




Nexus 1000V on OpenSource Hypervisors

Han Yang
Product Manager, Data Center Group

May, 2013



Cisco ONE/N1K Public Webinars, 1H 2013

| Date/Time | Topic |
|-----------------------------|---|
| Thur, Feb 21st at 0900 PST | Cisco Open Network Environment (Cisco ONE) – Next Phase of Network Programmability and SDN |
| Thur, Feb 28th at 0900 PST | Cisco One Platform Kit (onePK): Technical Deep Dive and key use cases |
| Wed, Mar 6th at 0900 PST | Nexus 1000V for Hyper-V with Microsoft SCVMM integration |
| Wed, Mar 13th at 0900 PST | Cisco ONE Controller: Technical Deep Dive and key use cases |
| Wed, Mar 20th at 0900 PST | 5000 Seat VDI Reference Architecture: Cisco UCS & Nexus 1000V, Citrix XenDesktop, and EMC VNX |
| Wed, Mar 27th at 0900 PST | Nexus 1000V v2.2 for vSphere: More scale, Multicast-less VXLAN, and VXLAN Gateway |
| Wed, April 3rd at 0900 PST | Cloud Services Router (CSR 1000V) - technical deep dive and key use cases |
| Wed, April 10th at 0900 PST | Cloud Security with ASA 1000V and Virtual Security Gateway v2.1 (VSG) |
| Wed, April 17th at 0900 PST | Secure Hybrid Cloud solution with Nexus 1000V InterCloud & VNMC InterCloud |
| Wed, April 24th at 0900 PST | Nexus 1100 Series Cloud Services Platform: new services & ecosystem including VXLAN-to-VLAN GW and Imperva's SecureSphere WAF |
| Wed, May 1st at 0900 PST | Cloud Networking Services: vWAAS and vNAM |
| Wed, May 8th at 0900 PST | Virtualized Multiservice Data Center (VMDC) with Virtualized Services |
| Wed, May 15th at 0900 PST | Nexus 1000V for KVM (with OpenStack and VXLAN) |

Register and view recordings/presentations here:

www.cisco.com/go/1000vcommunity

Physical | Virtual | Cloud Journey


PHYSICAL WORKLOAD

- One app per Server
- Static
- Manual provisioning




VIRTUAL WORKLOAD

- Many apps per Server
- Mobile
- Dynamic provisioning



CLOUD WORKLOAD

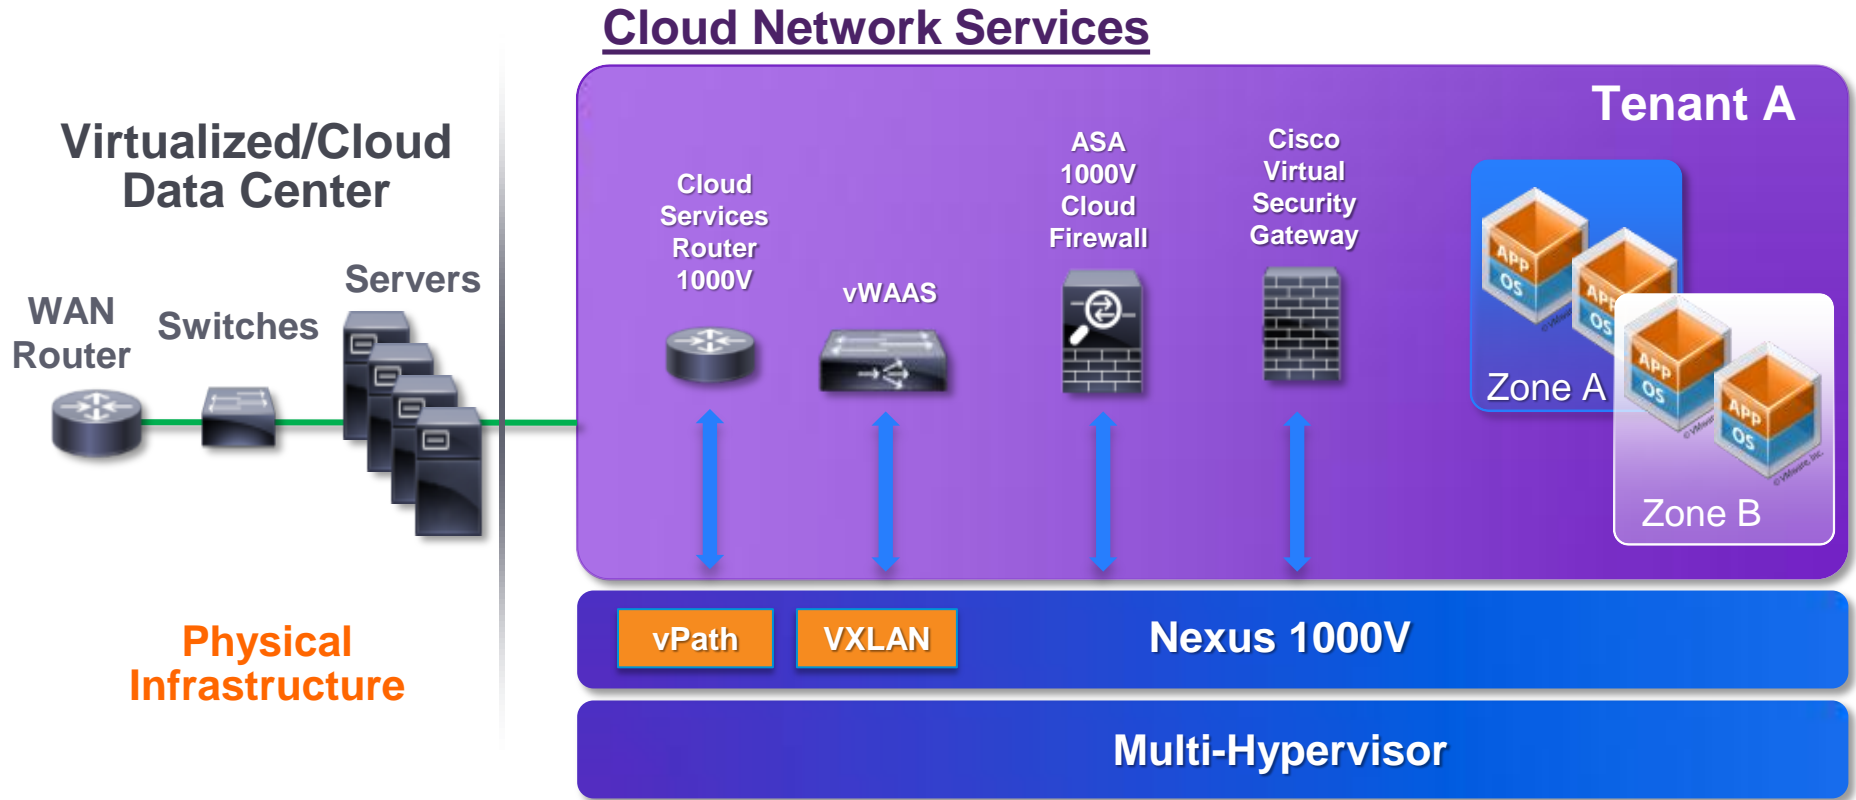
- Multi-tenant per Server
- Elastic
- Automated Scaling



CONSISTENCY: Policy, Features, Security, Management, Separation of Duties

| | | |
|------------------|--------------------|-----------------------------------|
| Switching | Nexus 7K/5K/3K/2K | Nexus 1000V, VM-FEX |
| Routing | ASR, ISR | Cloud Services Router (CSR 1000V) |
| Services | WAAS, ASA, NAM | vWAAS, VSG, ASA 1000V, vNAM** |
| Compute | UCS for Bare Metal | UCS for Virtualized Workloads |

Cisco Virtual Networking and Cloud Network Services



Nexus 1000V

- Distributed switch
- NX-OS consistency

VSG

- VM-level controls
- Zone-based FW

ASA 1000V

- Edge firewall, VPN
- Protocol Inspection

vWAAS

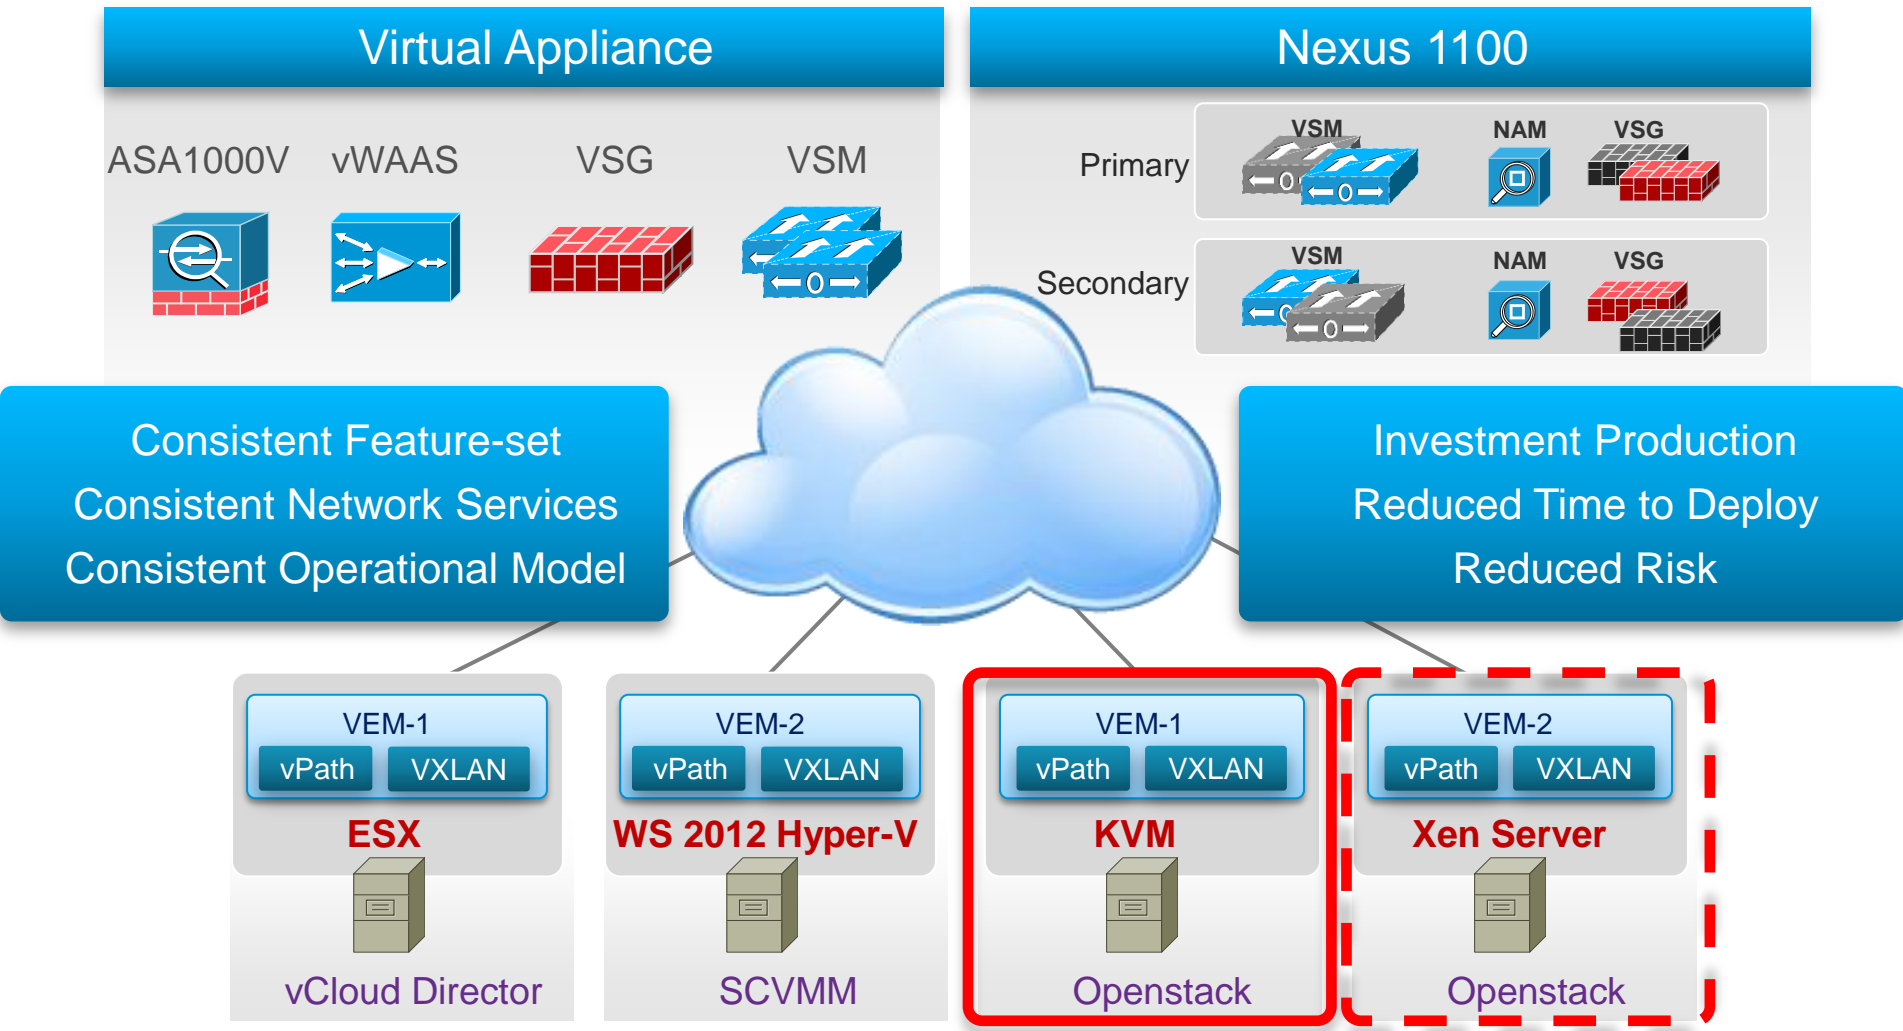
- WAN optimization
- Application traffic

CSR 1000V (Cloud Router)

- WAN L3 gateway
- Routing and VPN

Cisco Nexus 1000V

Multi-service, multi-hypervisor, multi-cloud platform



New Nexus 1000V Freemium Go-To-Market Model

No-Cost Version

\$695 per CPU MSRP

Nexus 1000V Essential Edition

The world's most advanced virtual switch

- Full Layer-2 Feature Set
- Security, QoS Policies
- VXLAN virtual overlays
- Full monitoring and management capabilities
- vPath enabled Virtual Services

Nexus 1000V Advanced Edition

Adds Cisco value-add features for DC and Cloud

- All Feature of Essential Edition
- VSG firewall bundled (previously sold separately)
- Support for Cisco TrustSec SGA policies
- Platform for other Cisco DC Extensions in the Future

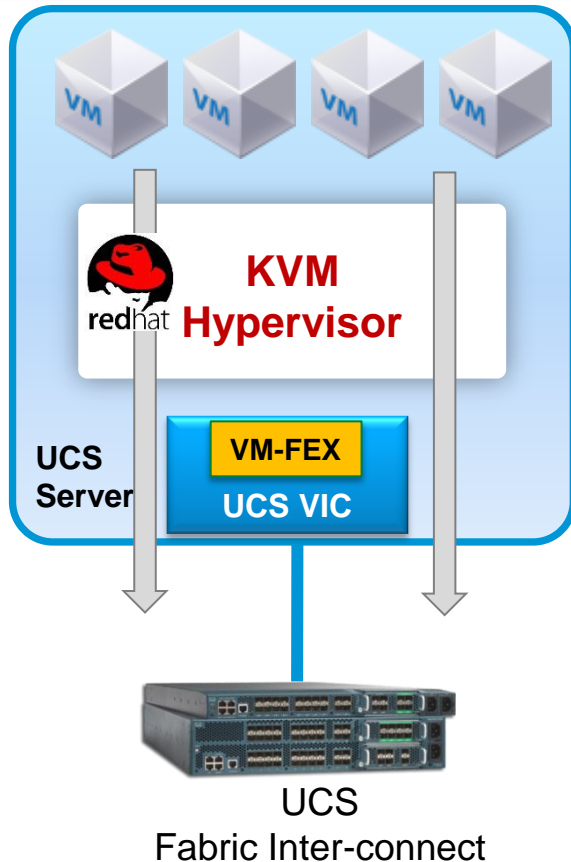
Freemium Pricing Model Offers Flexibility for Customers to Deploy Cisco Virtual Data Center




Nexus 1000V on KVM

Cisco VM Networking on KVM

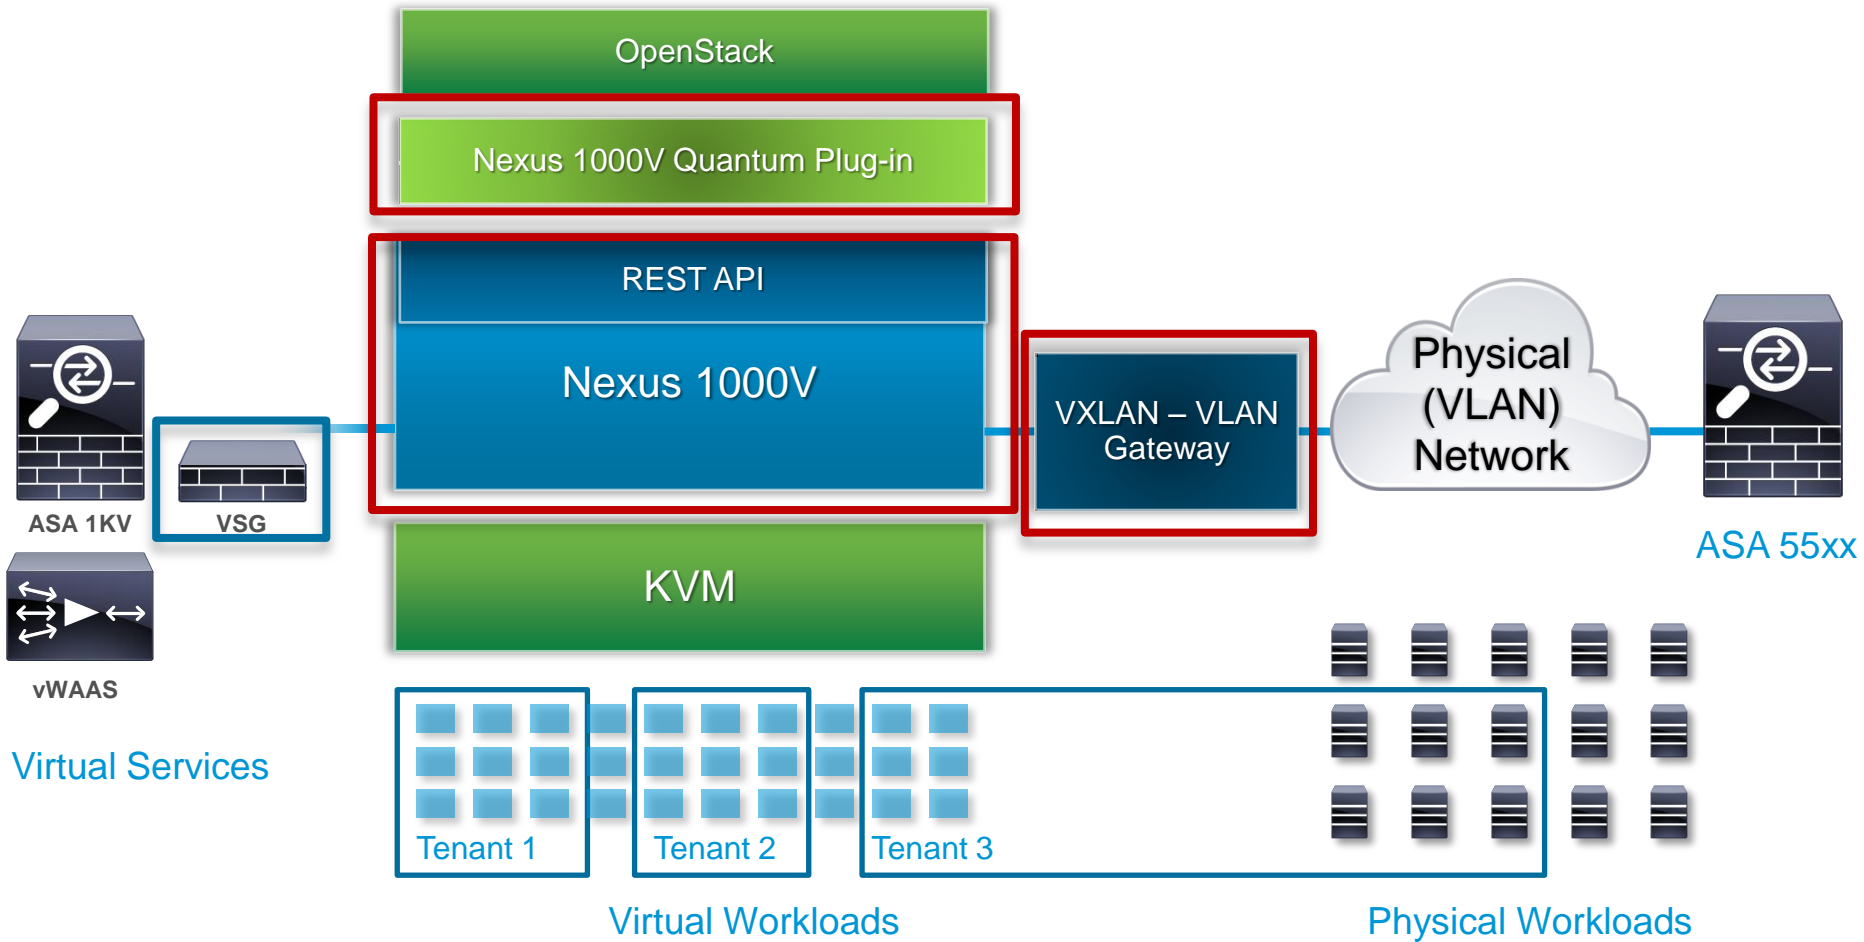
Bring VM awareness to physical network
(Cisco UCS VM-FEX)



Bring network to the hypervisor
(Cisco Nexus 1000V Switch)



Nexus 1000V on KVM Architecture




Nexus 1000V on KVM


- Expand Nexus 1000V hypervisor support to KVM
- VXLAN Support
 - Without IP multicast
 - Ease VXLAN deployment
- Highly Scalable
 - 16-32k ports on 256-512 servers with 300+ veth/server
- Initially with limited features when compared to Nexus 1000V on vSphere
 - Targeting service providers who need scale
- REST API

Cisco Nexus 1000 Portfolio

Virtual Appliance



Nexus 1010



- VSM:** Virtual Supervisor Module
- VEM:** Virtual Ethernet Module
- vPath:** Virtual Service Data-path
- VXLAN:** Scalable Segmentation
- VSG:** Virtual Security Gateway
- vWAAS:** Virtual WAAS
- ASA 1000V:** Tenant-edge security

Virtual Service Blades

- Virtual Supervisor Module (VSM)
- Network Analysis Module (NAM)
- Virtual Security Gateway (VSG)
- Data Center Network Manager (DCNM)



vPath

- Service Binding (Traffic Steering)
- Fast-Path Offload



VXLAN

- 16M address space for LAN segments
- Network Virtualization (Mac-over-UDP)

Increasing Scale

Rearchitected for Resiliency & Scale

- Migrating workload from VSM to VEM
- Loosely Coupled between VSM and VEM
- vMotion even when VSM is disconnected from VEM

Target Scale

- Veths/VSM: 16-32k
- VEMs/VSM: 256-512
- Veths/VEM: 300+
- Active VLANs: 4,096
- Active VXLANs: 16,384

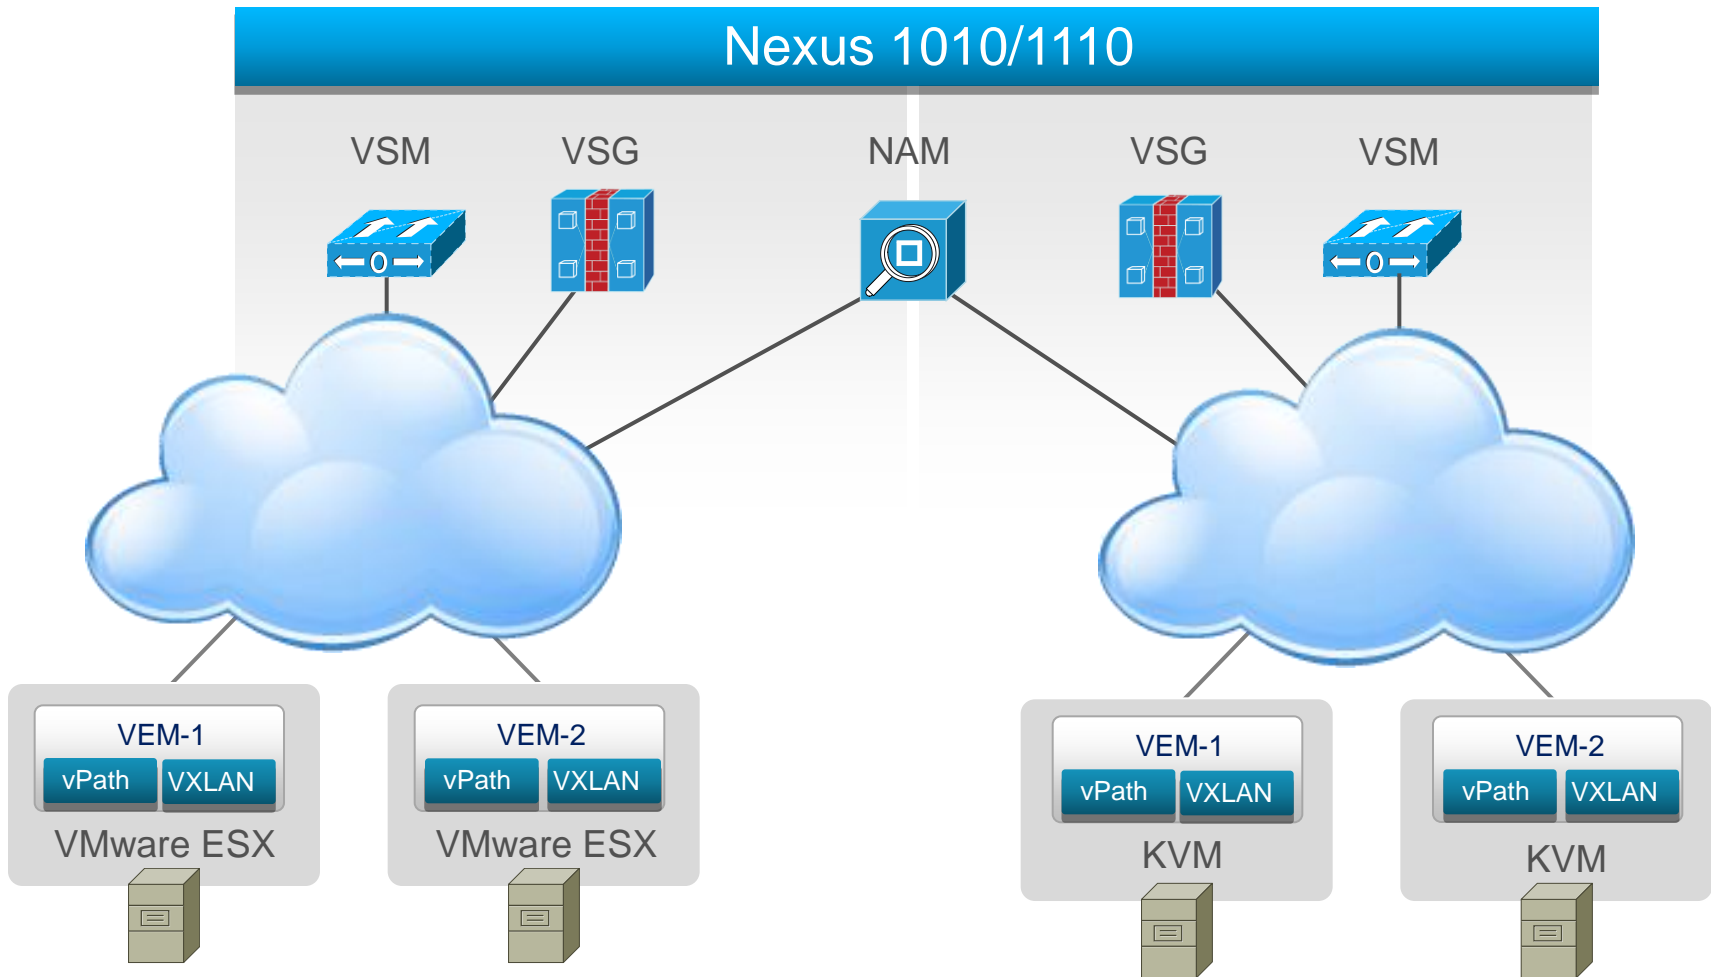
Cisco Nexus 1000V REST API

- vm-network-definition – info about VMND (id, vlan, ip-pool)
- fabric-network-definition- info about LN (name, id, connected-ports)
- virtual-port-profile – info about vEth (type, id, maxports, switch-id)
- uplink-port-profile – info about PNIC (state, type, id, maxports, switch-id)
- ip-address-pool – info about ip-pool (name, dhcp-server, range etc.)

For Use with Quantum Plugin

Cisco Nexus 1010/1110

Consistent Architecture across Hypervisors



Existing Nexus 1010 will support multi-hypervisor environment



VXLAN

Virtual Extensible Local Area Network (VXLAN)

- Ethernet in IP overlay network
 - Entire L2 frame encapsulated in UDP
 - 50 bytes of overhead
- Include 24 bit VXLAN Identifier
 - 16 M logical networks
 - Mapped into local bridge domains
- VXLAN can cross Layer 3
- Tunnel between VEMs
 - VMs do NOT see VXLAN ID
- IP multicast used for L2 broadcast/multicast, unknown unicast
- Technology submitted to IETF for standardization
 - With VMware, Citrix, Red Hat, and Others

← Ethernet Frame →



← VXLAN Encapsulation →

Enhanced VXLAN

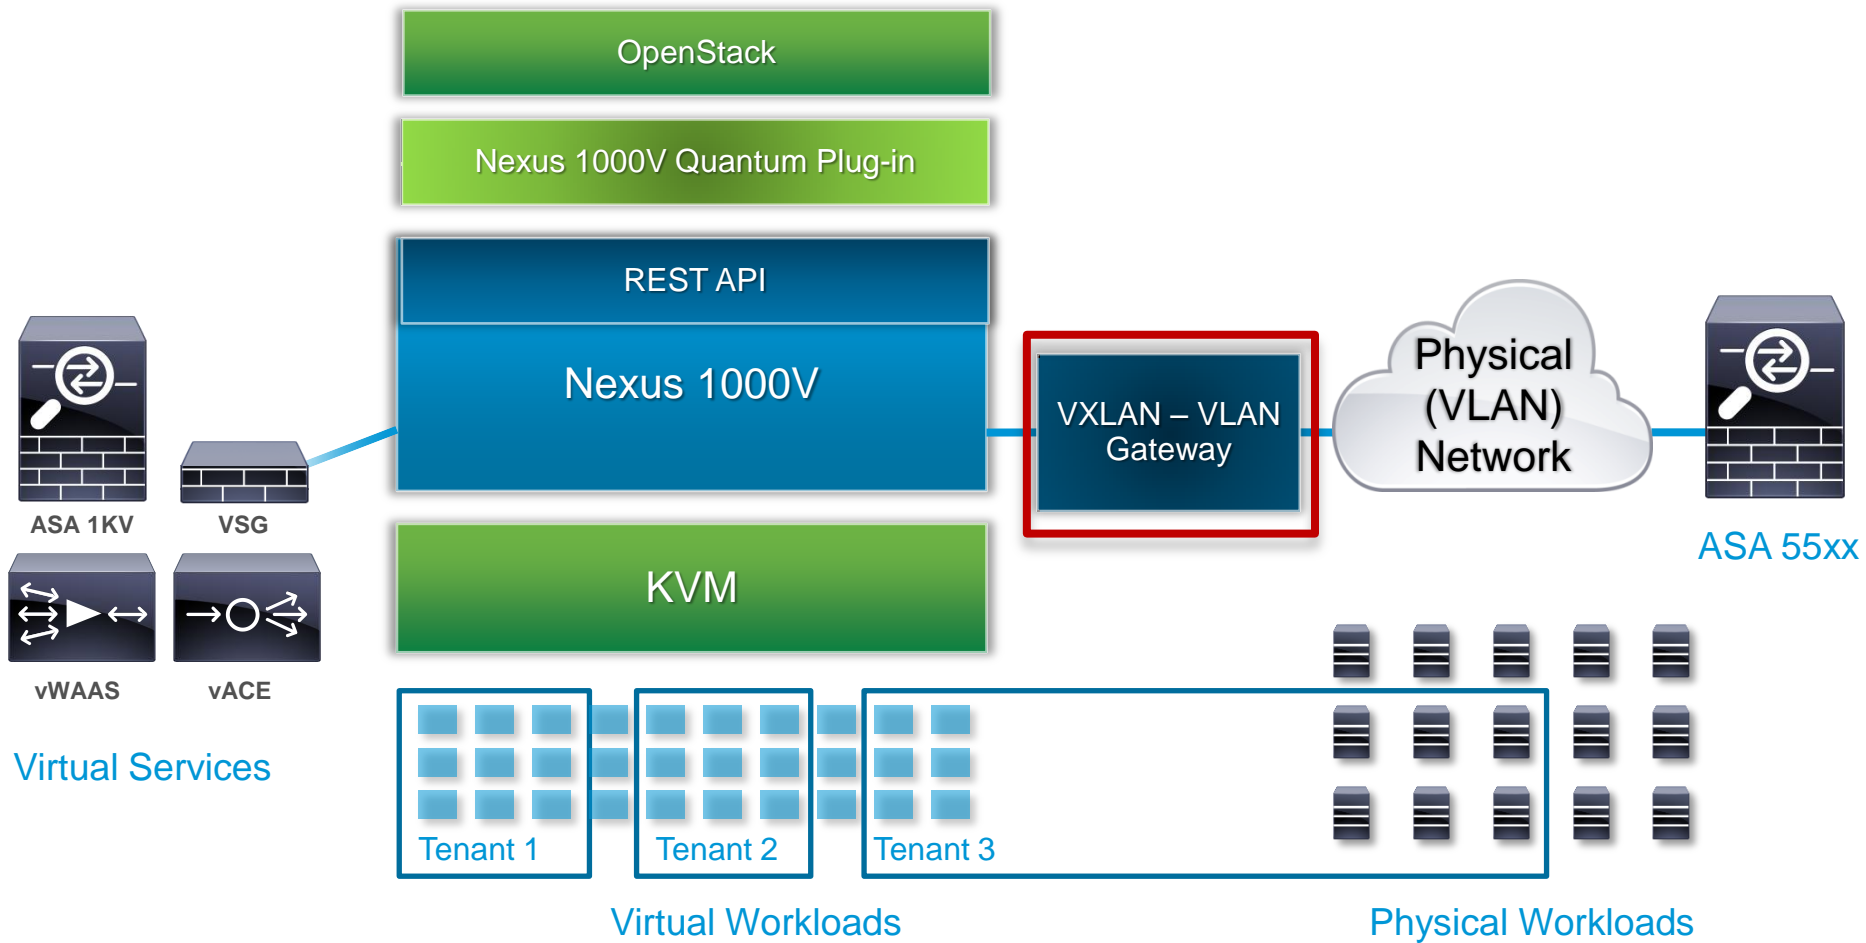
Current VXLAN

- Requires IP Multicast
- Requires large number of multicast groups
 - Ideally 1 multicast group for each VXLAN
- Multicast configuration
- Difficult to use multicast between data centers

Enhanced VXLAN

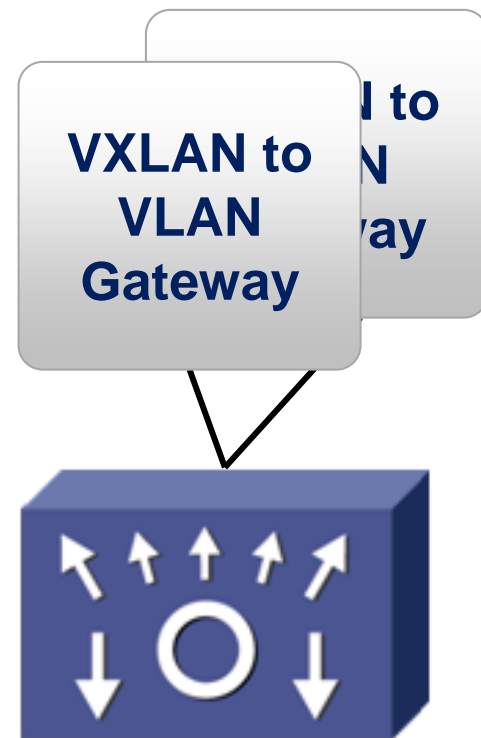
- VSM to know location of all VXLAN
- Head End Replication
- Restricted to 1 VSM active/standby pair (up to 500+ servers)

VXLAN to VLAN Gateway



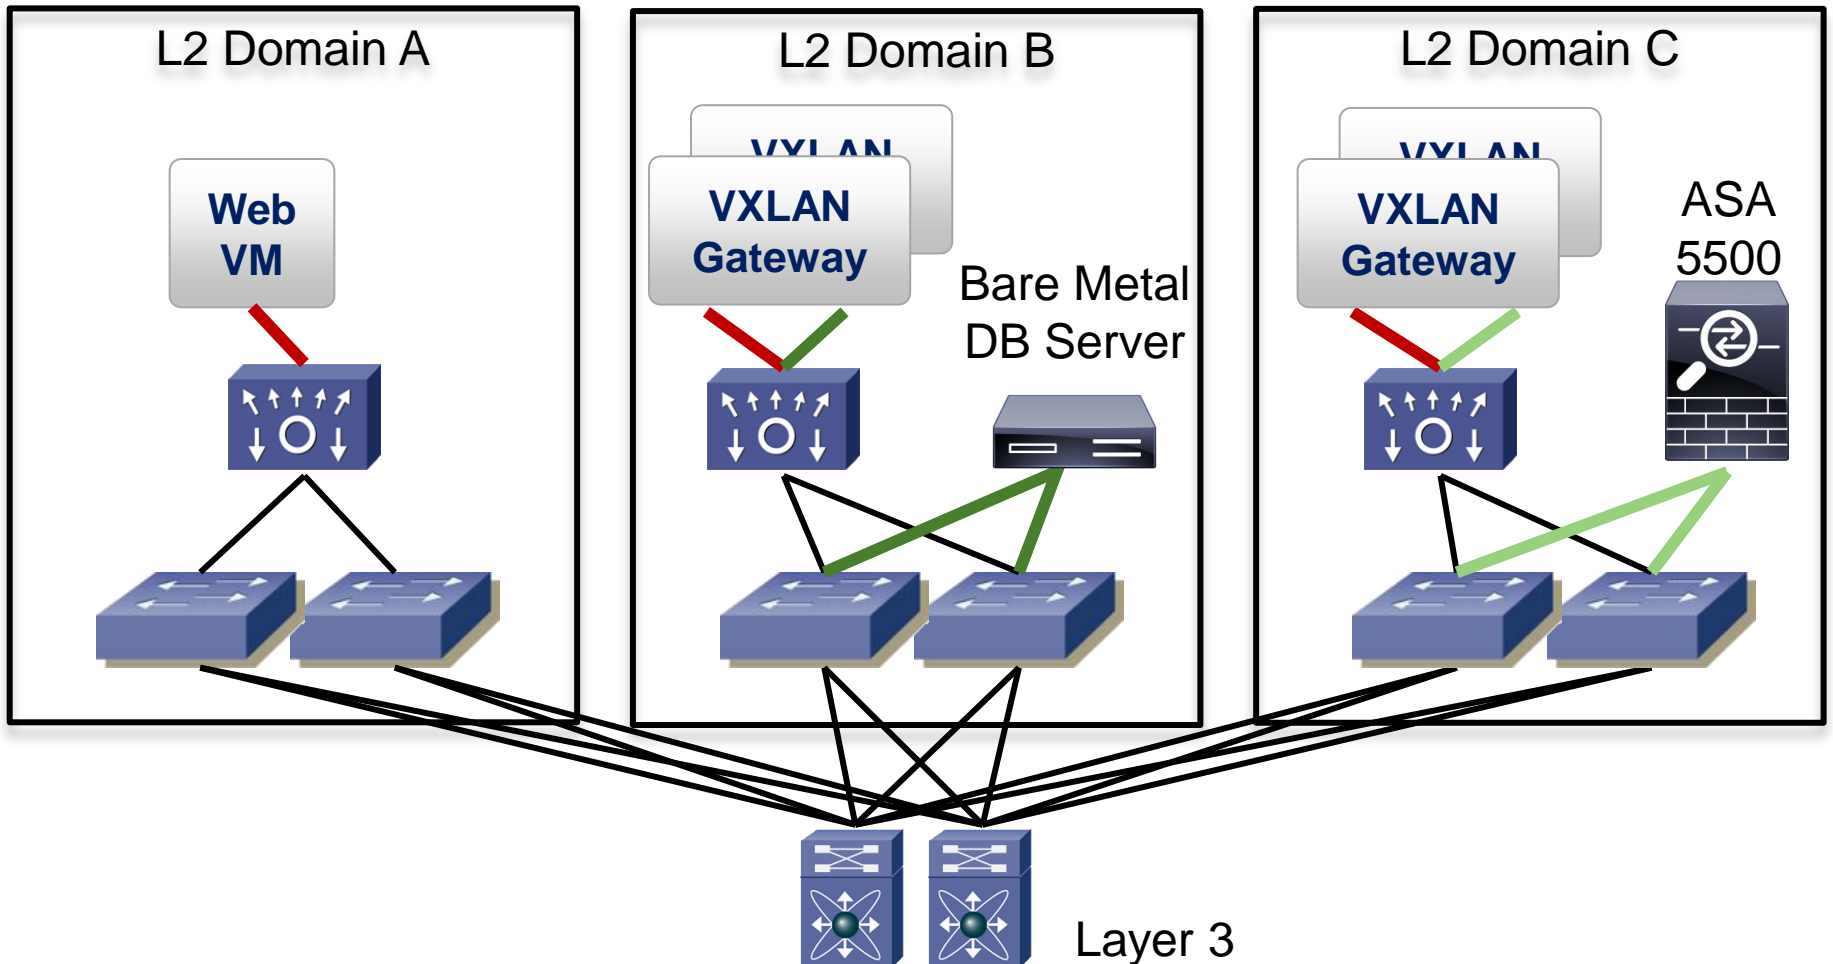
VXLAN to VLAN Gateway

- Hosted on local hypervisor as virtual machine connected to Virtual Ethernet Module
- Managed as a module from VSM
- Active/Standby VXLAN Gateway
- Integrated with OpenStack
- Scale:
 - 4 VXLAN Gateway per VSM
 - 2k Active VXLAN
 - 2k Active VLAN

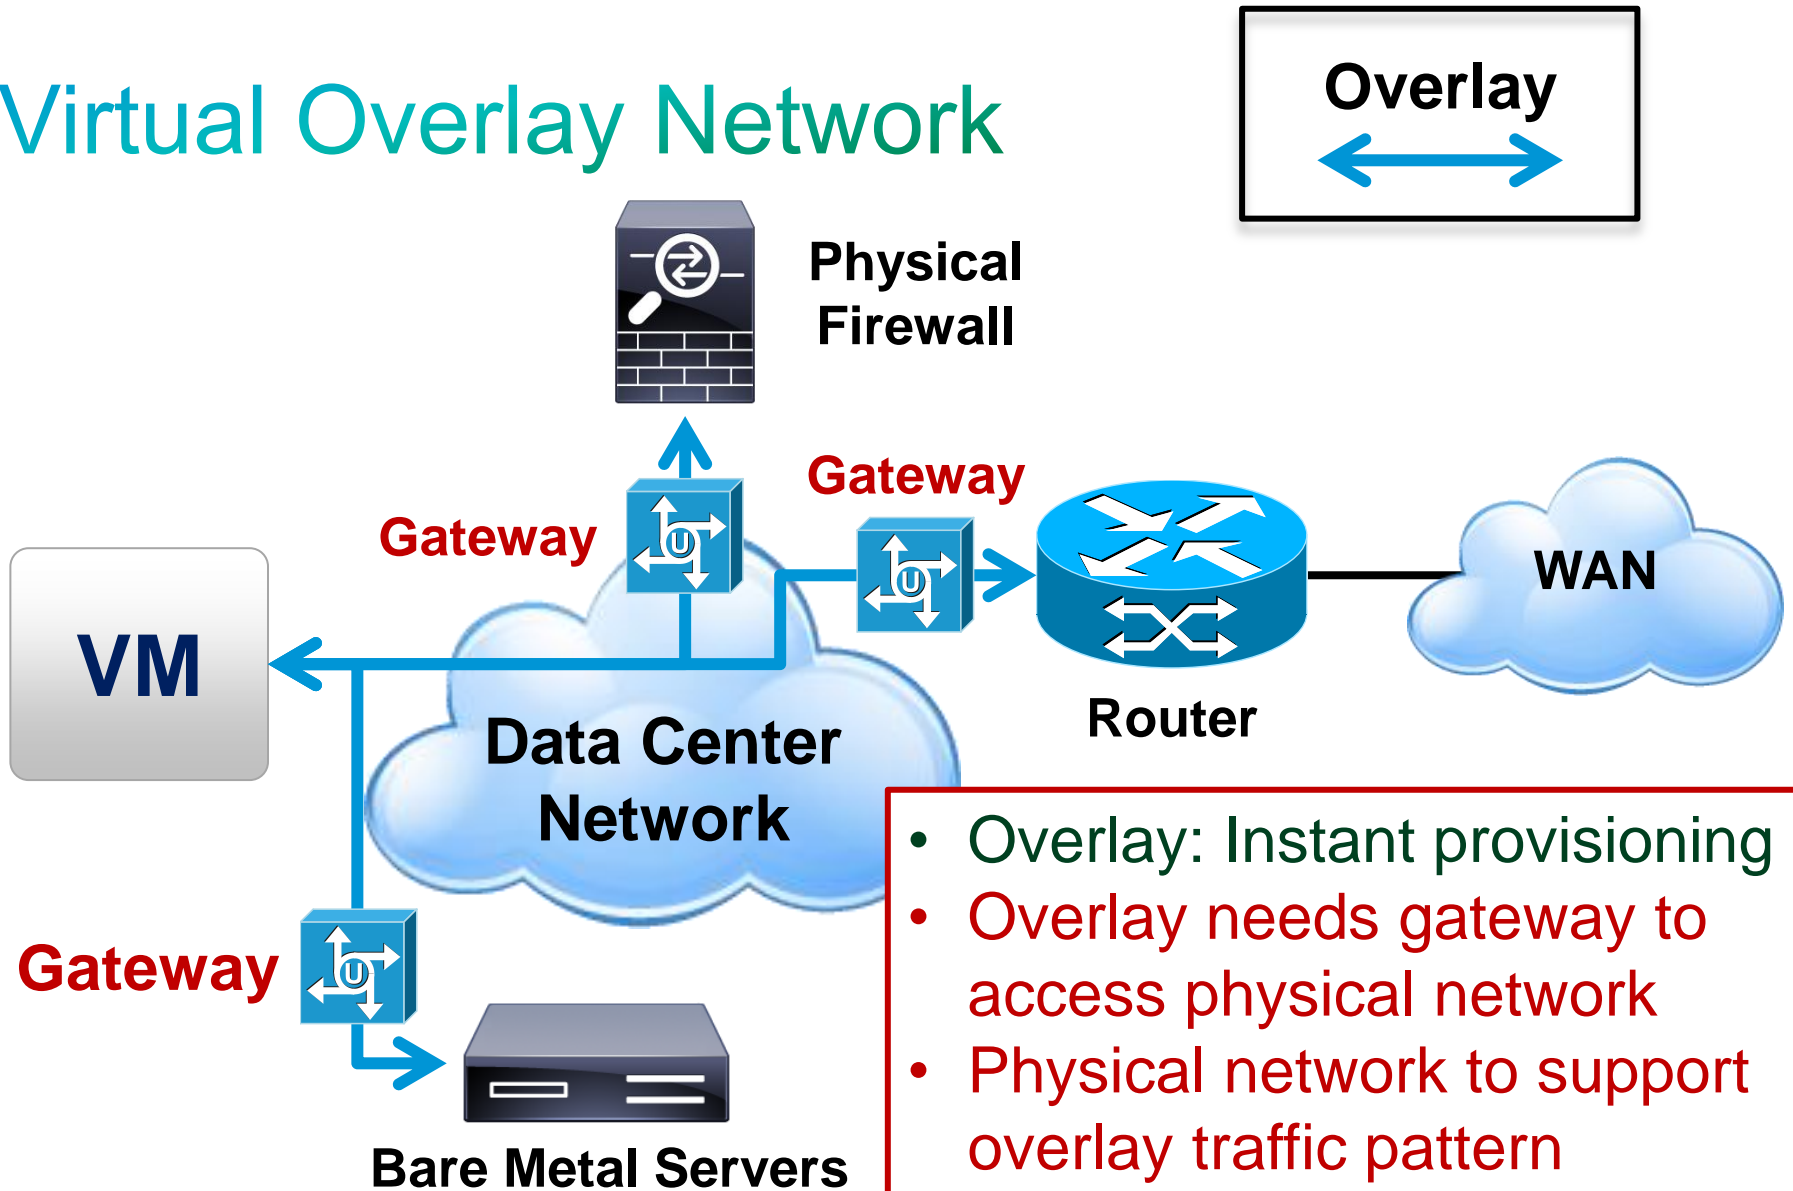


VXLAN to VLAN Gateway

- VXLAN 5500
- VLAN 100
- VLAN 200




Virtual Overlay Network




Broaden Mobility Domain with VXLAN

Unprecedented Infrastructure Flexibility

Rack-Wide VM Mobility



DC-Wide VM Mobility

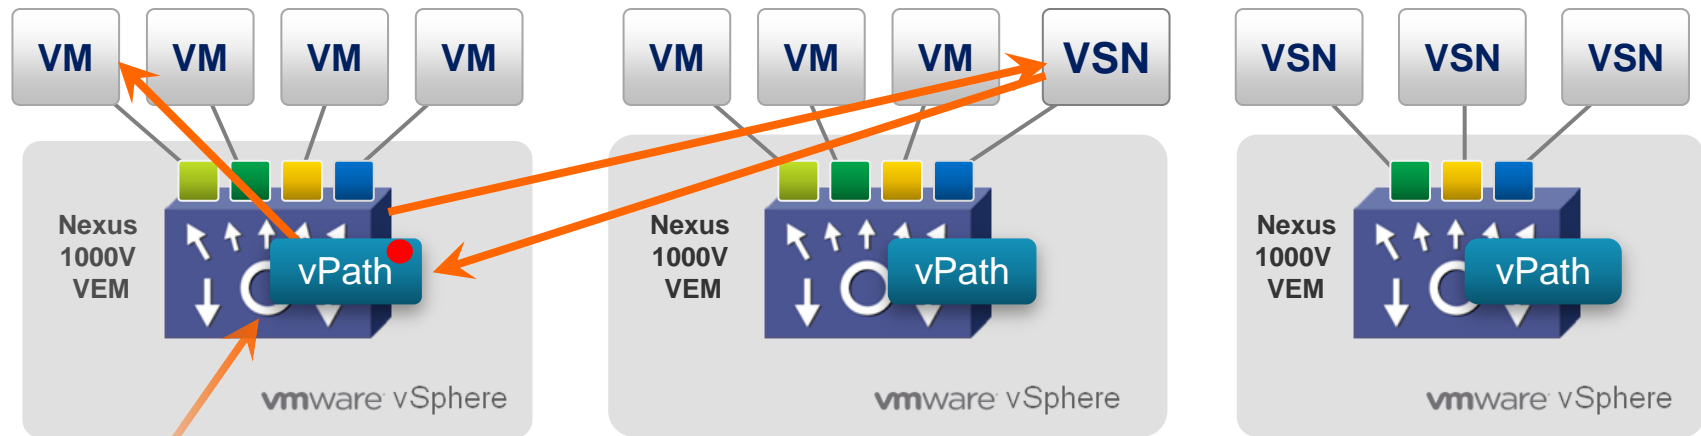




vPath

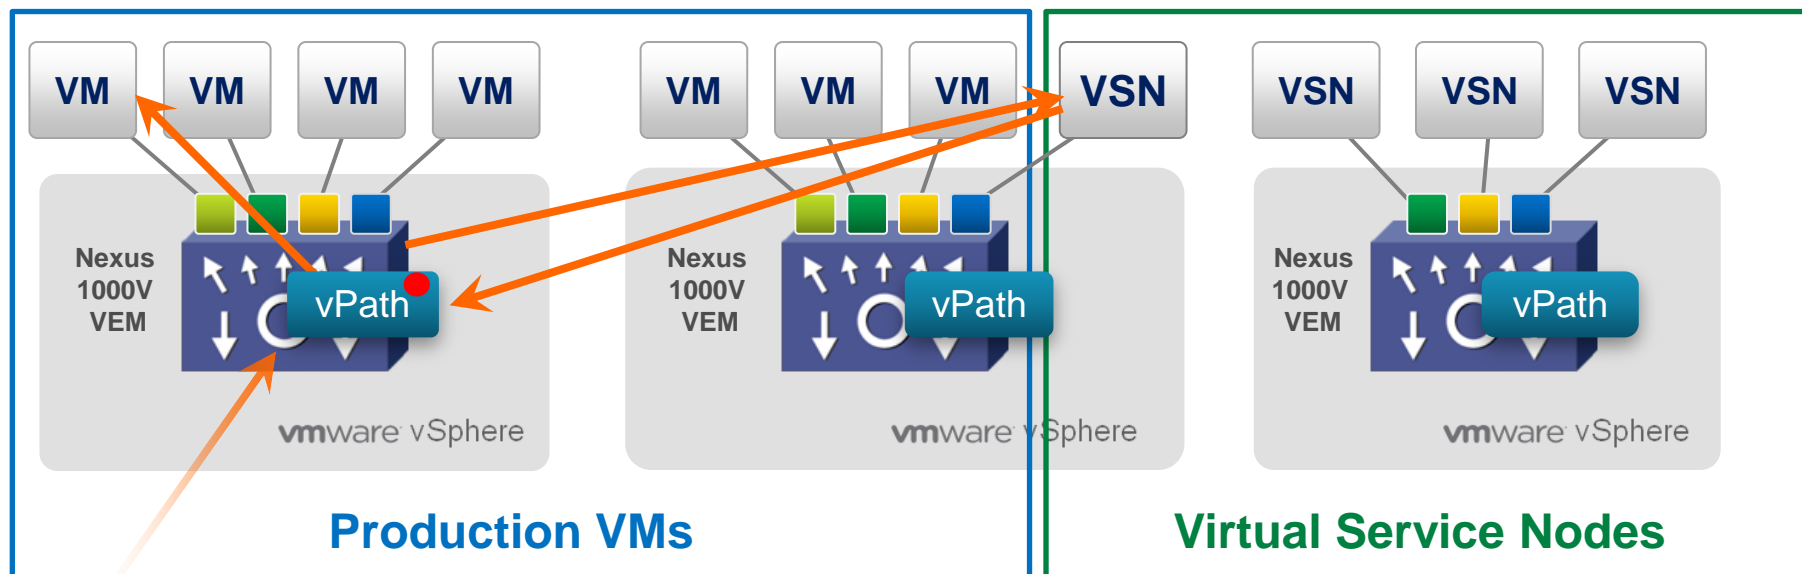
Cisco vPath: Intelligent Traffic Steering

Virtual Service Nodes (VSN)



- New flow is classified for VSN re-direction
- Initial packet(s) re-directed to VSN
- VSN installs a flow entry into vPath

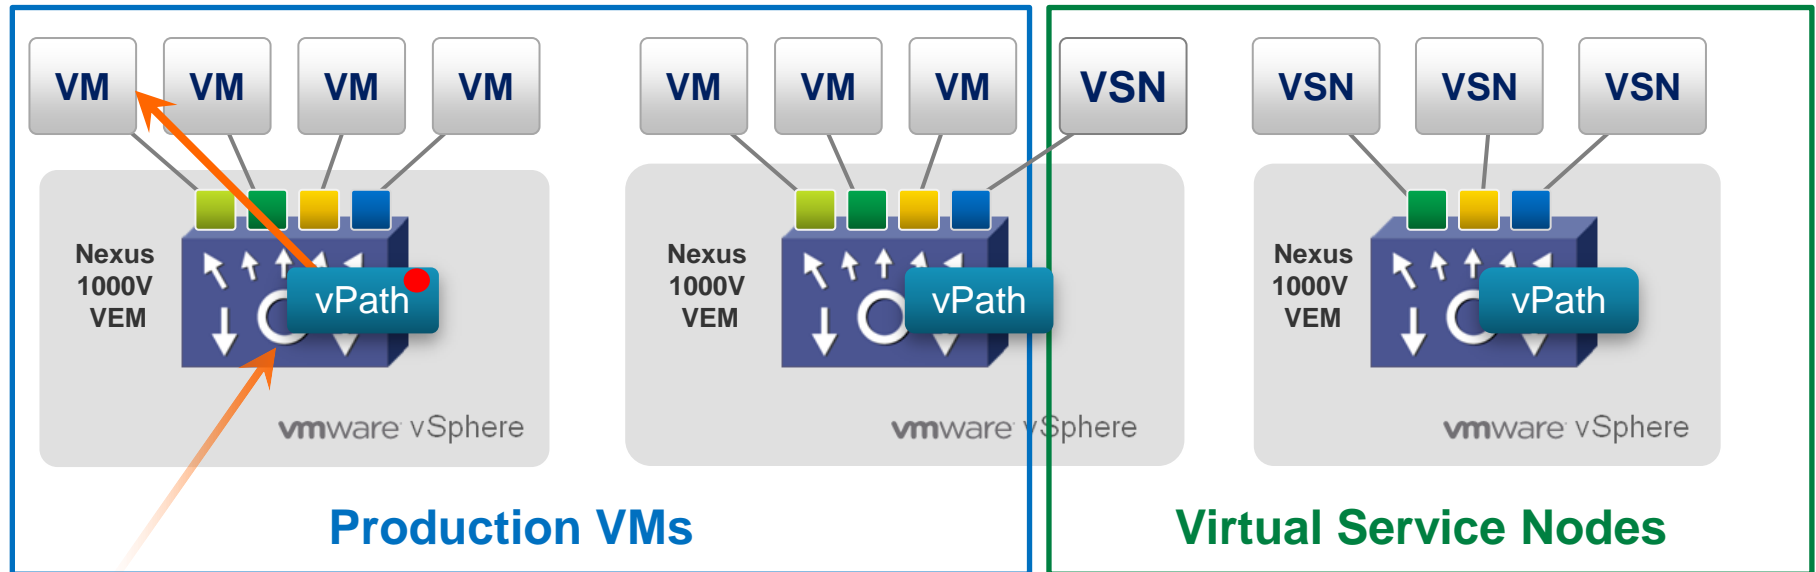
Cisco vPath: Flexible Deployment



- Service VMs placed with or separated from production VMs
- VSN can provide network service to multiple vSphere servers

Cisco vPath: Performance Acceleration

Scalable Acceleration in Virtual Ethernet Module



- Network service policy for subsequent packets in the flow are enforced in VEM
- Reduces traffic steering
- VEMs are part of the network service: Scalable Acceleration in hypervisor kernel




OpenStack & Quantum

OpenStack Core Projects



OpenStack Compute (Nova)

Software to provision virtual machines on commodity hardware at massive scale



OpenStack Object Storage (Swift)

Software to reliably store billions of objects distributed across commodity hardware



OpenStack Image Service (Glance)

Services for discovering, registering, and retrieving virtual machine images

OpenStack Core Projects Cont..




OpenStack Dashboard (Horizon)

A self-service web portal to allow administrators and users to manage OpenStack resources



OpenStack Identity (Keystone)

Provides “unified authentication” across all OpenStack projects and integrates with 3rd party authentication systems



OpenStack Network Service (Quantum)

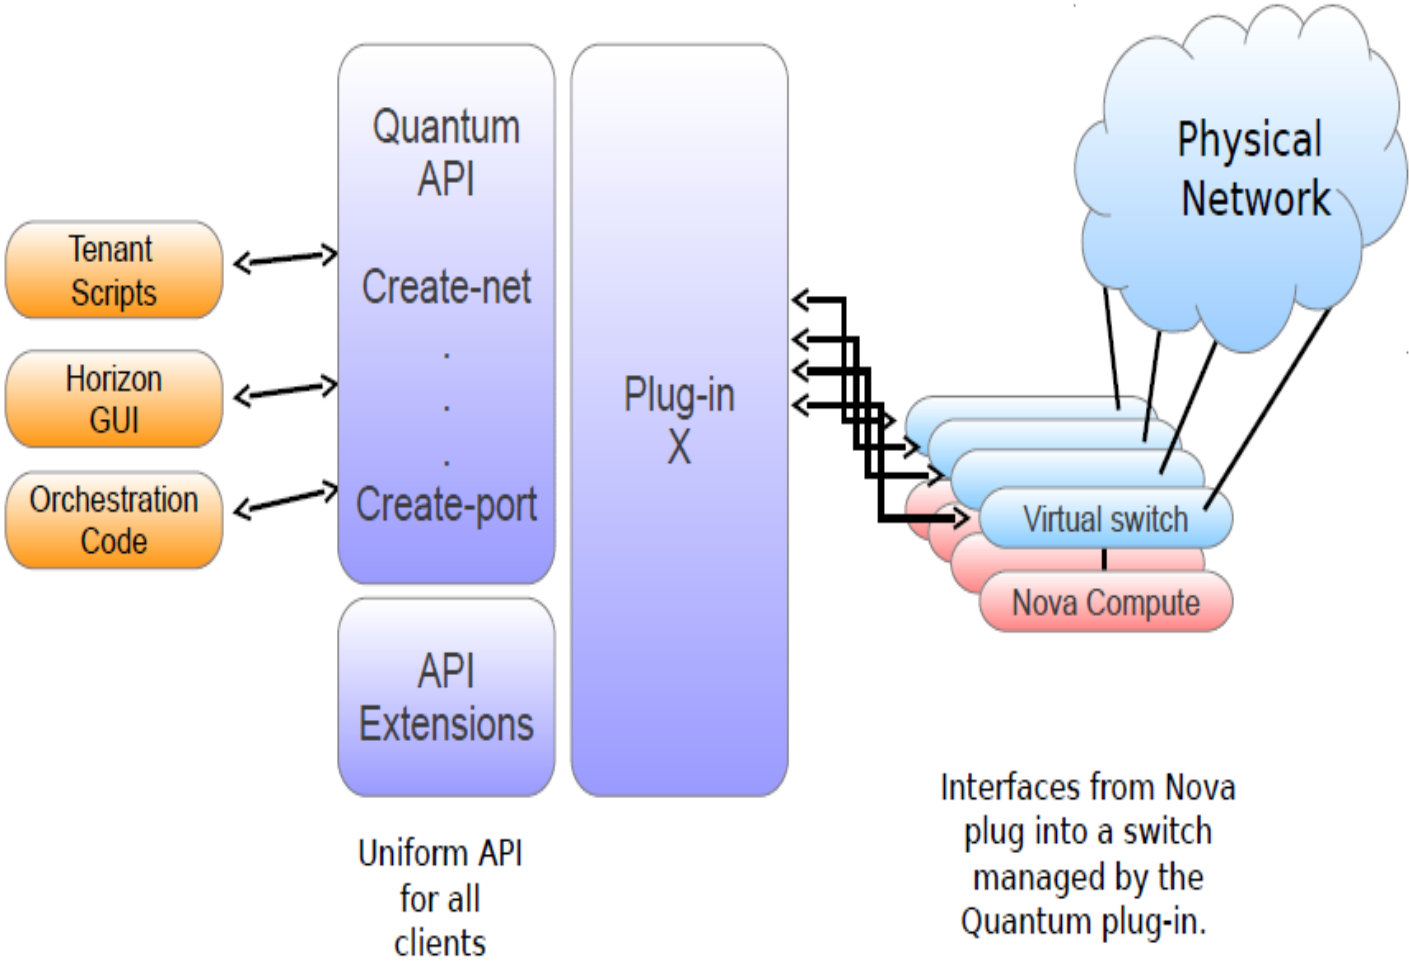
Provides “network connectivity as a service” between devices managed by other OpenStack services

Quantum Architecture

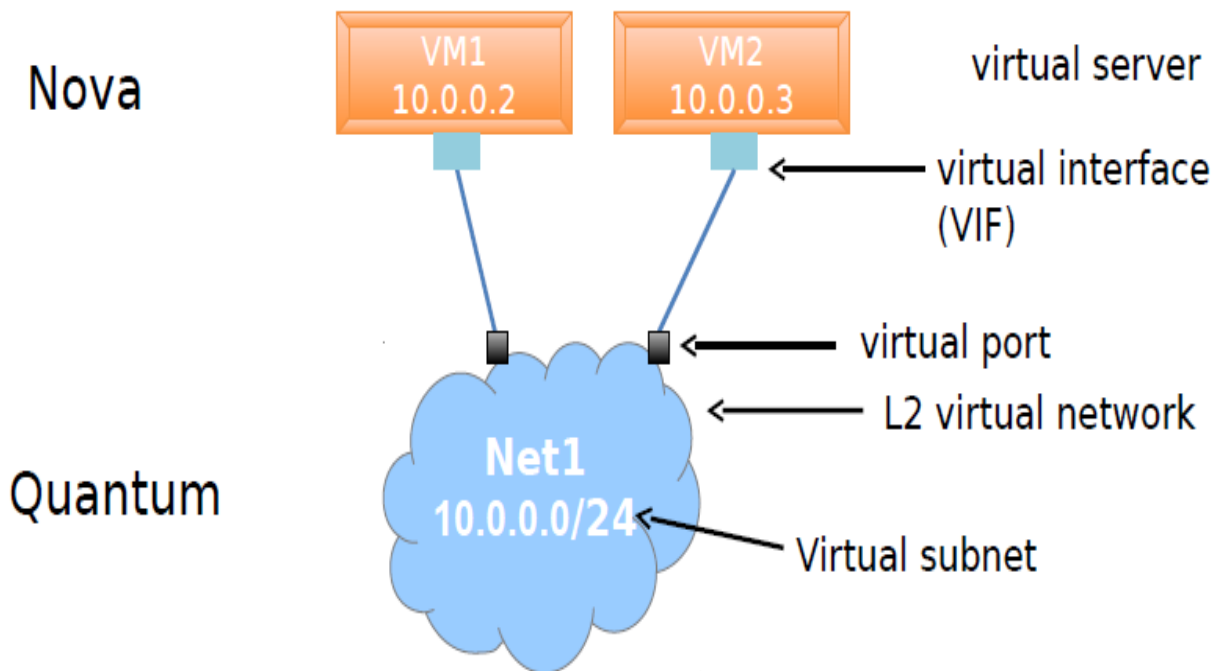
Clients

Quantum Service

Backend Networks Physical and Virtual



Basic Quantum API Abstractions



“virtual networks” and virtual subnets are fundamentally multi-tenant, just like virtual servers (e.g., overlapping IP's can be used on different networks)

Quantum API v2.0

- **Networks**

Create, Delete, Update
List, Show

- **Subnets**

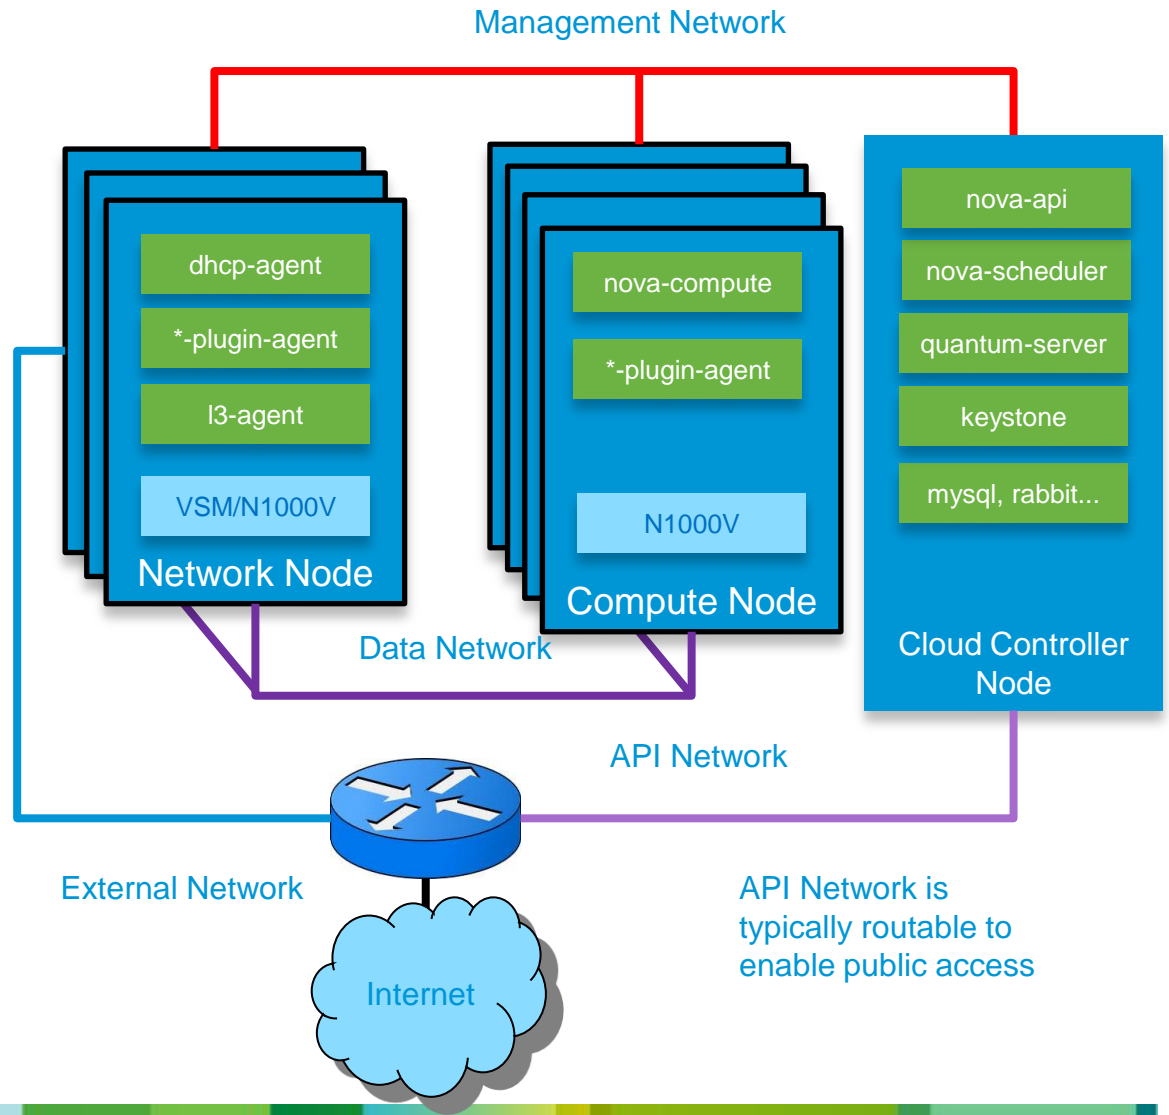
Create, Delete, Update
List, Show

- **Ports**

Create, Delete, Update
List, Show

OpenStack with Nexus 1000V

- Foundation of Virtual Services Architecture
 - vPath Service Insertion/Chaining
 - VXLAN Overlay Networking
 - CSR, VPN, Firewall, etc.
- Leverage Nexus 1000V REST API





Integrated Solution

Virtual Overlay Networks

Example: Virtual Overlay Networks and Services with Nexus 1000V

- Large scale L2 domains:

Tens of thousands of virtual ports

Hundreds of Servers

- Common APIs

Incl. OpenStack Quantum API's for cloud automation/orchestration

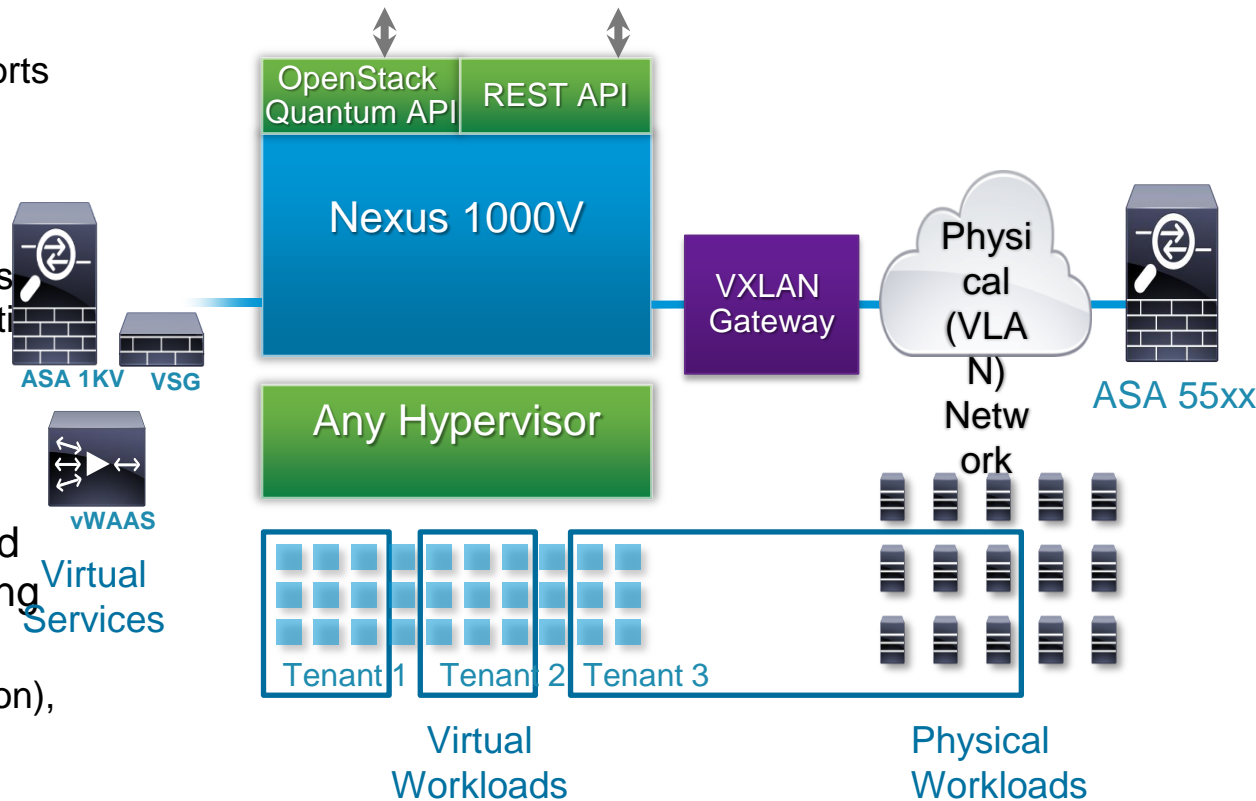
- Scalable DC segmentation and addressing: VXLAN

- Virtual service appliances and service chaining/traffic steering

VSG (cloud-ready security),
vWAAS (application acceleration),
vPATH

- Multi-hypervisor platform support: ESX, Hyper-V, OpenSource Hypervisors

- Physical and Virtual: VXLAN to VLAN GW



Tenant 1: virtual workloads protected by virtual firewall

Tenant 2: virtual workloads protected by physical firewall (via VXLAN)

Tenant 3: virtual & physical workloads in same L2 domain (via VXLAN)

Summary

- Expanding Nexus 1000V to open source hypervisors
- Increasing Nexus 1000V scale
- VXLAN: Easier to deploy without IP Multicast requirement
- VXLAN-to-VLAN Gateway for virtual to physical connectivity
- Leveraging OpenStack as primary management tool

Thank you.

