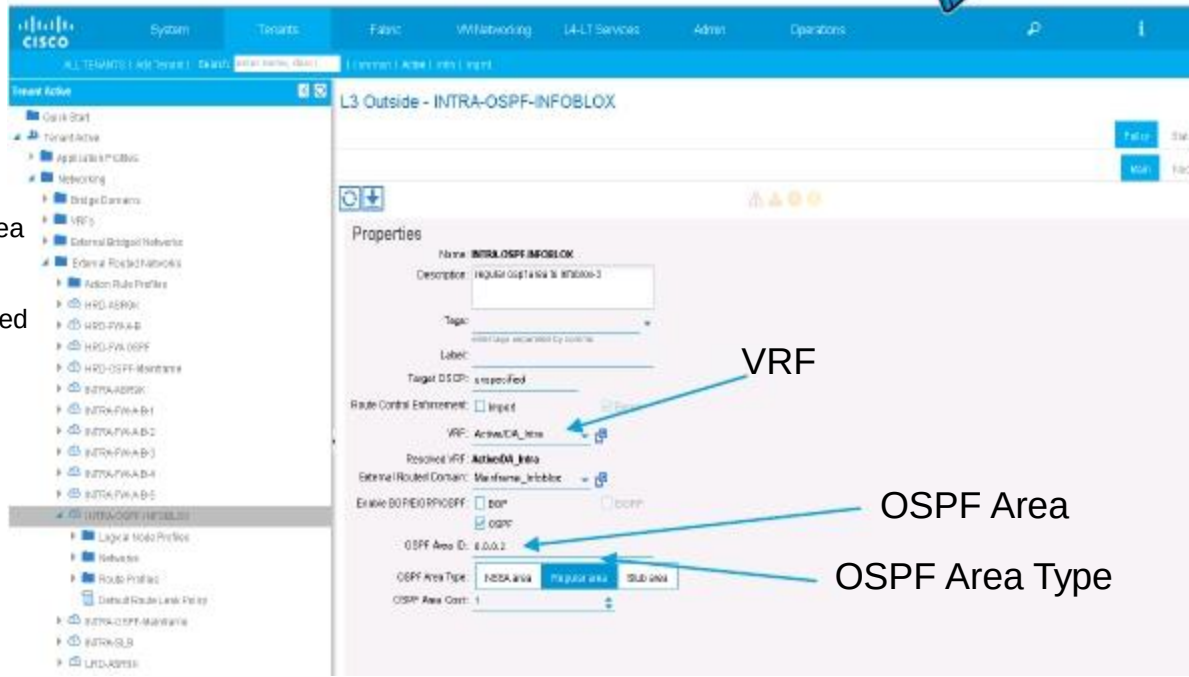

3S TECHNOLOGY CITY

Mainframe OSPF Integration

Mainframe: L3-Out ACI Deployment

Mainframe L3-out is a regular OSPF Area.

Defined external network instance for Export Route Control Subnet for 0.0.0.0/0 (make sure un-check

"Aggregate Export“)

Trying to “treat” as OSPF Stub Area.

Type 5 LSA(s); leaf(s) are OSPF ASBR(s)

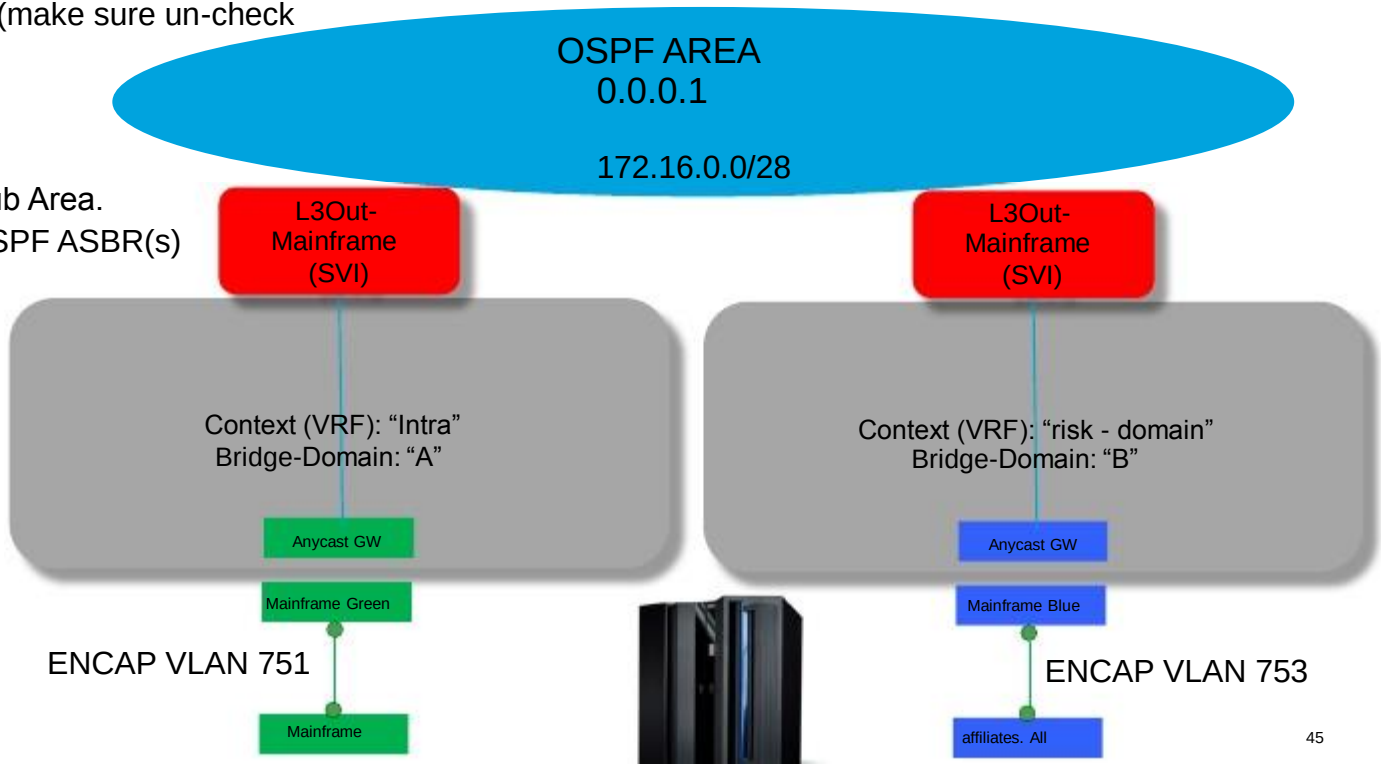

3S TECHNOLOGY CITY

ACI Configuration: External Networks

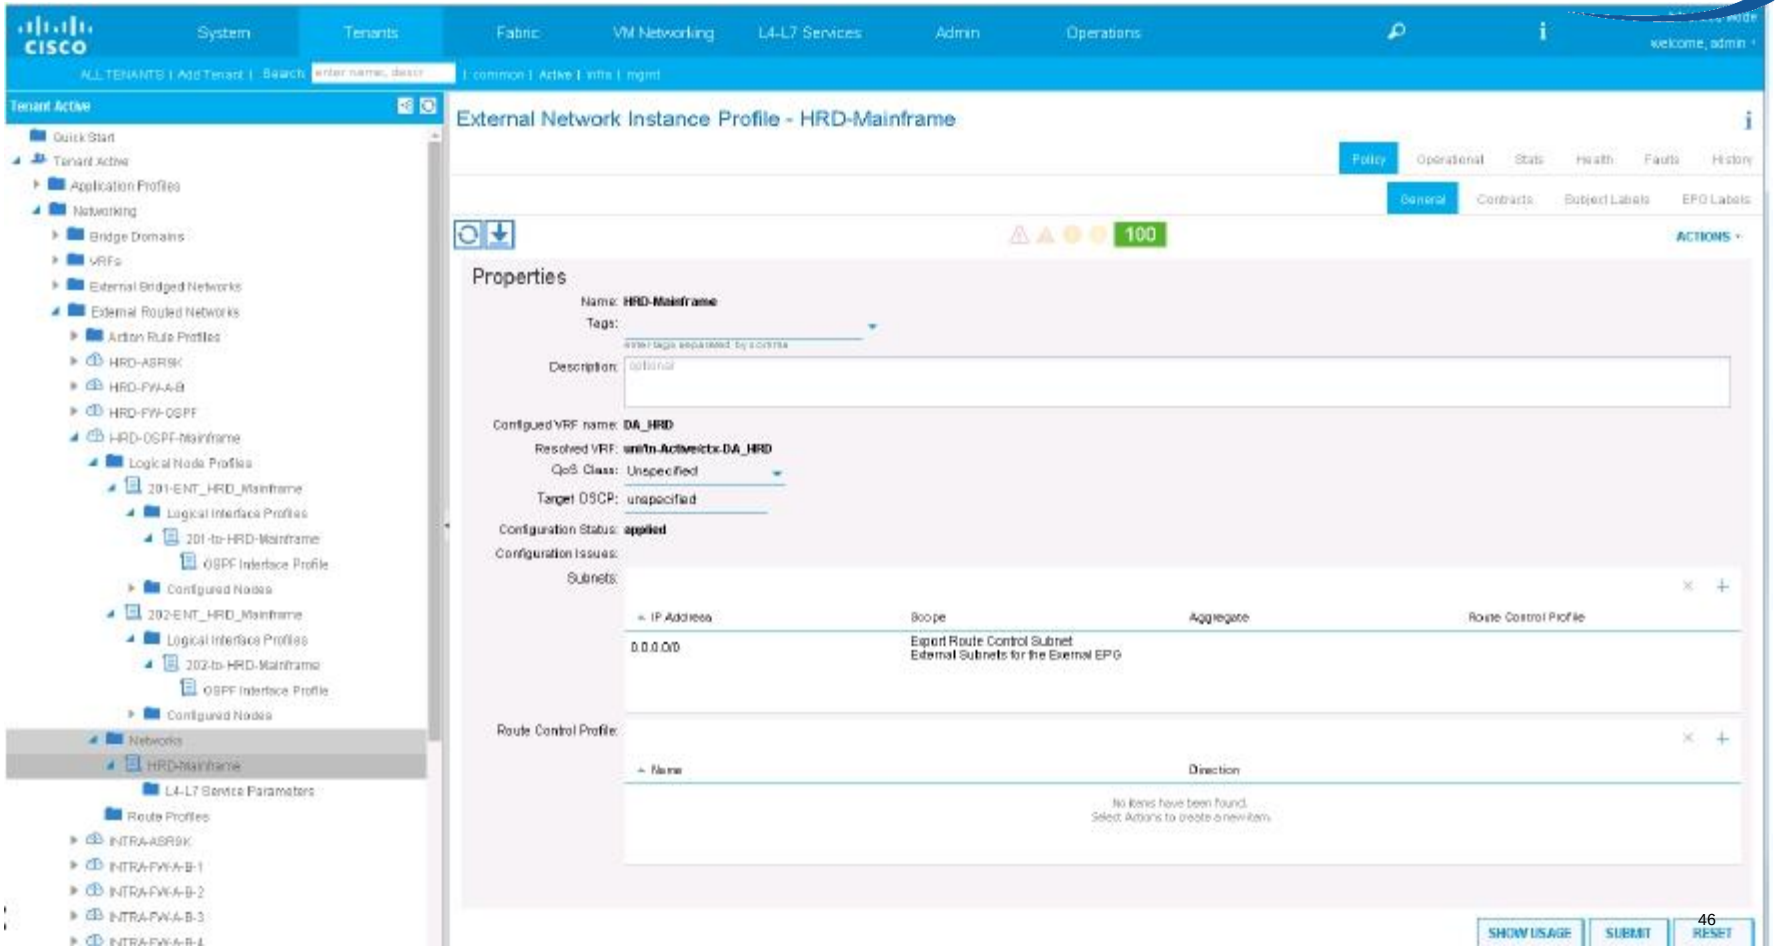

3S TECHNOLOGY CITY

Load-balancers Integration: Citrix/F5

Citrix 2-arm Load-balancer: Static-Bindings

External-arm (VLAN) for the VIP / Client

Static route for LB servers pointing to VIP

VLAN 10 SVI on L3out

VIP: 20.20.20.20/32

1) External-arm: VIP / Client
2) Internal-arm: server default-gateway is on the load-balancer.

VLAN 400 (Bridge-domain same for Servers)
192.168.50.100

Server(s) default-gateway

Static route to Servers

Internal-arm (VLAN) is the Server default-gateway on the load-balancer

ACI: Configuring the Server-side bridge-domain

Enabled Flooding (ARP) as this L2 Only

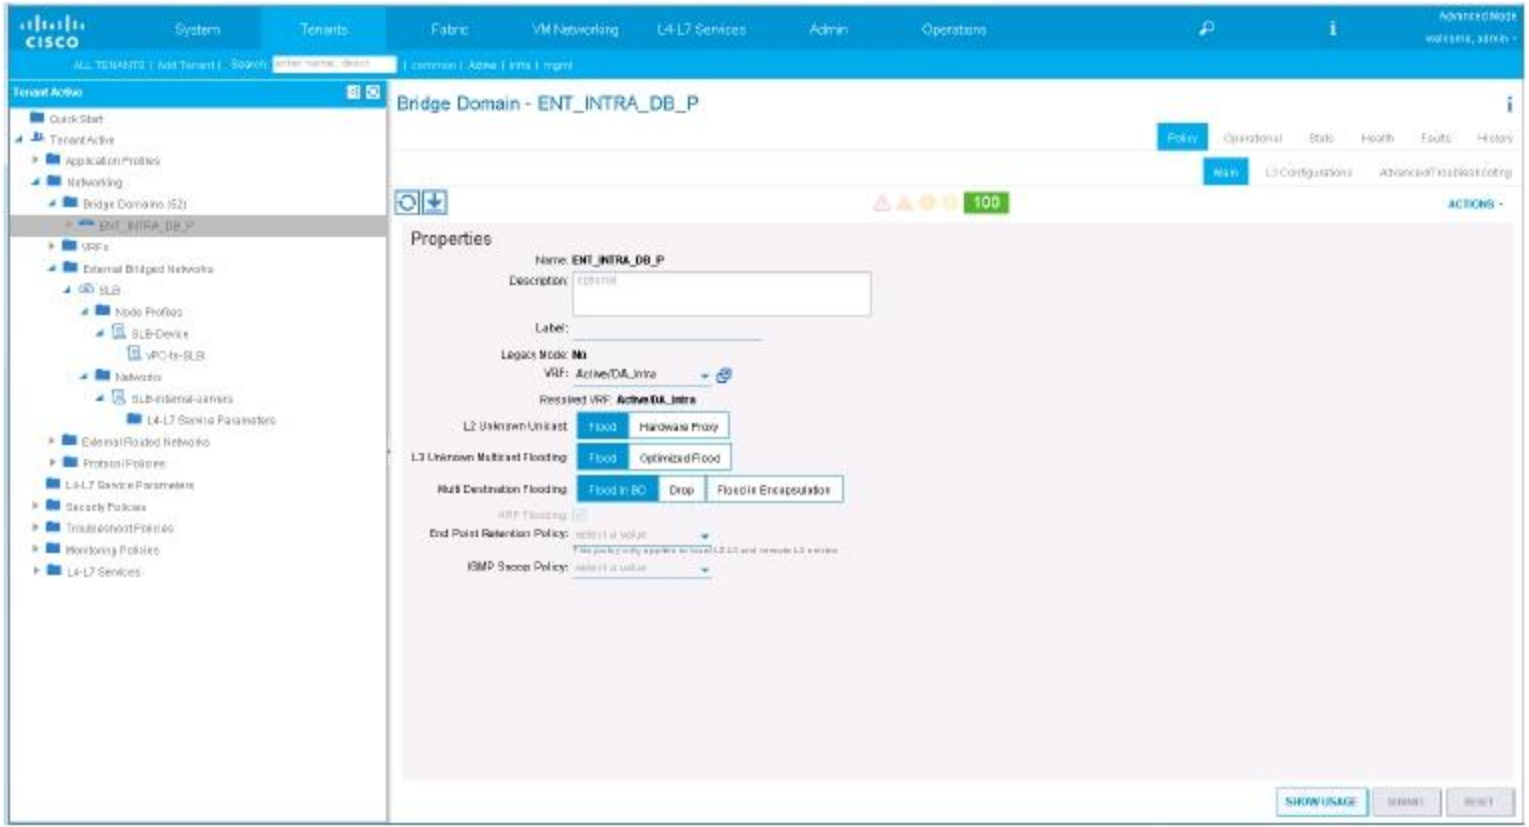

3S TECHNOLOGY CITY

External Connectivity

ACI Interaction with STP

- No STP running within ACI fabric
- BPDU frames are flooded within EPG. No Configuration required
- External switches break any potential loop upon receiving the flooded BPDU frame fabric
- BPDU filter and BPDU guard can be enabled with interface policy

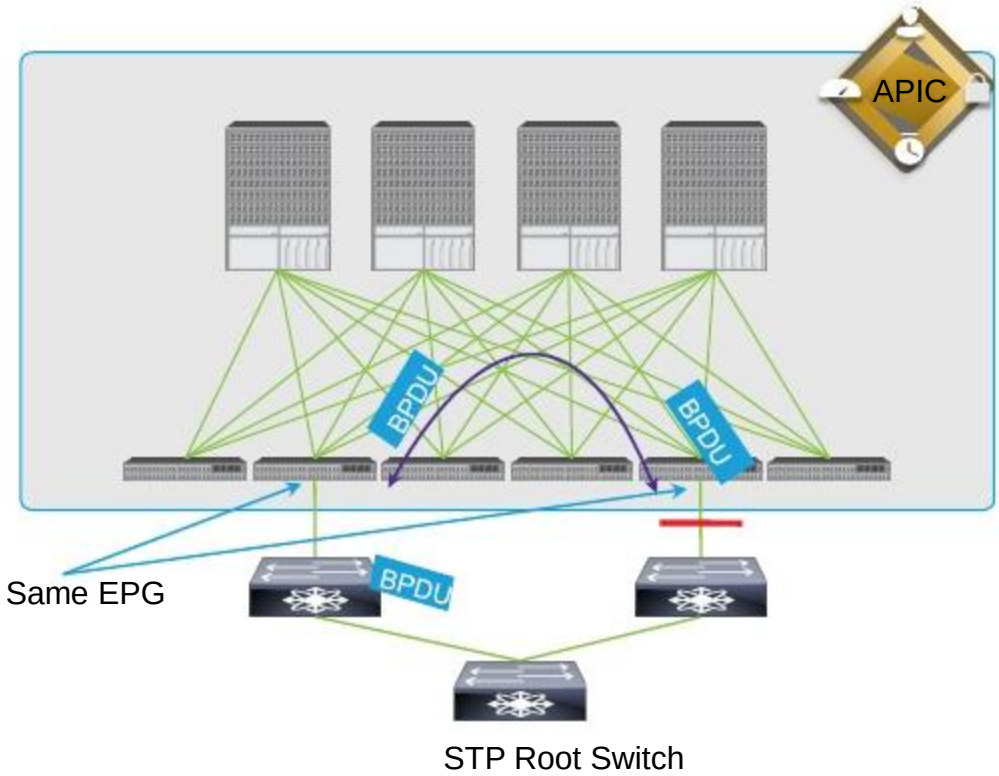

ASR9000 External L3out OSPF via SVI and vPC

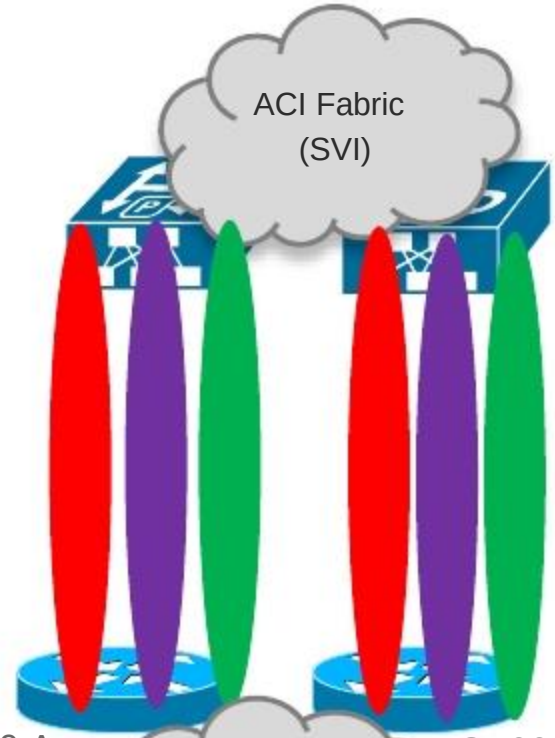

VRF: risk-domain
VLAN 902
172.18.159.64/29
OSPF Area 0

VRF: Intra
VLAN 900
172.18.0.64/29
OSPF Area 0

VRF: risk-domain
VLAN 904
172.18.181.0/29
OSPF Area 0

VRF: risk-domain
VLAN 903
172.18.159.72/29
OSPF Area 0

VRF: Intra
VLAN 901
172.18.0.72/29
OSPF Area 0

VRF: risk-domain
VLAN 905
172.18.181.8/29
OSPF Area 0

3S TECHNOLOGY CITY

Firewall Integration : ASA / Checkpoint

Extranet: Routing Between Contexts

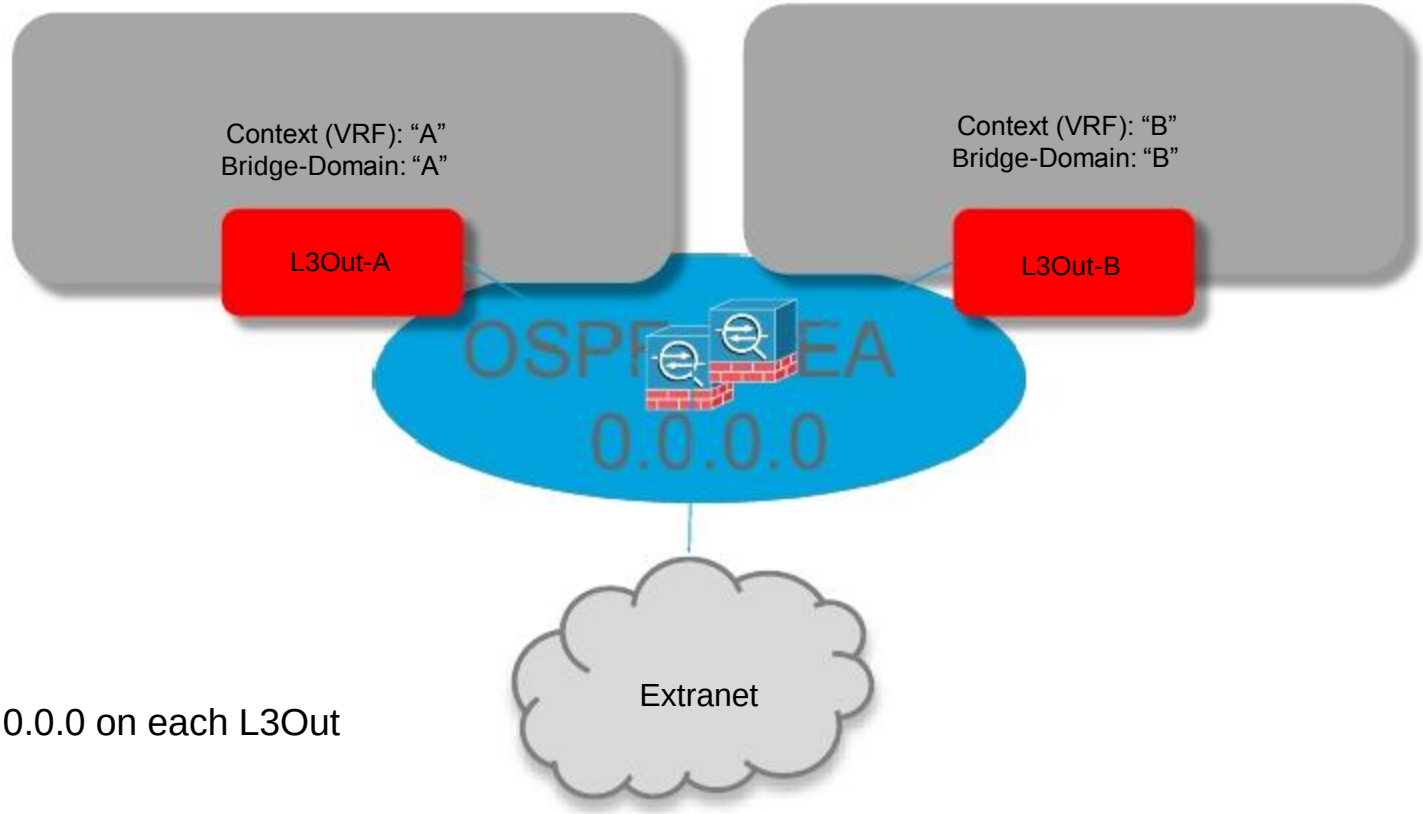

L3Out OSPF Area 0.0.0.0 on each L3Out

Local-Internet: Logical view

1) Intra-VRF default routes from ASR9k to Fabric to Internet Only
2) Other VRF(s) will have default-route point to Firewall and Firewall will route to Intranet; based on FW policy

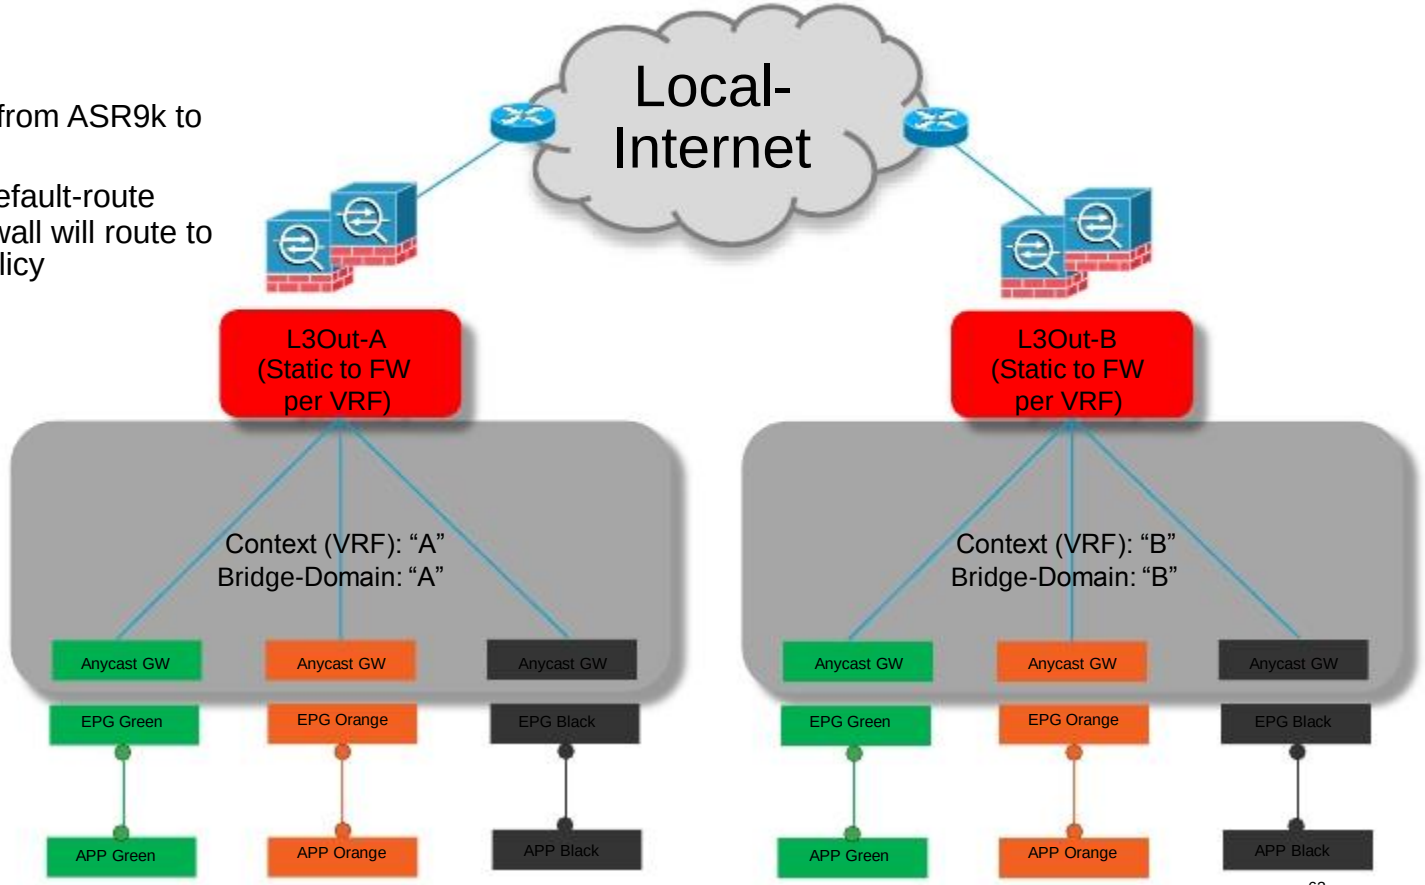

Intra-VRF and Intera-VRF Traffic Flows

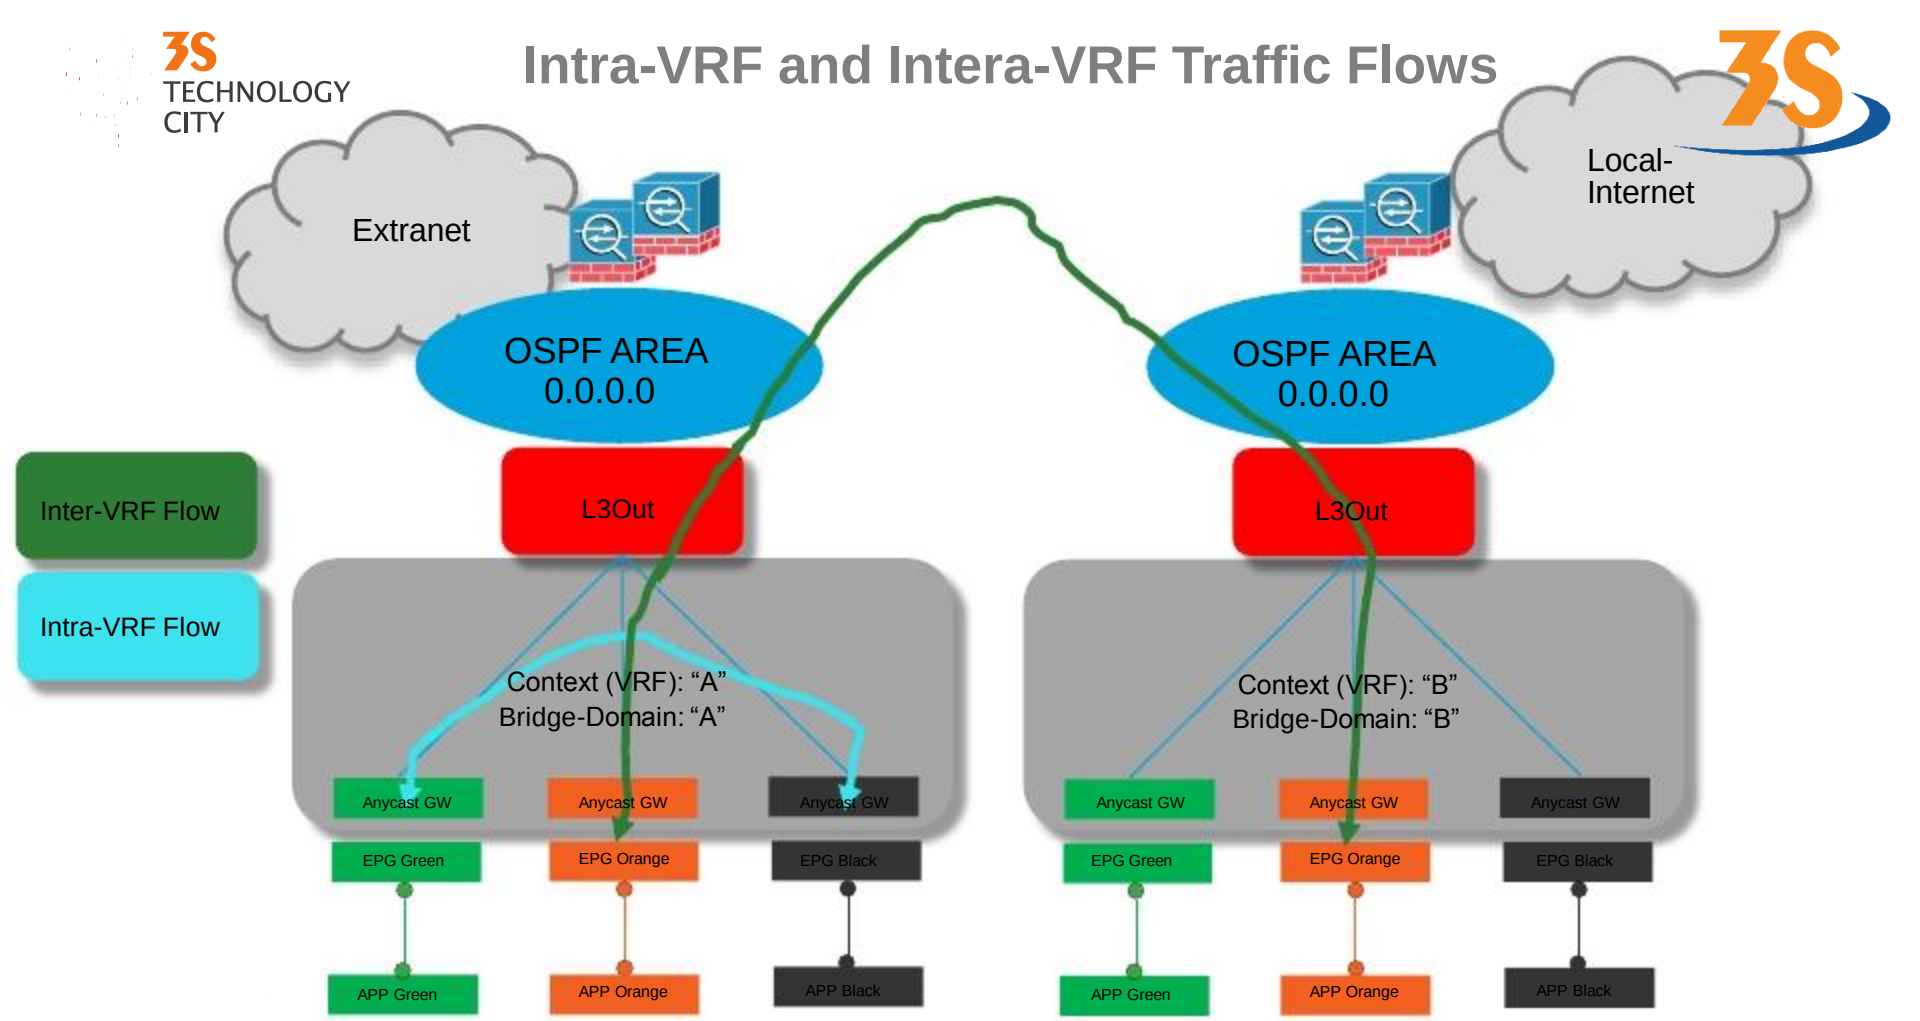

ACI and VMM vCenter Integration

- Cisco APIC integrates with the VMware vCenter.
- Ability to transparently extend the Cisco ACI policy framework to VMware vSphere workloads.
- APIC uses Application Network Profiles (ANPs) to represent the Cisco ACI policy.
- APIC creates a virtual distributed switch (VDS) in VMware vCenter for virtual networking.
- APIC manages all application infrastructure components. The network administrator creates EPGs and pushes them to VMware vCenter as port groups on the DVS.
- Server administrators can then associate the virtual machines and provision them accordingly.

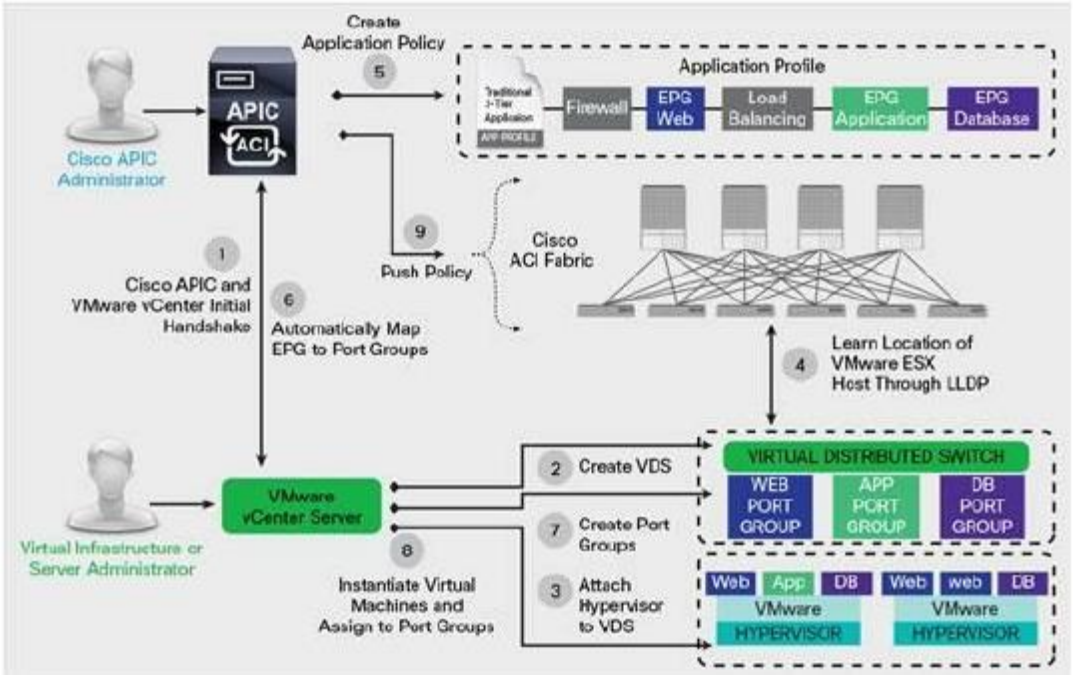

ACI and VMM vCenter Integration

- Show configured VMware VMM vCenter
- Focusing on vCenter 6 instances

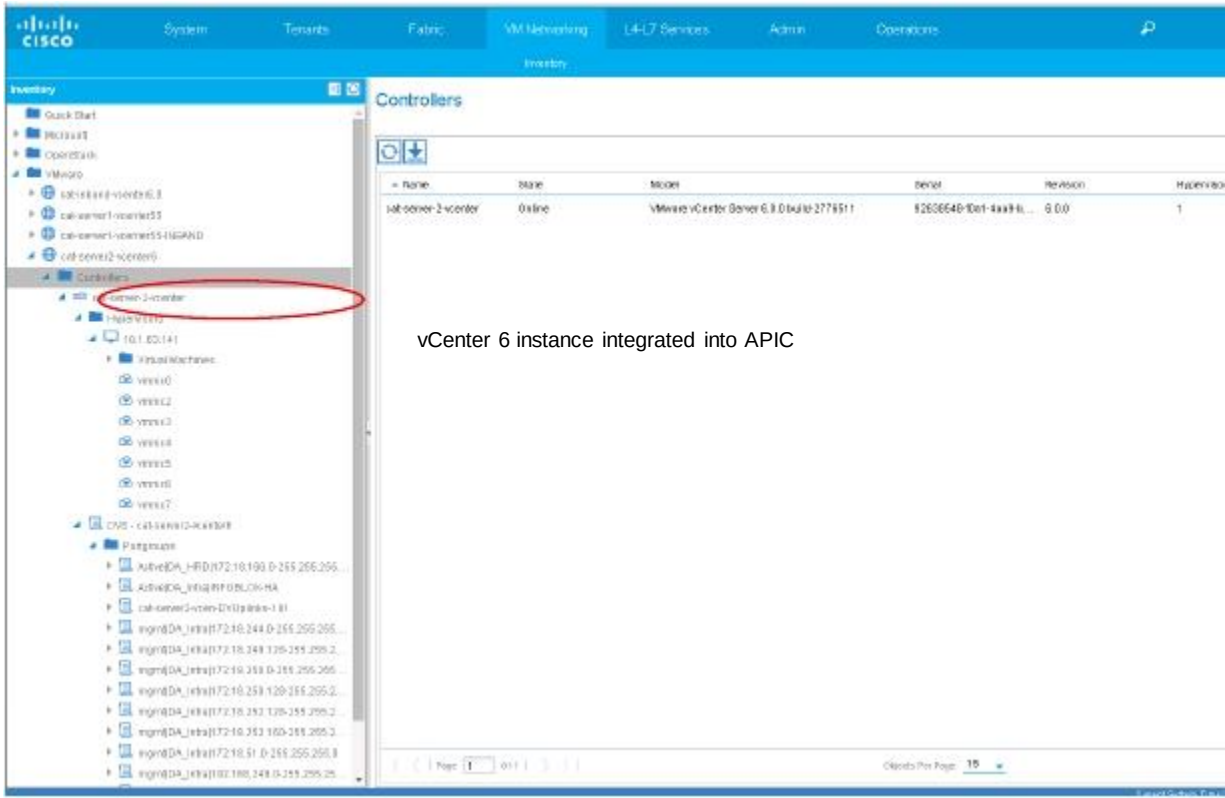

UCS Director workflows

- Provision new server
- Decommission server
- ACI - Create Context
- ACI - Create Bridge Domain
- ACI - Create EPG
- ACI - Create Application Profile
- ACI - Create Contract
- ACI - Assign EPG to PortChannel/Alias
- ACI - Unassign EPG from PortChannel/Alias
- ACI Combined Provisioning Workflow
- ACI Combined De-provisioning Workflow
- Create a data LUN (array based on 'class') for presentation via VPLEX
- Expand LUN and volume
- Remove LUN and volume
- Present virtual volume to a host
- Present virtual volume to a RP cluster

3S TECHNOLOGY CITY

Driving Innovation

The Path to Agility in an App-Centric World

3S TECHNOLOGY CITY

Policy-Driven Integrated Infrastructure Answers Customers' Request

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

Policy Driven Automation for a Cloud Model

Foundational Challenges

ACI Solves

Simplification ✓
Drive Agility & Automation

Optimize Operations > TCO ✓
Open & Programmable

Security to Meet Compliance ✓
Segment with Multi-tenancy

Elastic Model ✓
Develop Hybrid Cloud Model

Enhancing Bi-Modal IT with ACI and UCS

Tenant 1 = Mode 1

Manufacturing

IT Department

Recently Acquired Company

Marketing Department

Tenant 2 = Mode 2

Pivotal

openstack CLOUD SOFTWARE

apprenda

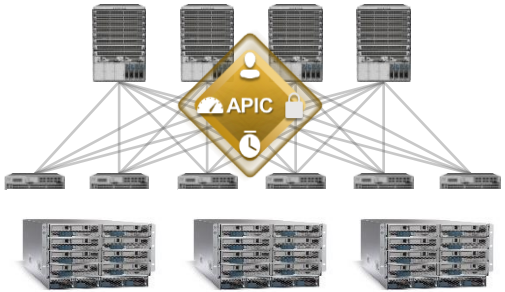

Common Infrastructure, and Management With Secure, Stable Separation

3S TECHNOLOGY CITY

Enhancing Bi-Modal IT with ACI and CliQr

ACI, UCS, UCSD = Mode 1

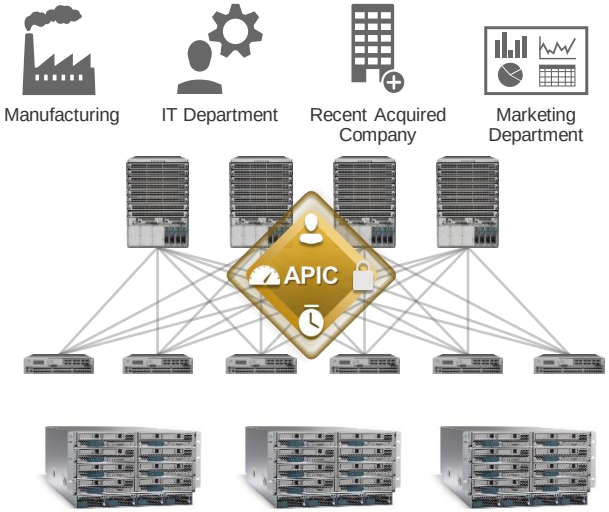

CliQr

Public Cloud = Mode 2

amazon web services™

Microsoft Azure

SOFTLAYER an IBM Company

Google Cloud Platform

Stability With on Premises Model 1. Agility with Public Model 2.

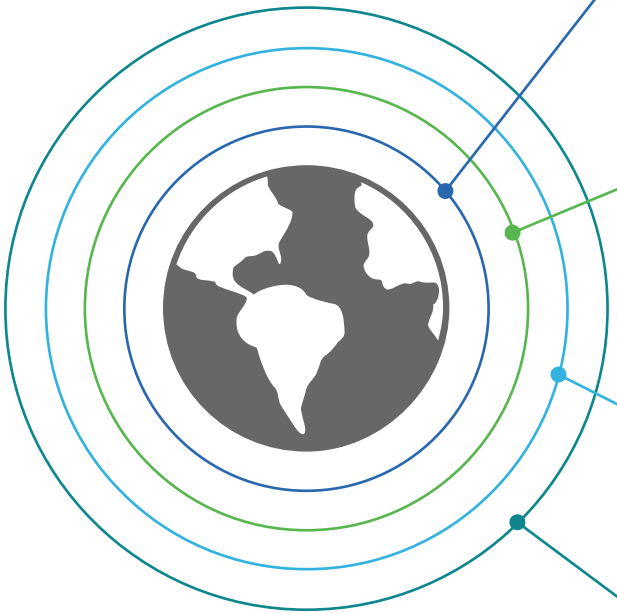

SW Overlay Apps

cisco · PLUMgrid · midokura · vmware NSX · nuagenetworks

Hypervisor and/or Container

Microsoft · vmware · docker · redhat.

Bare Metal

SAP BusinessObjects · Microsoft Exchange · ORACLE · hadoop

Orchestration/Automation

cisco · openstack · CliQr

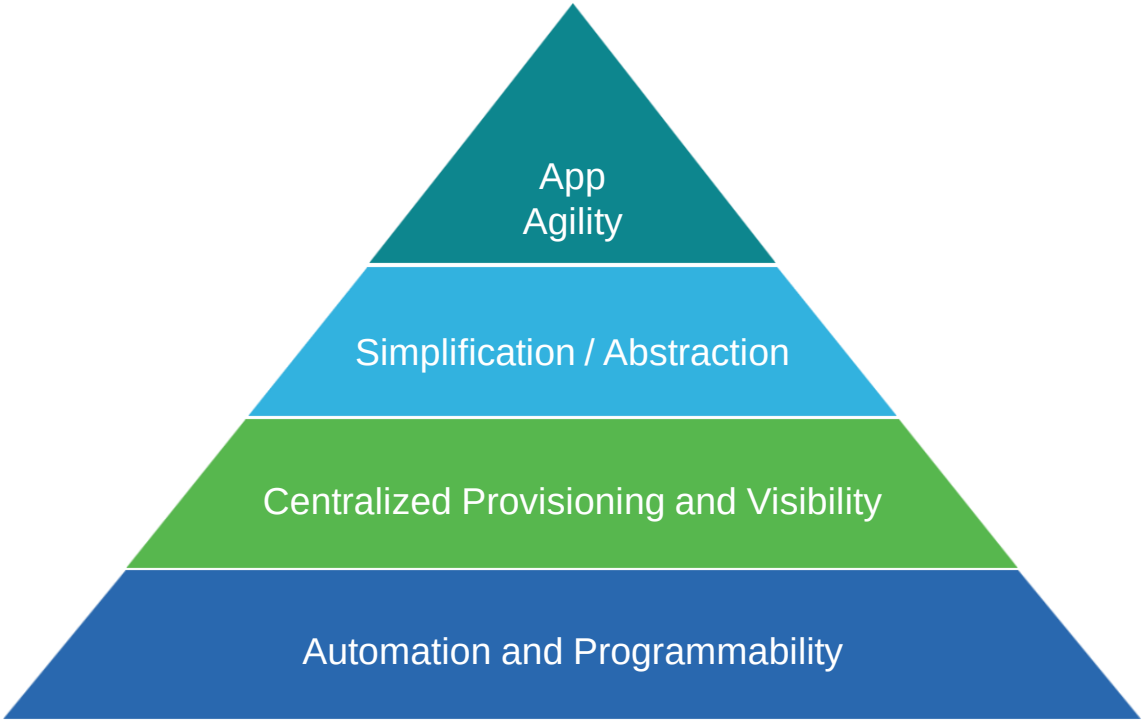
App
Agility
Simplification / Abstraction
Centralized Provisioning and Visibility
Automation and Programmability


Traditional
3-Tier
Application
APPLICATION
NETWORK PROFILE
F/W
ADC
WEB
ACC
APP
DB


ACI

Different Teams
Different Languages

Application Language

Security – SLA – Dependency – Performance – Compliance – Tenants – Geo-dependency

Network Language

VLAN – IP Addressing – Subnet – Firewalls – QoS – ACL – Load Balancer

3S TECHNOLOGY CITY

Our Vision for ACI: Scale, Security and Full Visibility

Tenant

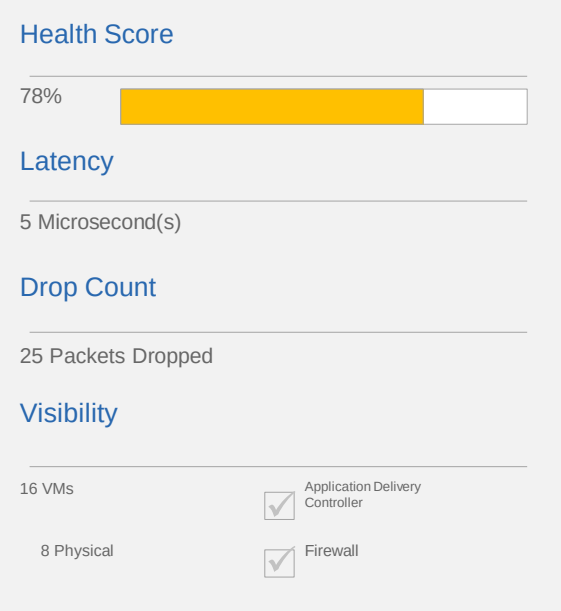

Application

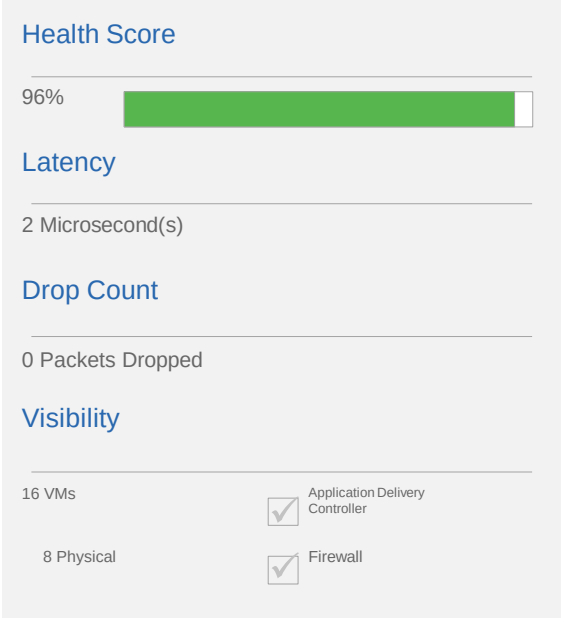

Enabled By Physical and Virtual Integration

3S TECHNOLOGY CITY

Implement Granular Security Groups

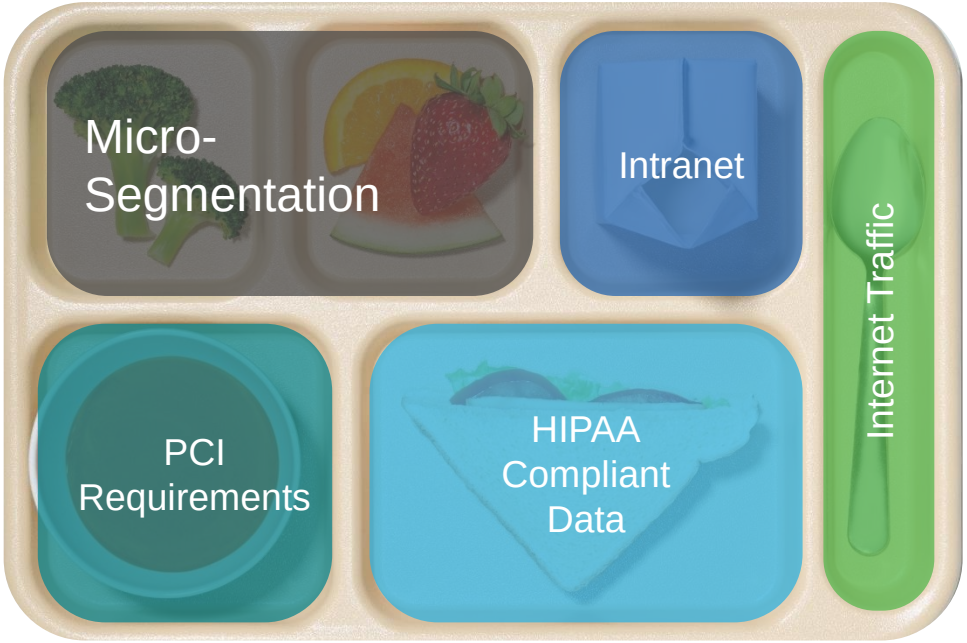

3S TECHNOLOGY CITY

ACI Security

With Focus on Simplification, Multi-Tenancy and Scalability

Integrated in ACI

Policy–
Physical and
Multi-Cloud

Automated
Audit, Detect,
Mitigate

Stateless Firewall
and Micro
Segmentation*

* State-full with Cisco
AVS

Investment Protection

Symantec. FORTINET radware intel Security Check Point SOFTWARE TECHNOLOGIES LTD. SOURCEfire A10 Networks CISCO CITRIX Infoblox CONTROL YOUR NETWORK AVI Networks f5

Validated for Deployment in PCI Compliant Networks

ACI Visibility
Across Physical, Virtual and Containers

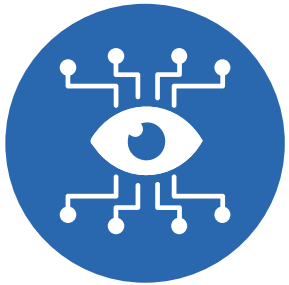

Integrated
Overlay/Underlay

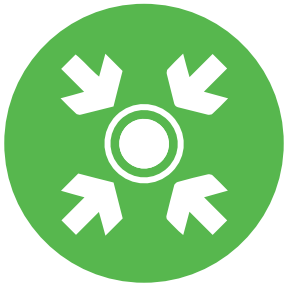

Centralized
Management
and Open APIs

Health
Scores

3S TECHNOLOGY CITY

Hybrid Cloud Orchestration

3S TECHNOLOGY CITY

CliQr CloudCenter

Any App, Any Cloud, One Platform

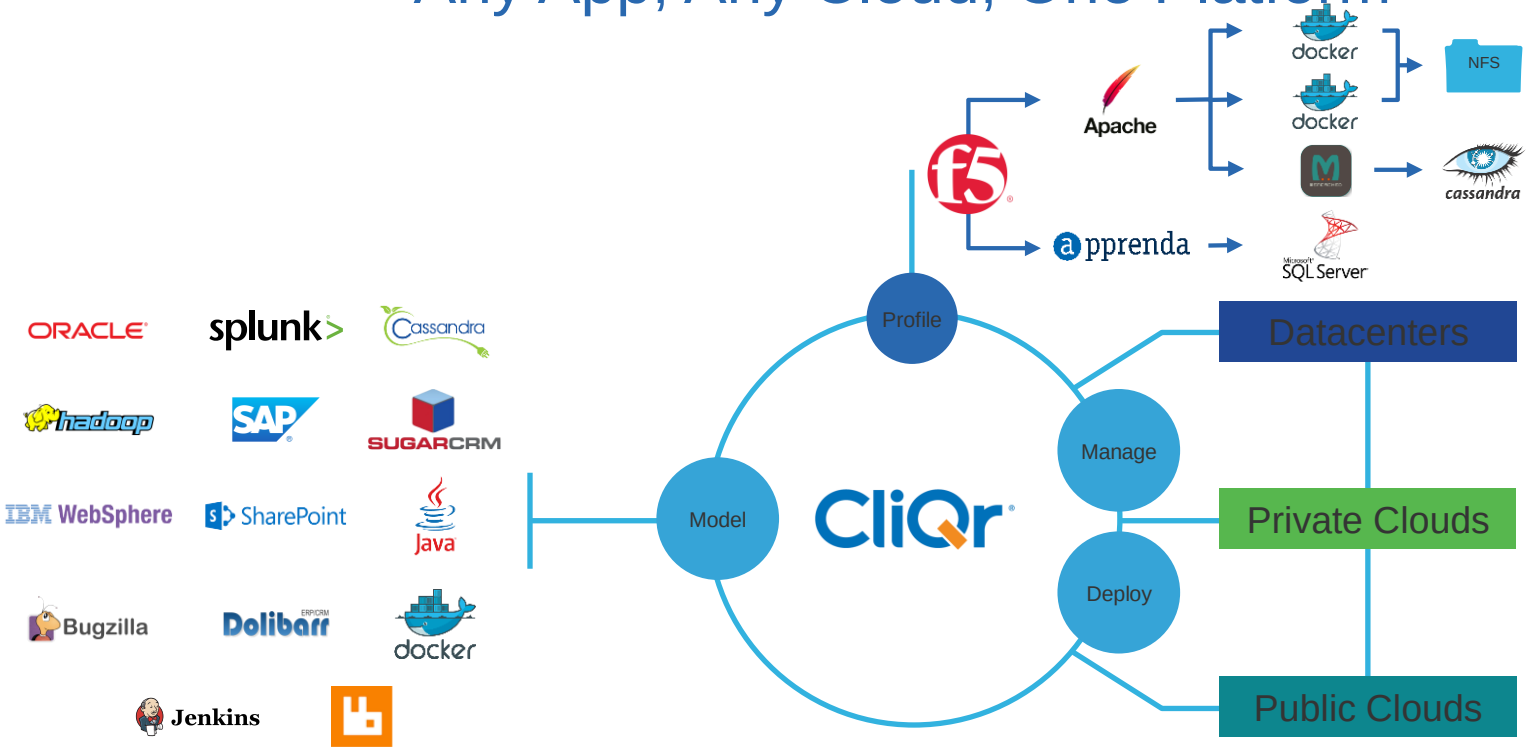

Working Together
End-to-End Orchestration

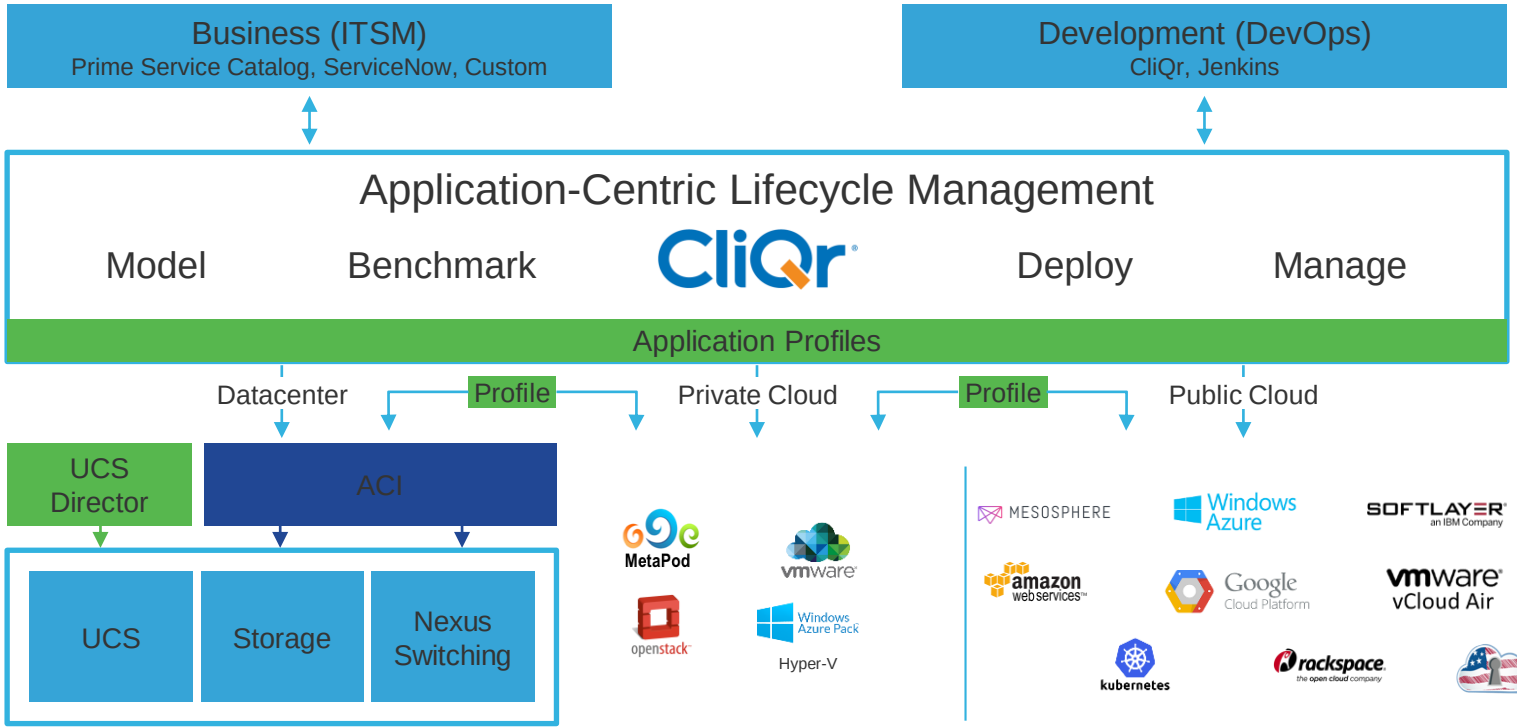

3S TECHNOLOGY CITY

Application Driven Datacenter & Cisco Virtual Topology System

3S TECHNOLOGY CITY

Application Driven Data Center

Today's App

Scale: Connectivity, Tenancy, Secure Access

+

Integration: APIs, IOT, M2M, Cloud

+

Data: Volume, Streaming & Predictive Analytics

+

Agile: CI/CD, Devops, Scale Out, Containers, Microservices

Traditional Applications

Monolithic Model

Multi-tier Apps

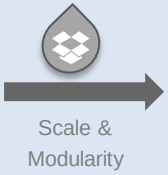

Scale & Modularity

Cloud-native applications

Business Agility with cloud model

Micro-services / Bi-Modal IT / DevOps

Focus on Products

Disjoint approaches to solve technical demands. Cohesiveness as "after thought"

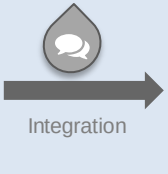

Integration

Focus on Integrated Solutions

Data Center is the foundation for business agility. Delivered as a solution and / or as a service

Manual Interaction

IT Silos based approach

Configuration driven

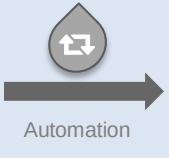

Automation

Policy Driven Automation

Enterprise-wide policy, hyper-convergence and cross-domain automation. Consumption driven with analytics and programmability

What is Cisco Virtual Topology System (VTS)?

- Overlay Provisioning and Management System
- Automates Overlay provisioning across Cisco Datacenter Top of Rack Nexus switches (Nexus 2000- Nexus 9000), Virtual Switches & DCI routers
- Automates fabric provisioning for both virtual and bare metal workloads.
- Programmable using North Bound REST APIs
- Tighter Integration with Orchestration systems such as Openstack, vCenter and Cisco NSO

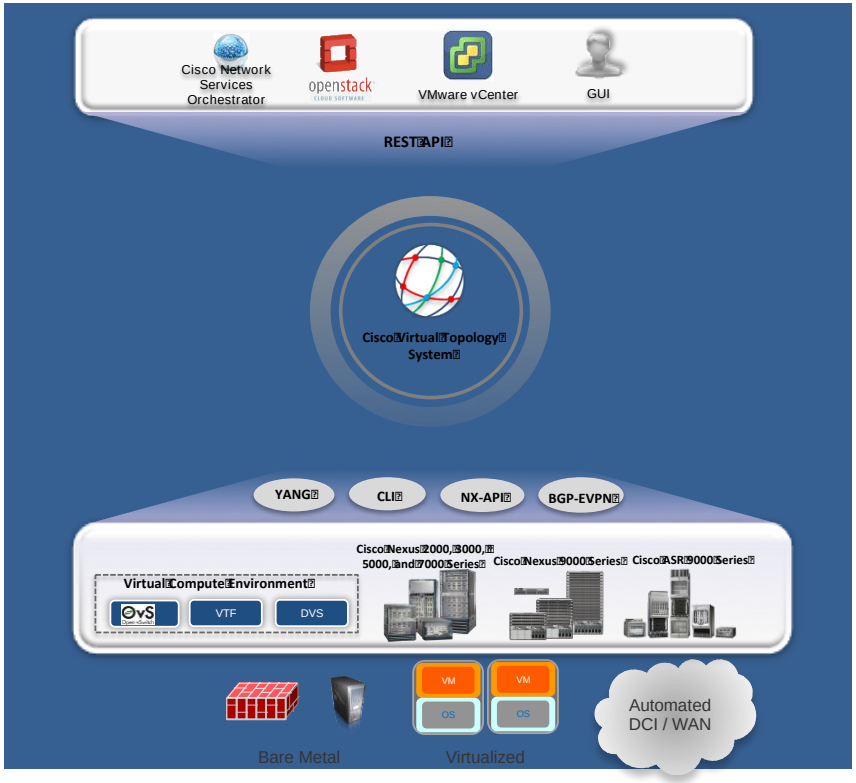

Simplified Management for Ease of Operations

VTS : Addressing Application Needs

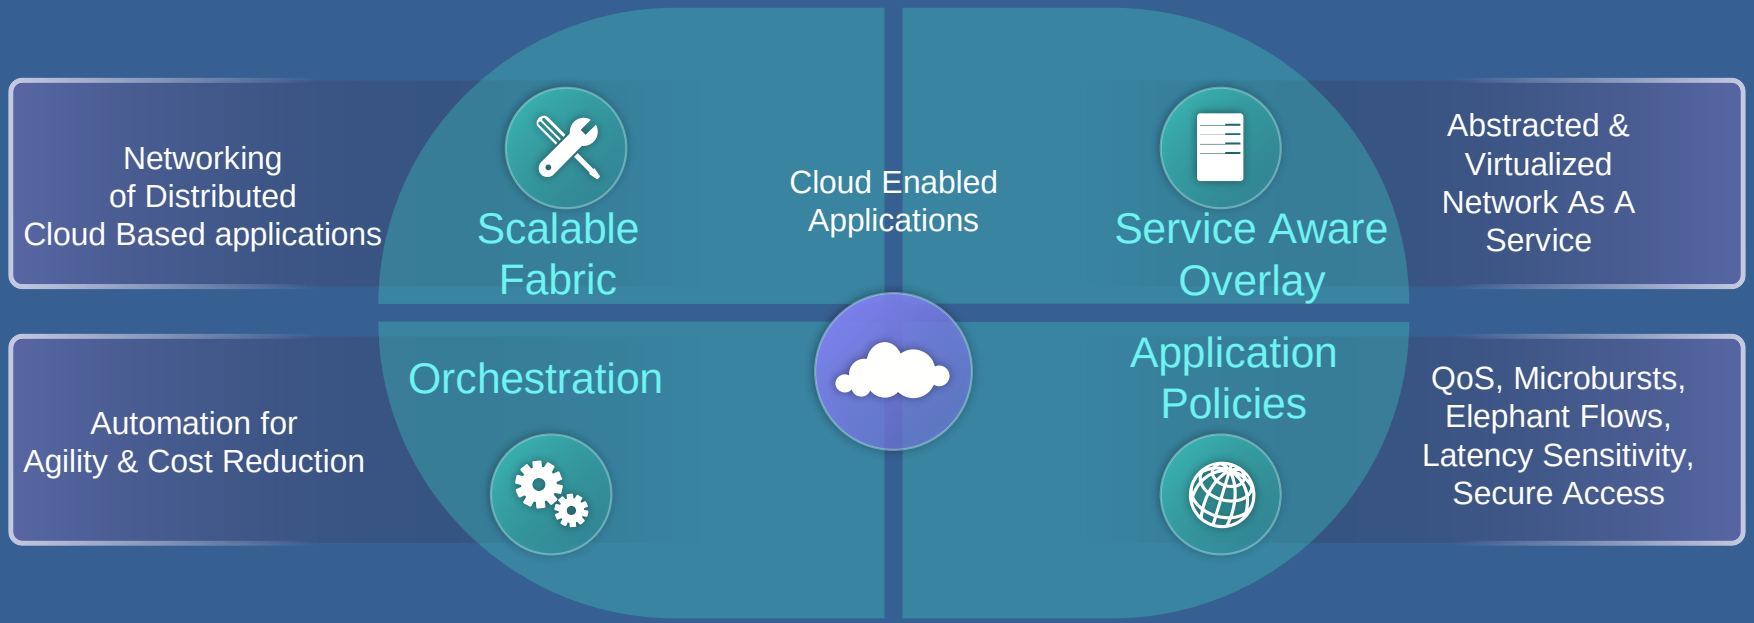

Scale & Automation Drives Business Success

3S TECHNOLOGY CITY

Why Cisco VTS?

Open

- Control & Data Plane
- Programmable Architecture (NB & SB)
- Interoperability (MPLS/VPN, OTV)

Agility and Automation

- Network as a Service
- Integration with Orchestrators
- Automated DCI/WAN
- Multi-Tenancy

Seamless Integration

- Multi-Hypervisor Multi-VMM
- Heterogeneous Workloads
- Custom NB Integration
- Services Integration (P&V)

Scale Performance Efficiency

- Scale-Out PODs
- Fabric Efficiency
- Multi-POD & Multi-DC

Investment Protection

- Host Based Overlays
- N2k-N9k, ASR Support
- Bare metal Apps/Services
- Interoperability

Policy Driven

- Infrastructure
- Network Connectivity
- Group Based Policies
- Service Assurance

3S TECHNOLOGY CITY

Scalable Fabric, Network Virtualization & Overlays

Scalable Fabric to meet Application Demands

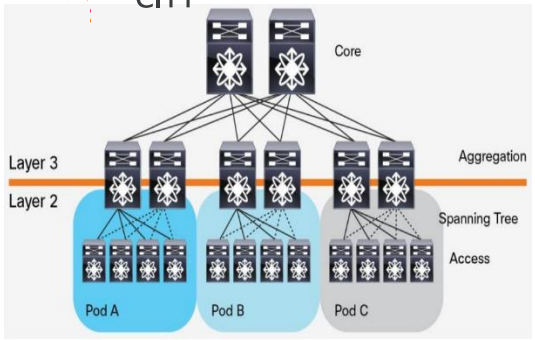

Three tier DC with STP

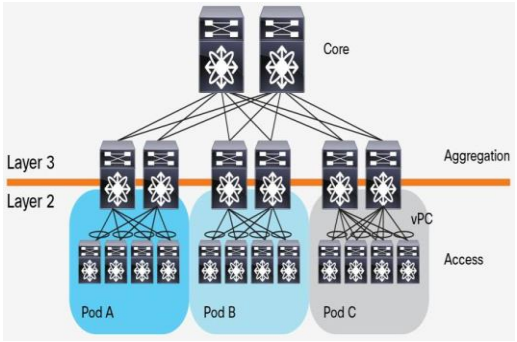

Three tier DC with VPC

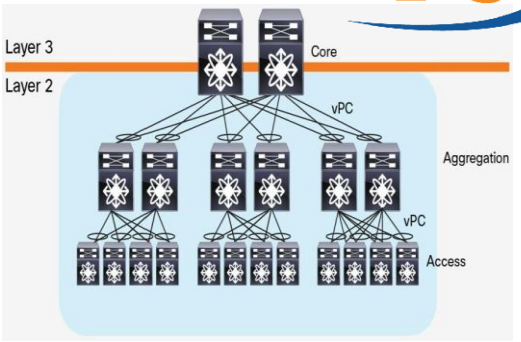

Three tier DC with L2 Extended

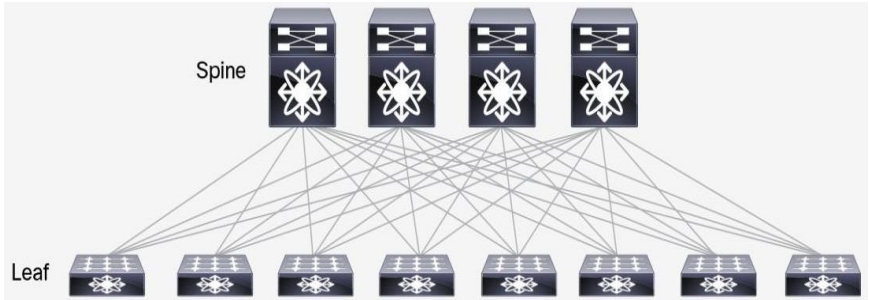

Spine-Leaf, L3 Fabric

VTS

Spine Leaf Evolution

Fabric Path + BGP → VXLAN (Flood & Learn) → Fabric Path + BGP + Automation → VXLAN + EVPN → VXLAN + EVPN + Automation

Network Virtualization and Multi-Tenancy using Overlays

Network virtualization: ability to separate, abstract and decouple the physical infrastructure & topology from a 'logical' topology or infrastructure typically by creating overlay networks.

Network overlays helps disassociates applications from physical networks infrastructure & topology, allowing a transition to cloud based multi-tenanted & scalable networks.

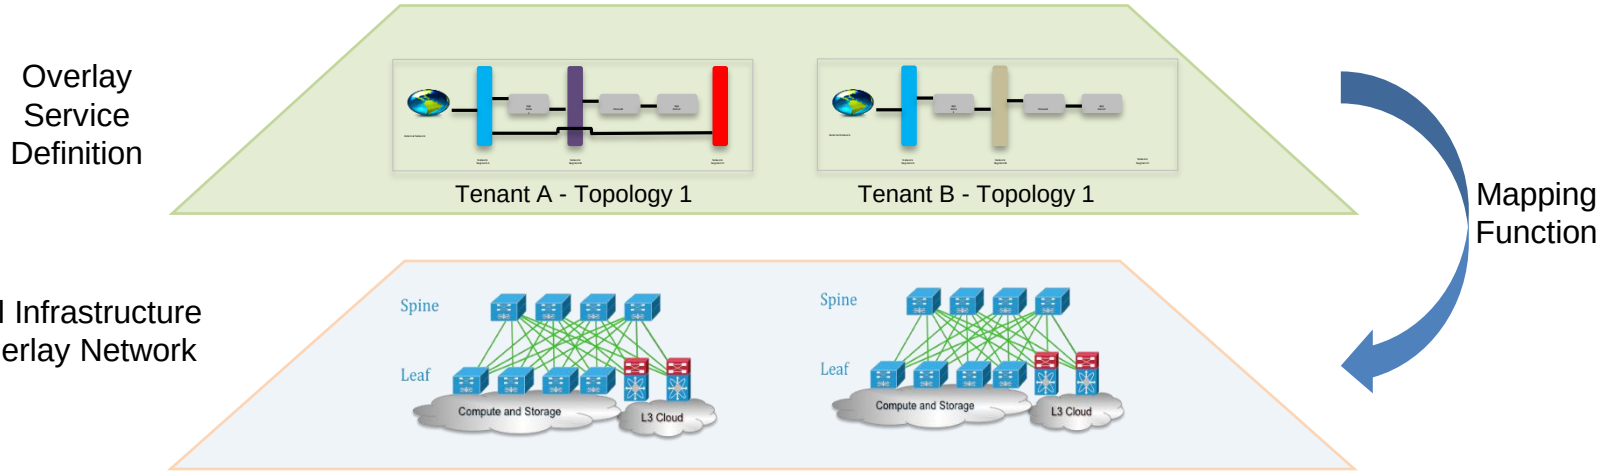

Routing and Forwarding Requirements For Overlays

Must Have Requirements for CP & DP:

- Underlay Topology Agnostic
- IP Only Underlay
- Open Standards Based
- Scalable multi-tenancy
- Optimal forwarding of L2 and L3
- Unified CP/DP for inter & intra POD

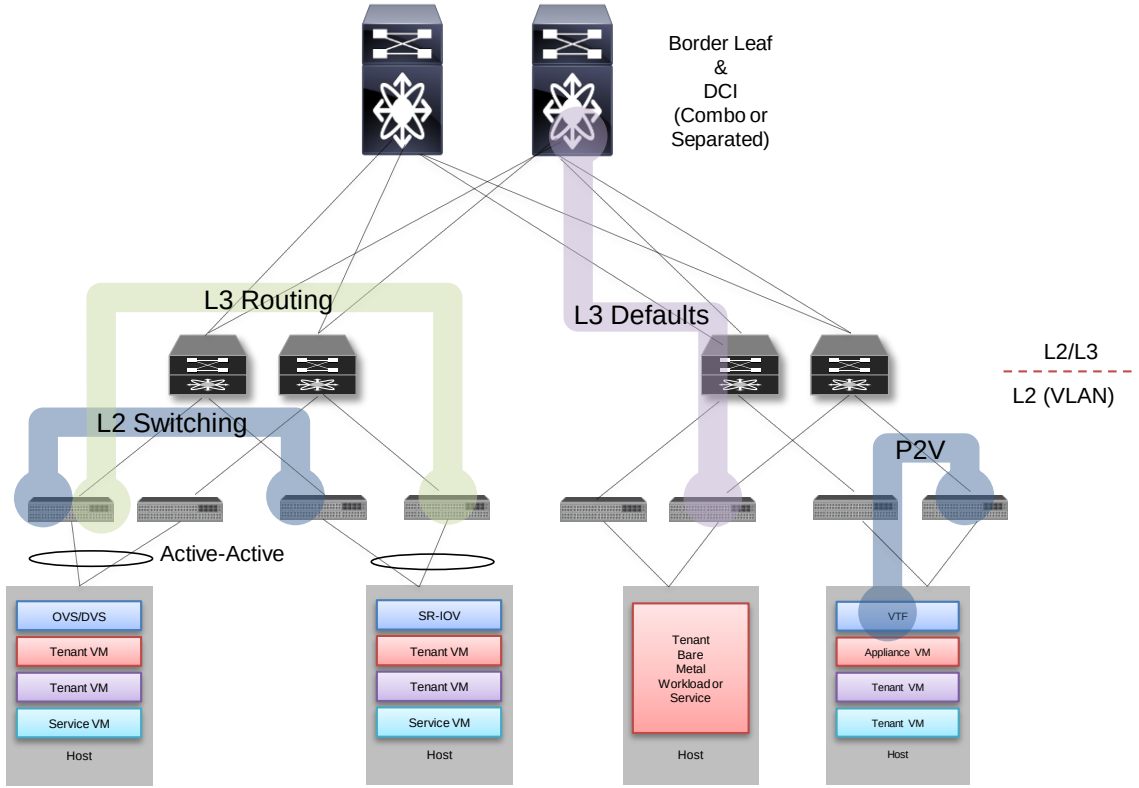

MPBGP-EVPN/VXLAN based Overlays

EVPN

Overlay Forwarding Table		
T1,S1	MAC, IP Address	P1/2
T1,S2	MAC, IP Address	VTEP2
T2,S3	MAC, IP Address	VTEP3
T2,S4	MAC, IP Address	VTEP4

Layer-2 MAC and Layer-3 IP information distribution by Control-Plane (BGP)

- Built in multi-tenancy (at scale)
- Integrated Routing/Bridging (IRB) for Optimized Forwarding
- Minimize flooding through ARP suppression
- Fast convergence upon network failures and host movements
- Security through VTEP peer-authentication

VXLAN

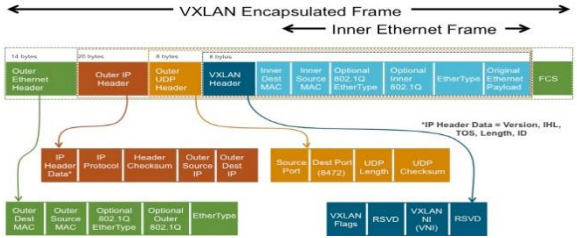

- IP routing – proven, stable, scalable
- ECMP – utilize all available network paths
- Flexible placement of multitenant segments
- Better utilization of network paths
- Scalable network domain (16M VNI vs. 4K VLANs)

EVPN Control Plane & VXLAN Data Plane in Action

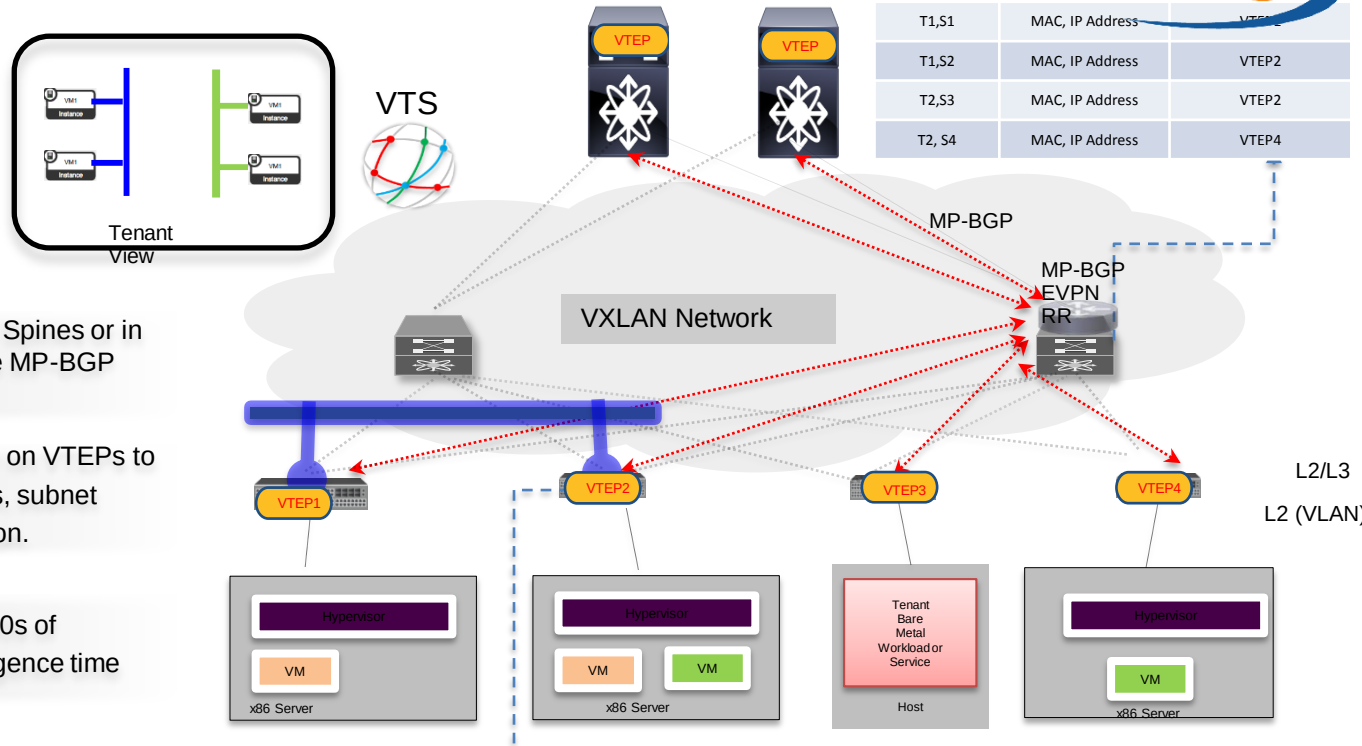

T1,S1	MAC, IP Address	
T1,S2	MAC, IP Address	VTEP2
T2,S3	MAC, IP Address	VTEP2
T2, S4	MAC, IP Address	VTEP4

VTS can be used to deploy the RR at the Spines or in VTS Control plane and then automate the MP-BGP peering on all the VTEPs in the fabric

Use MP-BGP with EVPN Address Family on VTEPs to distribute internal host MAC/IP addresses, subnet routes and external reachability information.

MP-BGP enhancements to carry up to 100s of thousands of routes with reduced convergence time

Overlay Forwarding Table		
T1,S1	MAC, IP Address	VTEP1
T1,S2	MAC, IP Address	P1/2
T2,S3	MAC, IP Address	P1/2
T3,S4	MAC, IP Address	VTEP4

BGP Update
Host-MAC
Host-IP
Internal IP Subnet
External Prefixes

VTEP – VXLAN Tunnel End-Point

VTS Architecture: Hardware VTEPs

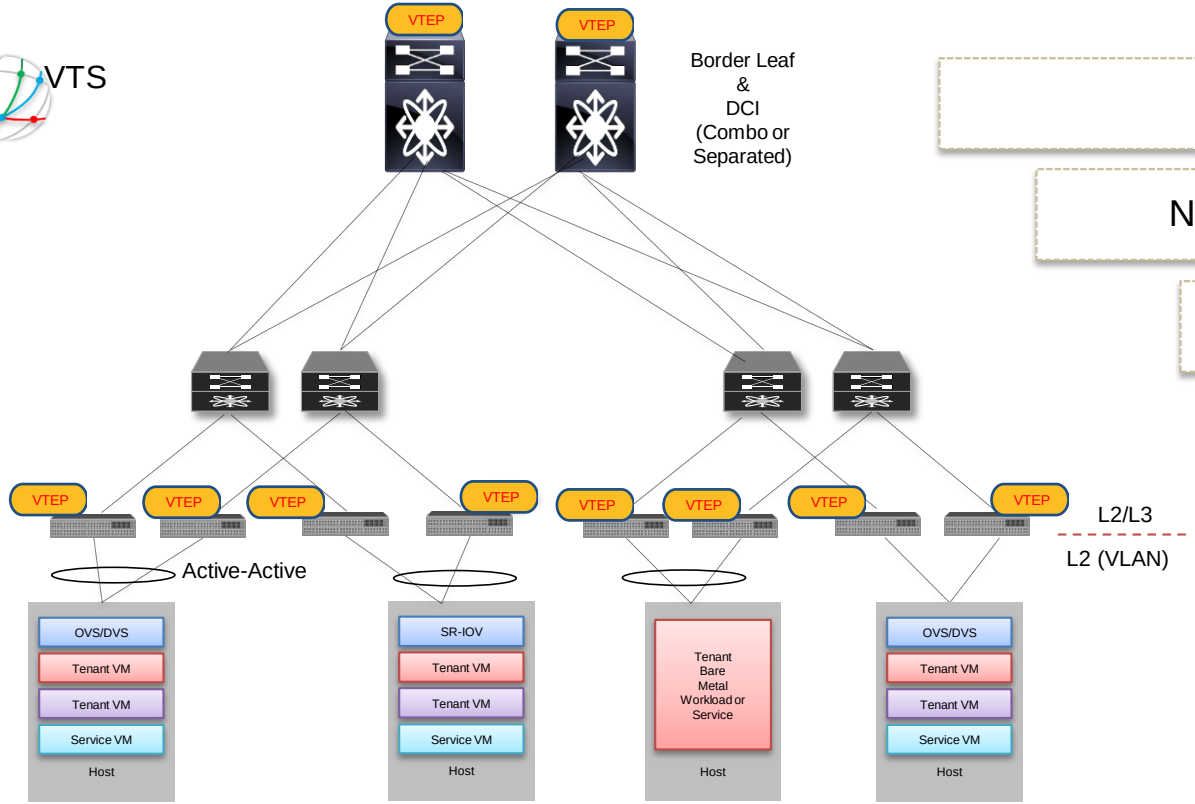

Nexus 9200/9300/5600 – ToR

Nexus 9x00/5000/7x00– Spine (RR)

ASR9000, Nexus 7x00– DCI

VTS Architecture: Software VTEPs

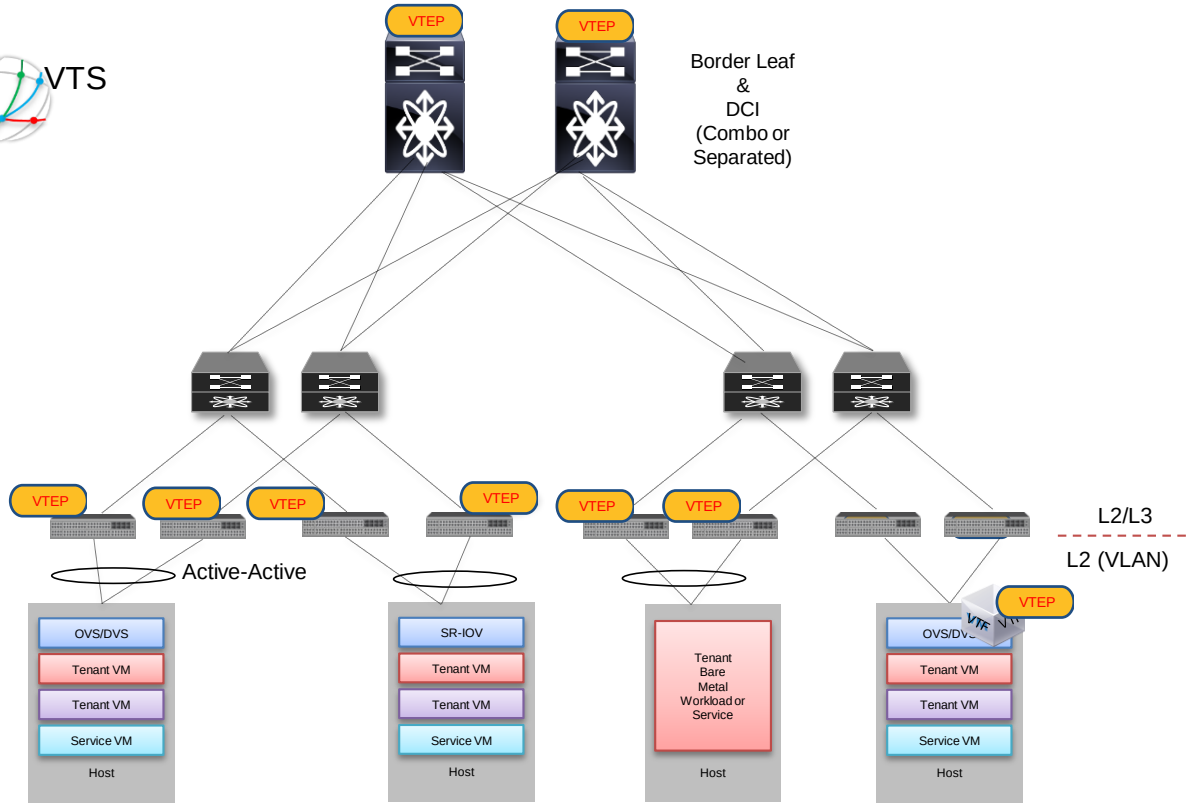

3S TECHNOLOGY CITY

VTS Use Cases

3S TECHNOLOGY CITY

VTS Use Cases

Multi-Tenant Data Centers

Network-Function Virtualization

3S TECHNOLOGY CITY

VTS – Multitenant Data Center
A case study

- Cloud Management: OpenStack
- Host OS: RHEL, Hypervisor: KVM
- Sites: Multi-POD, Multi-DC
- Core: MPLS Core
- Servers Connected as VPC
- Services: Firewall & Load Balancer
- Management and IP Storage Network
- L2/L3 Connectivity
- Selectively allow L2 outside POD
- Remote access for branches
- Integrated DCI/BL
- N9k within POD and ASR9k as DCI

VTS Architecture – Single POD

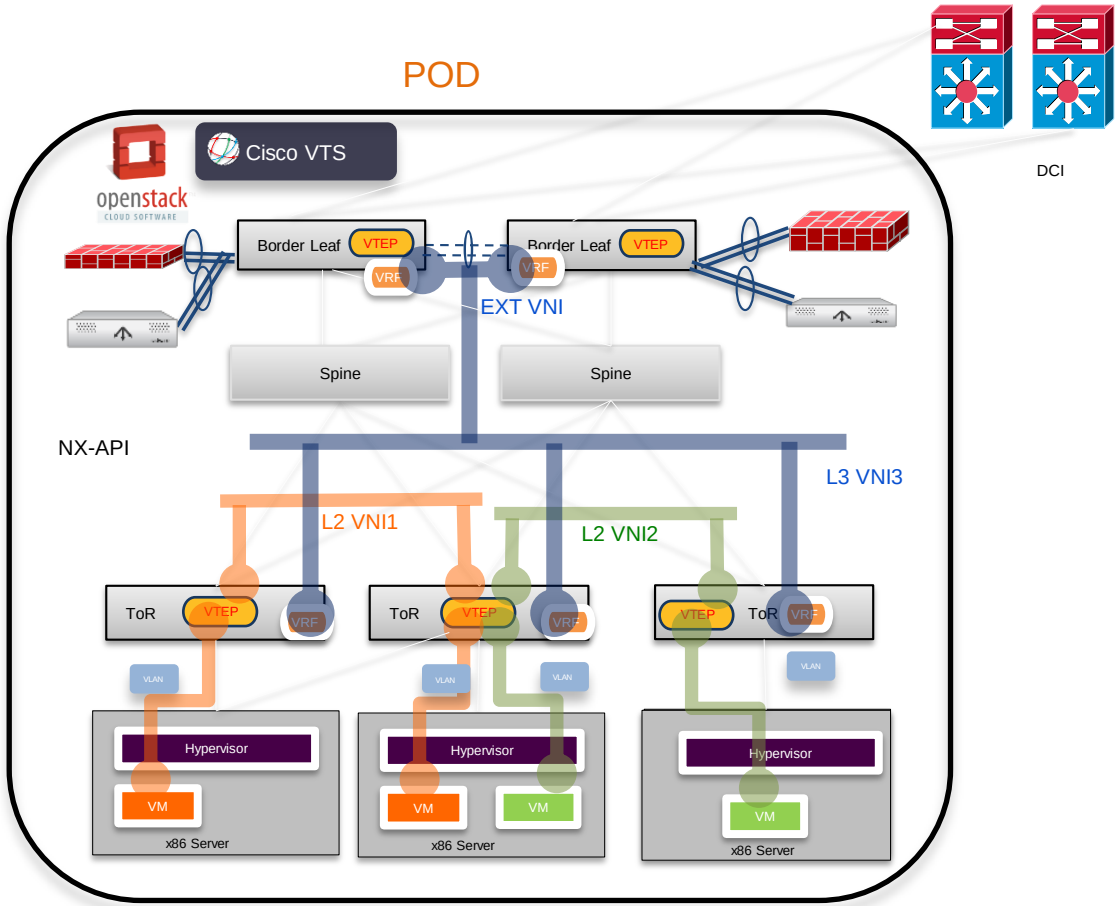

VTS Architecture Multi-Site

- For Disaster recovery, High Availability
- Integrate EVPN/VXLAN to MPLS-L3VPN

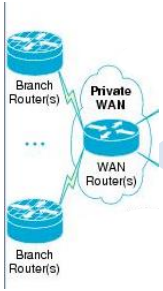

IP/MPLS WAN

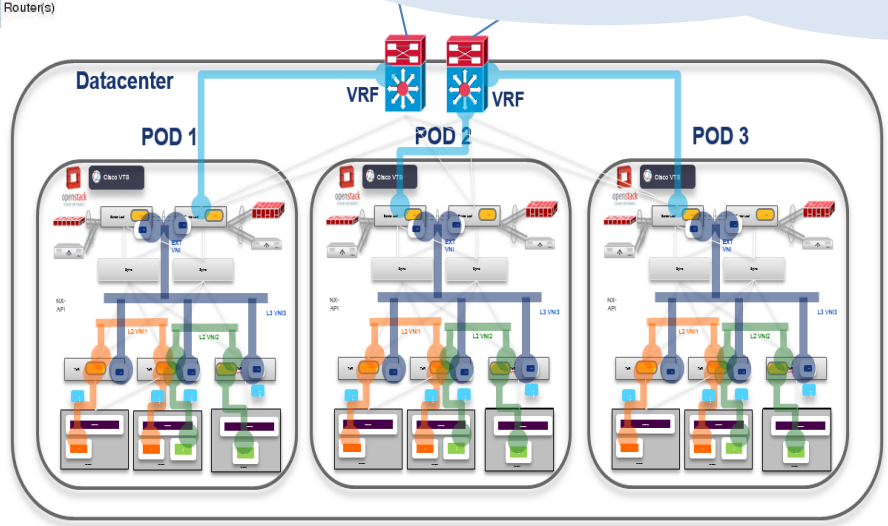

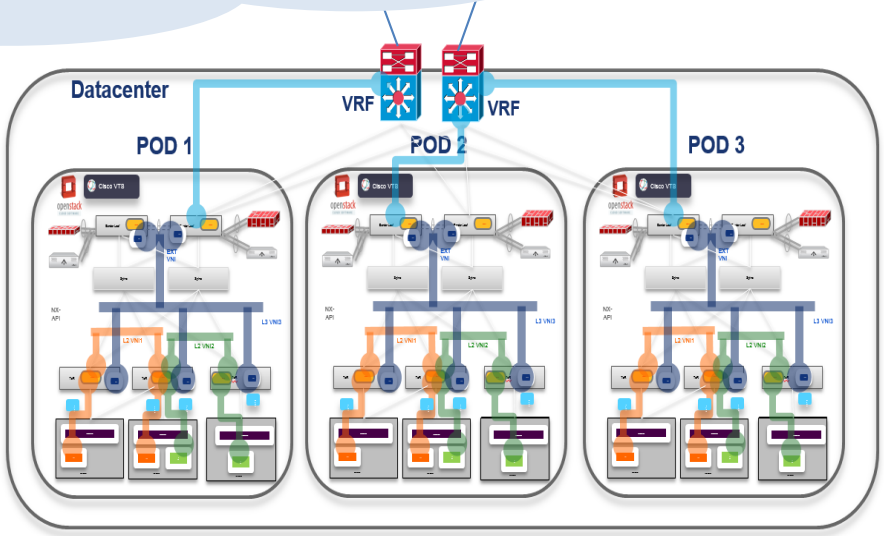

Cisco NFV Integration with VTS

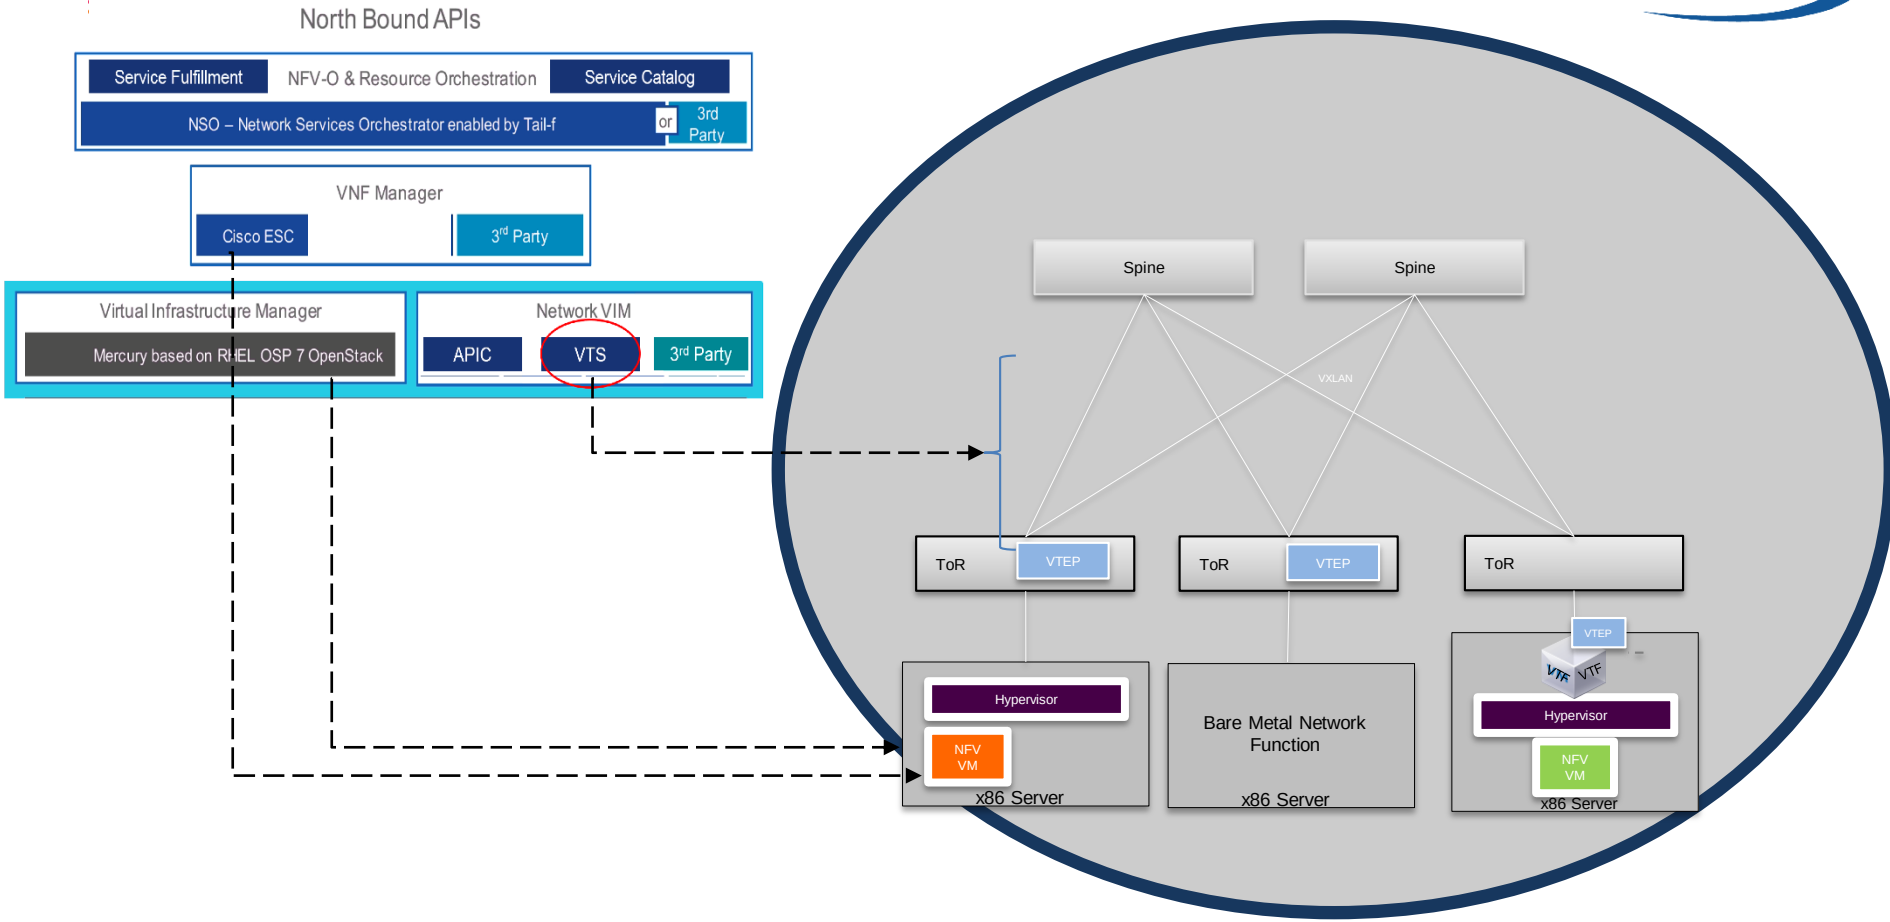

Innovation for a Better Future