

Introduction to Cisco Virtual Topology System

Speaker Name: Linda Wang (lindawa@cisco.com)

Speaker Title: Solution Support Engineer

Date: September 7, 2016


Agenda

- Cisco Data Center SDN Strategy
- VTS Architecture Overview
- VTS Functionality
- VXLAN & BGP EVPN
- VTS and OpenStack
- VTS and VMware
- Use Cases
- Roadmap
- Demo

Cisco SDN: Updates on Nexus Portfolio Offerings

Extended NX-API Support Across Nexus 2K-9K


Application Centric Infrastructure



NEW! ACI Release for Nexus 9000 (Shipping June 2015)

- Microsoft Azure and System Center Integration
- Programmability examples: vCenter plugin, ACI toolkit etc.
- Simplified operations
- Stretched fabric, multiple destinations from 30KMs to 150KMs
- Group-based policy on Openstack
- New ACI ecosystem partners (CliQr)


Programmable Fabric



NEW! Virtual Topology System (VTS) for software overlay provisioning and management across for Nexus 2K-9K (2H 2015)

- Standards-based fabric support on Nexus 5600/7x00 with VXLAN BGP EVPN (shipping with Nexus 9000 today)

Programmable Network



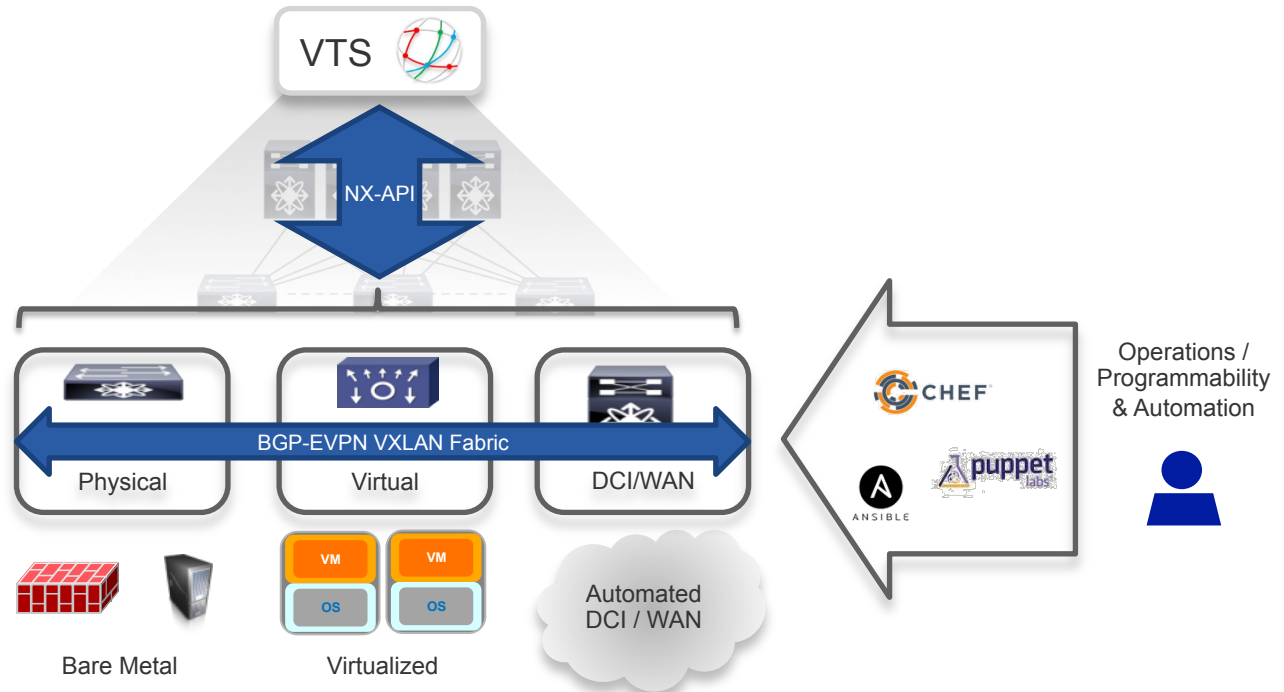
NEW! Unified Open NX-OS Release for Nexus 3000 and Nexus 9000 (Q3 2015)

- Enhancements to NX-API – object store and model driven
- Native 3rd party RPM applications integration (tcollector, Nagios, Ganglia, Puppet / Chef etc.)
- Linux utilities support for seamless tool integration across compute and network
- SDK for custom application integration

NEW! Common NX-API across N2K-N9K (2H 2015)

Programmable Fabric

NX-API, VXLAN BGP EVPN Fabric, and Virtual Topology System (VTS)

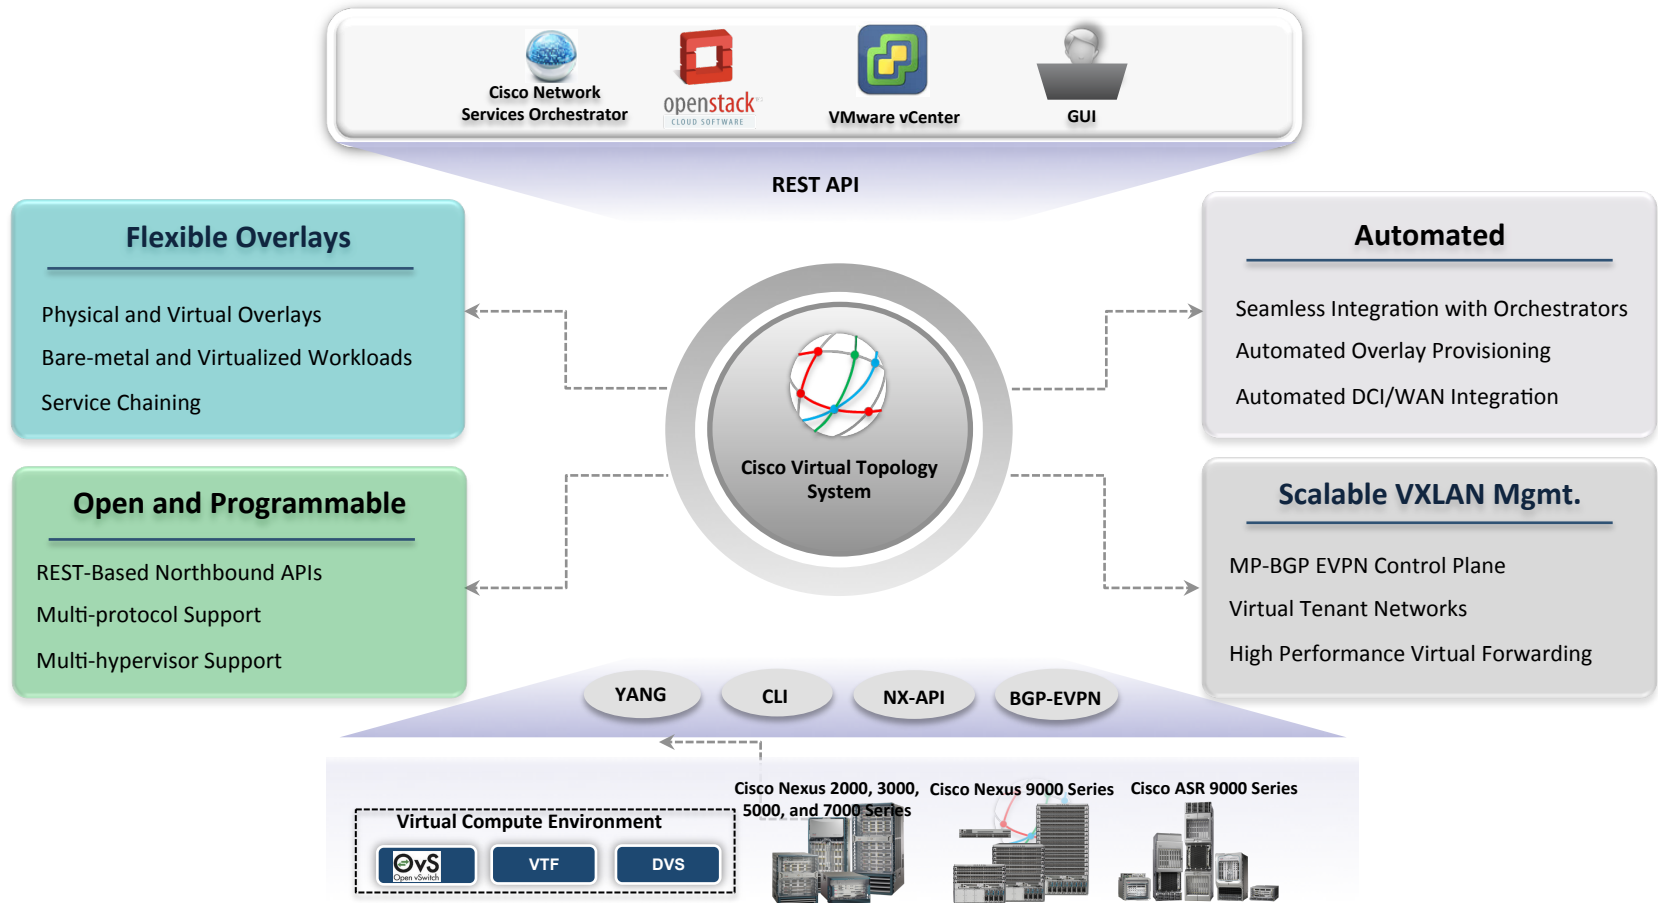


VTS for overlay provisioning and management across Nexus 2000 – Nexus 9000 (2H 2015)

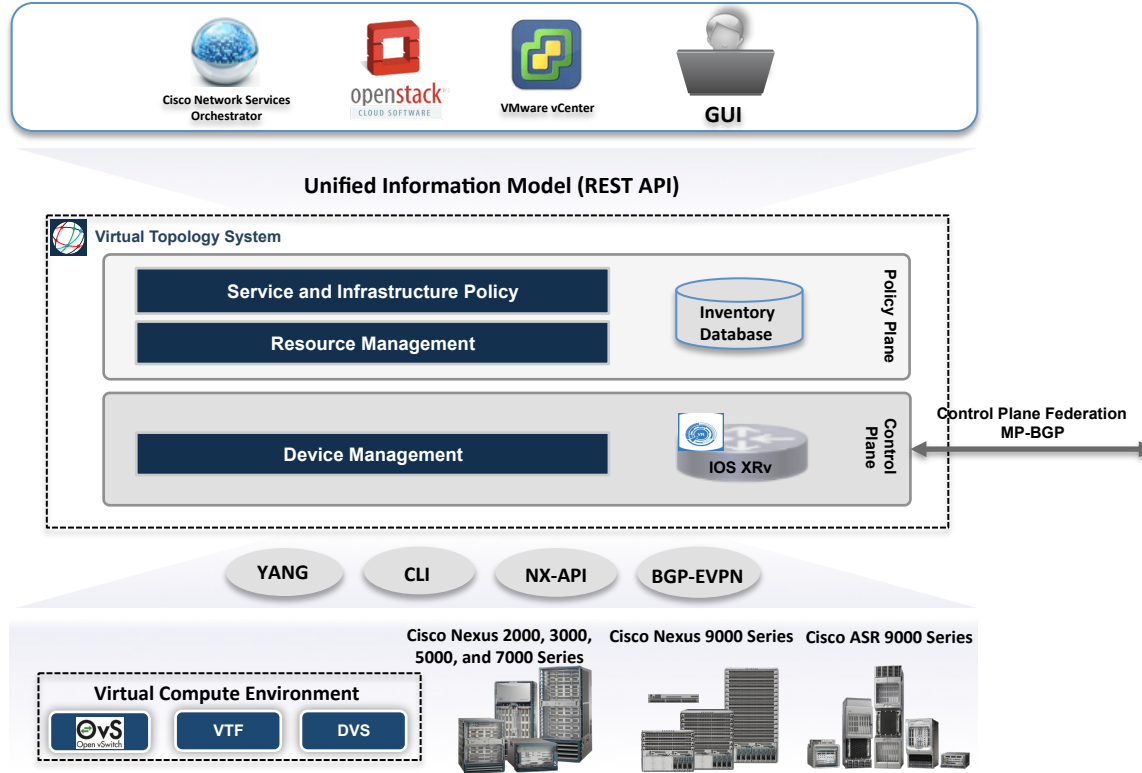
VTS Architecture Overview

Cisco Virtual Topology System (VTS)

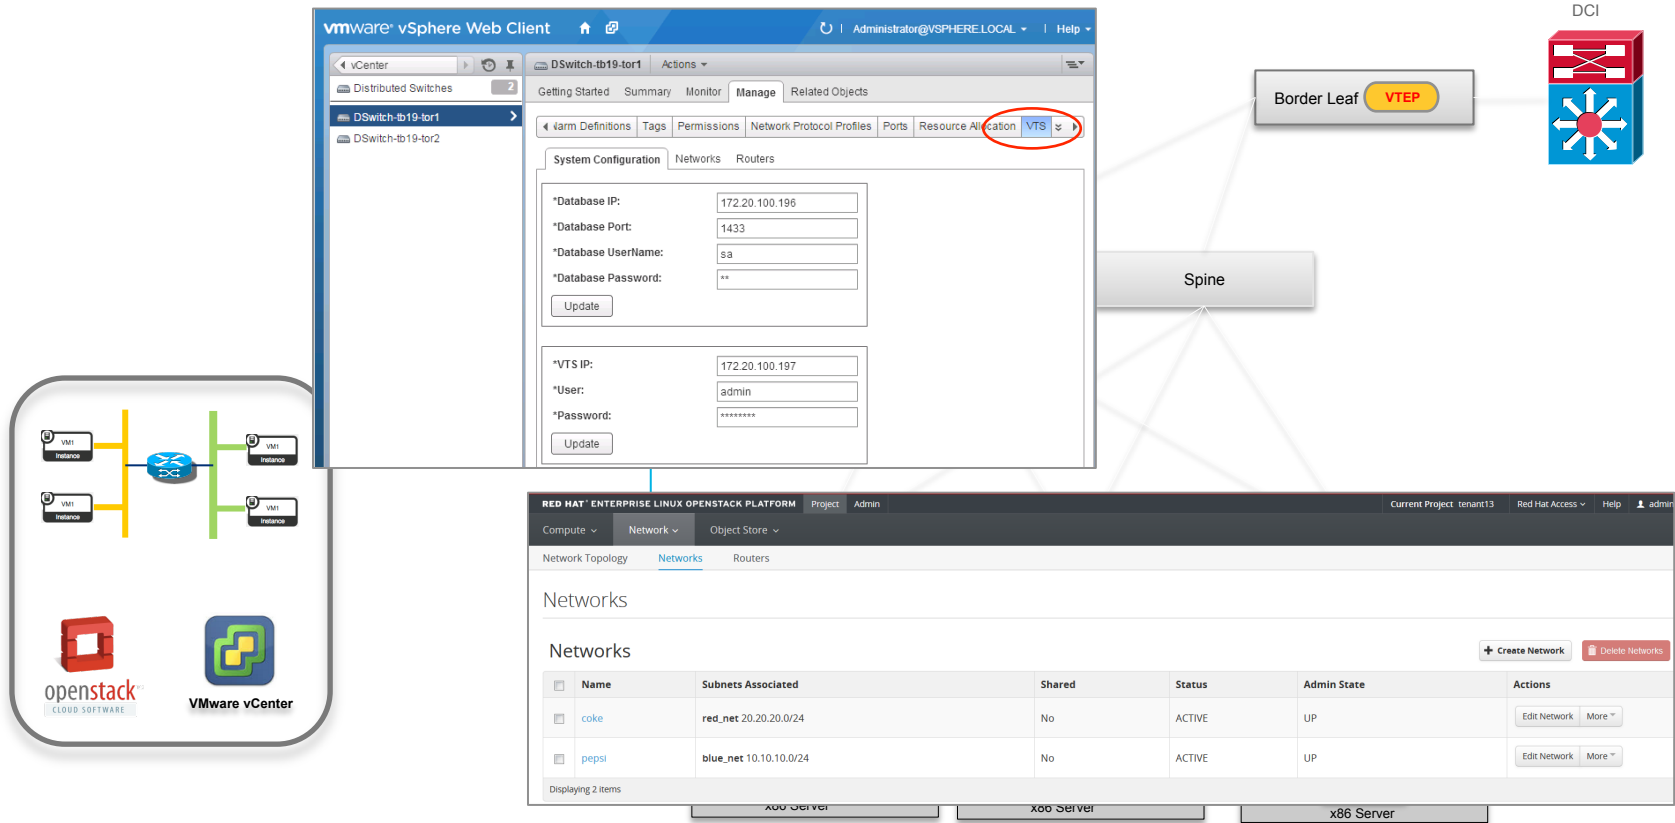
Overlay Provisioning & Management System



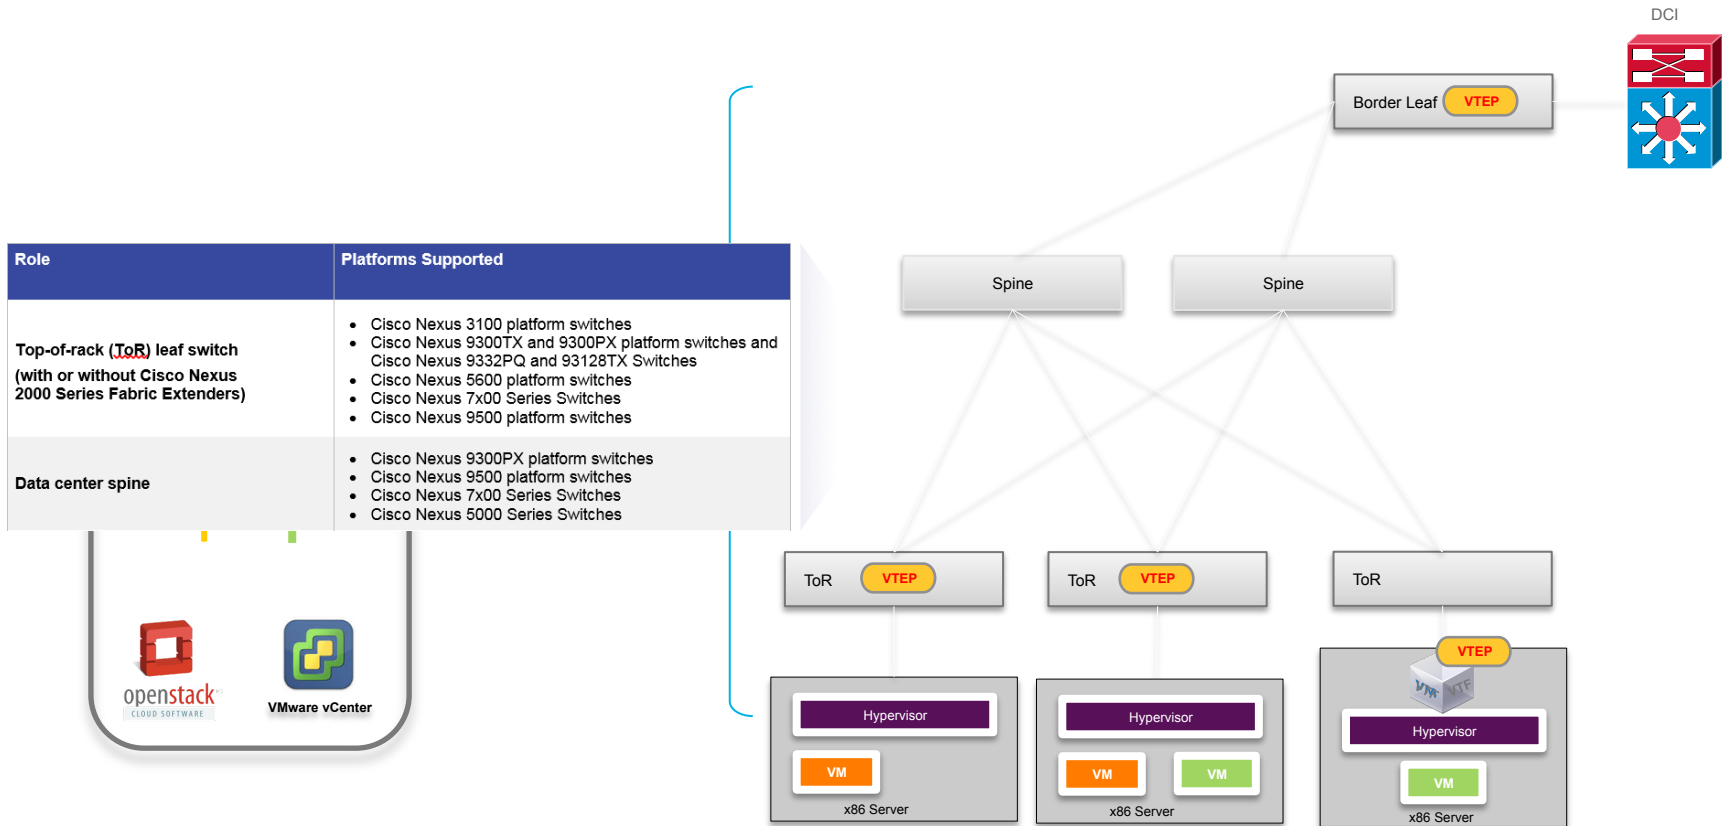
VTS Architecture



VTS Architecture



VTS Architecture – Hardware Switches



VTS Architecture - VTF

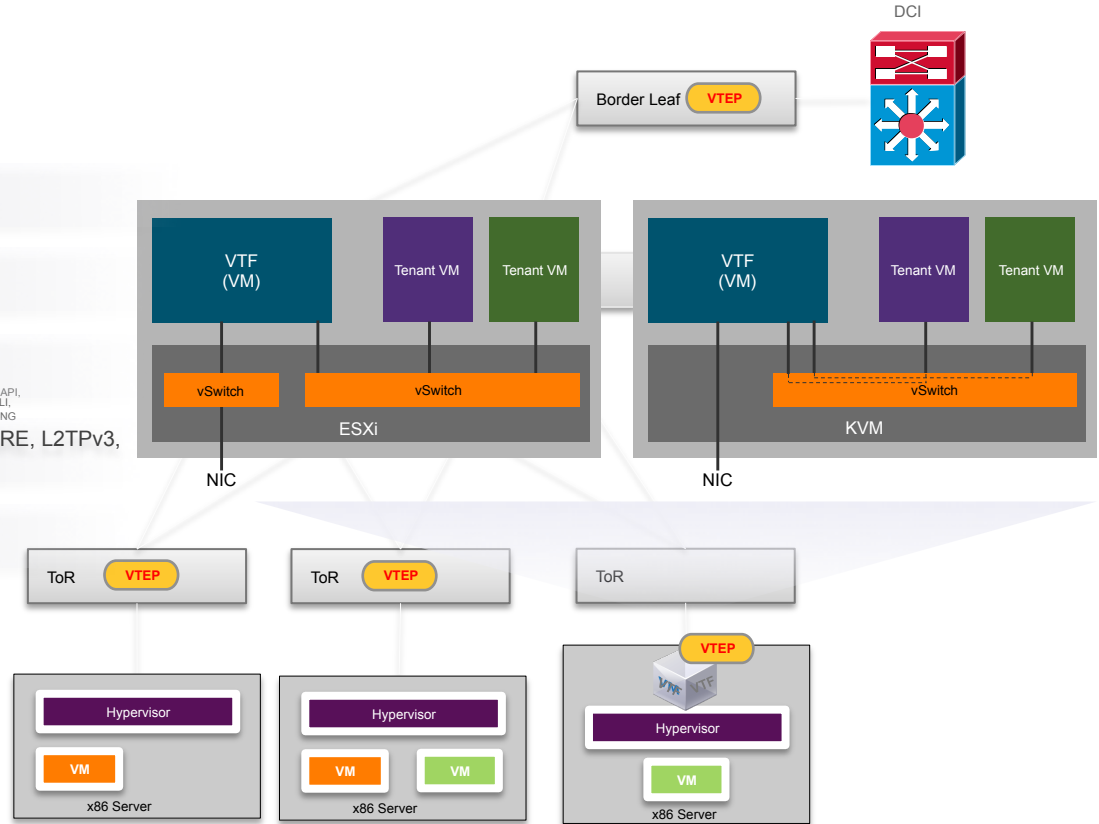
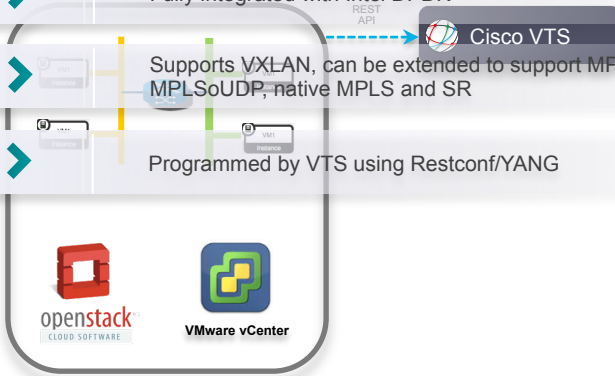
➤ User space, Multi-tenant, line rate packet forwarder

➤ Uses Vector Packet Processing technology

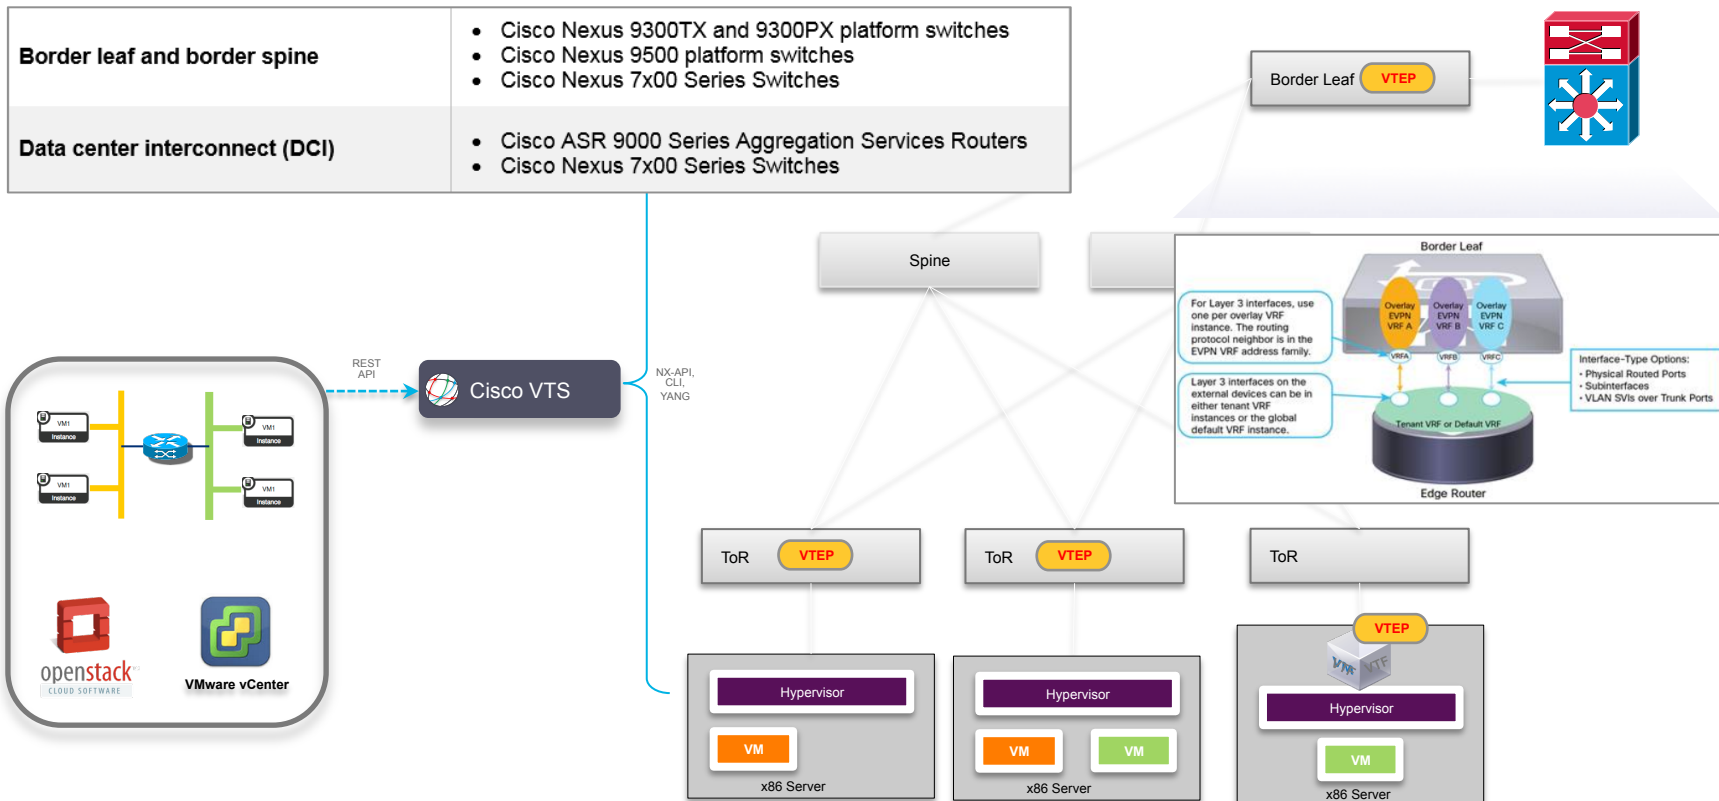
➤ Fully integrated with Intel DPDK

➤ Supports VXLAN, can be extended to support MPLSoGRE, L2TPv3, MPLSoUDP, native MPLS and SR

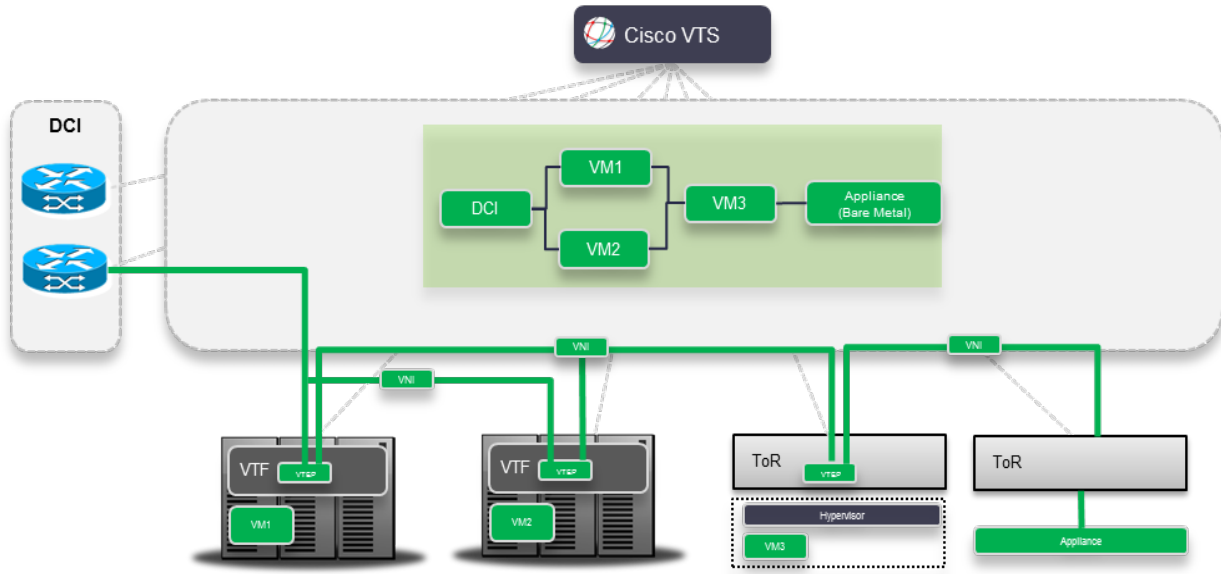
➤ Programmed by VTS using Restconf/YANG



VTS Architecture – Border Leaf and DCI



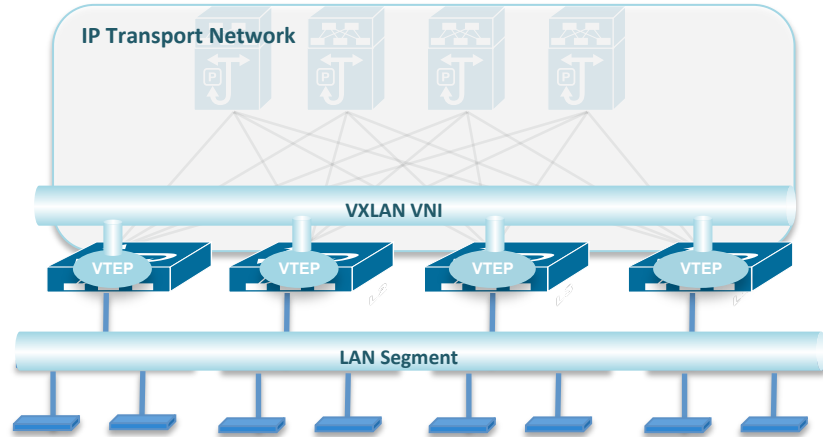
Flexible Overlays using Virtual Topology Forwarder



- Hardware based Overlays
- Software (VTF) based Overlays
- Hybrid Overlays

Virtual Overlays for Brownfield Deployment

VXLAN as Data Center Overlay technology



Underlay Network:

- IP routing – proven, stable, scalable
- ECMP – utilize all available network paths

Overlay Network (VXLAN):


- Standards-based overlay
- Flexible placement of multitenant segments
- Better utilization of network paths
- Scalable network domain (16M VNI vs. 4K VLANs)

Modes of Operation:

- Multicast based flood and learn (No control plane)
- BGP EVPN (BGP control plane with MP-BGP Extensions)

EVPN Control Plane

S1	MAC, IP Address	VTEP1
S2	MAC, IP Address	VTEP2
S3	MAC, IP Address	VTEP3
S4	MAC, IP Address	VTEP4



Industry standard protocol for multi-vendor support



Built in multi-tenancy support



Scalable, protocol driven control plane architecture



Fast convergence upon network failures and host movements



Minimize flooding through ARP suppression



Security through VTEP peer-authentication

Overlay Forwarding Table

S1	MAC, IP Address	P1/2
S2	MAC, IP Address	VTEP2
S3	MAC, IP Address	VTEP3
S4	MAC, IP Address	VTEP4

VTS Setup Steps

- Install Operating System on Bare Metal Server (e.g. Centos 7.1)
- Install OpenStack (e.g. Liberty)
- Setup VTC VM
- Setup XRv VM*
- Install OpenStack Plugin + Host Agent
 - a. Modify Credentials File (Automated Install)
 - b. Install OpenStack Plugin + Host Agent (auto Install)
- Configure Day0 config on TOR and Spine Switch

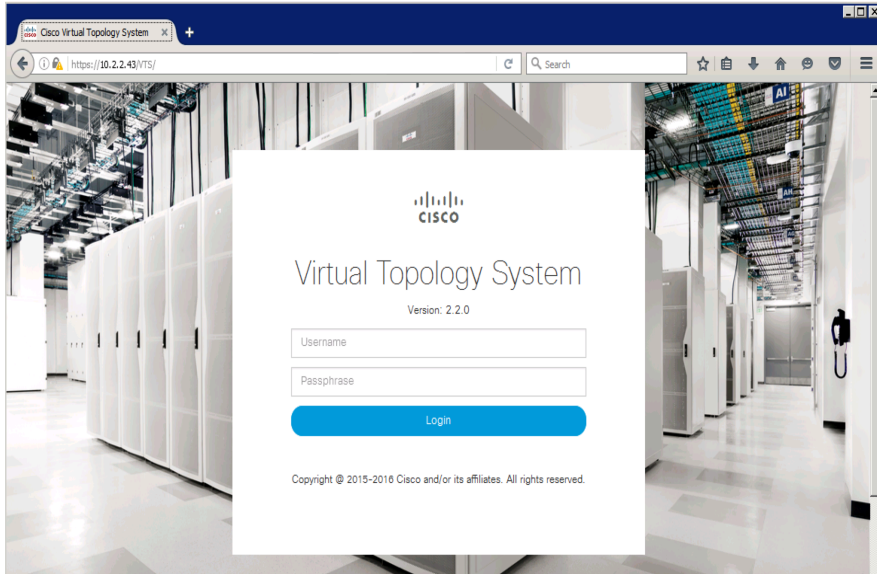
*Only if VTS L3 HA is required, or XRv VM is to be used as BGP-RR

Post-Installation Tasks

- Create Authorization Group
- Device Discovery > Export Inventory
- Edit Inventory > Import Inventory
- Check Network Inventory > Edit Device Loopback/ASN
- Check Host Inventory
- Manage Resource Pool (Assign VNI, VLAN, Multicast Range)
- Overlay provisioning (via VIM or GUI)

VTS Functionality

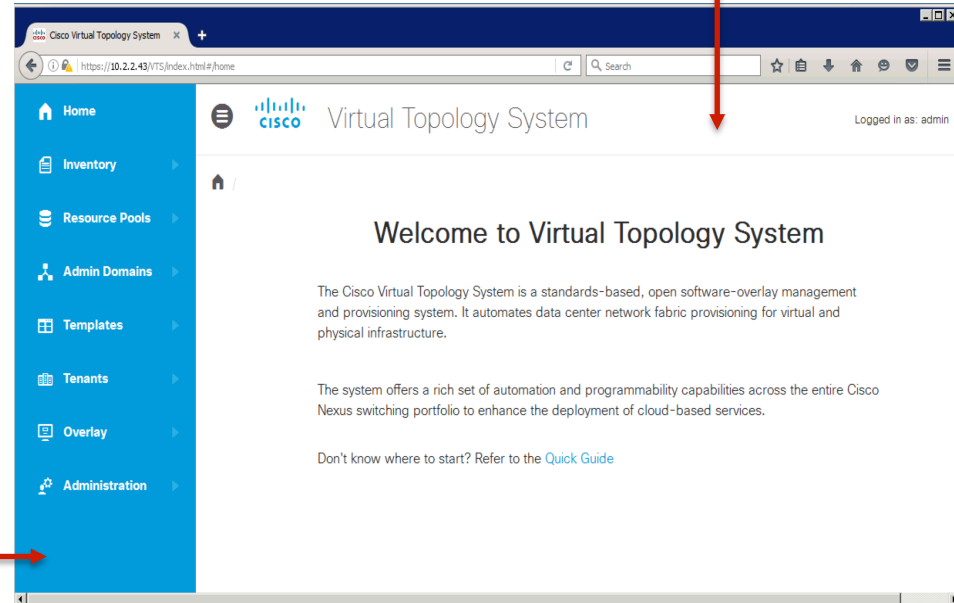
VTS Functionality: VTS Initial Configuration



Run through initial system configuration

Login to VTS WebGUI

Work
Pane



Navigation
Pane

VTS Functionality: Discovery & Topology View

Cisco Virtual Topology System

Virtual Topology System

Logged in as: admin

/ Inventory / Topology

Select Tenant: --ALL--

Show Hosts: Yes

Show Standalone Devices: No

Fabric Import & Discovery

1. Performing Topology Auto Discovery using LLDP
2. Import devices using inventory CSV file
3. Adding Device (Host and Network Devices) using VTS WebUI

VTS Functionality: Programming the Fabric

→ <https://10.2.2.42:8443/VTS/index.html#/overlay-network>

- Home
- Inventory
- Resource Pools
- Admin Domains
- Templates
- Tenants
- Overlay
- Administration



Virtual Topology System

Logged in as: admin



Home / Overlay / Network



Select Tenant:

admin

Search



<input type="checkbox"/>	Network Name	Tenant	Zone	Admin State	Status	Subnet	Network Type	External Network	VNI
<input type="checkbox"/>	Net-1	admin	admin		active		vlan	false	10421
<input type="checkbox"/>	Net-2	admin	admin		active		vlan	false	10422

Showing 1 to 2 of 2 entries

Previous

1

Next

Programming the Fabric

1. OpenStack /VMware
2. VTS GUI
3. NSO

VTS Functionality: Fabric Synchronization

The screenshot shows the Cisco Virtual Topology System (VTS) web interface. The browser address bar displays `https://10.2.2.42:8443/VTS/index.html#/network-inventory`. The page title is "Virtual Topology System" and the user is logged in as "admin". The navigation menu on the left includes Home, Inventory, Resource Pools, Admin Domains, Templates, Tenants, Overlay, and Administration. The main content area shows the "Network Inventory" page with tabs for "Device" and "Fabric Connection". A table lists three network devices:

	Device Name	Admin State	IP Address	Auth Group	Device Platform	Device Role	Group Tag	BGP-ASN	Loopback Interface Number	Loopback Interface IP/Mask	Sync
<input type="checkbox"/>	N7k-DCI	unlocked	10.2.1.37	VTC-POD2	N7K	dci		65100	0	1.1.3.1/32	Check Sync
<input type="checkbox"/>	N9k-BLeaf-1	unlocked	10.2.1.35	VTC-POD2	N9K	border-leaf		65001	0	1.1.2.3/32	Check Sync
<input type="checkbox"/>	N9k-Leaf-1	unlocked	10.2.1.33	VTC-POD2	N9K	leaf		65001	0	1.1.2.1/32	Check Sync

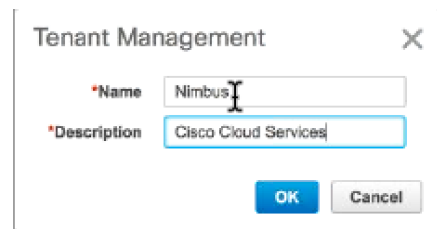
The "Sync" column is highlighted with a red border, and each row has a "Check Sync" link. The "Device Name" column contains information icons (i) for the three devices.

Fabric Synchronization

1. Sync before programming
2. Multiple VTS per fabric

VTS Functionality – Overlay Configuration through GUI

- Create a Tenant
- Add Network to the Tenant
 - Add Subnet
 - Select the ToR and host facing ports
- VLAN and VNID is automatically allocated
- Layer-2 VXLAN segment is configured on the ToR switches

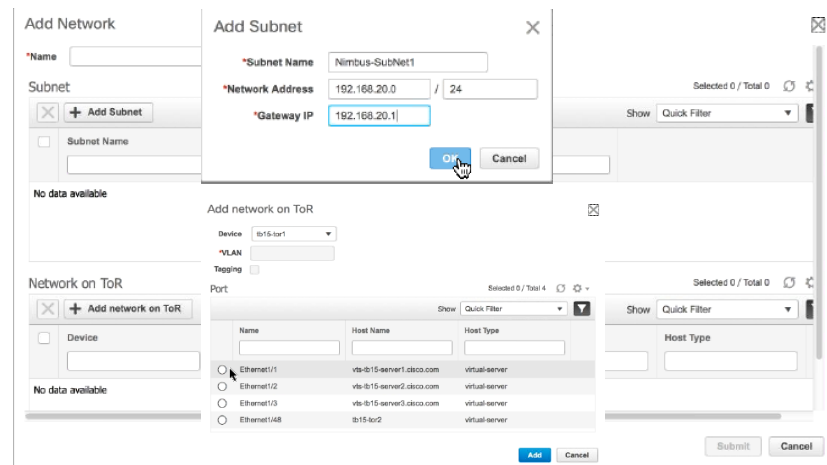


Tenant Management

*Name: Nimbus

*Description: Cisco Cloud Services

OK Cancel



Add Network

Add Subnet

*Subnet Name: Nimbus-SubNet1

*Network Address: 192.168.20.0 / 24

*Gateway IP: 192.168.20.1

OK Cancel

Add network on ToR

Device: b15-sw1

VLAN

Tapping

Network on ToR

+ Add network on ToR

Device

No data available

Port

Name	Host Name	Host Type
<input type="radio"/> Ethernet1/1	vls-b15-server1.cisco.com	virtual-server
<input type="radio"/> Ethernet1/2	vls-b15-server2.cisco.com	virtual-server
<input type="radio"/> Ethernet1/3	vls-b15-server3.cisco.com	virtual-server
<input type="radio"/> Ethernet1/48	b15-sw2	virtual-server

Add Cancel





Submit Cancel

VTS Functionality: Overlay Visibility

- Home
- Discovery
- Inventory
- Resource Pools
- Admin Domains
- Tenants
- Overlay
- Administration

Home / Overlay / Network

   Select Tenant: 

<input type="checkbox"/>	Network Name	Tenant	Zone	Admin State	Status	Subnet	Network Type	External Network	VNI
<input type="checkbox"/>	client-net	admin	admin	<input checked="" type="checkbox"/>	active		vlan	false	5000
<input type="checkbox"/>	client-net-ext	admin	admin	<input checked="" type="checkbox"/>	active		vlan	true	5001
<input type="checkbox"/>	demo1-private	admin	admin	<input checked="" type="checkbox"/>	active		vlan	false	5008
<input type="checkbox"/>	INTERNAL-MGMT	admin	admin	<input checked="" type="checkbox"/>	active		vlan	false	5003

Showing 1 to 4 of 4 entries

Previous Next


VTS Functionality: Backup

- Cisco VTS uses NCS internally
- NCS Backup backs up the database (CDB), state files, config files and rollback files from the installation directory
- Command: ncs-backup
- Backup will be stored in following directory
- /var/opt/ncs/backups/ncs- VERSION@DATETIME.backup

VTS Functionality: Restore

- Need to stop NCS first before performing restore
- `sudo -u ncs /bin/bash -c '/etc/init.d/ncs stop'`
- Command: `ncs-backup --restore`
- We can then select the list of backups to restore from
- Once restore is done, we can start NCS
- `sudo -u ncs /bin/bash -c '/etc/init.d/ncs start'`

VTS Functionality: High Availability



IP Address	Role	Status
10.68.234.210	master	slave
10.68.234.211	slave	master

Cisco VTS leverages on *Pacemaker* and *Corosync* for HA

Pacemaker: Performs detection and recovery of machine and app-level failure

Corosync: Provides membership, messaging and quorum capabilities

L2 HA

L2 HA Requisites:

1. Two instances of VTC
2. Require 3x IP Addr for VTCs - 1x VIP, 1x VTC1, 1x VTC2
3. Setup HA through VTS CLI interface

VXLAN & BGP EVPN


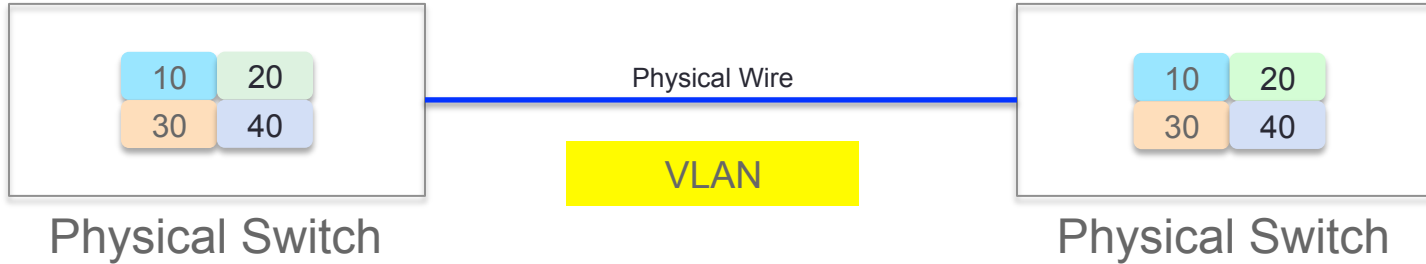
VXLAN Overview

- Limitations in Layer-2 Data Centers
 - Use of Spanning-Tree for loop prevention
 - Inefficient use of links
 - Slow convergence
- Limited VLAN space (12 bits = 4096 VLANs)
 - Very small number of virtual domains for today's tenant requirements

VXLAN Overview (CONT..)

- Uses Layer 3 Data Center Fabric
 - No need for Spanning-Tree Protocol
 - All ports are in forwarding state using Equal-Cost Multi-Path (ECMP)
 - Uses IP/UDP header (MAC in IP/UDP)
 - Divided into Underlay and Overlay(s)
- Uses Larger Layer-2 Segment Space (24 bit = 16M)
 - VXLAN Network Identifier (VNI) or VXLAN Segment ID
 - Scales to accommodate large number of tenants

VLAN vs VXLAN – In Simplest Form



1 VNI is 1 Broadcast Domain

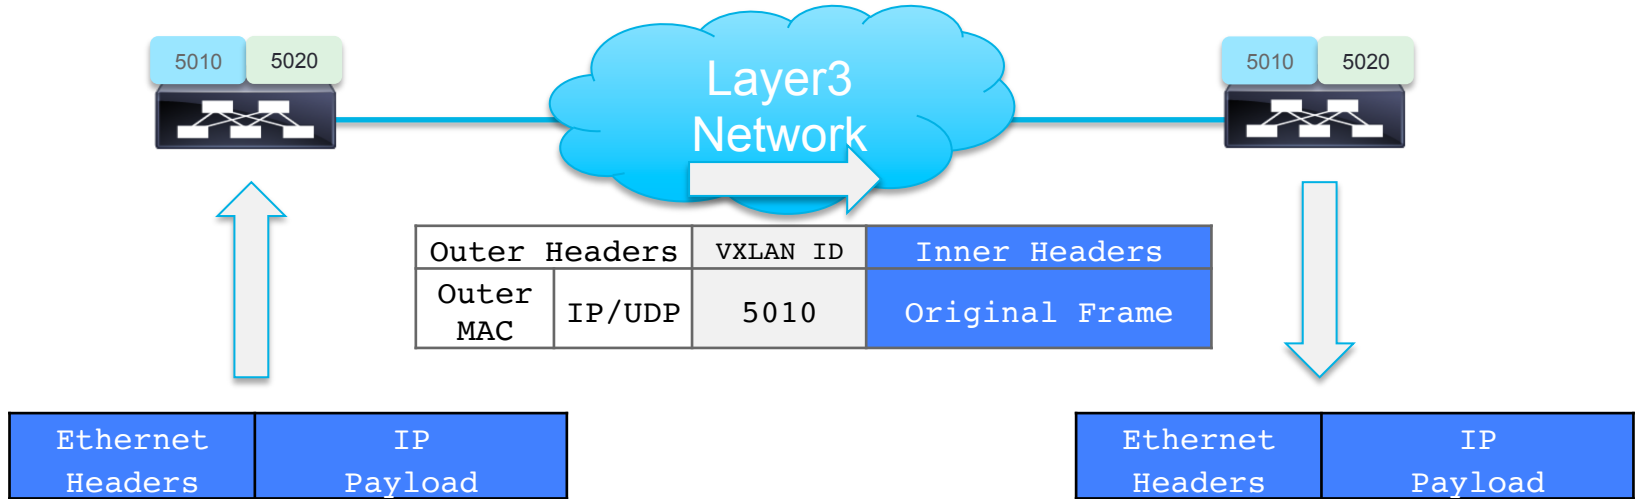
VXLAN Overview and Terminology – Control Plane

- MP-BGP EVPN is the Standardized Control Plane
- MP-BGP EVPN used to distributed NLRI (EVPN Routes = MAC + IP)
- Integrated Bridging and Routing (IRB)
 - Cisco follows Symmetric IRB (becoming the standard IRB method)
- Leafs = VTEPs (Virtual Tunnel Endpoints)
 - Perform the VXLAN encap/decap function
 - Uses logical interface called an NVE interface
 - VLANs are mapped to VNIs
- NVE Interface uses loopback as TEP address
 - Network Virtualized Endpoint
 - Loopback advertised into Underlay for reachability


VXLAN Overview and Terminology – Overlays

- L2 VNI used for VXLAN Bridging
- L3 VNI used for VXLAN Routing
 - Special VNI used as the Tenant VRF identifier
 - Ingress VTEP routes to Layer-3 VNI
 - Egress VTEP routes to destination Layer-2 VNI
- Multi-tenancy using VRF construct
 - Extended Community string in MP-BGP EVPN
 - Route-Distinguishers/RD – VNI+BGP Router-ID
 - Route-Targets/RT – BGP AS:VNI
- MP-BGP EVPN allows for use of Anycast Gateway
- MP-BGP EVPN makes use of ARP Suppression

VXLAN – Simple Packet Flow




VXLAN – Life of a Packet




VXLAN – BUM Traffic

- IP Multicast is used for Broadcast, Unknown Unicast, & Multicast (BUM) Traffic from Endpoints and VM's



VTS and OpenStack

OpenStack Integration



Controller/Neutron Node:

- OpenStack Neutron ML2 Plugin
- Sync Agent (neutron-vts-sync-agent)
 - Keeps VTS in-sync with OpenStack's Neutron database

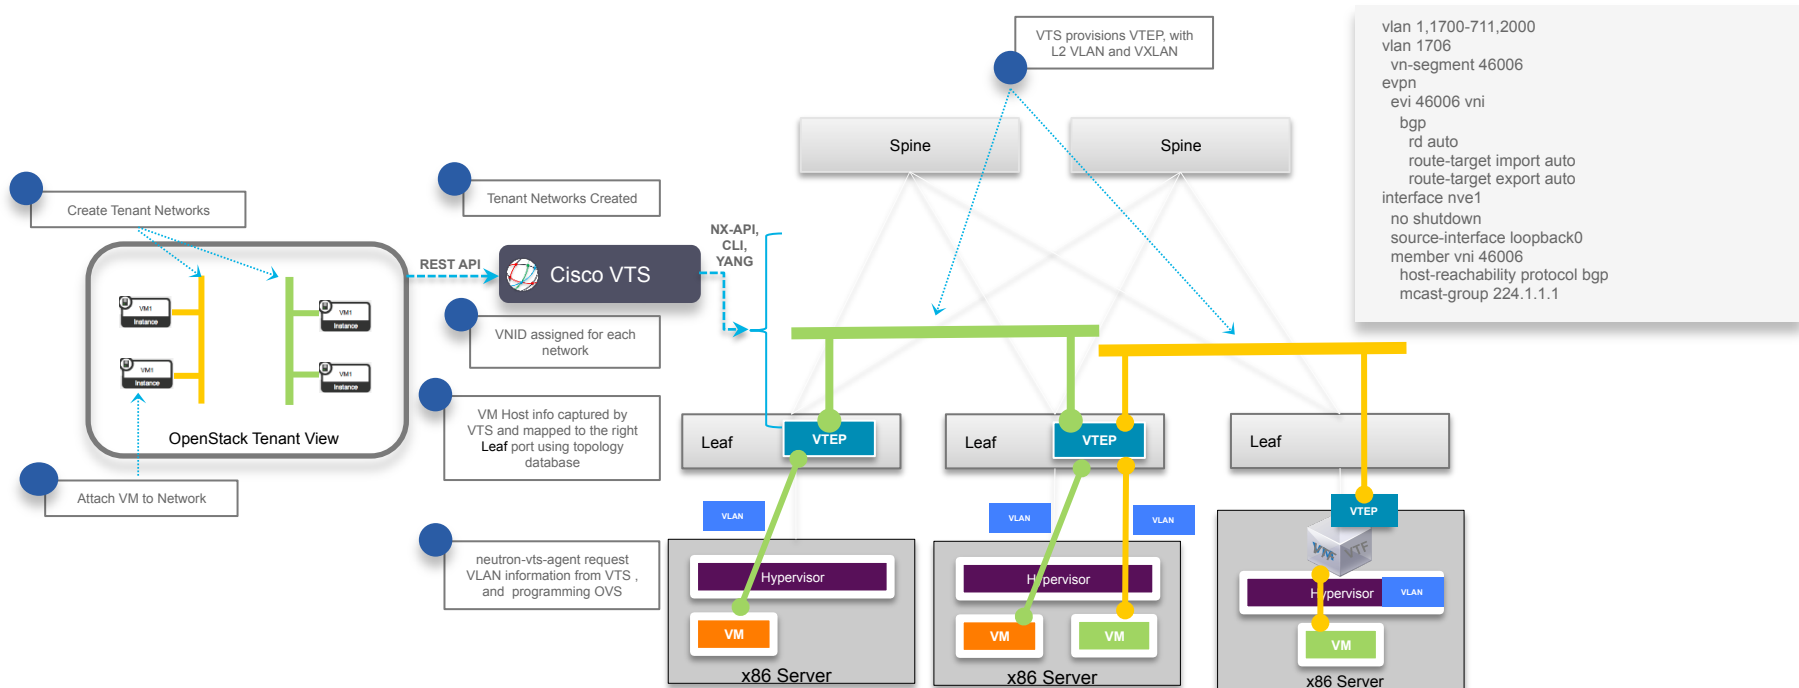
Compute Node:

- Host Agent (neutron-vts-agent)
 - Replace neutron-openvswitch-agent
 - Listens for port create/delete/update events from OpenStack
 - Queries VTS for port VLAN information, and programs OVS

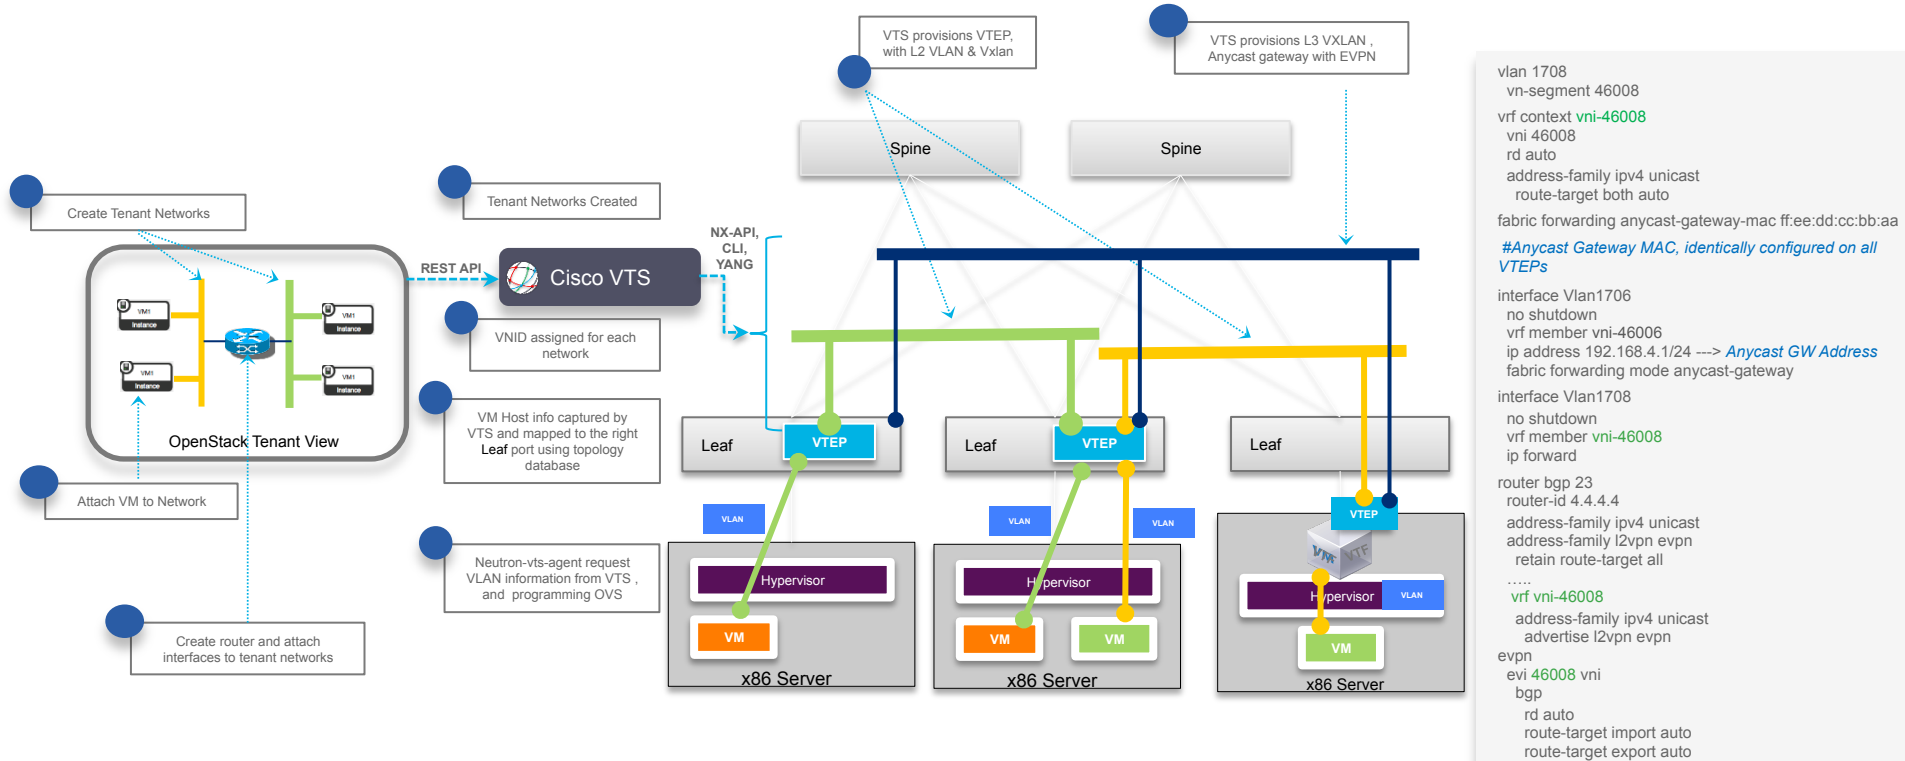
Support:

- Juno, Kilo, & Liberty

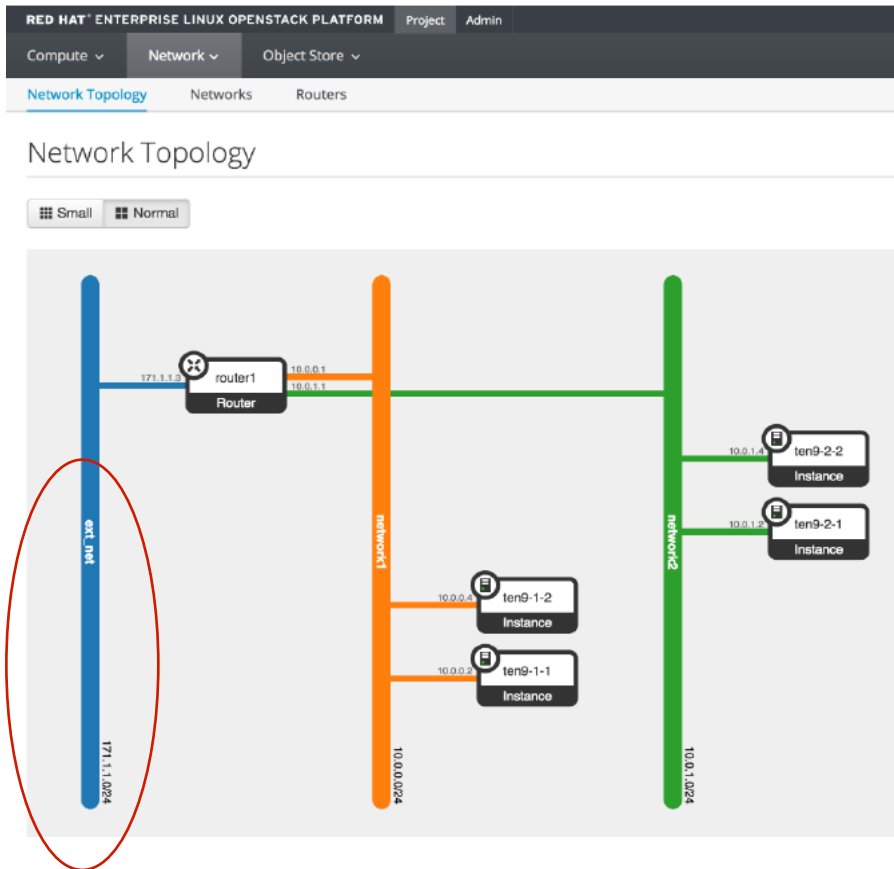
VTS: OpenStack Workflow – Layer 2 Networks




VTS: OpenStack Workflow – Layer 3 Networks



VTS: OpenStack Workflow – External Networks



Border-Leaf to DCI Connectivity



OpenStack Configuration

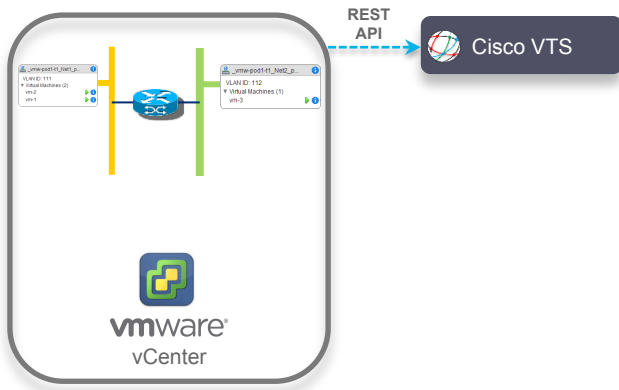
1. Create an External Network in OpenStack
2. Attach a Subnet to the External Network
3. Set External Network as Gateway for Router

VTS will provision Border Leaf and DCI router

1. VRF on the L3 interface (physical/subinterface/VLAN)
2. BGP peer between Border Leaf and DCI
3. "default originate" under BGP peer on DCI

VTS and VMware

VMware Integration



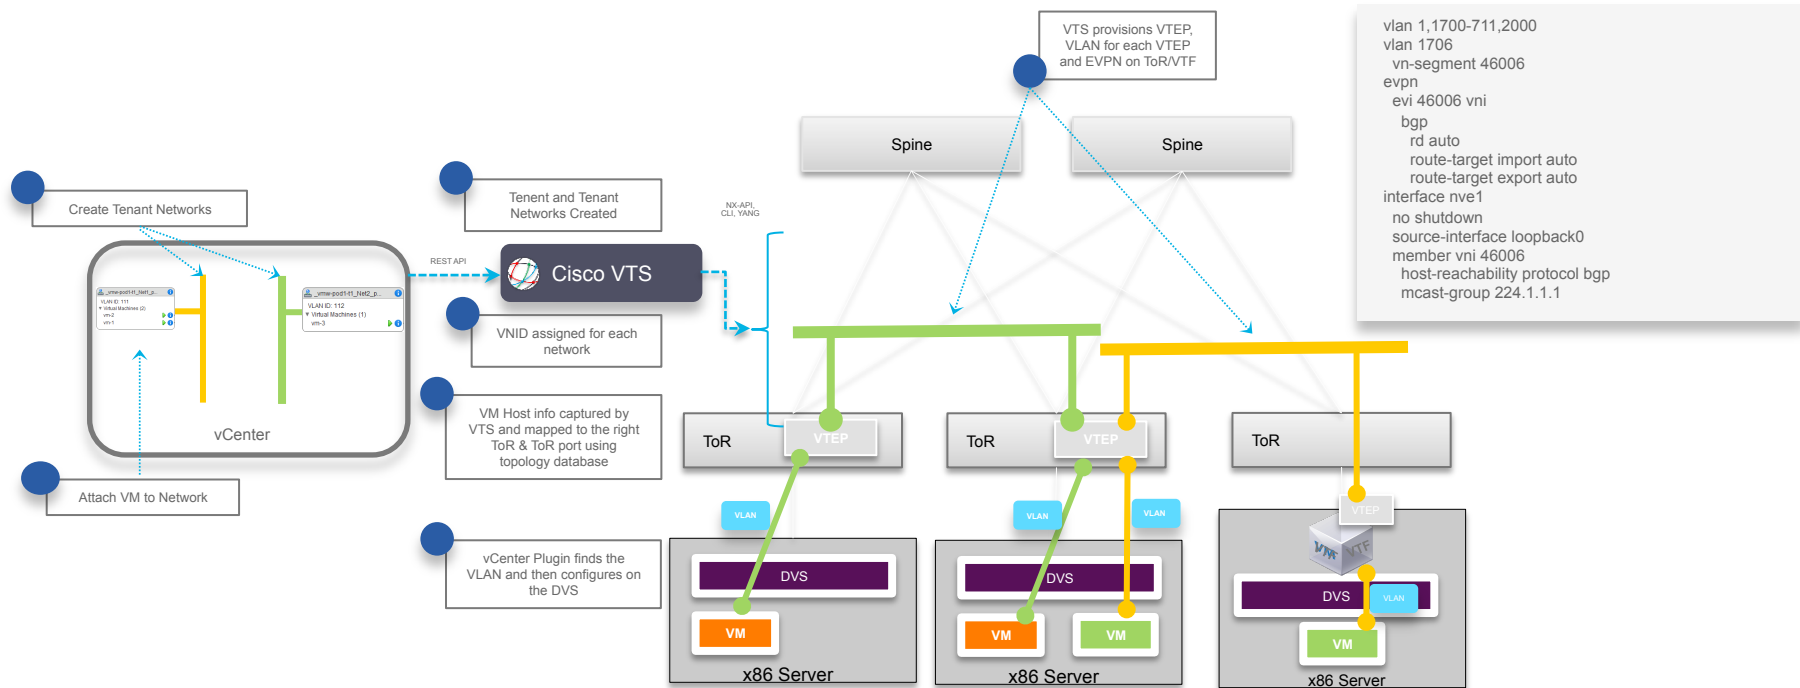
vCenter

- vCenter Extension Plugin (MOB)
- Extension to Distributed Virtual Switch
 - VTS tab added to vDS Manage navigation bar
 - Used to create networks that equate to port-groups

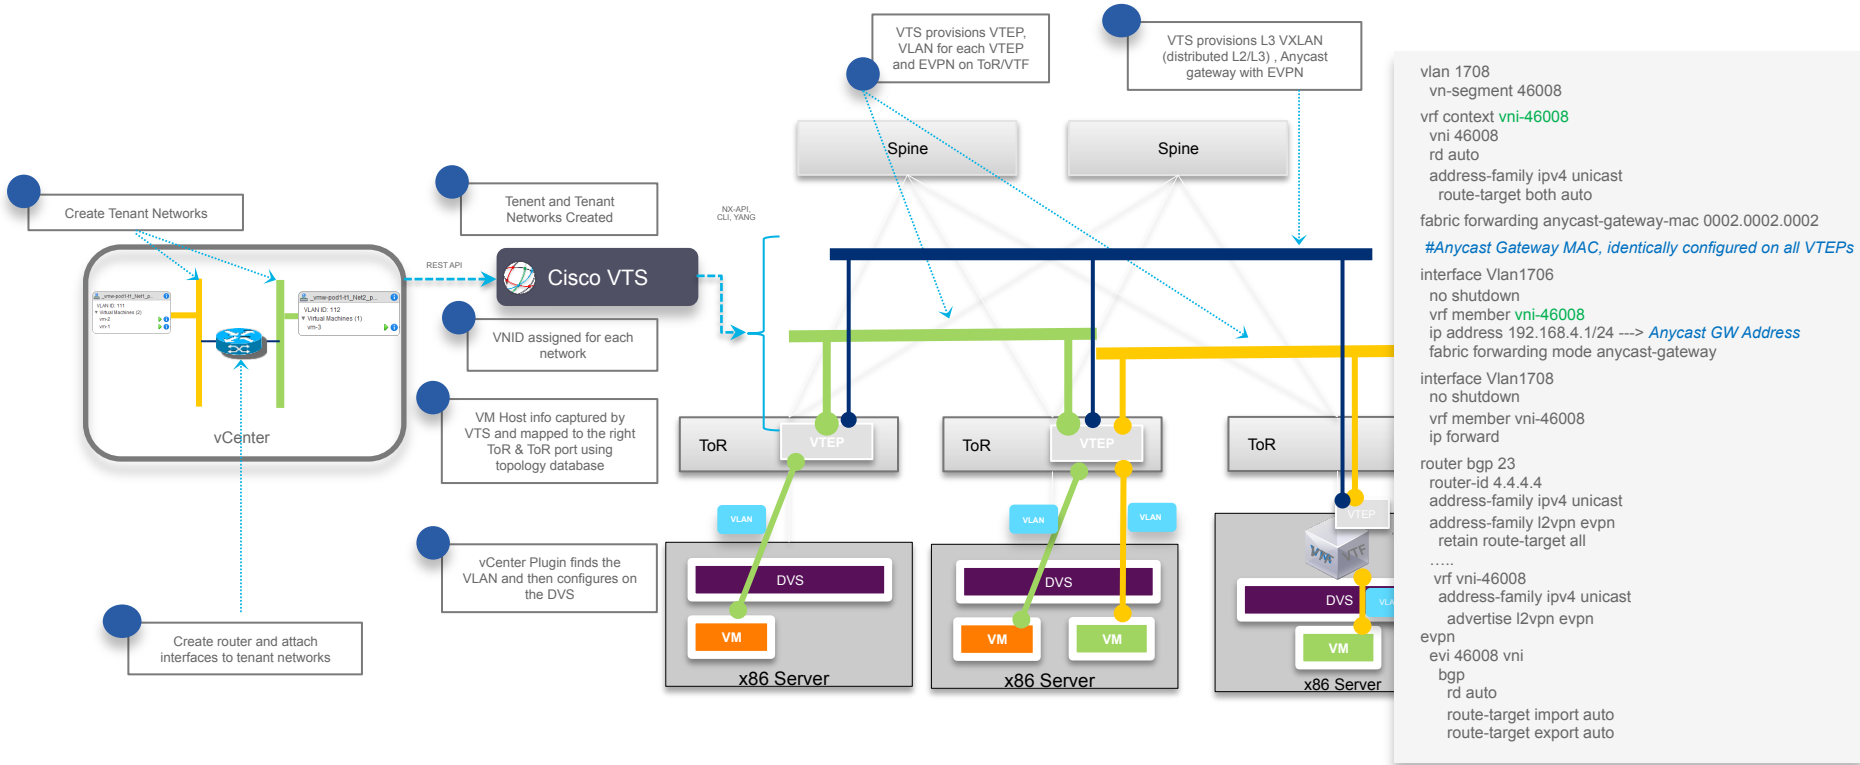
Support:

- Windows vCenter 5.5 & vCenter Server Appliance 6.0

VTS: VMware Workflow – Layer 2 Networks



VTS: VMware Workflow – Layer 3 Networks



VTS: VMware Workflow – External Networks

Router Details

Router: router

Networks: Net1


Subnets: Subnet1

External Networks: ExtNet

👉

Subnet	Tenant Name	Network Range	Gateway IP	Detach
Subnet2	vmw-pod1-t1	172.17.1.0/24	172.17.1.1	<input type="button" value="Detach"/>
Subnet1	vmw-pod1-t1	172.16.1.0/24	172.16.1.1	<input type="button" value="Detach"/>

Border-Leaf to DCI Connectivity



VMware Configuration

1. Create an External Network on vCenter vDS plugin
2. Attach the External Network to the Router

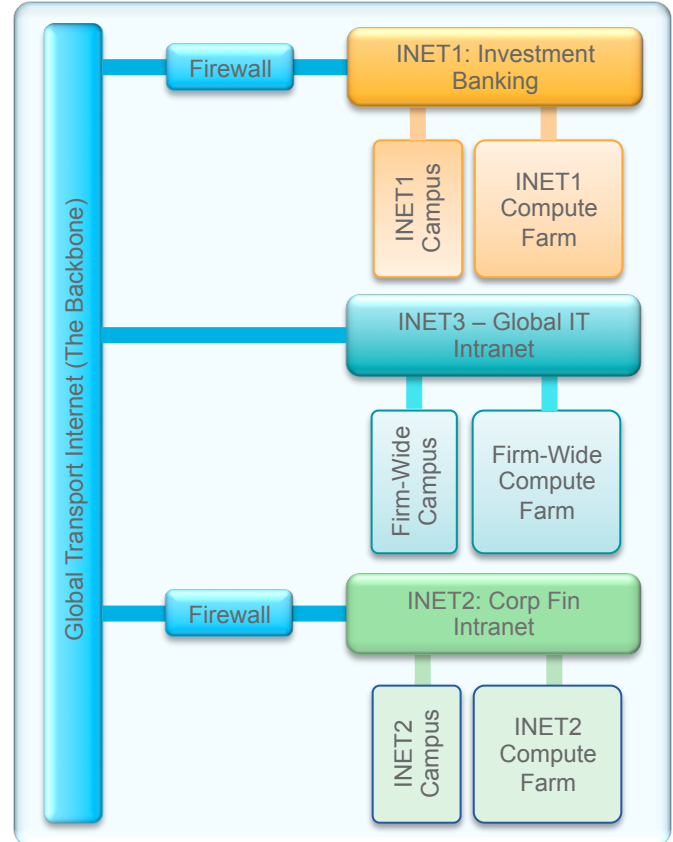
VTS will provision Border Leaf and DCI router

1. VRF on the L3 interface (physical/subinterface/VLAN)
2. BGP peer between Border Leaf and DCI
3. “default originate” under BGP peer on DCI

Use cases

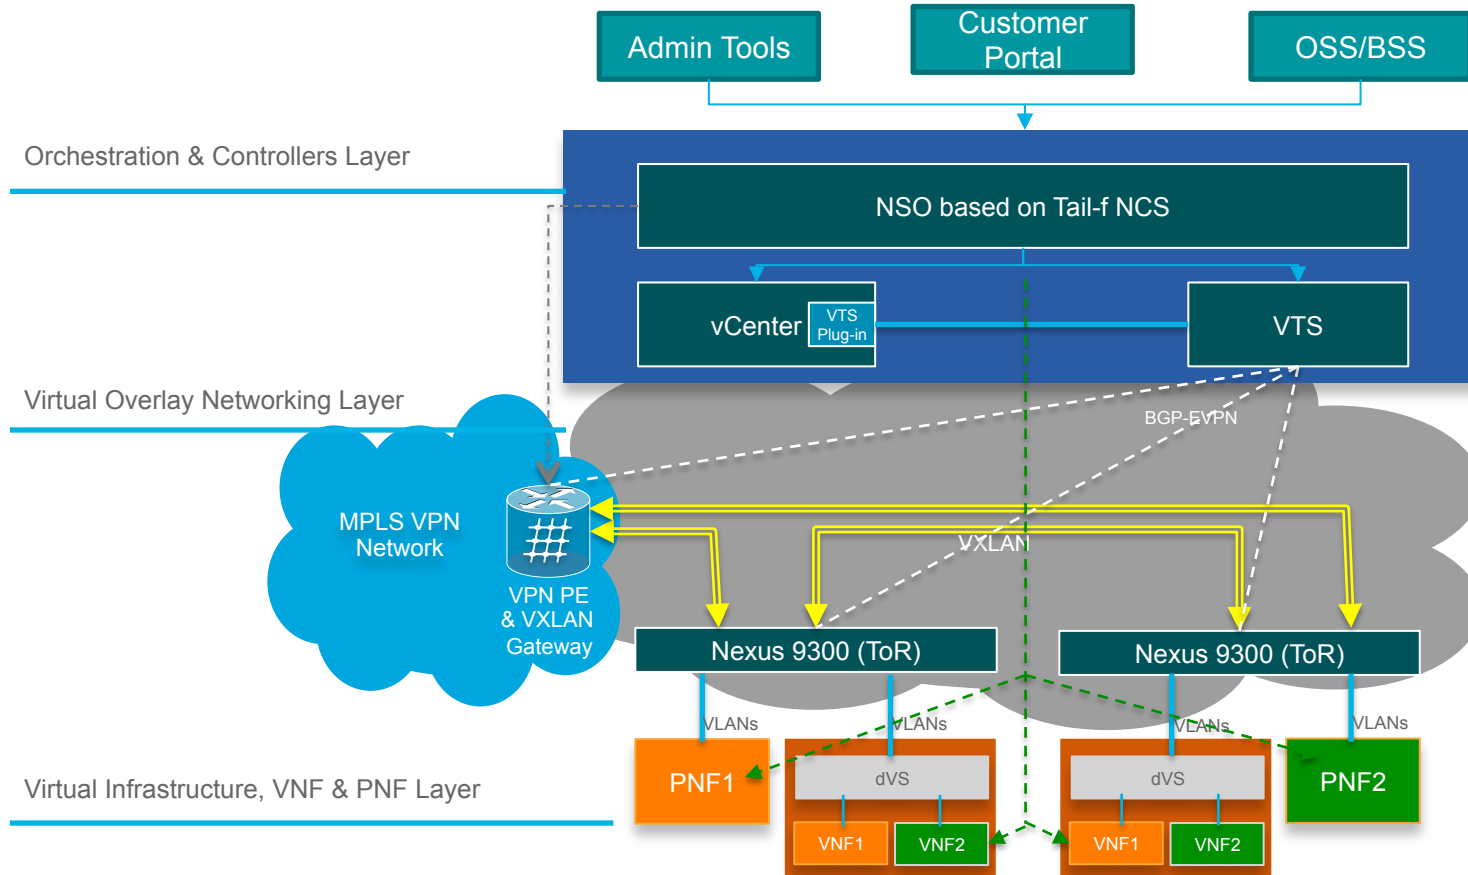
Multi-tenant cloud use case

- In FSI, customers have used the separated or isolated infrastructures for the sake of compliances.
- For example, a large US FSI has the following three infrastructures:
 - Investment Banking Intranet (INET1)
 - Corporate Financing Intranet (INET2)
 - Global IT Intranet (INET3)
- The hardware resource waste resulted from three sets of infrastructures is horrendous.
- Leveraging Cisco's software overlay SDN will allow them to tap into the unused compute capacity and avoid hardware acquisition.



NFV Use Case

NSO: Network Service Orchestrator
VTS: Virtual Topology System
VTS: Virtual Topology System
ToR: Top of Rack switch
PNF: Physical Network Function
VNF: Virtual Network Function
dVS: distributed Virtual Switch
PE: Provider Edge



VTS Roadmap

Q1 2016 2.2	Q2CY2016 2.3	Q3CY2016 2.4	Q4CY2016 2.5	1HCY2017
<ul style="list-style-type: none"> • Static and dynamic allocation of VLAN resources • DCI/WAN integration enhancements Open Stack Kilo • Static Route provisioning • vCenter plugin enhancements • VTS Upgrade 	<ul style="list-style-type: none"> • VTF vhost_user for higher VTF performance Bonded Interface • VTS GUI based Static/Dynamic router extensions • ESI based multi-homing • NFVI integration • Static Route Target • Open Stack Liberty support –Nexus 9200 Support. • L2VNI IPV6 support • User specified VRF naming. 	<ul style="list-style-type: none"> • Resource Pool Enhancements /Static VNI • Multiple VNI to multicast address mapping. • L2VNI-DCI integration with ASR9k • IPv6 routing P2P • VTS GUI based External Router extensions • Route-Target GUI enhancements • Nexus 9300-EX Support. • Collectd on VTF. • Dual Interfaces on VTF 	<ul style="list-style-type: none"> • AAA Support • VTS GUI based Underlay extensions • VTS GUI based Network Extensions for ACL/QoS • Multi-VMM • Multicast/Ingress Replication in VTF • Overlay Network visibility • NED retries • EVPN VRF to Global route leaking 	<ul style="list-style-type: none"> • Full/Port-local VLAN significance • Service assurance: Overlay statistics & log collection • Security Policies (Security Groups, GBP) • vRealize support • Tenant based RBAC • L3 Service Chaining • DCNM integration • N3k ToR/flood&learn • MPLS, Segment routing. • L4-L7 Service Integration • NAT/ACLs/QoS on VTF • Backup/Restore/ISSU • Third Party device Support • IPAM Integration • 3rd Party VNF integration. • FW in VTF • Containers • vMotion across DVS • N7k as Leaf • N1Kv integration. • Centralized NAT • Infra OAM for install and upgrade

▪ *Roadmap is subject to change and is not in priority order*

Demo

- VTS Web UI demo
- OpenStack+ VTS Working Case Demo

More Information On CCO

- Cisco VTS White Paper:
 - <http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/virtual-topology-system/white-paper-c11-734904.html>
- VTS Data Sheet:
 - <http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/virtual-topology-system/datasheet-c78-734877.html>
- VTS 2.3 User Guide:
 - http://www.cisco.com/c/en/us/td/docs/net_mgmt/virtual_topology_system/2_3/user_guide/Cisco_VTS_2_3_User_Guide.html
- VTS 2.3 Installation Guide:
 - http://www.cisco.com/c/en/us/td/docs/net_mgmt/virtual_topology_system/2_3/install_guide/Cisco_VTS_2_3_Installation_Guide.html
- Cisco VTS 2.3 Day Zero Configuration Examples:
 - http://www.cisco.com/c/en/us/td/docs/net_mgmt/virtual_topology_system/2_3/day_zero/VTS_Day_Zero_Configuration.html

Thank you



CISCO

TOMORROW starts here.