
CISCO
TOMORROW
starts here.
Cisco live!

Architecting an OpenStack Based Cloud with Cisco Infrastructure

BRKVIR-2601

Errol Roberts, DSE

Agenda

- Trends
- Introduction to OpenStack
- Infrastructure Consideration
- GBP and OpenStack
- Scaling OpenStack Deployments
- Conclusion

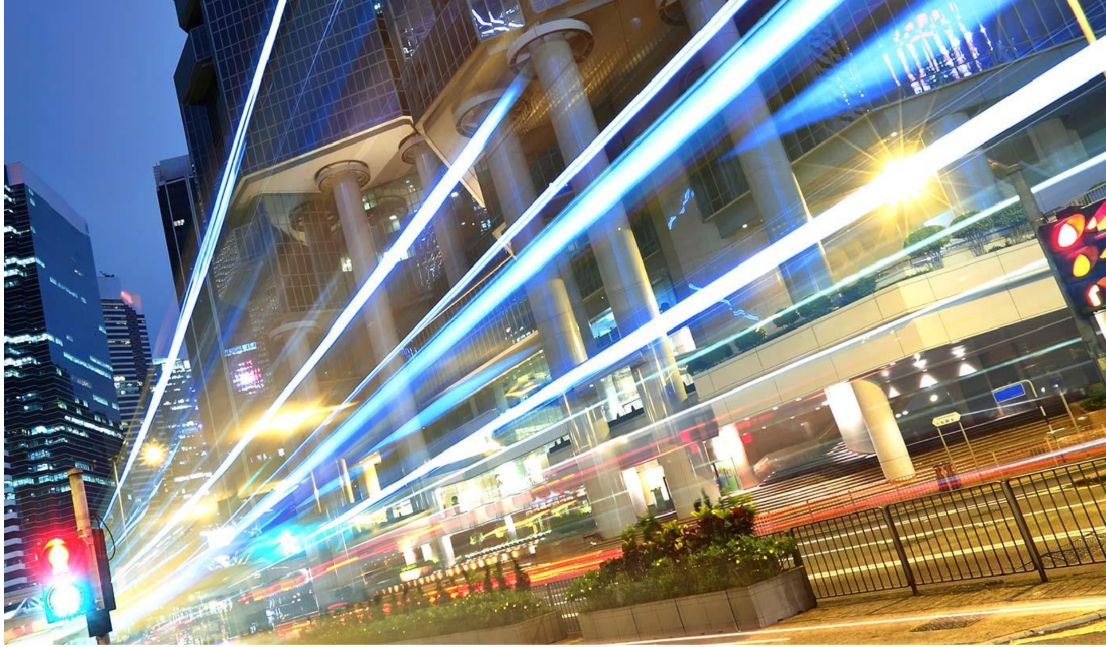

Market Trends

- **IT Spending is Shifting to the 3rd Platform**, comprised of **Cloud**, Mobile, Social, and Big Data. Gartner predicts that in 2015, the 3rd Platform will account for 30% of all IT spending and 100% of growth.
- **Web Scale IT.** Gartner notes that in the next few years, companies will need to think, act, and build applications and infrastructure in the same way Amazon, Google and Facebook do, using **cloud** to quickly deploy replicatable hardware on-demand.
- **Open Source.** Open source is leading the way in technology development, as developers seek to leverage the innovative solutions of others and concentrate their efforts on new services or applications.
- **Multi/Hybrid Cloud Management**. Few large companies want to put all their eggs in one basket and will be looking for ways to efficiently manage deployments across multiple clouds.
- **Software Defined Everything**. Agile development methods are essential to delivering application and service flexibility. Software defined networking, storage, data centres and security will finally make computing dynamic.

Orchestration in the IT World

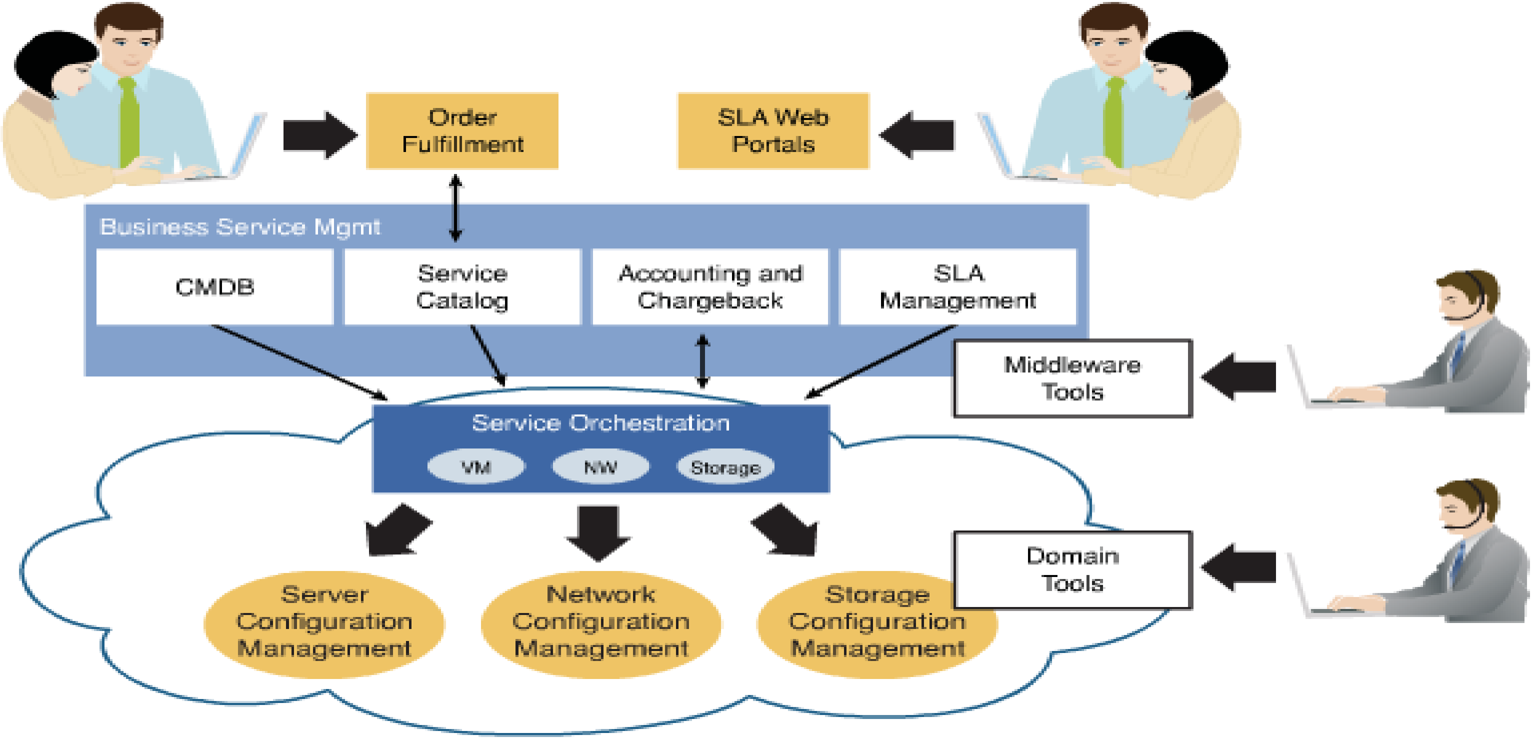

What is OpenStack?

“OpenStack is a global collaboration of developers and cloud computing technologists producing the ubiquitous open source cloud computing platform for public and private clouds. The project aims to deliver solutions for all types of clouds by being simple to implement, massively scalable and feature rich. The technology consists of a series of interrelated projects delivering various components for a cloud infrastructure solution.”

- openstack.org

Translated, OpenStack is software to run cloud services and the community behind that software.

OpenStack

Overview

Open source Cloud Computing Platform for Private and Public Clouds

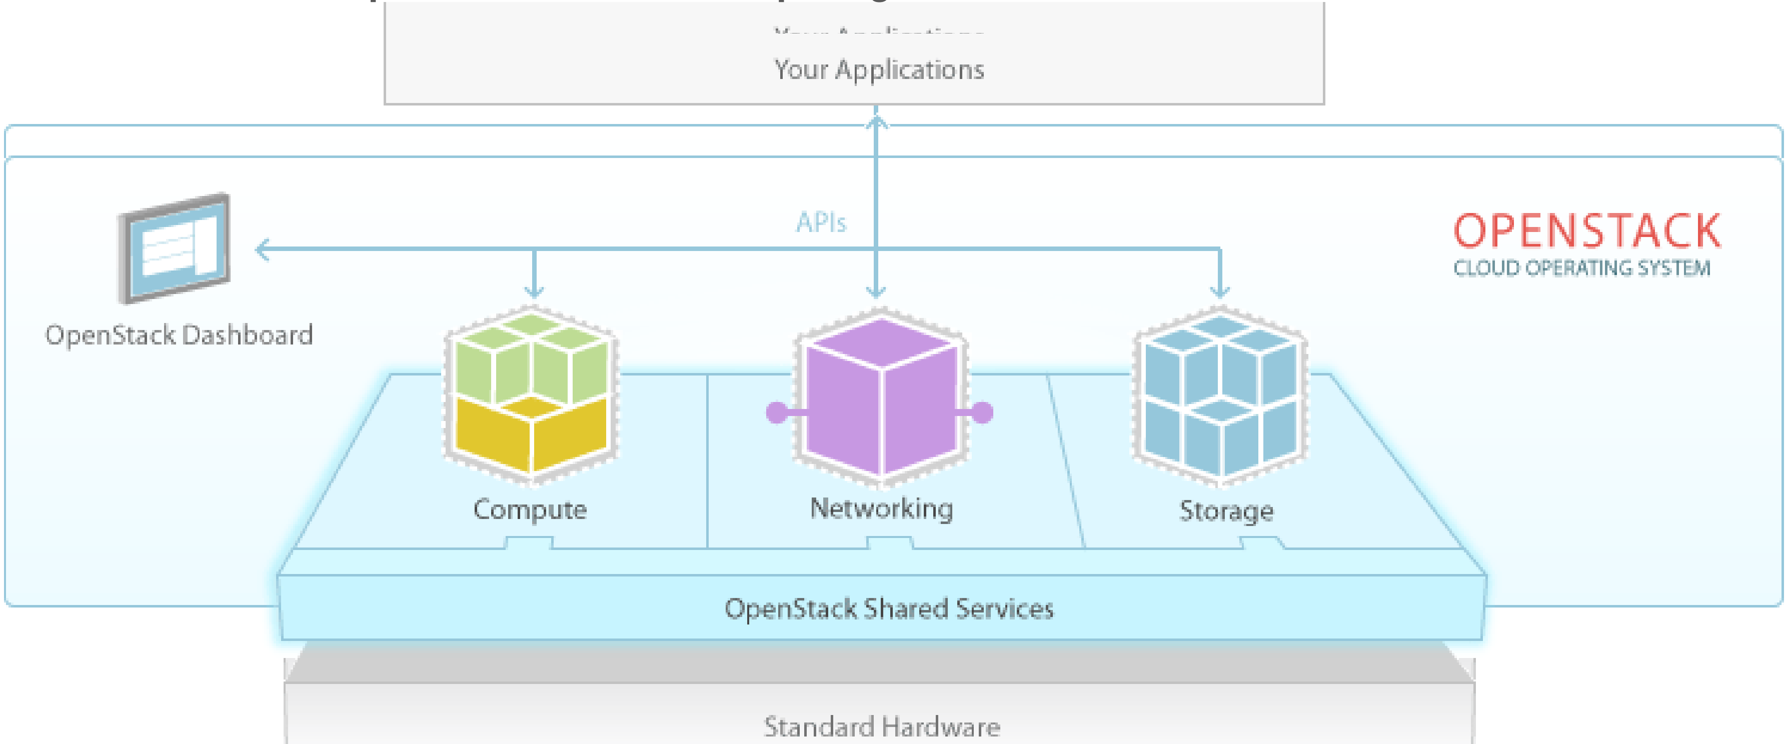

Design tenets – scale & elasticity, share nothing & distribute everything

OpenStack Community History

- Founded in July 2010 by Rackspace Hosting, NASA and partners
  – NASA and Rackspace contributed the initial code
- Code has gone through eight releases
  – OpenStack has a 6-month time-based release cycle
- Over 169 companies have joined the community
  – OS/Hypervisor vendors
  – Public Cloud/service providers
  – Equipment manufacturers
  – OpenStack software and services

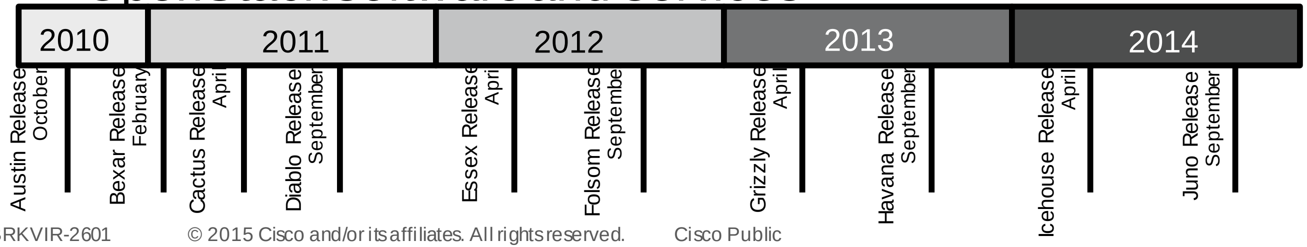

OpenStack is Transforming Cloud Deployment

Cost Reduction	Investment Protection	Access to Innovation	Increased Agility	Faster ROI
Greater Control	Choice of Vendors	Community Engagement	Modify & Scale on Demand	Accelerated Deployment

openstack™ Cisco live!

OpenStack is Transforming Cloud Development

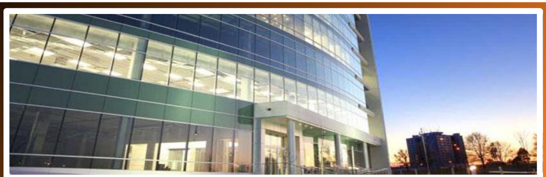

Enterprise/Public Sector

Application deployment speed in a highly dynamic IT environment

Service Provider

End-to-end cloud delivery that is automated and tenant aware

84% of RedHat users indicate OpenStack part of future plans

Typical Use Cases

- Development and Testing
- Proof of Concept Implementations
- Applications in a Private IaaS Environment
- Cloud Scale Applications
- Small to Medium Data Processing Applications

PRIVATE AND HYBRID CLOUDS

- Big Data
- Mobile Applications
- Multi Data Centre Deployments
- Seamless Hybrid Clouds Deployments
- PaaS SaaS

Cisco and OpenStack

- Code contributions across several services – Network. Compute, Dashboard, Storage
- Foundation Board member

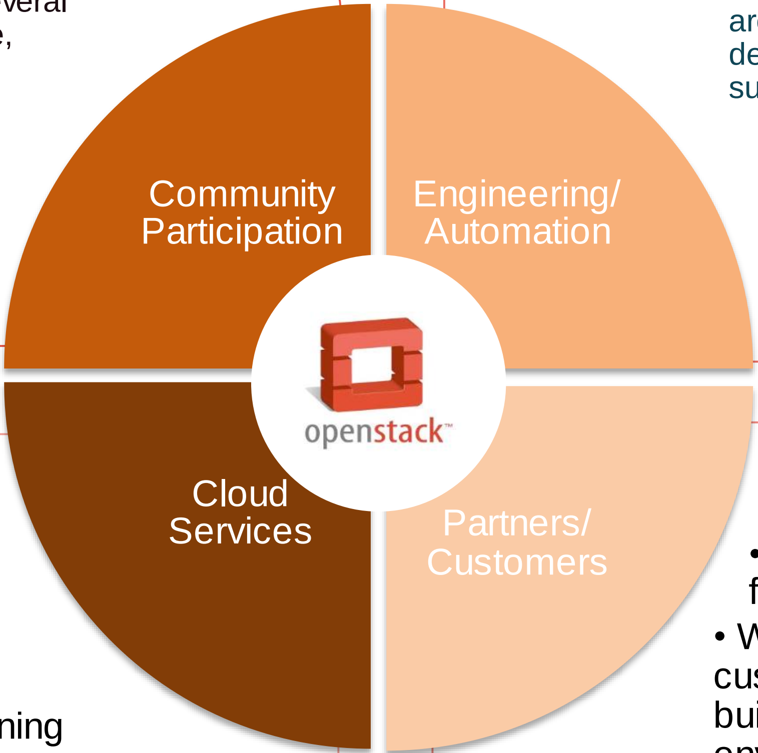

•Automation (Puppet) and architectures (HA) for production deployment and operational support

- •Neutron/Nova Plug-ins for Cisco product lines – Nexus, DFA, APIC, UCS, CSR/ASR
- •Co-developed solutions (Red at, Canonical, SUSE

metacloud

- OpenStack based Global Intercloud hosted across Cisco and partners data centres
- Cisco Webex Service running on OpenStack

•Cisco Validated Designs for production deployments

- Work closely and jointly with customers to design and build their OpenStack environment

Cisco Top Vendor in Network Contributions

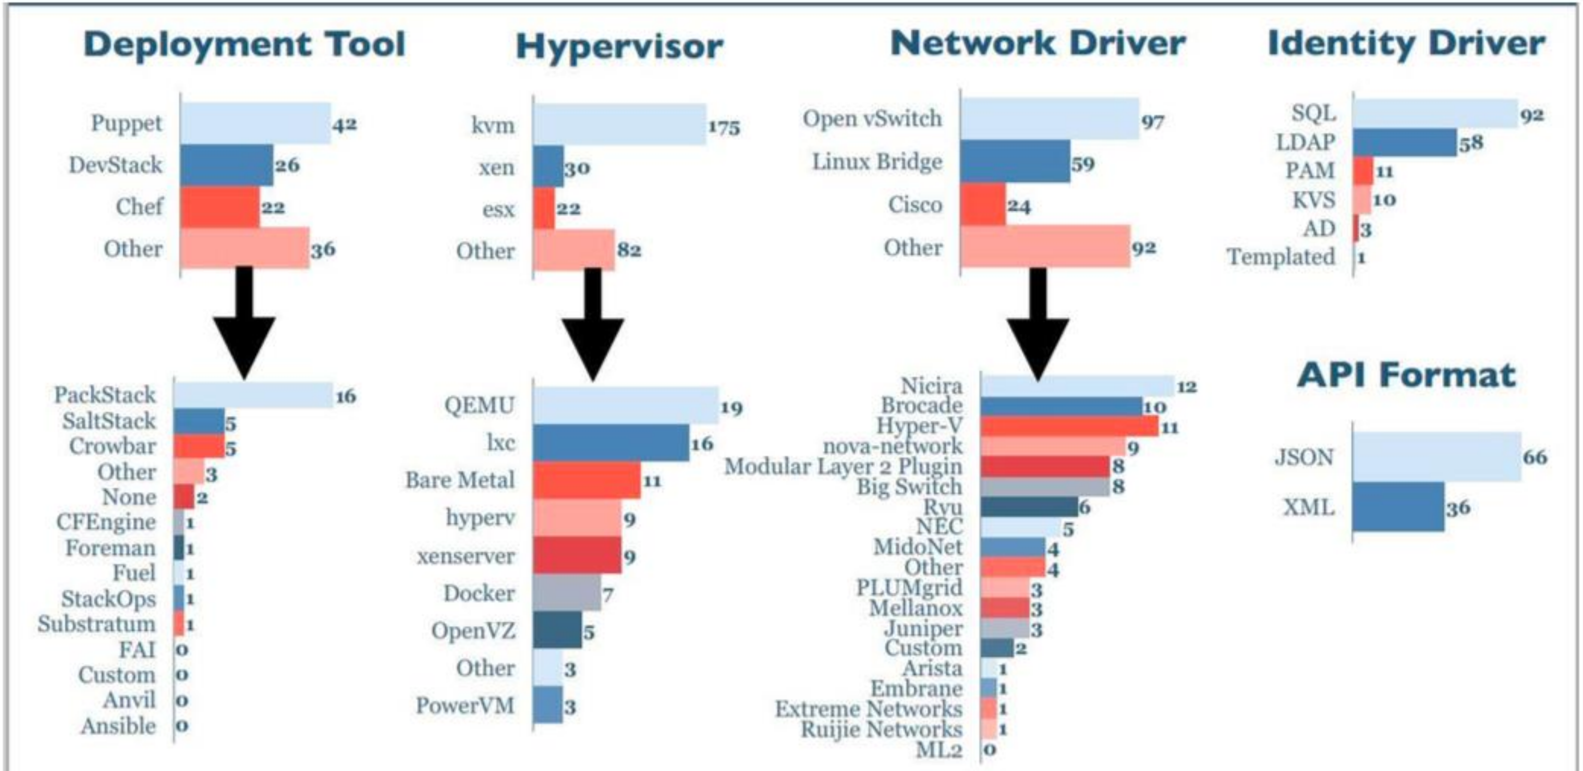

OpenStack Atlanta User Survey

OpenStack Basics

OpenStack is "Project" Based

Core Projects Shown

Compute
"Nova"
- Houses VMs
- API driven
- Support for multi-hypervisors

Storage
Image, Object, Block
"Glance, Swift, Cinder"
- Instance/VM image storage
- Cloud object storage
- Persistent block level storage

Dashboard
"Horizon"
- Web app for controlling OpenStack resources
- Self-service portal

Identity
"Keystone"
- Centralised policies
- Tenant mgmt.
- RBAC
- Ext. integration (LDAP)

Networking
"Neutron"
- Networking as a service
- Multiple models
- IP address mgmt.
- Plugins to external HW

Telemetry
"Ceilometer"
- Central collection point
- Metering and monitoring

Orchestration
"Heat"
- Template-based orchestration engine
- More rapid deployment of applications

Database
"Trove"
-DBaaS
-Single-tenant DB within instance

Data Processing
"Sahara"
- Fast provisioning of Hadoop clusters

New!

Architecture

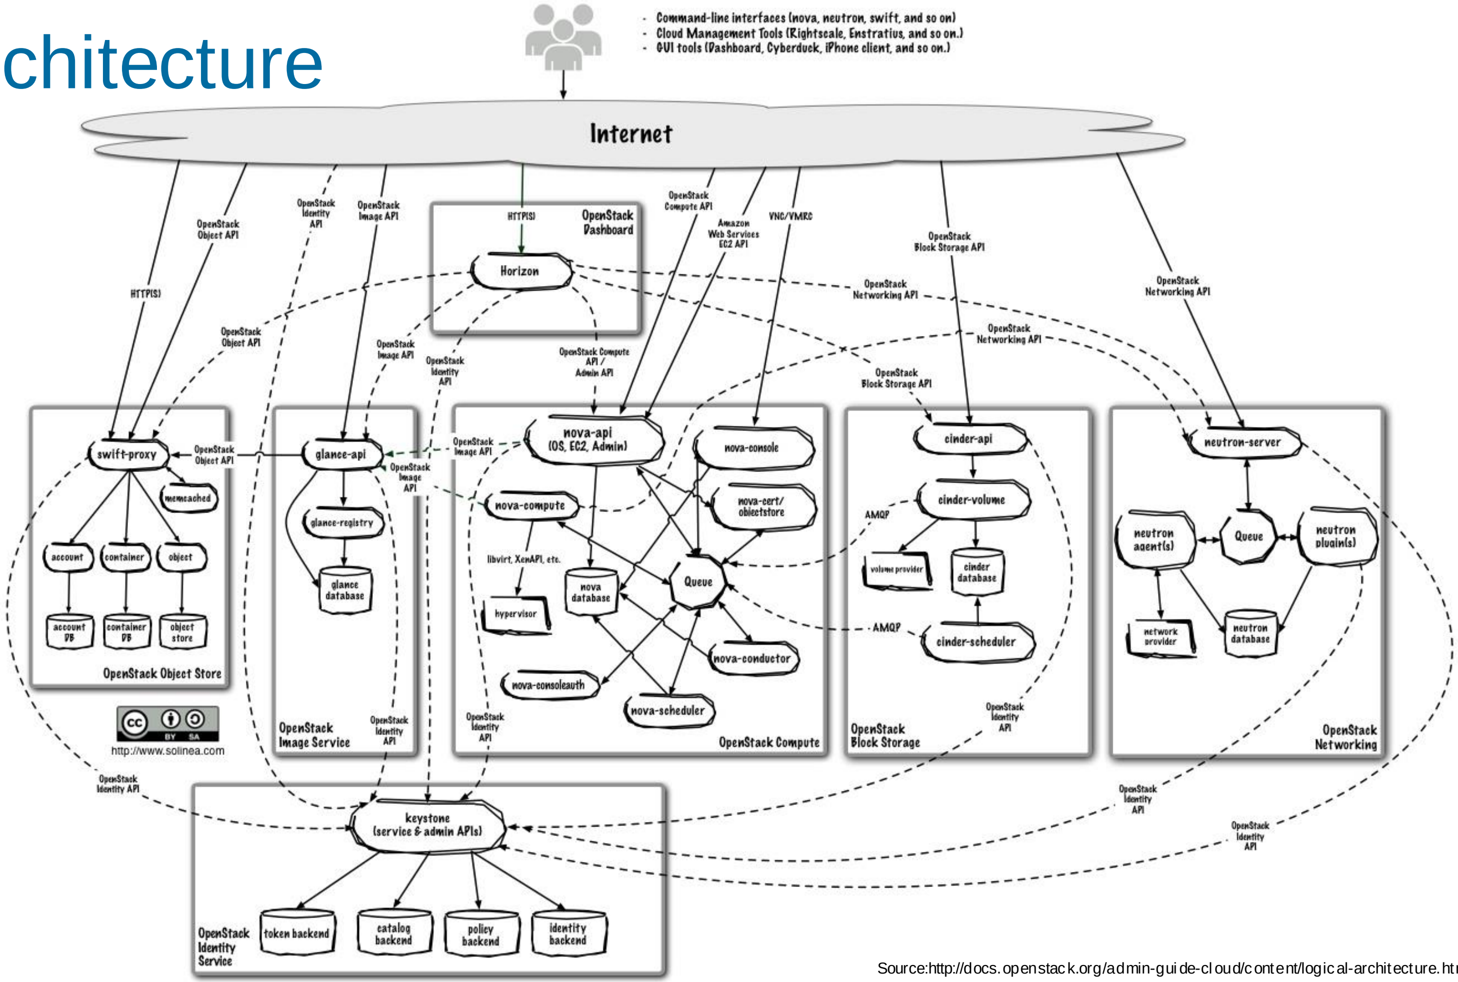

Source:http://docs.openstack.org/admin-guide-cloud/content/logical-architecture.html

REST API and Messaging Between Components

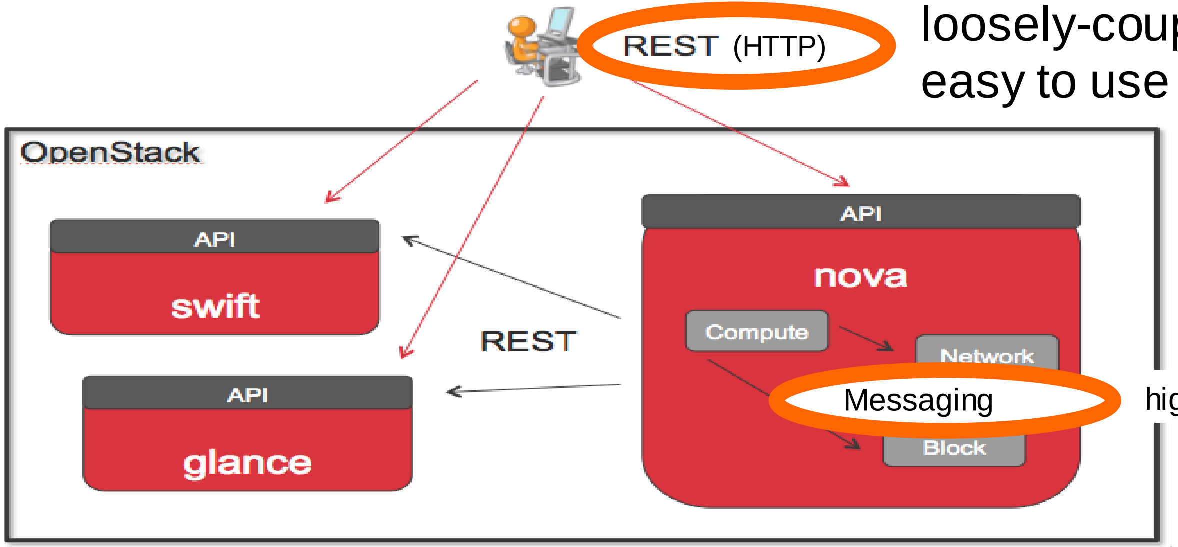

loosely-coupled
easy to use

high performance

- **Representational state transfer (REST)**
- **Advanced Message Queuing Protocol (AMQP)**

OpenStack Software Architecture

Rapidly evolving set of open API's and services for cloud applications

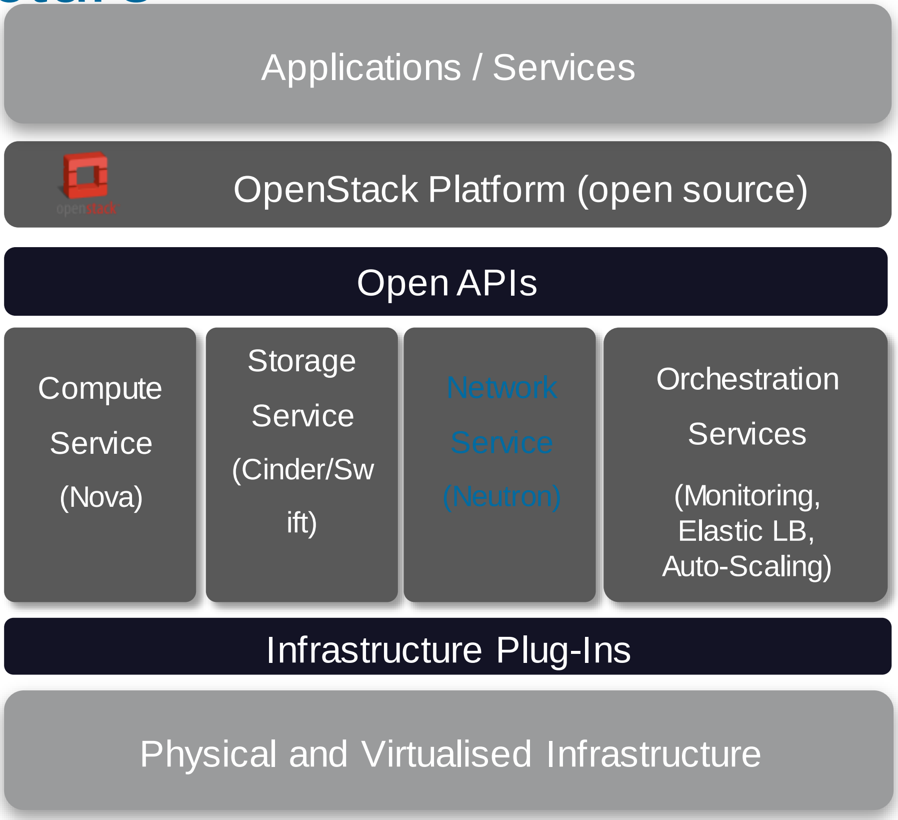

OpenStack Architecture

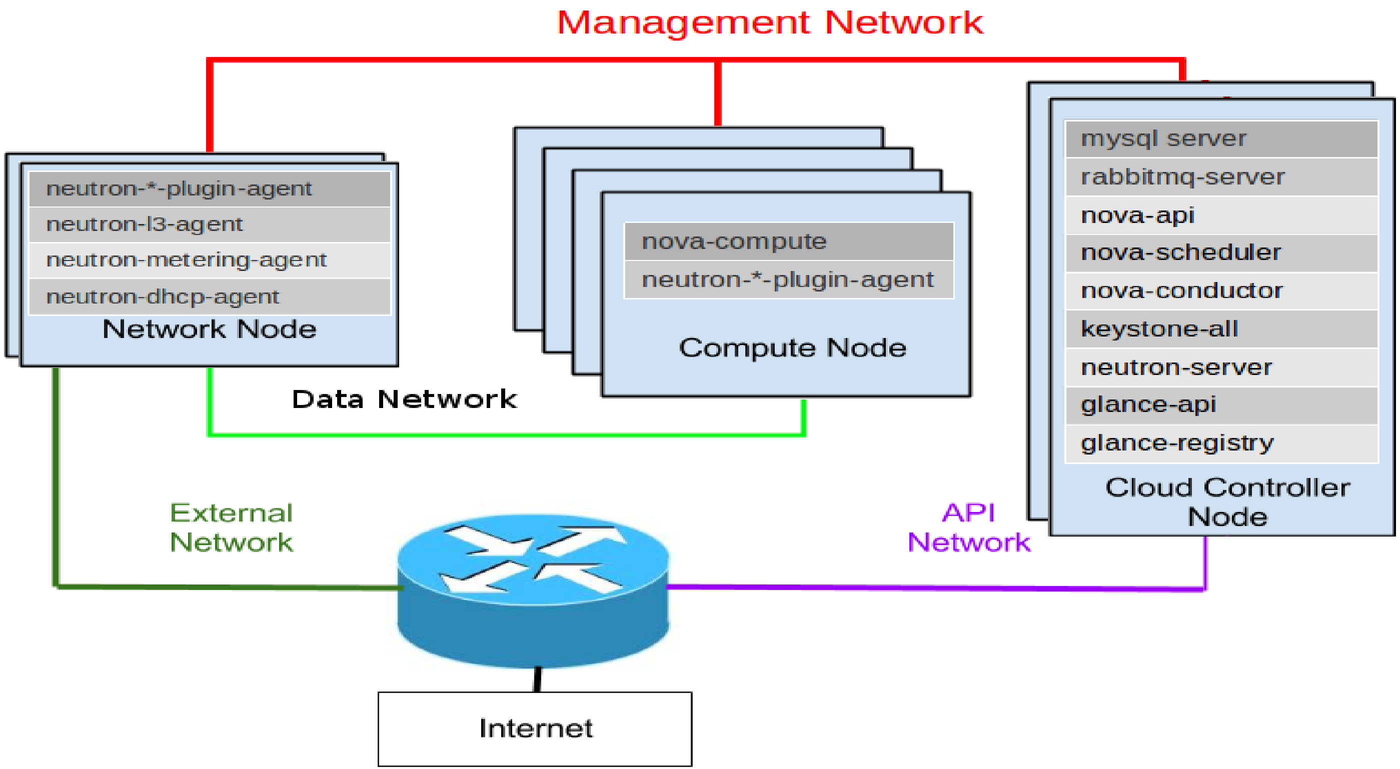

OpenStack Plugin Model

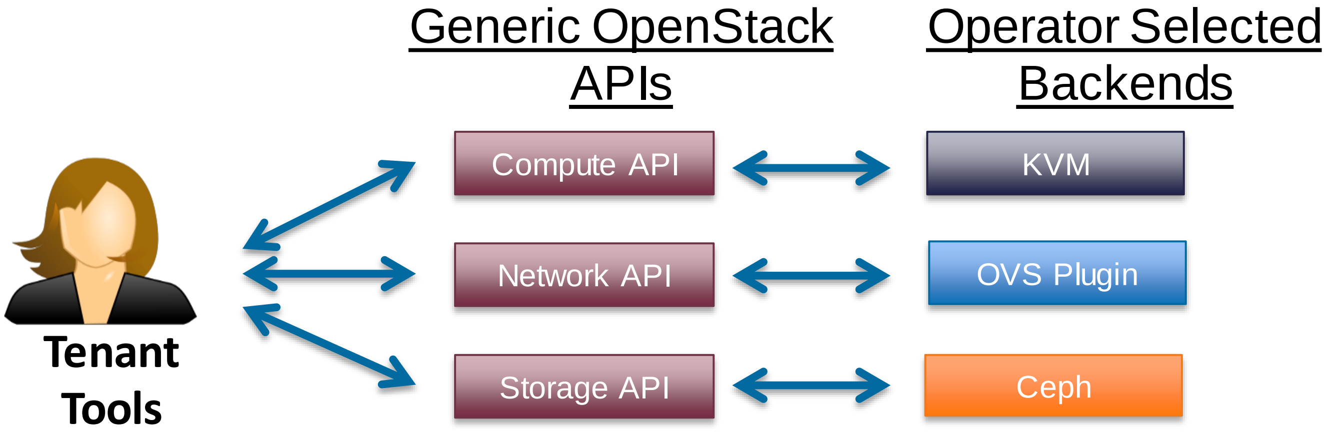

- Cisco plugin supports multiple sub-plugins
- Modular L2 (ML2) evolution of Neutron
- Allow multiple plug-ins to exist as sub-plugin drivers

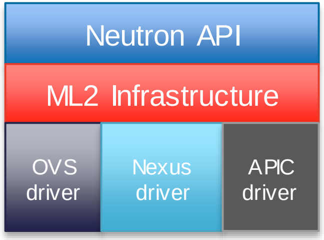

OpenStack Neutron Architecture

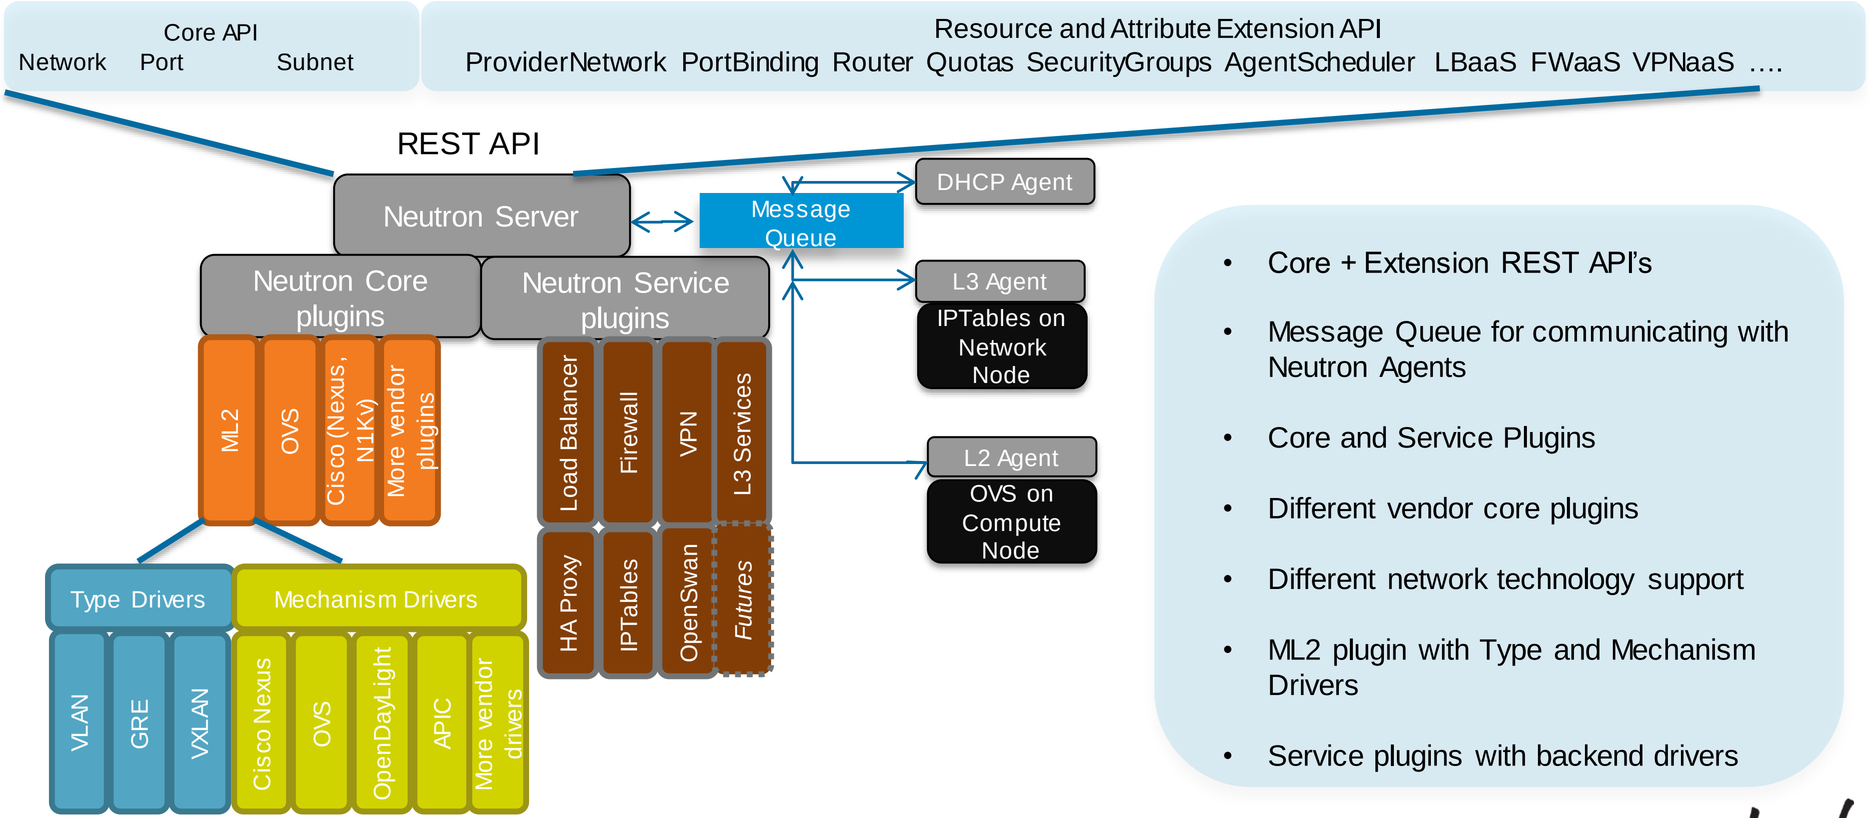

- Core + Extension REST API's
- Message Queue for communicating with Neutron Agents
- Core and Service Plugins
- Different vendor core plugins
- Different network technology support
- ML2 plugin with Type and Mechanism Drivers
- Service plugins with backend drivers

Neutron Model

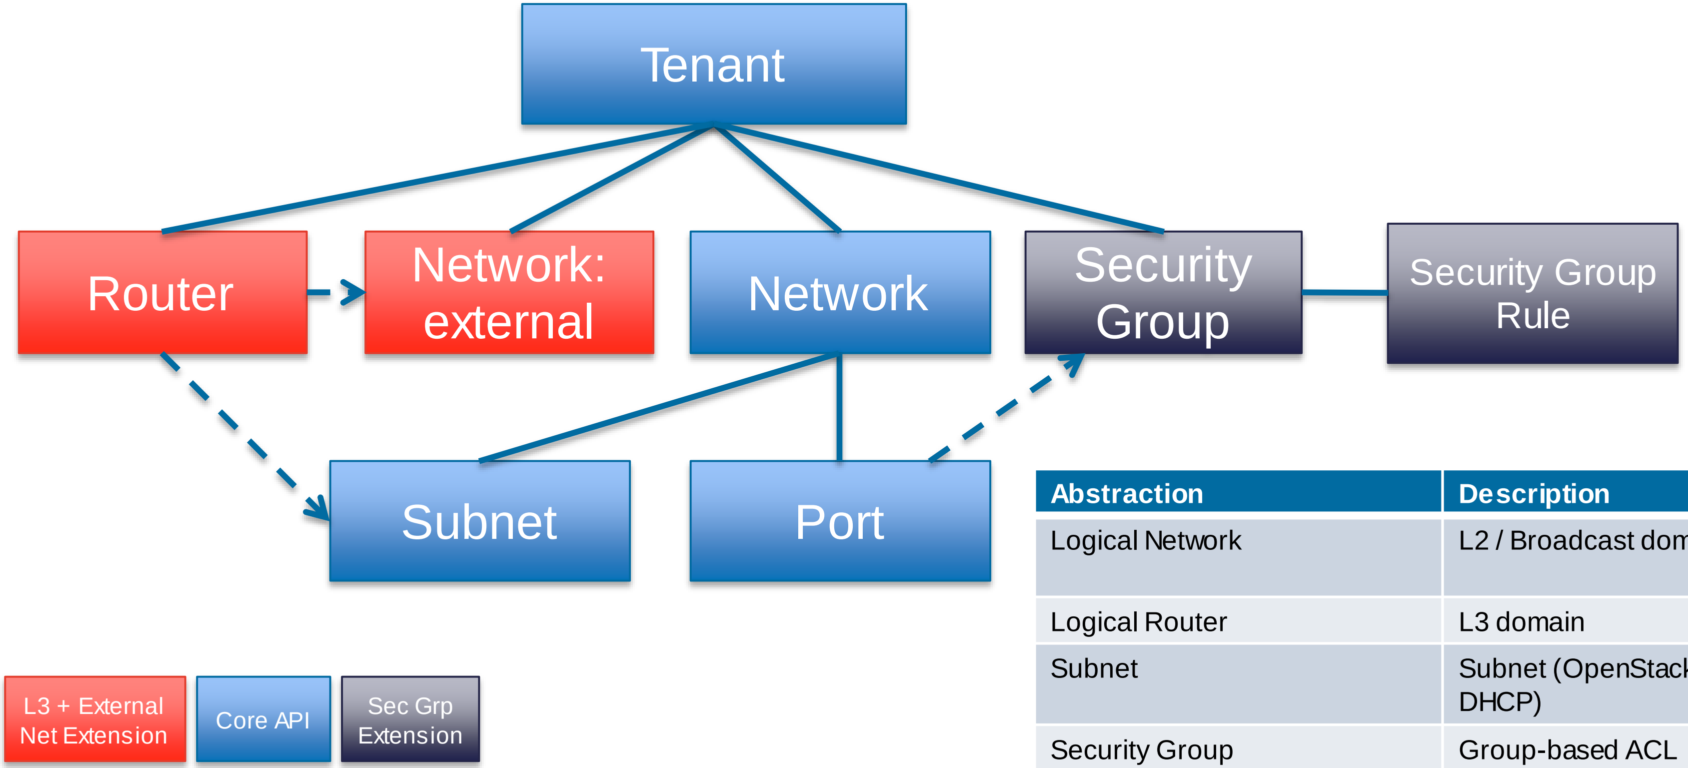

Abstraction	Description
Logical Network	L2 / Broadcast domain
Logical Router	L3 domain
Subnet	Subnet (OpenStack IPAM / DHCP)
Security Group	Group-based ACL

Neutron Integration with Other OpenStack Services

- Neutron relies on OpenStack Identity Service (KeyStone) for authentication and authorisation of all API requests
- OpenStack Compute Service (Nova) communicates with Neutron API for plugging VM onto a network through port creation
- Tenants and Administrators use the GUI based OpenStack Dashboard Service (Horizon) for managing Neutron networks

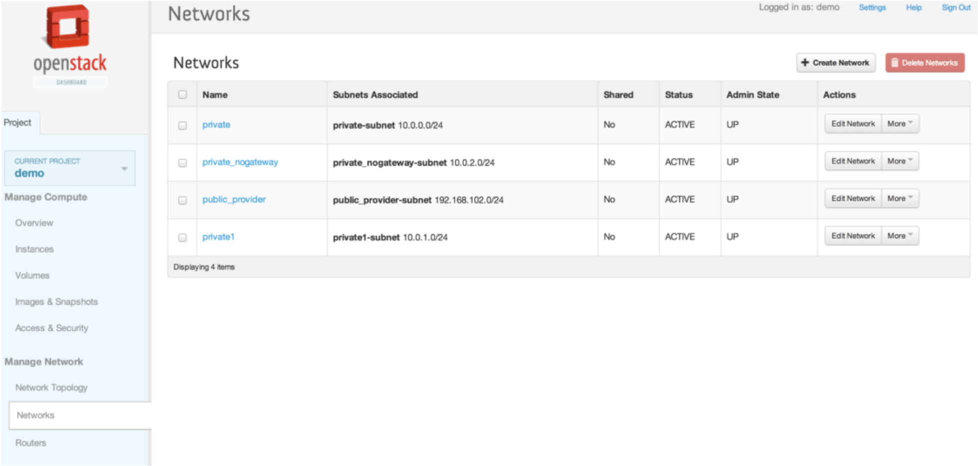

Neutron Networking for Tenant Isolation

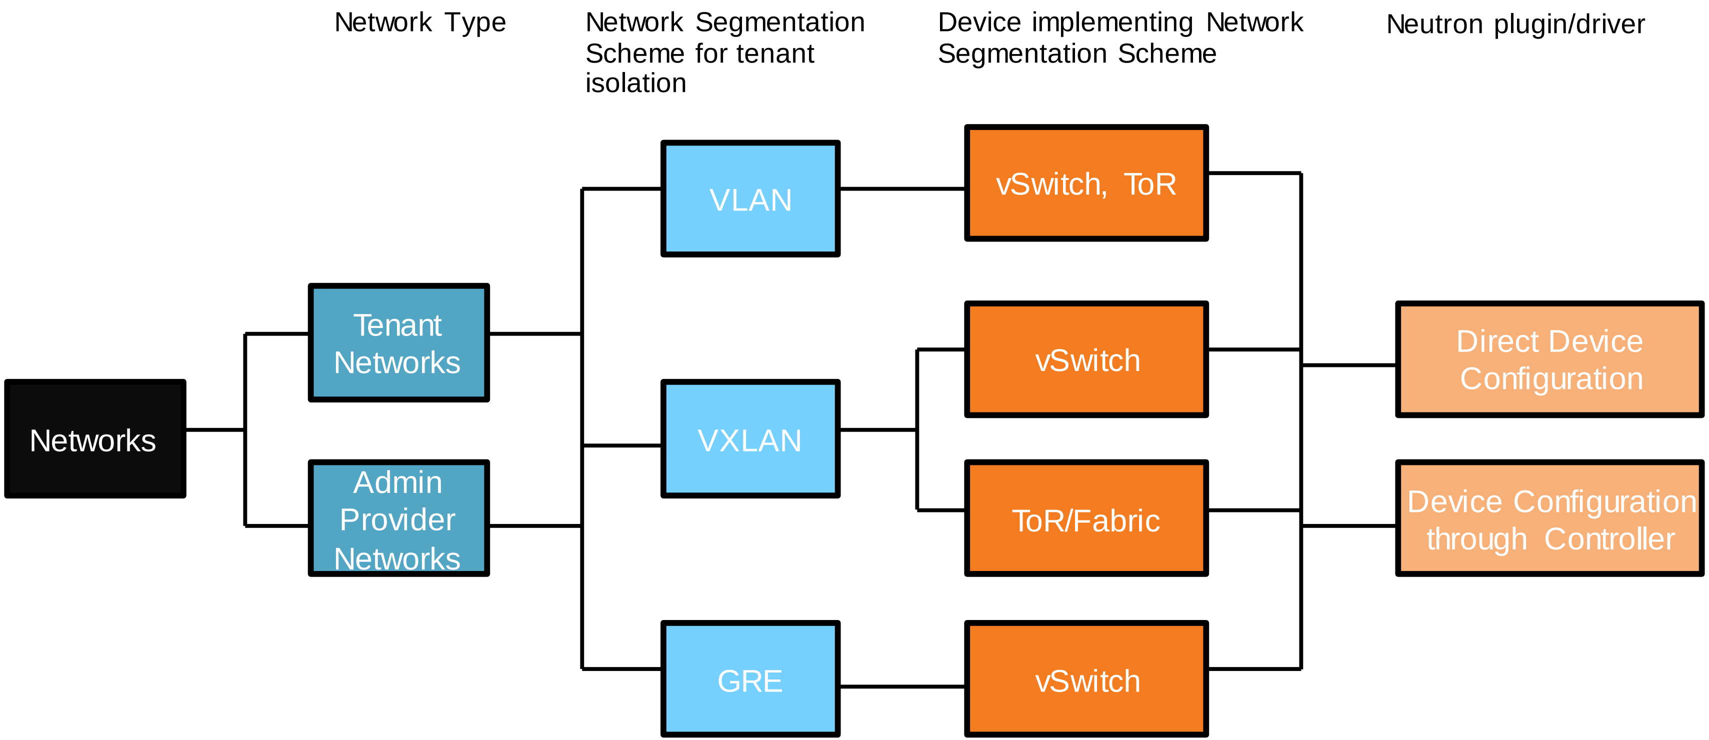

Neutron Networking for Layer 3 Services

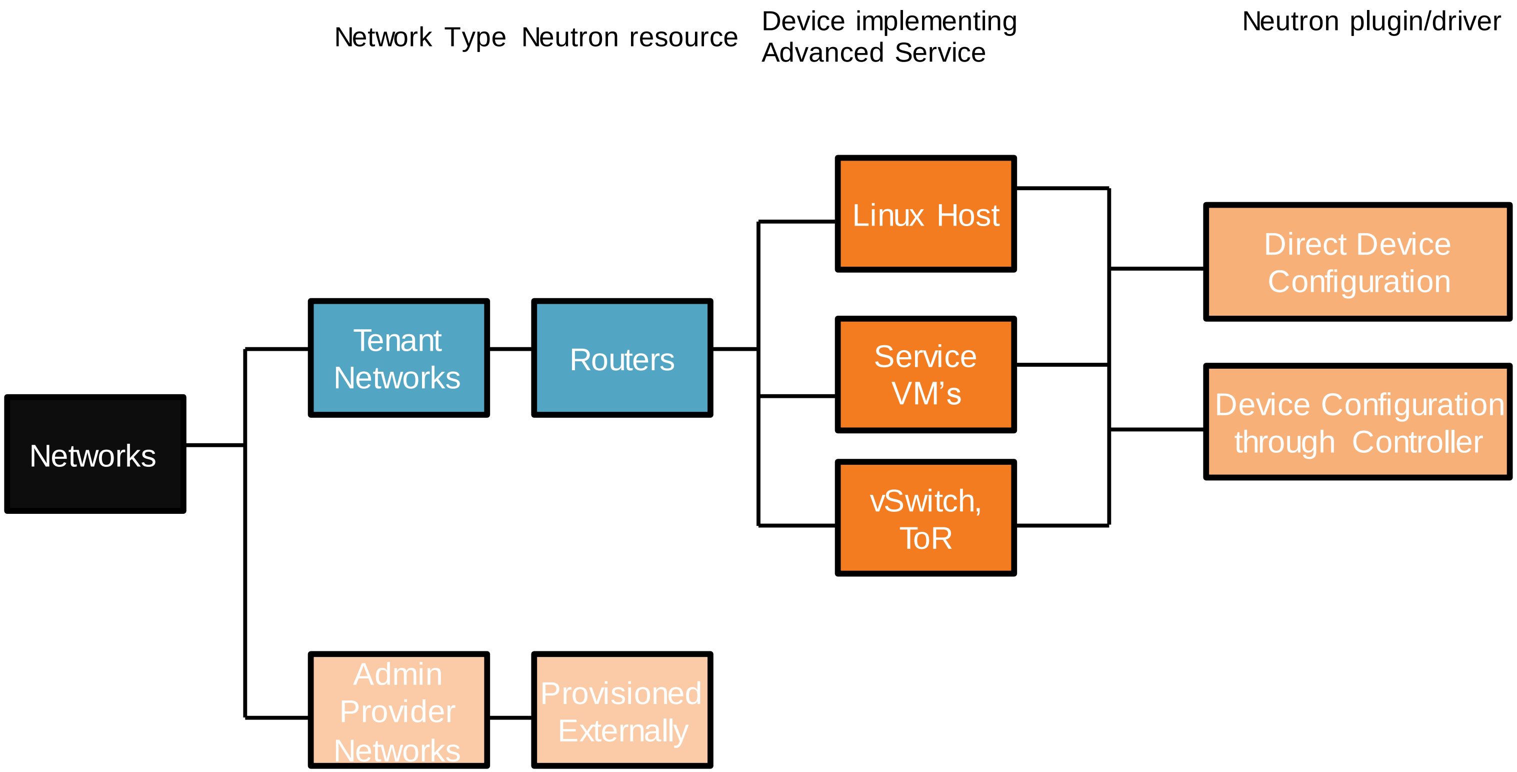

Linux Networking Devices on the Compute Host

- There are four distinct type of virtual networking devices: TAP devices, veth pairs, Linux bridges, and Open vSwitch bridges. For an ethernet frame to travel from eth0 of virtual machine vm01, to the physical network, it must pass through nine devices inside of the host: TAP vnet0, Linux bridge qbr*XXX*, veth pair (qvb*XXX*, qvo*XXX*), Open vSwitch bridge br-int, veth pair (int-br-eth1, phy-br-eth1), and, finally, the physica network interface card eth1.
- Vnet connects to the extra set of linux bridge device so that IPTables can be used for SecGroups
- Veth Pair acts like a “patch panel”
- qvo: veth pair openvswitch side
- qvb: veth pair bridge side
- qbr: bridge
- qr: l3 agent managed port, router side
- qg: l3 agent managed port, gateway side

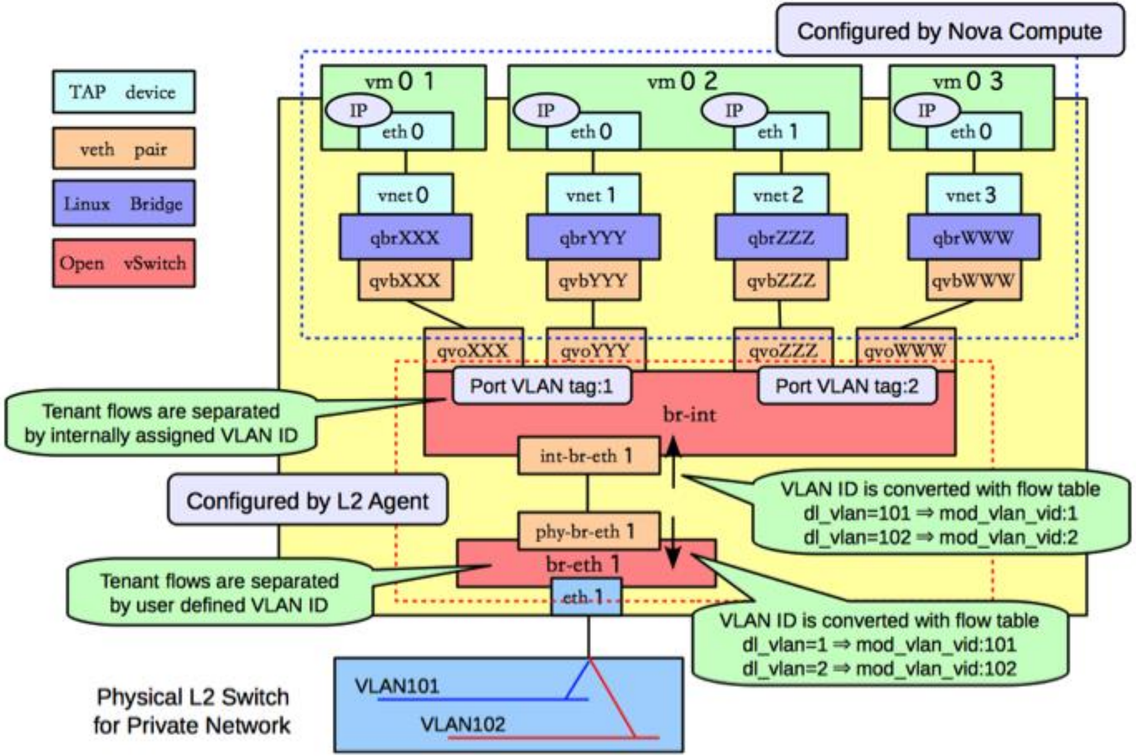

Linux Networking Devices on the Network Node

- Connection to public network
- Routers
- Floating IPs, SNAT
- DHCP

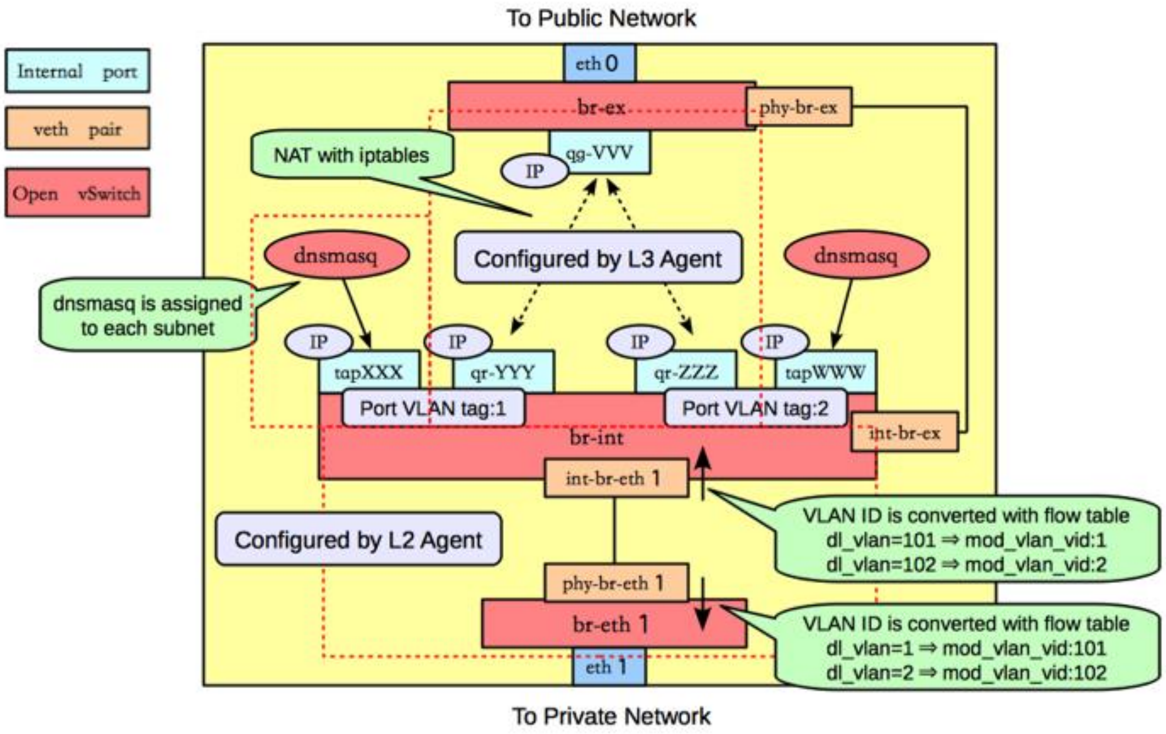

Typical Network

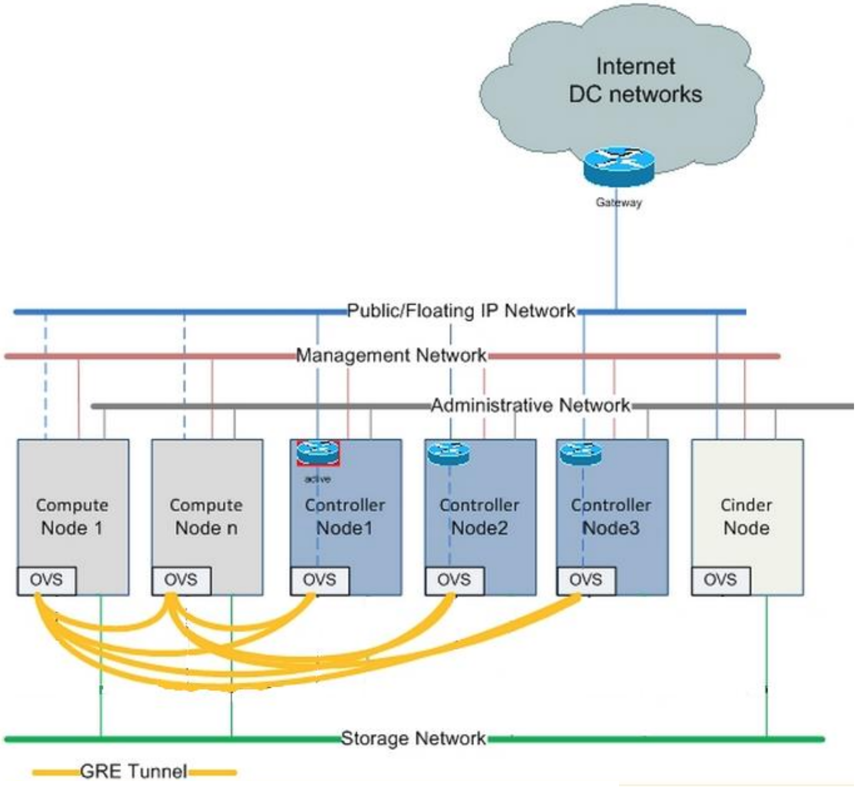

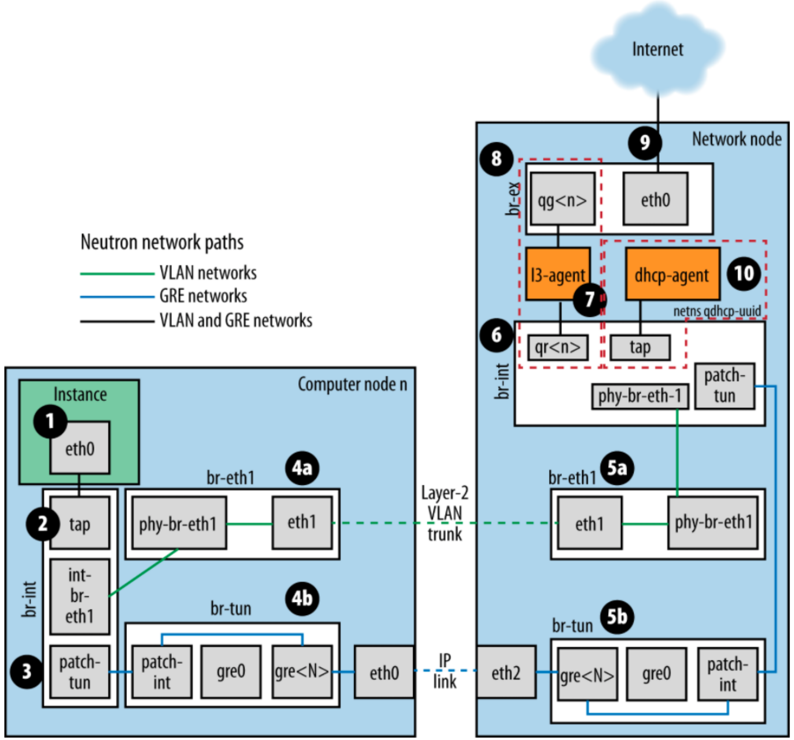

Image Source:
http://docs.mirantis.com/fuel/fuel-4.1/reference-architecture.html#advanced-network-configuration-using-open-vswitch

Openstack Network Troubleshooting:
http://docs.openstack.org/trunk/openstack-ops/content/network_troubleshooting.html

Neutron Advanced Services via Service Plugins

- Layer 3
  - Enables creation of router for connecting/attaching Layer 2 tenant networks that require L3 connectivity
  - Requires creation of Floating IP's for associating VM private IP address to public IP address
  - External Gateway for forwarding traffic outside of tenant networks
- LoadBalancer
  - Requires creation of a load balancer pool with members for a tenant
  - Enables creation of a virtual IP (VIP) that when accessed through the loadbalancer, directs the request to one of the pool members
  - Health Monitor Check for pool members
- VPN
  - Related to a specific tenant subnet and router
  - VPN connection represents the Ipsec tunnel established between two sites for the tenant
  - Requires creation of VPN – IKE – Ipsec - Connection
- Firewall
  - Provides perimeter firewall functionality on Neutron logical router for a tenant
  - Requires creation of Firewall – Policy – Rules

Neutron Router

- Run on network node
- L3 Agent in network node
- Uses Namespaces per tenant
- Icehouse proposed DVR
  - Can run on any compute node
  - Better E-W traffic

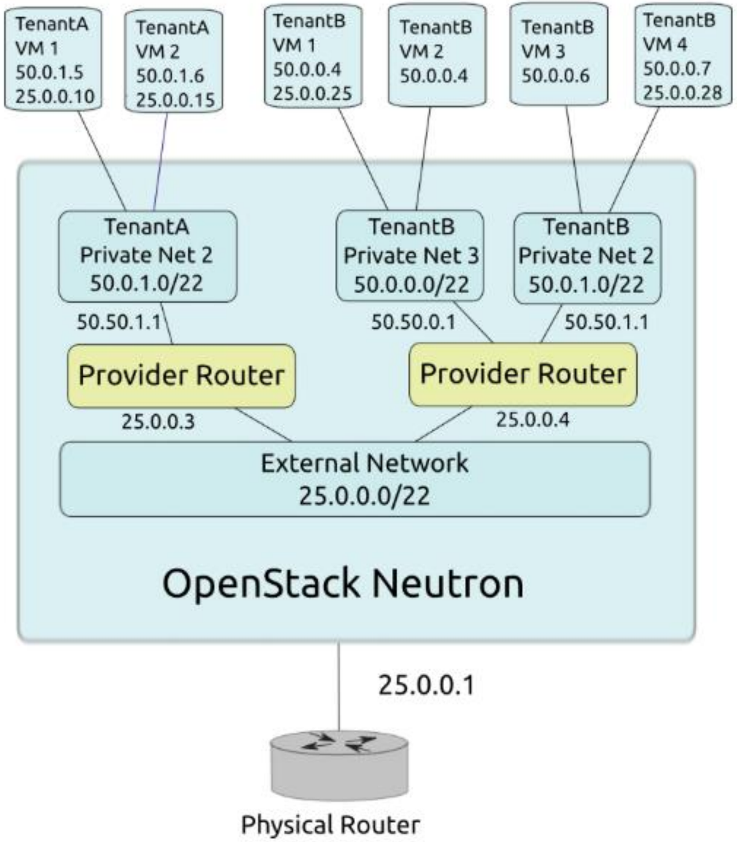

Source:https://developer.rackspace.com/blog/neutron-networking-l3-agent/

LBAAS

- Introduced in Grizzly
  - Backend - HAProxy Only
  - Round Robin, Least Connections, Source IP
  - Single Agent/node
- Havana
  - Multi-vendor support
  - HAProxy - Multiple agents/nodes, statistics, improved health monitor
- Features
  - Monitor – Ping,HTTP
  - Connection limits
  - Session Persistence

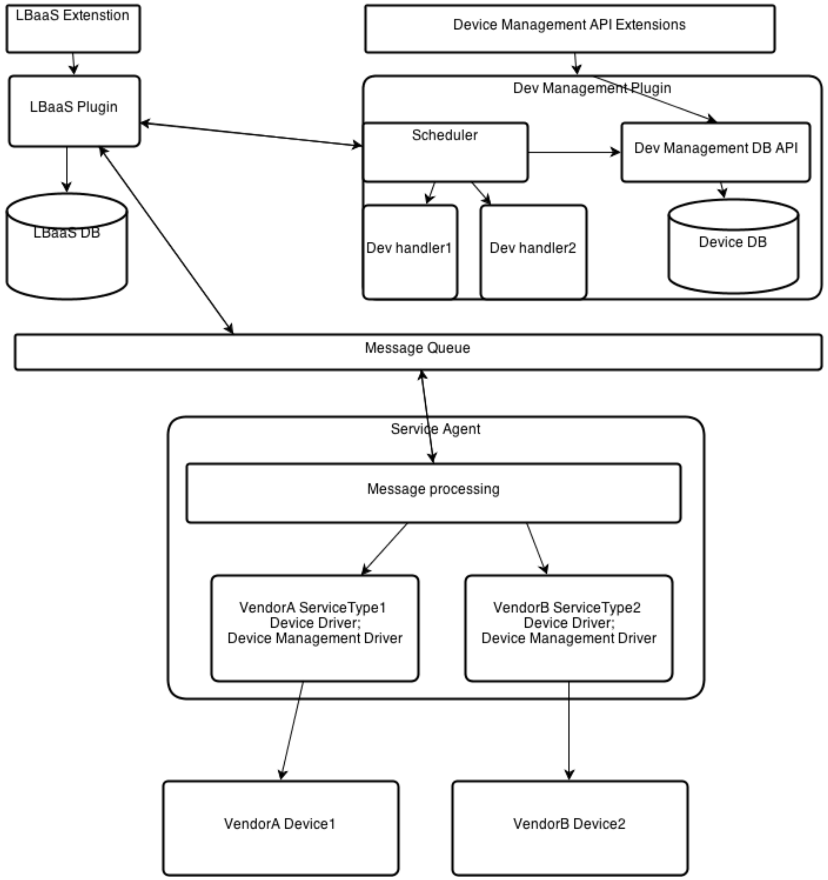

Source:https://wiki.openstack.org/wiki/Neutron/LBaaS/Architecture/Scheduler

VPNaaS

- Experimental introduction in Havana
- Based on OpenSwan
- Site to Site
- Pre-shared Keys

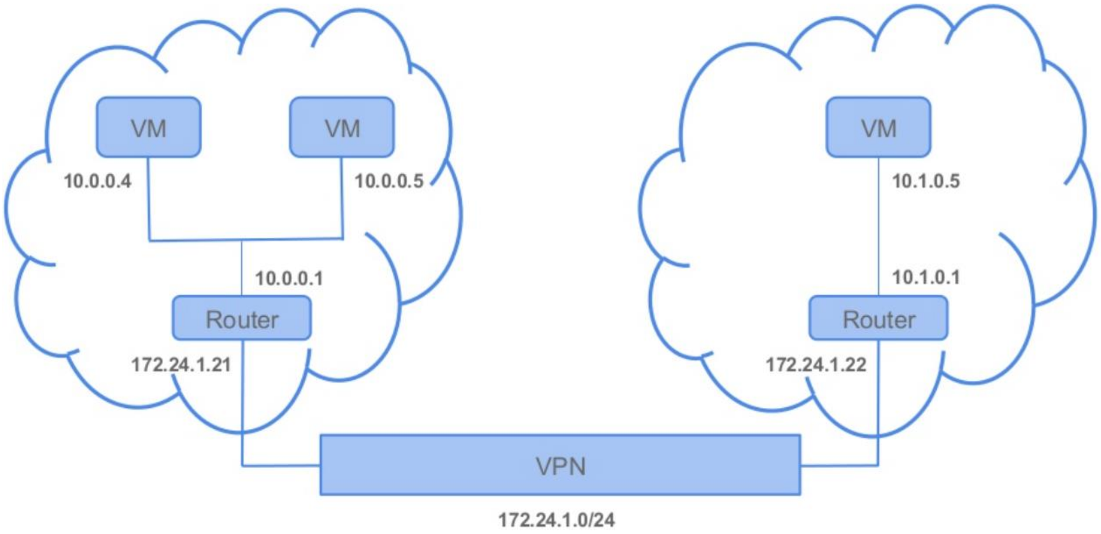

FWaaS

- Could be 3rd Party
- IPTables based
- Modeled after ASA and Checkpoint
- Default insertion is L3 GW

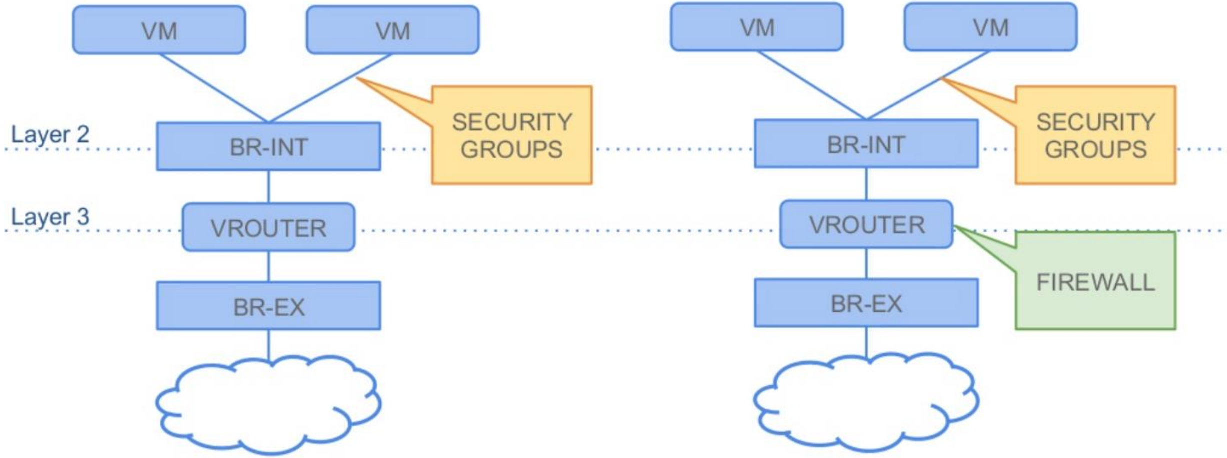

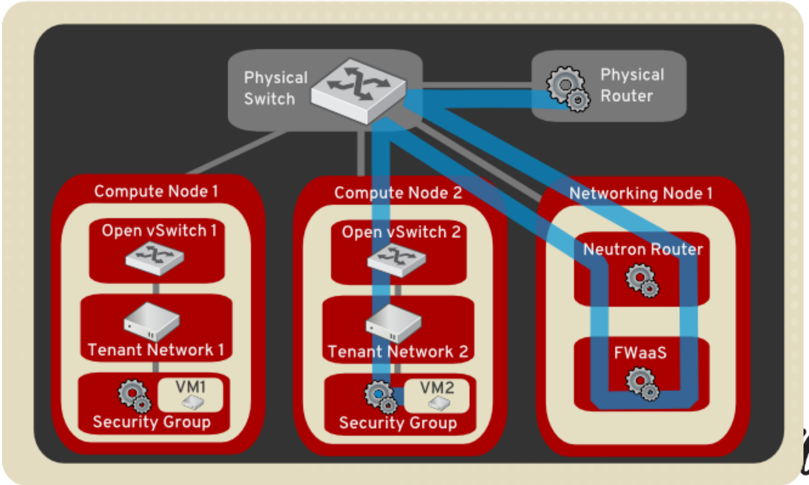

http://docs.openstack.org/admin-guide-cloud/content/install_neutron-fwaas-agent.html

Multi-Segment Networks

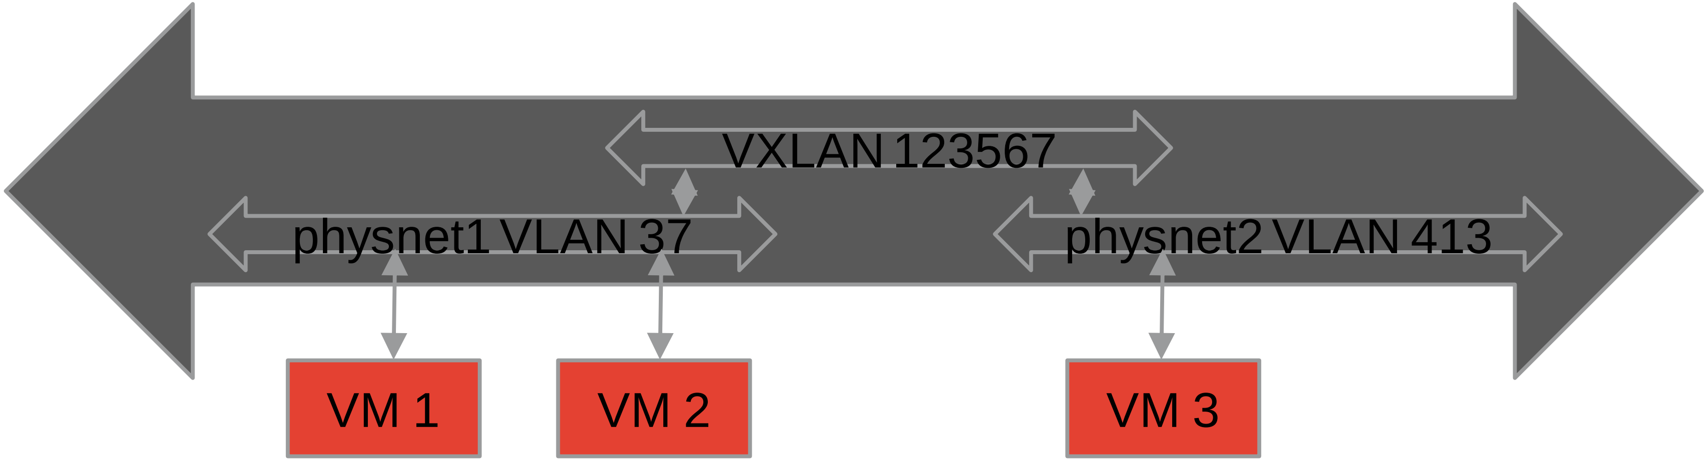

- Created via multi-provider API extension
- Segments bridged administratively (for now)
- Ports associated with network, not specific segment
- Ports bound automatically to segment with connectivity

Infrastructure Considerations & Solutions
Neutron

Applications Typically Start Like This

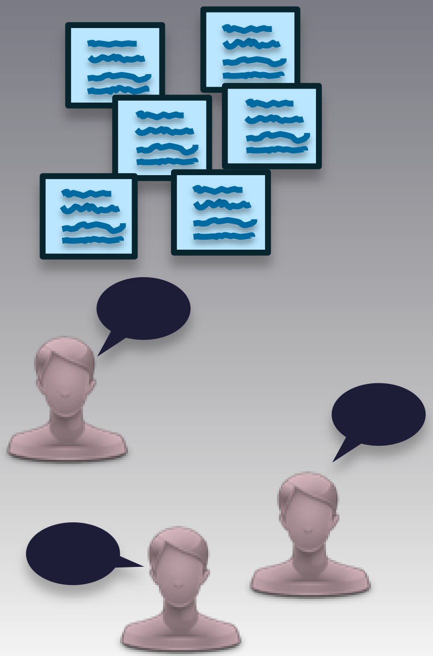

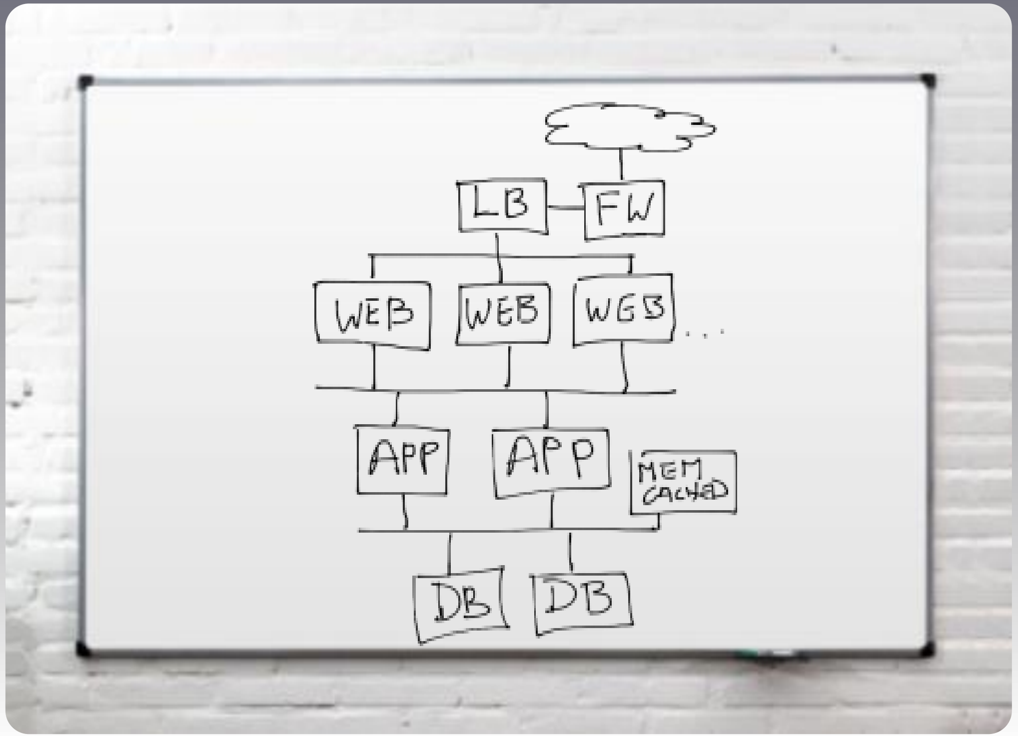

App Developers Define Their Own Network Topology

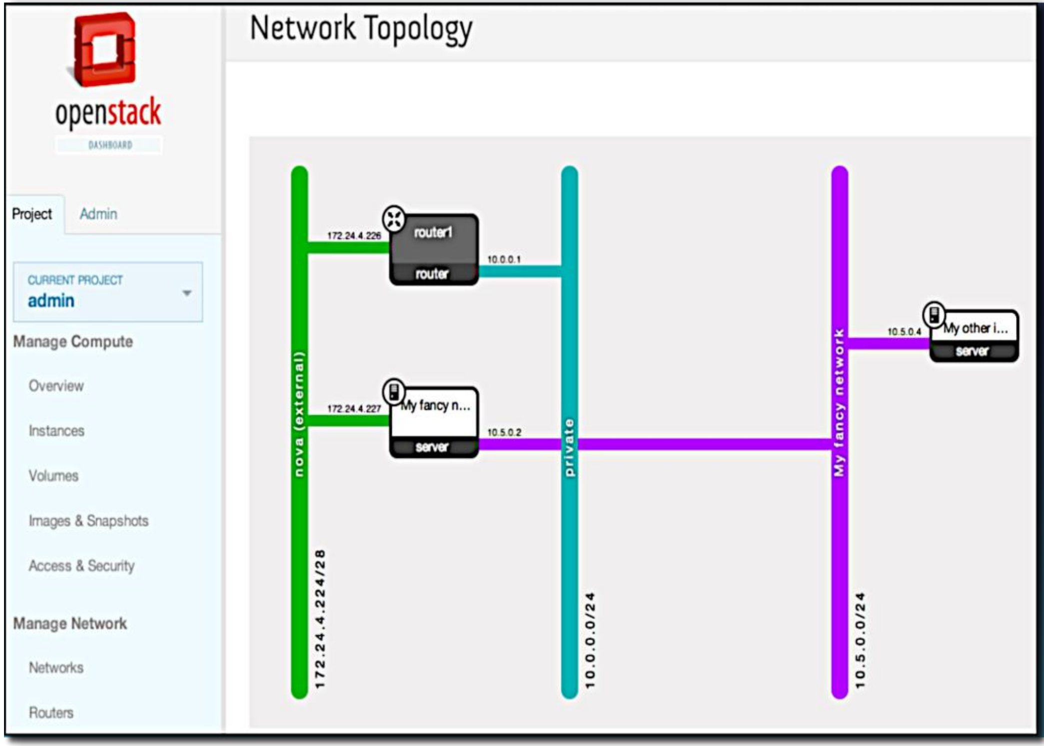

- Virtual, isolated networks
- Virtual routers
- Load-balancers
- Public, private addresses
- Access rights and security

OpenStack Design

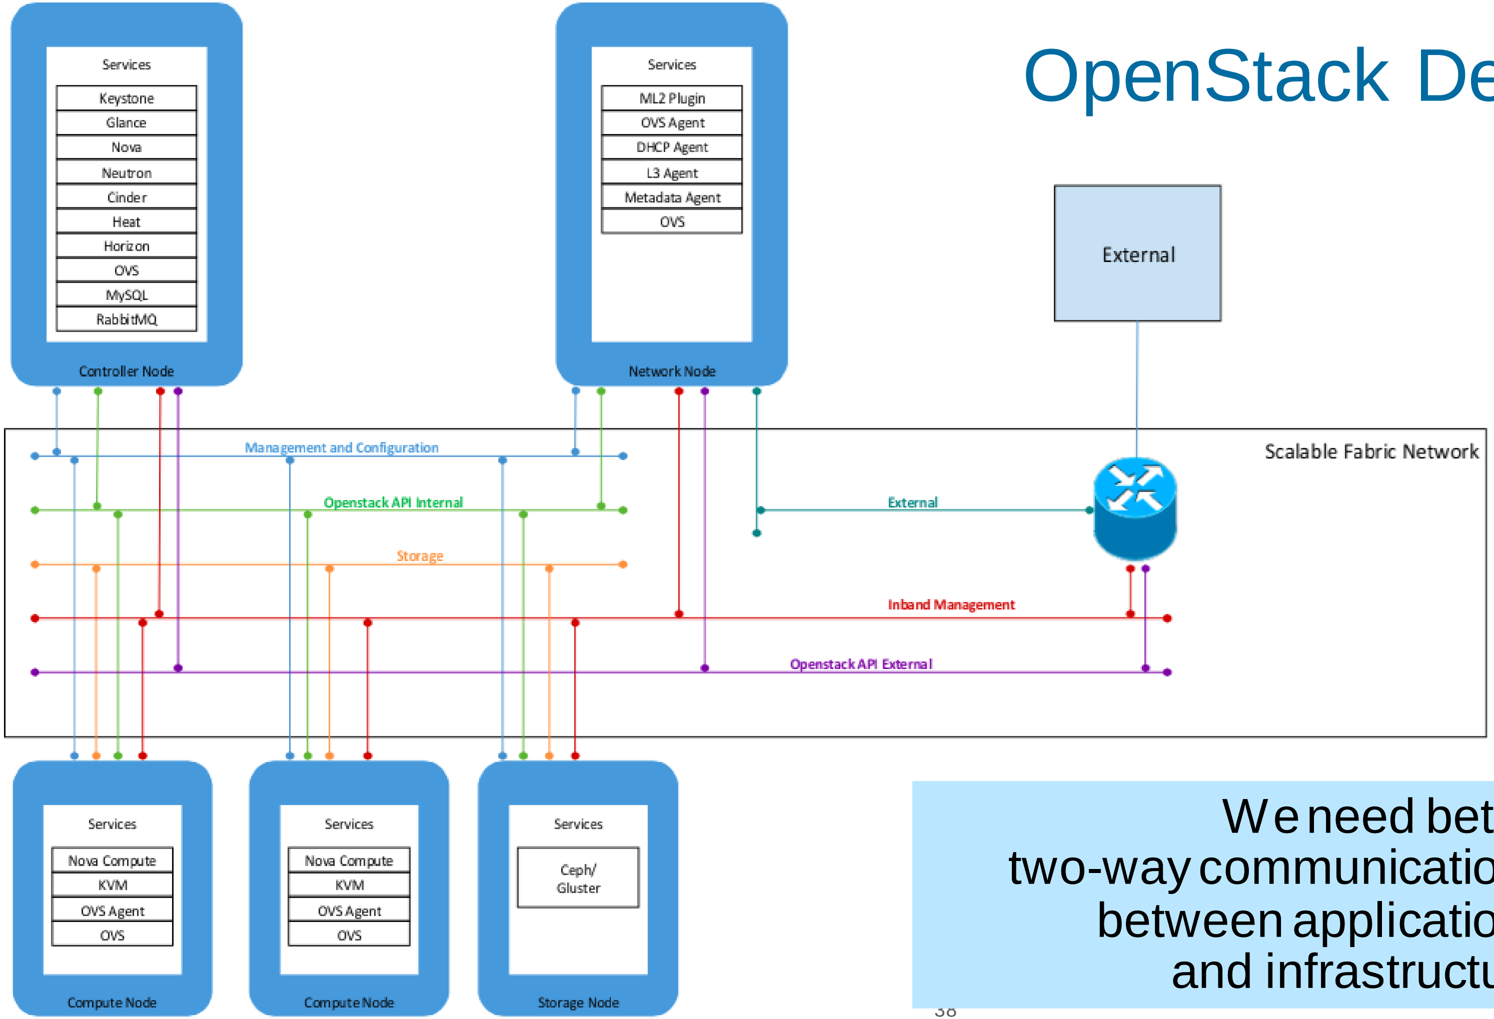

We need better two-way communications between applications and infrastructure

Infrastructure Components

Integrations	Release (IceHouse – May'14, Juno – Nov'14)
Cisco Physical Nexus Switches (N3K/5K/6K/7K/9K) Plugin and ML2 Driver	IceHouse release
Cisco Virtual Nexus 1000v Switch Plugin	Icehouse release
Cisco UCS VM-FEX ML2 Driver	Juno release
Cisco Virtual Cloud Services Router 1000v Service Plugins (L3)	Juno release
Cisco Virtual Cloud Services Router 1000v Service Driver (VPN)	IceHouse release
Cisco Dynamic Fabric Automation Fabric ML2 Driver	Juno release
Cisco Application Policy Infrastructure Controller ML2 Driver	Juno release

Neutron Cisco Nexus Plugin

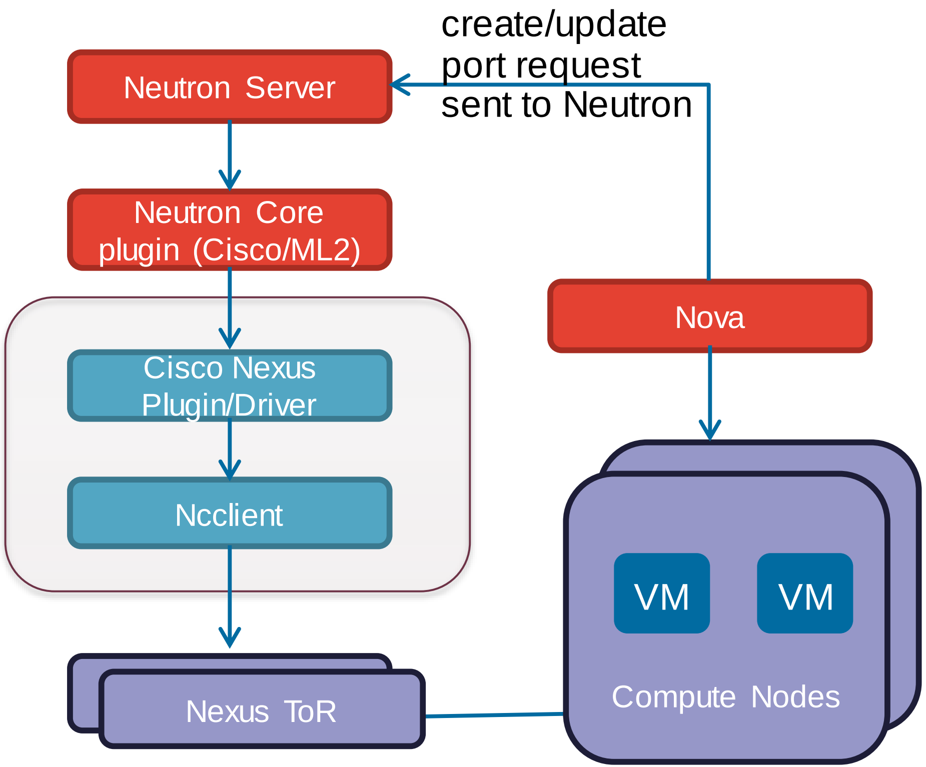

Demo of Nexus plugin at the end !

Benefits

- Works with Nexus 3k/5k/6k/7k/9k
- Support for Neutron Provider Networks
- Dynamic VLAN and SVI provisioning/deprovisioning on ToR
- Network based Overlays using VXLAN

Neutron Cisco Nexus1000v Plugin (KVM)

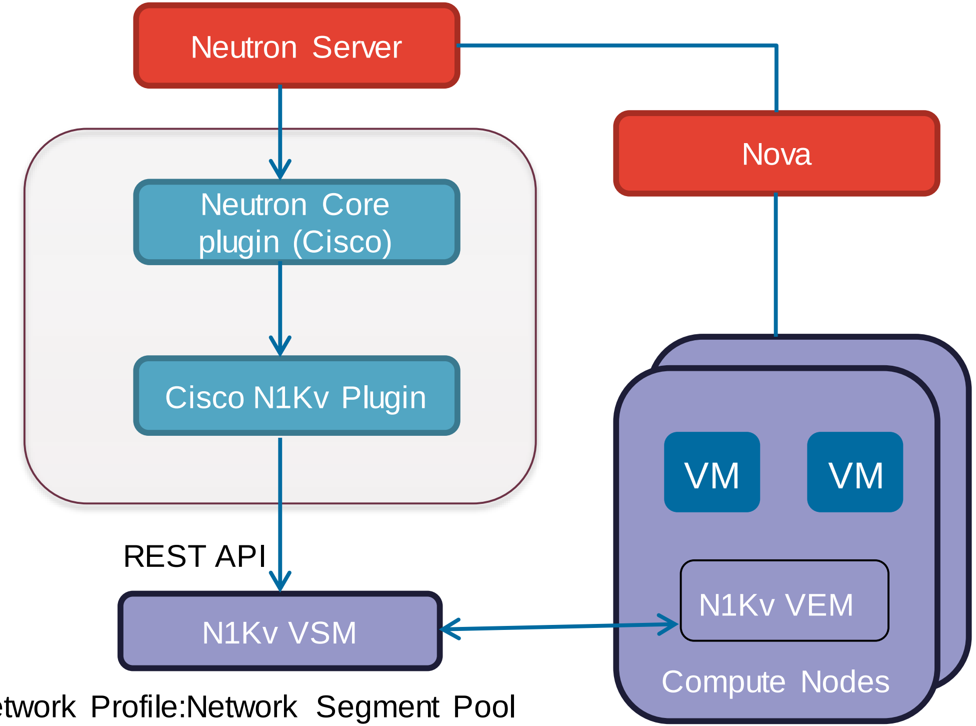

Network Profile:Network Segment Pool

Policy Profile:Port Profile,

Neutron N1Kv specific API extensions usage –

```
neutron network-profile-create PROFILE_NAME
vlan --segment_range 400-499
```

Network Profile (admin)

```
neutron net-create NETWORK_NAME --
n1kv:profile_id PROFILE_ID
```

```
neutron policy-profile-list
```

Policy Profile defined in VSM (periodic polling)

```
neutron port-create NETWORK_NAME --
n1kv:profile_id PROFILE_ID
```

Policy Profile

Benefits:

- Network Profiles – VLAN, VXLAN (multicast/unicast), Trunk
- Policy Profiles – ACLs, QoS
- VXLAN Gateway Service VM

Neutron Cisco UCS VM-FEX Driver (KVM)

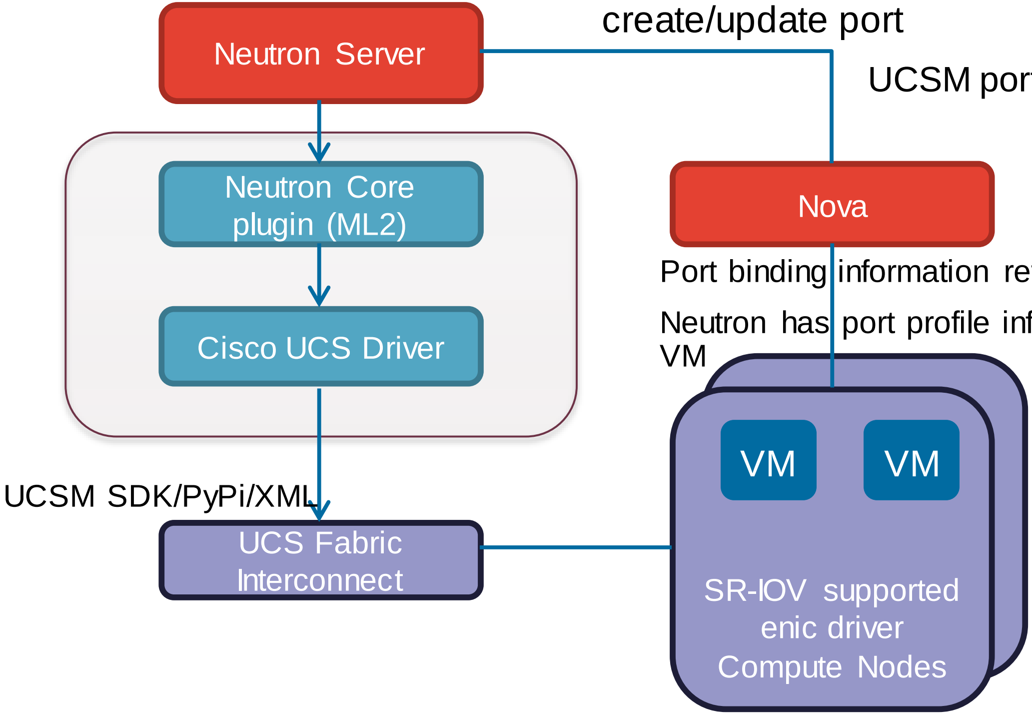

<interface type='hostdev' managed='yes'>
<mac address='fa:16:3e:f1:dd:e6'/>
....
<virtualport type='802.1Qbh'>
<parameters profileid='Net1Profile'/>
</virtualport>
</interface>

UCSM port profile

Benefits:

- Bypasses the vswtich
- Improves throughput

Neutron's Routing Reference Implementation

- Limitations
  - x86 box for L3 services
  - HA (Neutron logical router)

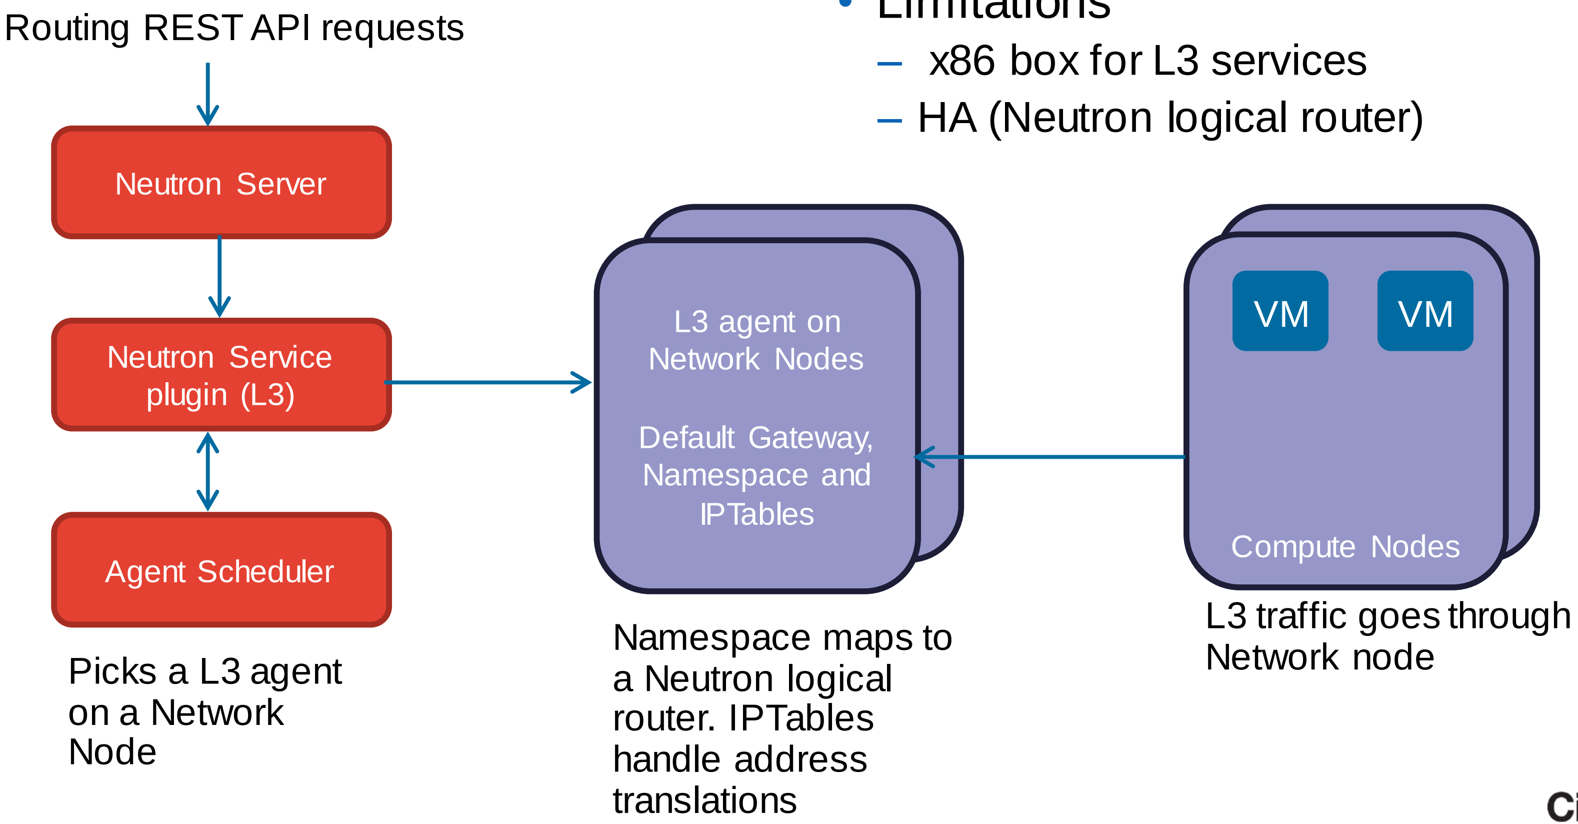

Neutron + NFV (Cisco Driven Architecture)

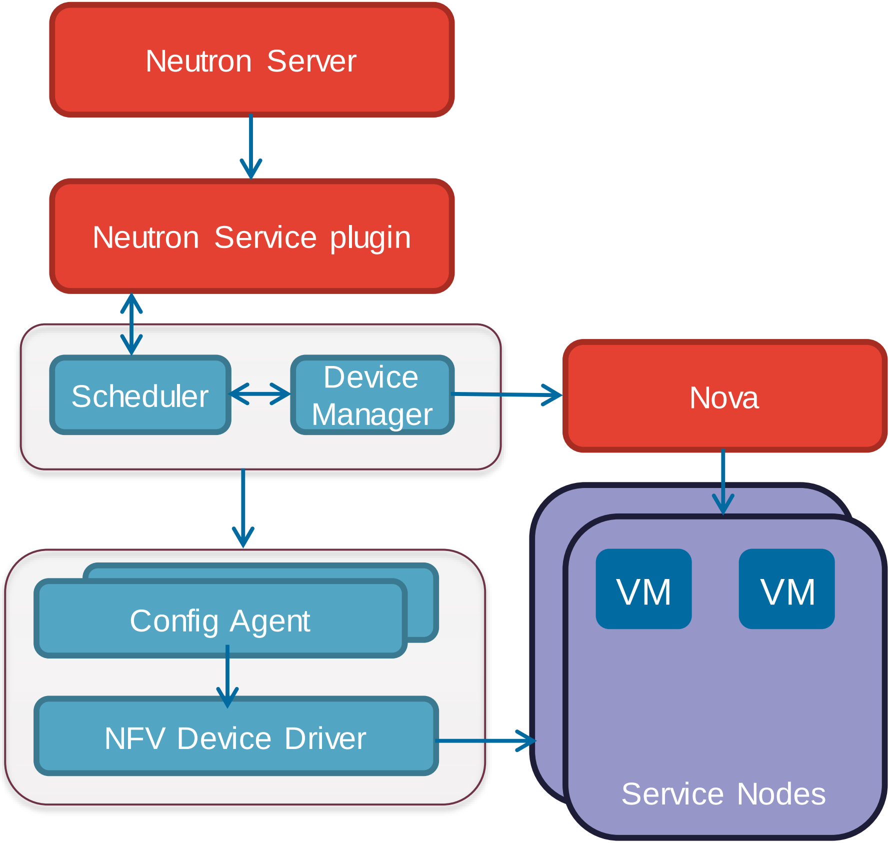

- Service Plugins
  - Management of logical resources
- Scheduler
  - Select Hosting device
- Device Manager
  - Lifecycle management of devices (Spinning up of NFV devices)
  - Book-keeping of processing capacity in devices (Avoid over allocation)
- Config Agent
  - Apply configuration to devices
  - Monitor health devices

Neutron Cisco CSR1000v for Neutron L3 Service

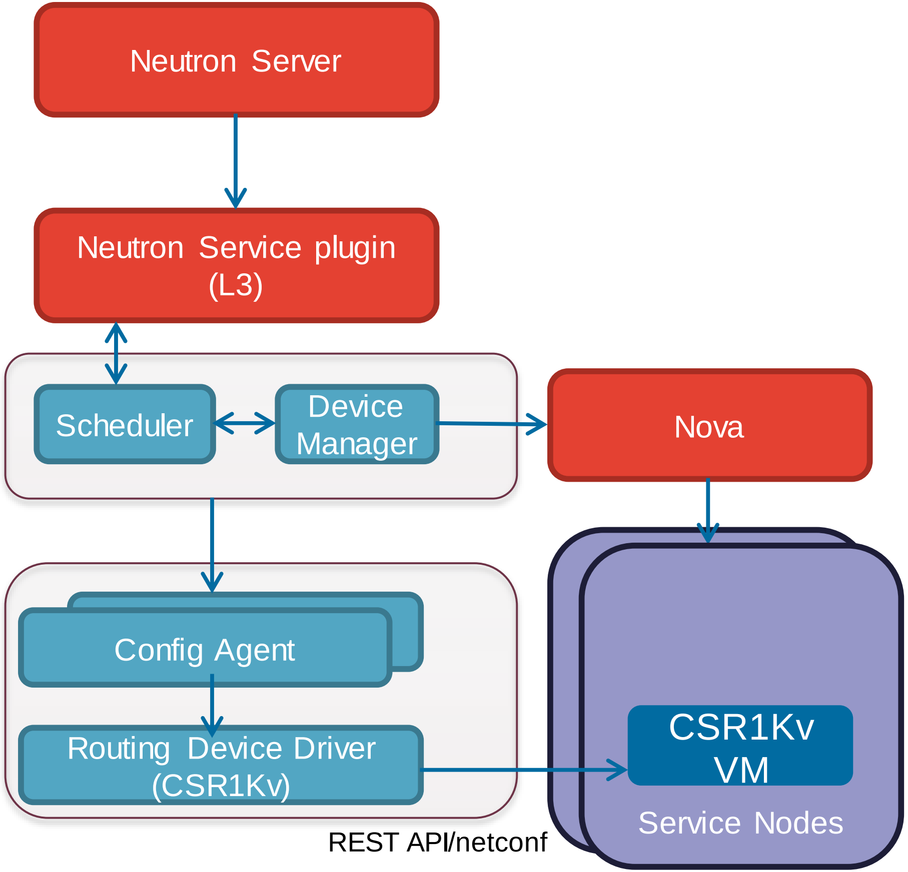

- Mapping of Neutron reference L3 implementation -
  - Linux namespaces - CSR1Kv VRF
  - Router ports (qr) on bridge – CSR1Kv VLAN sub interfaces
  - Gateway ports (qg) on bridge - CSR1Kv VLAN sub interfaces
  - Linux IPTables – CSR1Kv NAT
- Benefits
  - Available as NFV services
  - Scalable solution
  - Integrates with N1Kv

Example CSR1Kv Config for a Neutron Logical Model

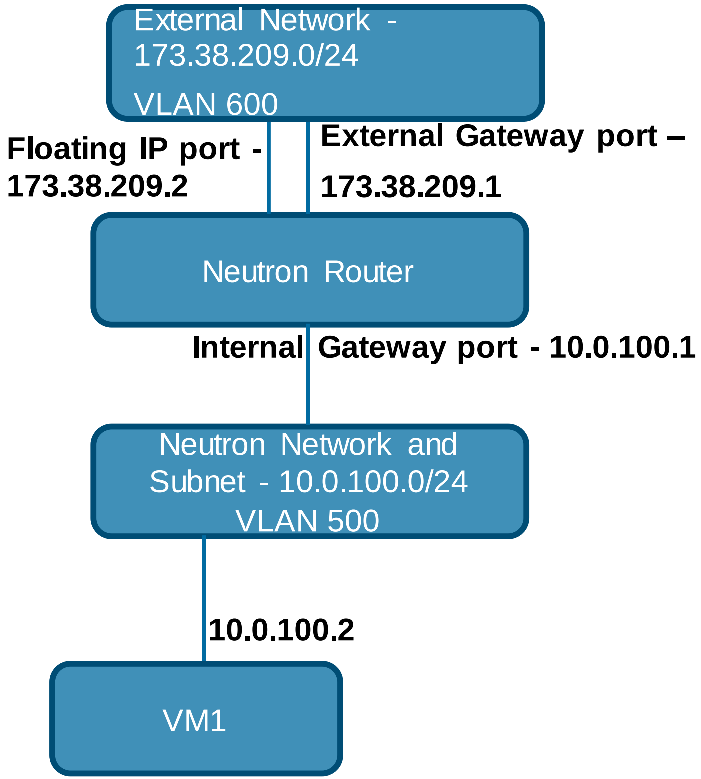


```
interface GigabitEthernet2.500
 encapsulation dot1Q 500
 ip vrf forwarding nrouter-462986b8
 ip address 10.0.100.1 255.255.255.0
 ip nat inside

interface GigabitEthernet2.600
 encapsulation dot1Q 600
 ip vrf forwarding nrouter-462986b8
 ip address 173.38.209.1 255.255.255.0
 ip nat outside

ip nat inside source static 10.0.100.2 173.38.209.2
 vrf nrouter-462986b8 match-in-vrf
```

Neutron Cisco CSR1000v VPN Service Driver (KVM)

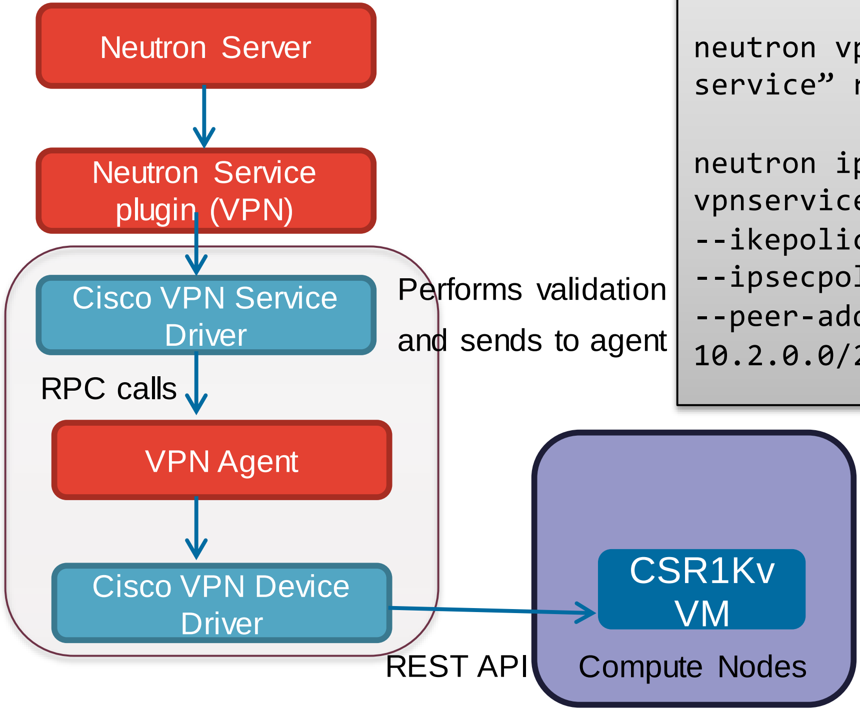


```
neutron vpn-ikepolicy-create ikepolicy1

neutron vpn-ipsecpolicy-create ipsecpolicy1

neutron vpn-service-create --name myvpn --description "My vpn
service” router1 mysubnet

neutron ipsec-site-connection-create --name vpnconnection1 \ --
vpnservice-id myvpn \
--ikepolicy-id ikepolicy1 \
--ipsecpolicy-id ipsecpolicy1 \
--peer-address 172.24.4.23 --peer-id 172.24.4.23 --peer-cidr \
10.2.0.0/24 --psk secret
```

Benefits

- CSR1Kv secure VPN qualified solution
- Unlock rich CSR1Kv features into OpenStack

Neutron Cisco Dynamic Fabric Automation(DFA) Driver

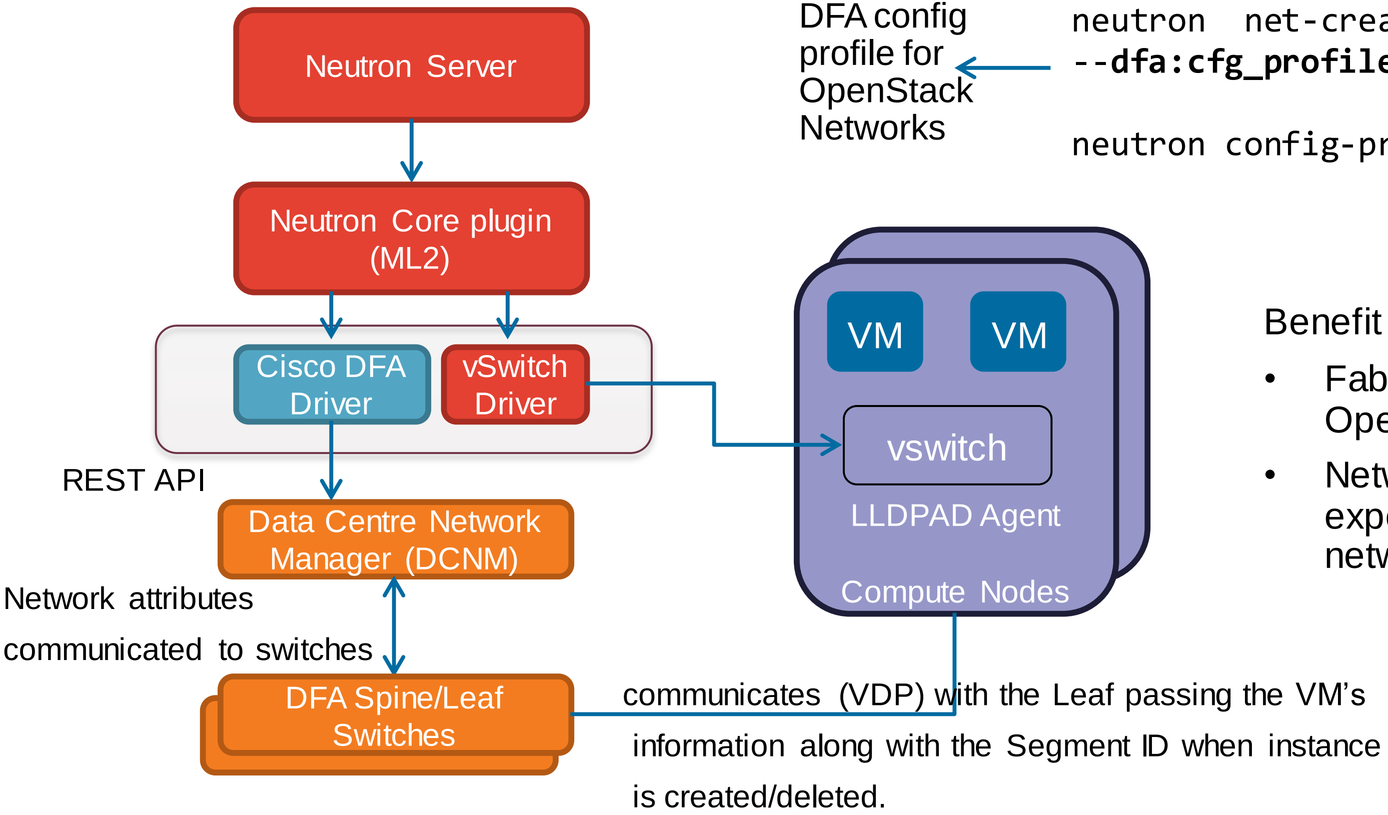

DFA config profile for OpenStack Networks

```
neutron  net-create NETWORK_NAME
--dfa:cfg_profile_id PROFILE_ID

neutron config-profile-list
```

Benefit

- Fabric based overlays with OpenStack
- Network Fabric Advantages exposed to OpenStack networks

Evolving the Neutron API

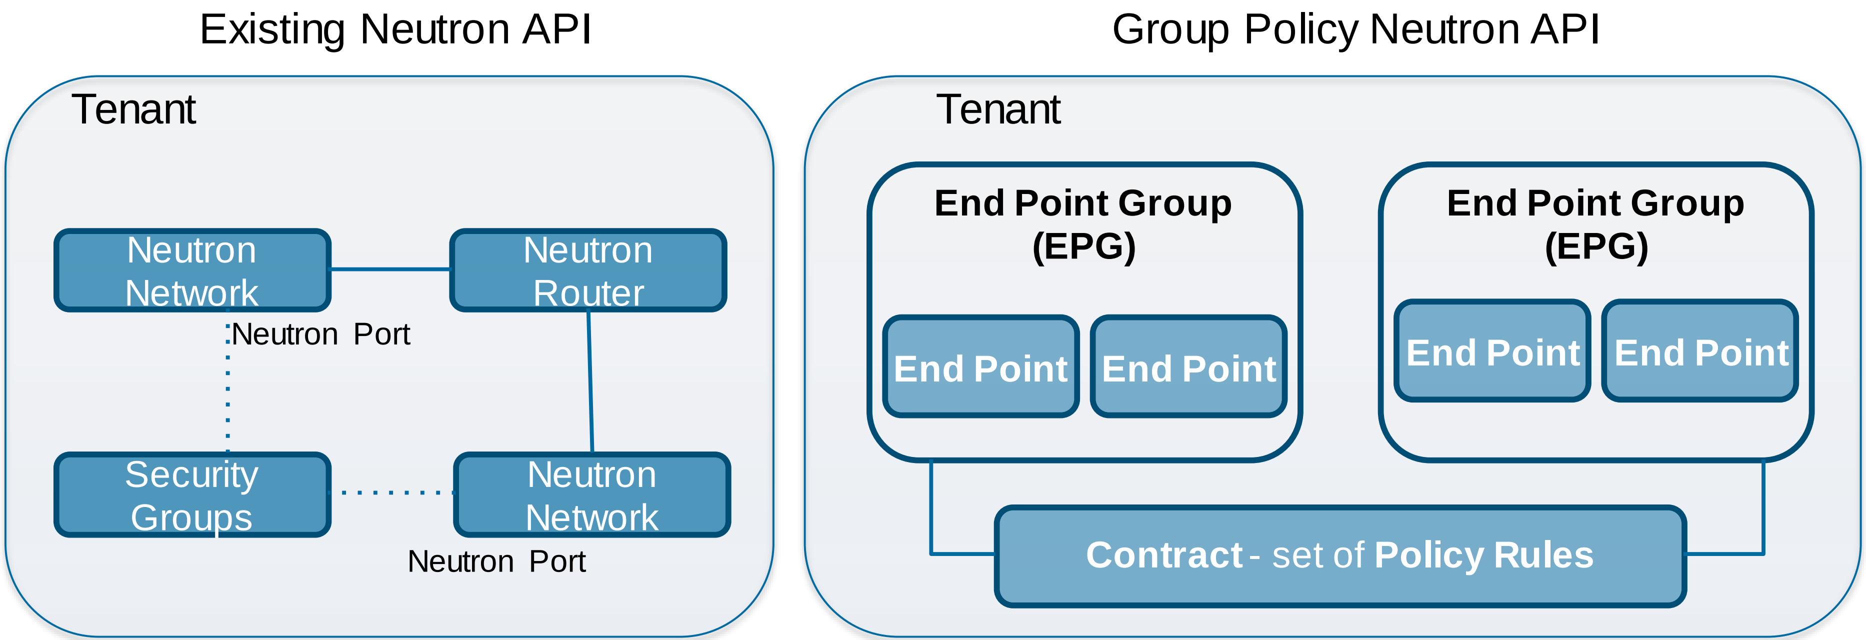

API to provide clear separation between Application developer and Infrastructure manager

- Application developer doesn't need to care about network centric resources such as Networks/Routers etc (existing Neutron API)
- Infrastructure Manager doesn't need to care about application requirements such as what ports requires to be opened for the applications

Neutron Cisco CSR1000v for Neutron L3 Service

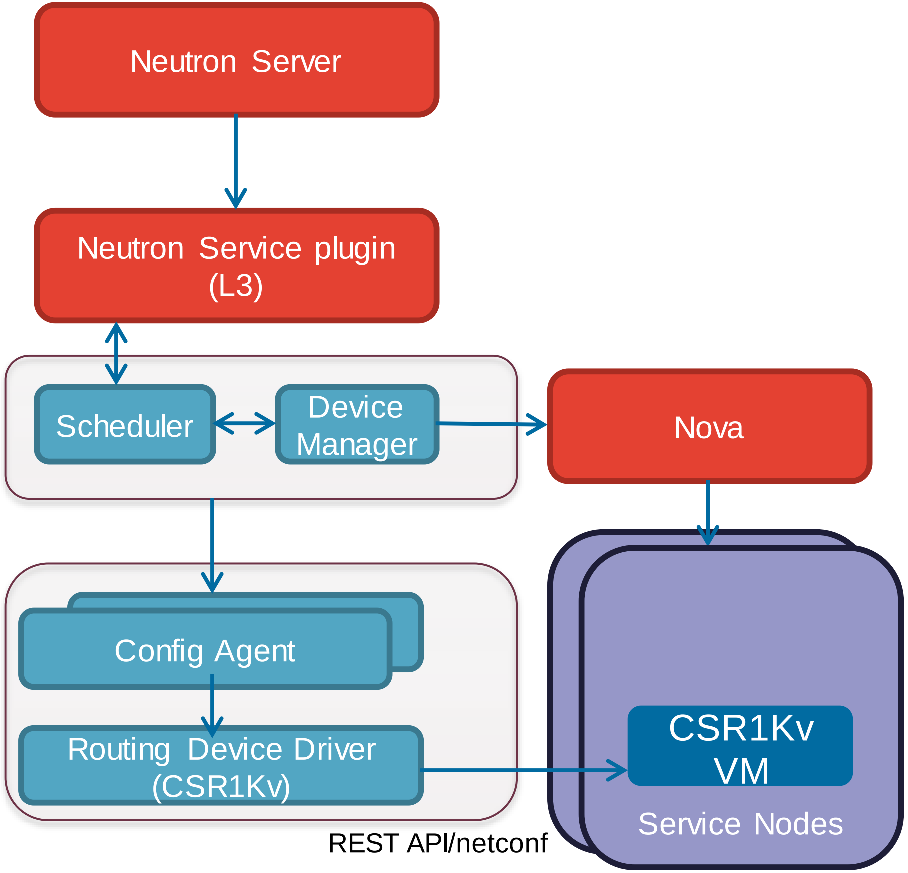

- Mapping of Neutron reference L3 implementation -
  - Linux namespaces - CSR1Kv VRF
  - Router ports (qr) on bridge – CSR1Kv VLAN sub interfaces
  - Gateway ports (qg) on bridge - CSR1Kv VLAN sub interfaces
  - Linux IPTables – CSR1Kv NAT

- Benefits
  - Available as NFV services
  - Scalable solution
  - Integrates with N1Kv

Neutron Cisco CSR1000v VPN Service Driver (KVM)

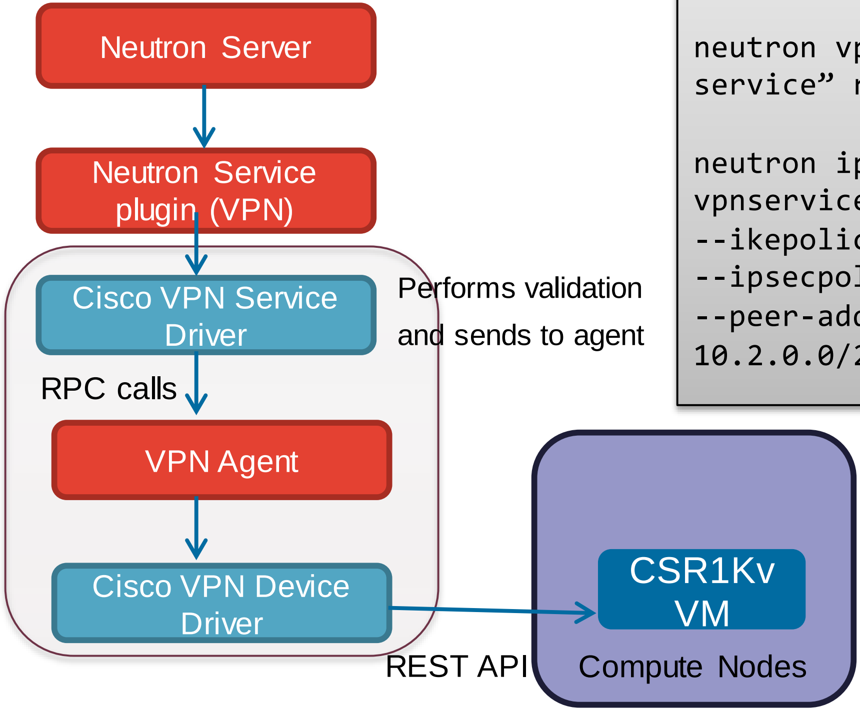


```
neutron vpn-ikepolicy-create ikepolicy1

neutron vpn-ipsecpolicy-create ipsecpolicy1

neutron vpn-service-create --name myvpn --description "My vpn
service" router1 mysubnet

neutron ipsec-site-connection-create --name vpnconnection1 \ --
vpnservice-id myvpn \
--ikepolicy-id ikepolicy1 \
--ipsecpolicy-id ipsecpolicy1 \
--peer-address 172.24.4.23 --peer-id 172.24.4.23 --peer-cidr \
10.2.0.0/24 --psk secret
```

Benefits

- CSR1Kv secure VPN qualified solution
- Unlock rich CSR1Kv features into OpenStack

Benefits of Cisco UCS Integrated Infrastructure

Foundation for Scalable Clouds

Cisco UCS
Unified, Programmable,
Rapid Provisioning, Scalable

Cisco Nexus
Scalable, Secure, Network Fabric

Storage Partners
Choice of: Direct Attach, NAS/SAN

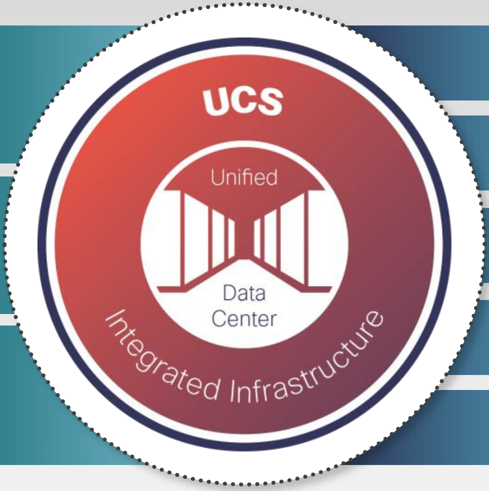

Standard "Building Blocks"

Performance, Scalability, Availability

Secure Multi-Tenancy

ACI Ready

Hybrid and Intercloud Enabled

Simplifies operations -- Maximises ROI -- Accelerates deployment

Cisco UCS - Powering Applications at Every Scale

Cisco UCS for Cloud Deployment

Adding Value and Innovation

Programmability

- Cisco UCS programmability is accomplished through published APIs and Python scripting

Centralised Management

- Racks of computing and storage nodes can be managed through a single GUI
- Quickly spin up new or replacement computing/storage assets

Standardisation and Consistency

- Service profiles enables standardisation across compute assets and speeds new installation

Cost Effective and Customisable

- High memory density supports more guest OS instances in fewer physical servers
- Flexibility of hardware options enables selecting the best-fit server for computing, storage, and controller nodes

Starter Edition Topology

- PortChannel from each Cisco UCS fabric interconnect to ToR
- Fabric interconnects in end-host mode, with dynamic pinning of server vNICs to uplink PortChannels

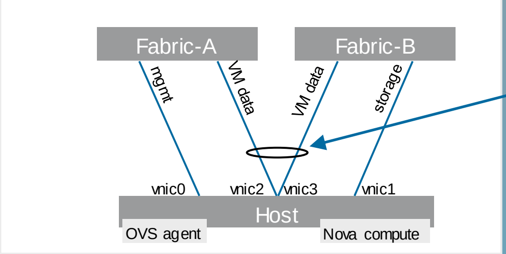

- Link aggregation on OVS
- Load balancing

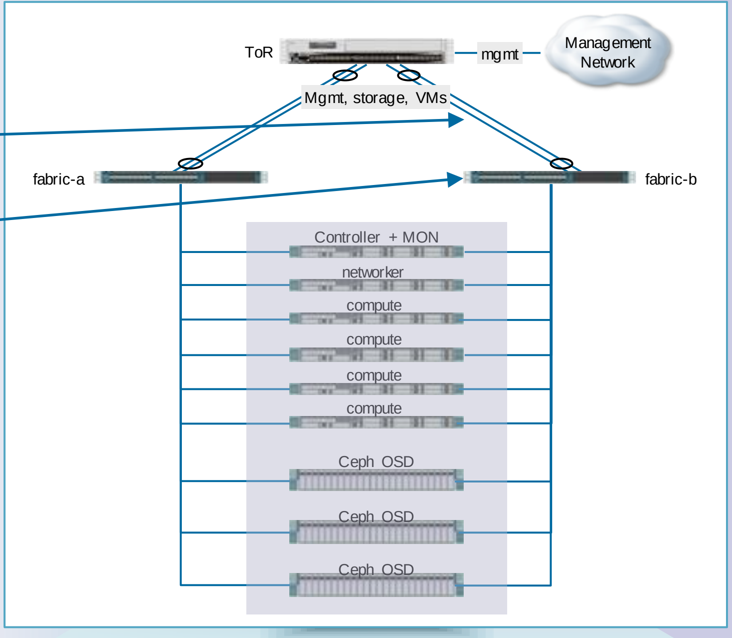

UCSO Topology

Starter Edition: Controller and Network Nodes

- OpenStack services on single controller node:
  - Horizon
  - Keystone
  - Glance API: as Ceph client (operating system images)
  - Nova scheduler
  - Neutron server
  - AMQP
  - mySQL
  - Cinder volume: as Ceph clients for persistent block storage on Ceph cluster for boot from volume and data
  - Heat: installed as foundation for subsequent Starter releases that involve Heat templates
- Packstack installer on controller node
- Networker node (other Neutron services, OpenvSwitch agent)
  - For scaling to production-level deployment
  - RHEL OpenStack installer allows networker co-location with controller

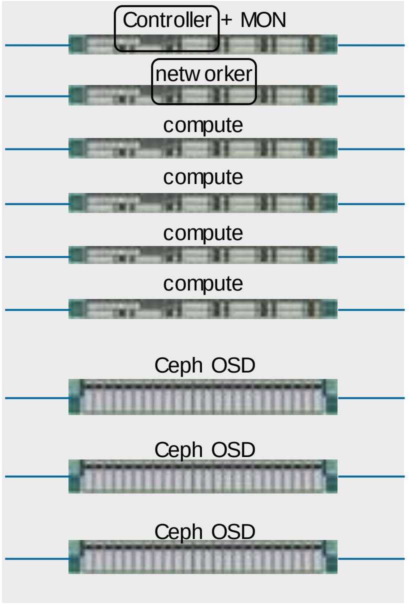

UCS C3160 as Ceph Storage Node

- C3160 is ideal for Ceph Object Stores and as well as block based Ceph deployments
- Optimised for high throughput workloads
- Power efficient server
- Petabytes of local storage in a standard 19inch rack
- Investment protection and reduced operational cost

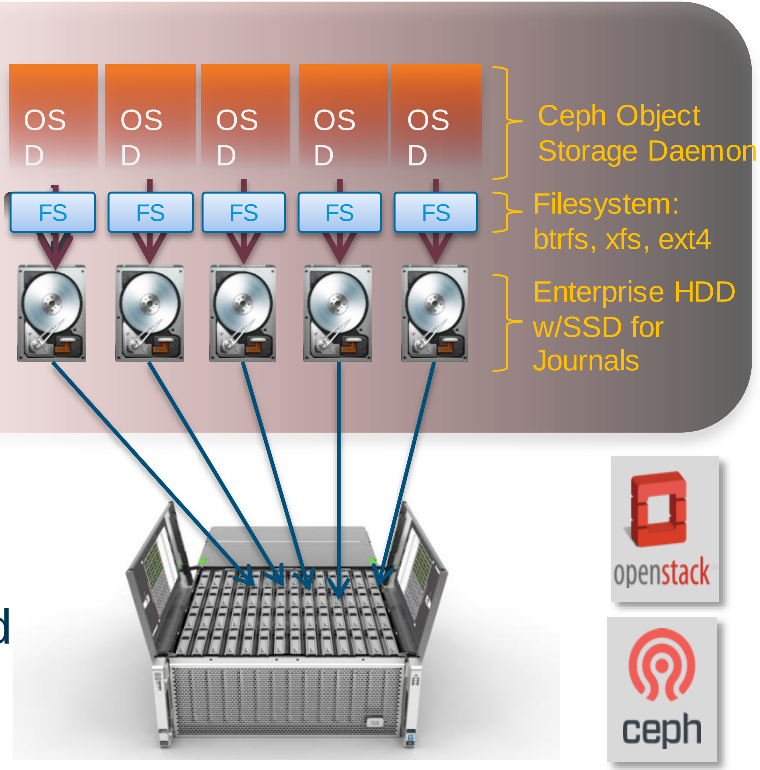

Policy w/ OpenStack

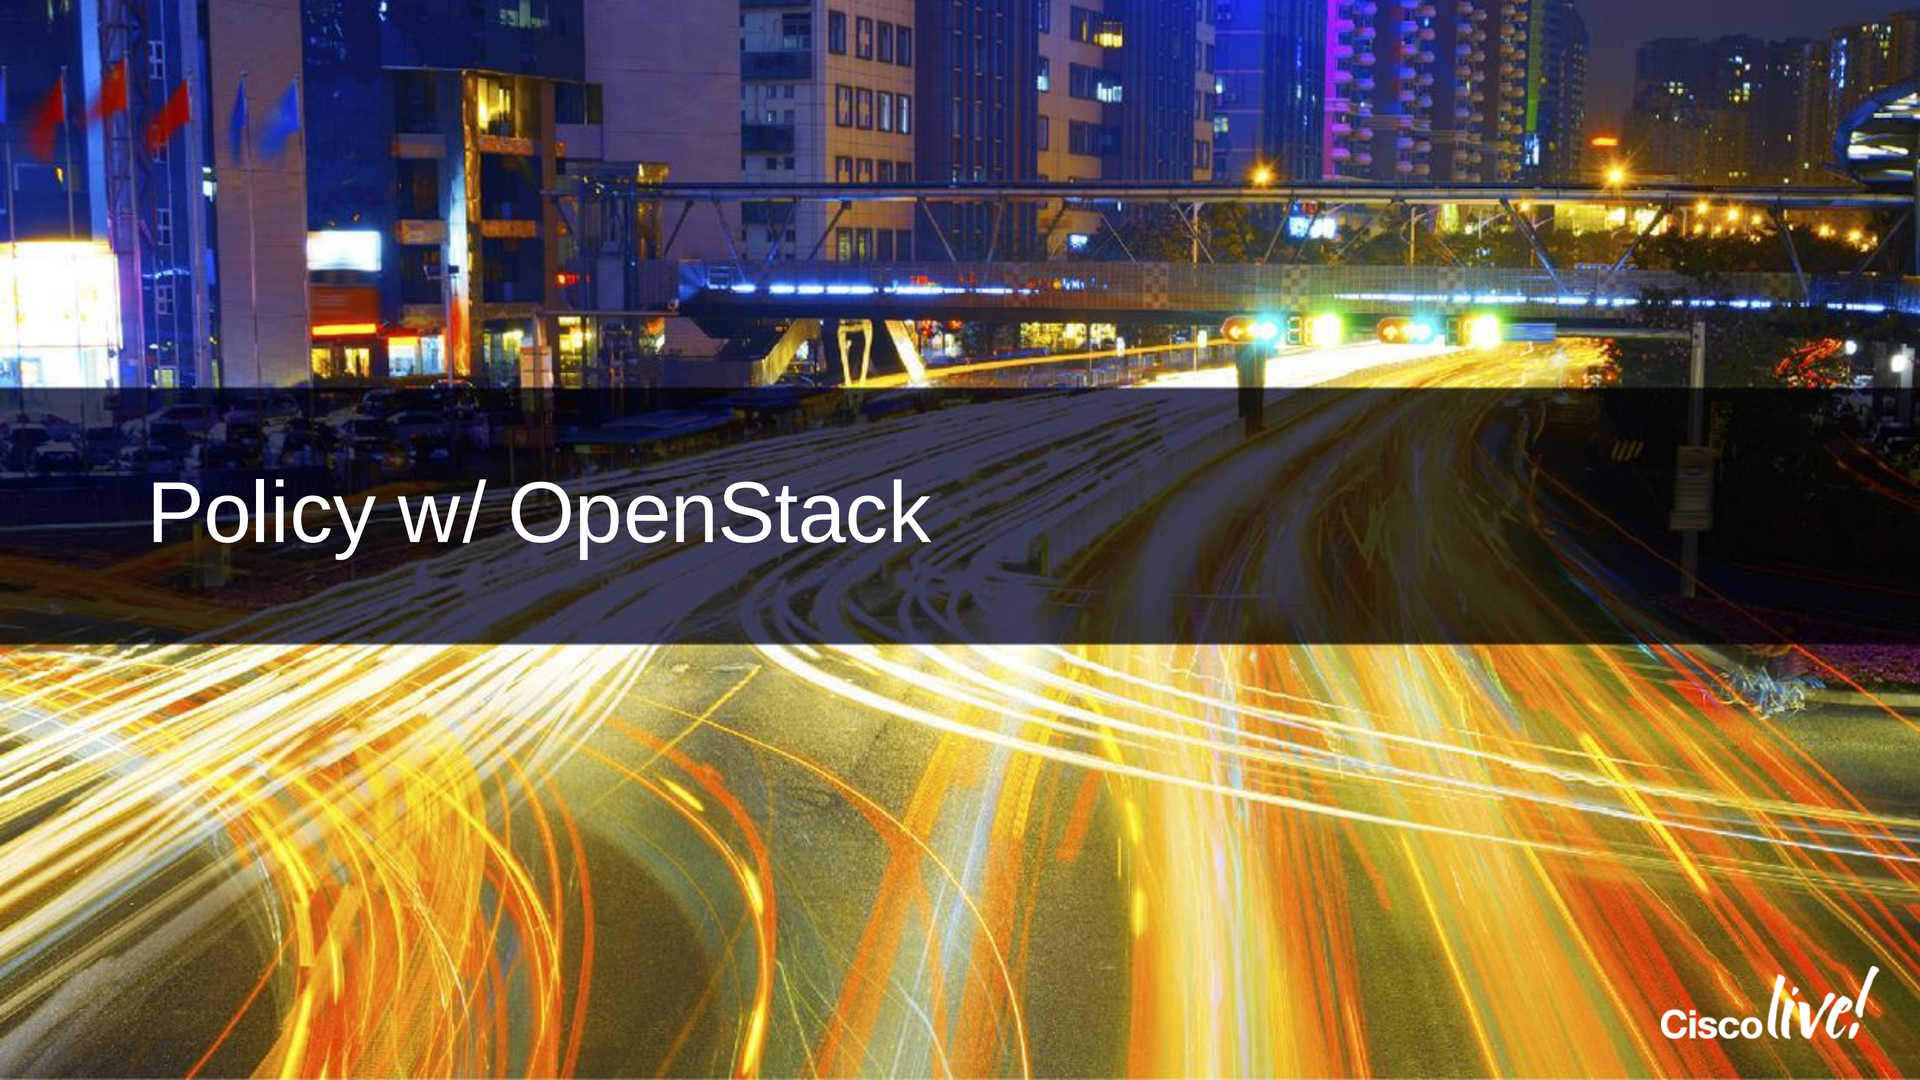

What's Wrong with OpenStack Networking Today?

Cloud Application Model

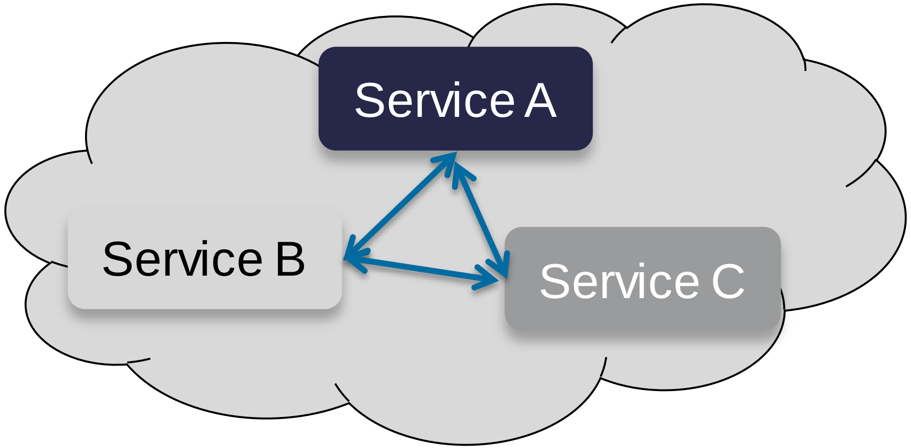

- No broadcast / multicast
- Resilient / Fault Tolerant
- Scalable Tiers
- Built around loosely coupled services
- Don't care about IP addresses

Neutron Model

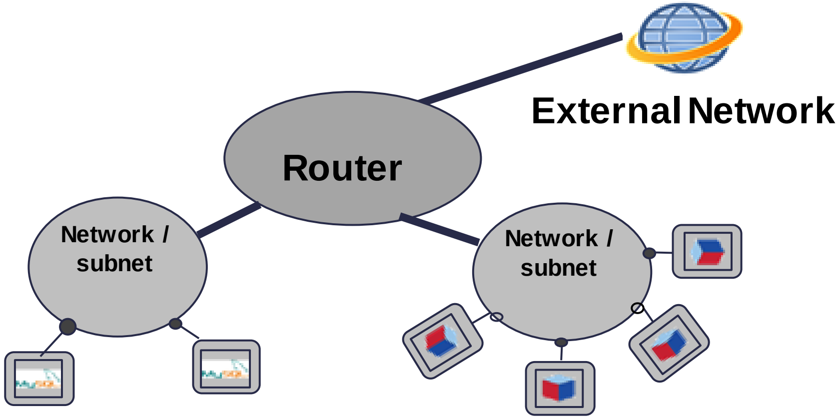

- L2 / Broadcast is the base API!
- Network / routers / subnets
- Based on existing networking models
- No concept of dependency mapping or intent

Where Can We Do Better

Separation of Concerns

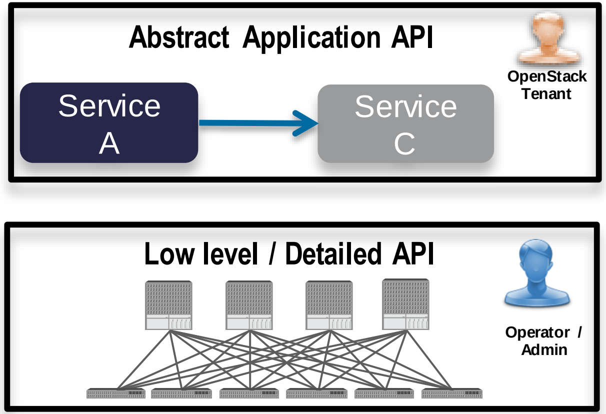

- Separate application requirements from low level APIs
- Separate tenant from operator

Dependency Mapping

Service A consumes service B and Service C

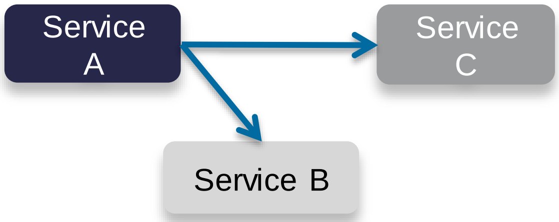

- Build self-documenting dependency maps of tiers of an application

Enable Network Services

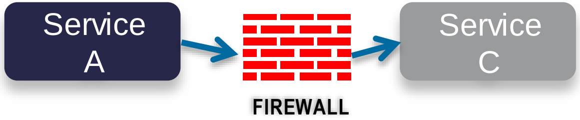

- Define network service chains between tiers of an application *without* low level configuration

Introducing Group-Based Policy

- Intent-based API for describing application requirements
- Separates concerns of tenants and operators
- Captures dependencies between tiers of an application
- Plugin model
  - Supports mapping to Neutron APIs
  - Supports “native” SDN drivers

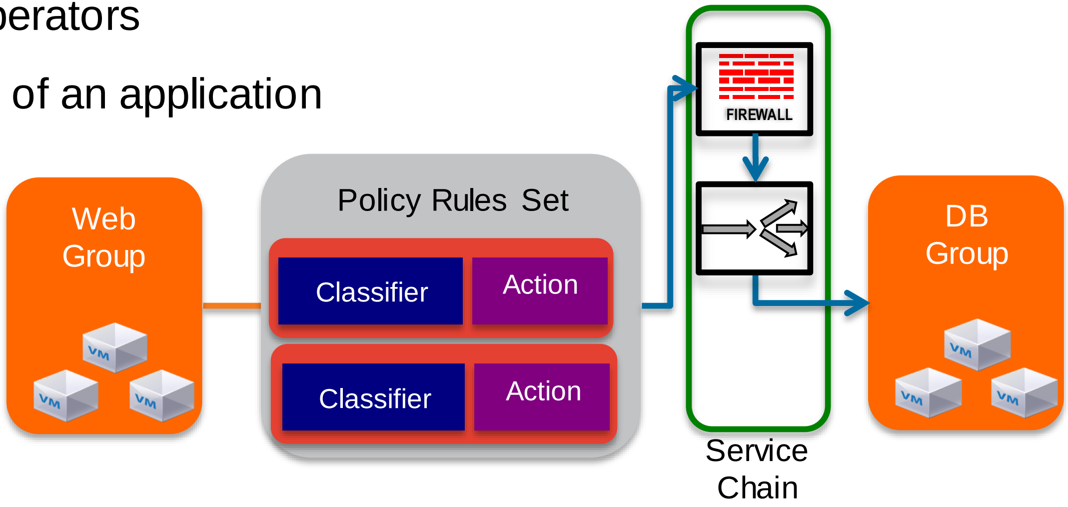

Group-Based Policy Model

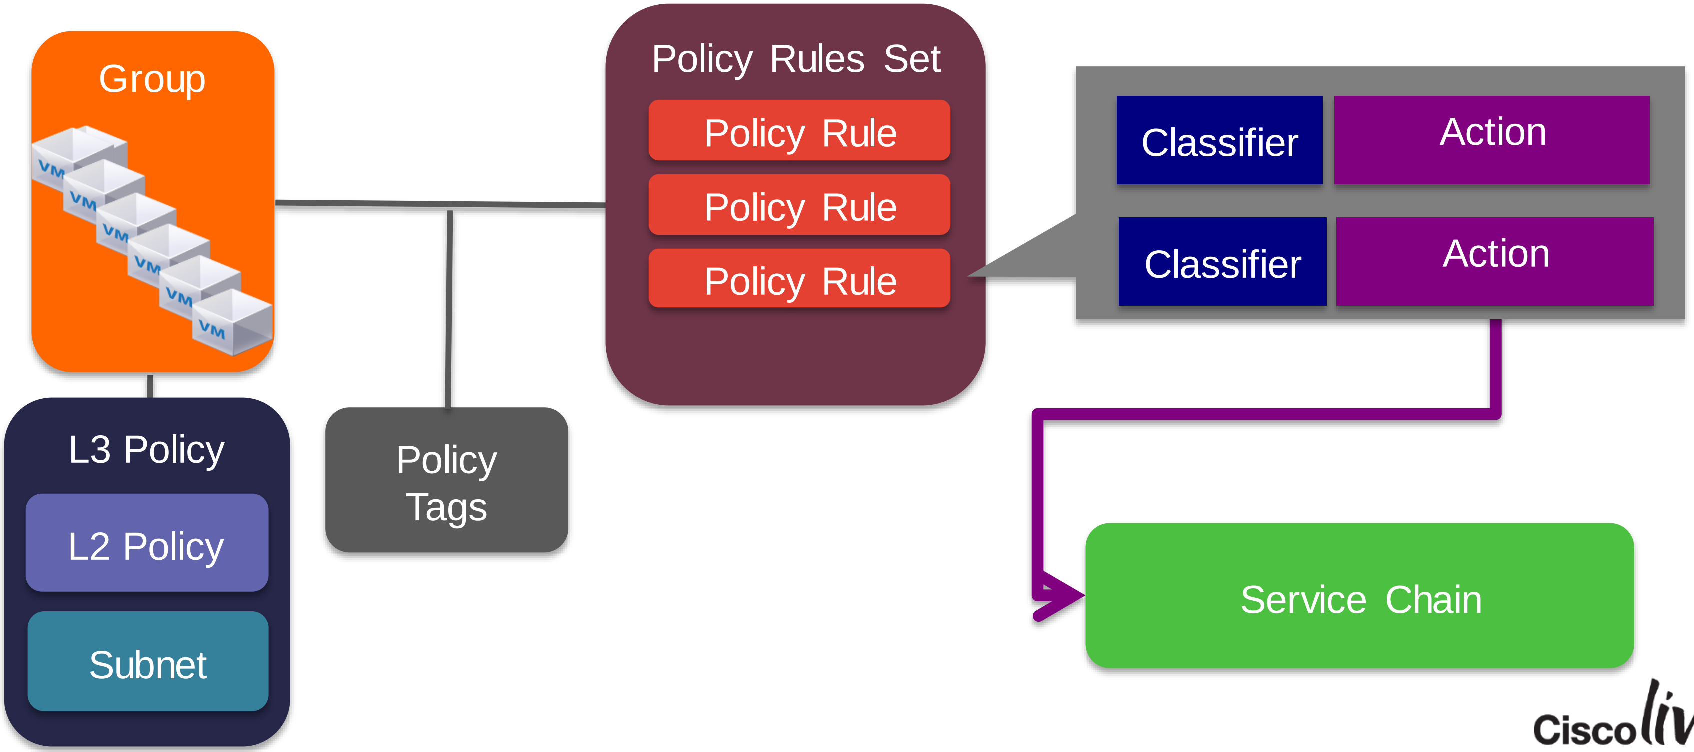

OpenStack GBP Architecture

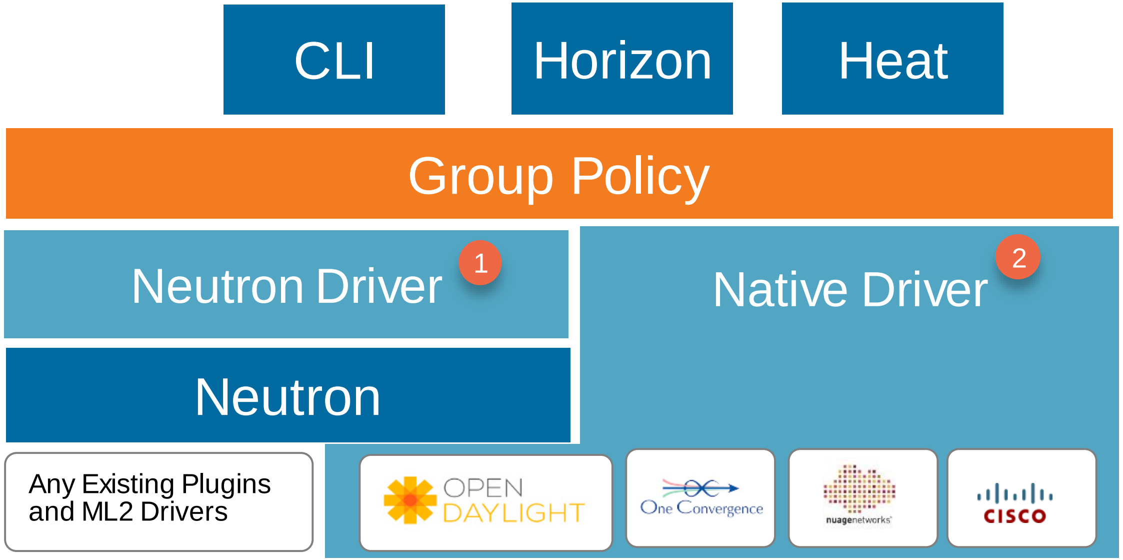

Open model that is compatible with ANY physical or virtual networking backends

1. Neutron Driver maps GBP to existing Neutron API and offers compatibility with any existing Neutron Plugin

2. Native Drivers exist for OpenDaylight as well as multiple vendors (Cisco, Nuage Networks, and One Convergence)

Group-Based Policy Model

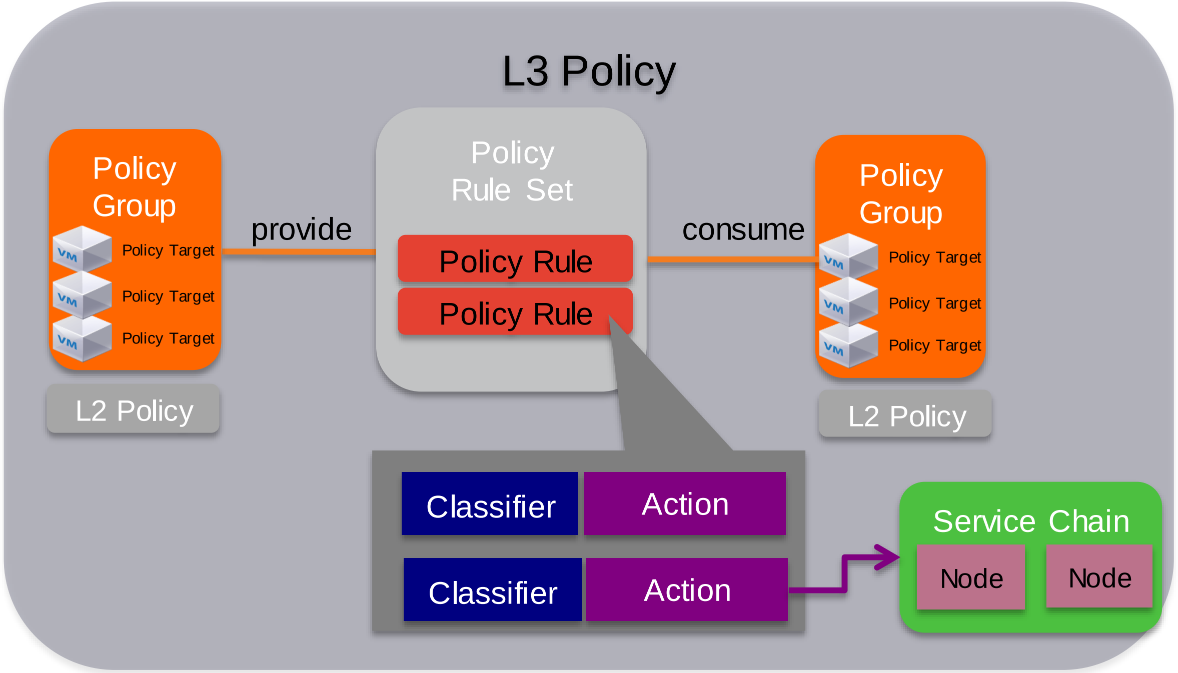

Policy Group: Set of endpoints with the same properties. Often a tier of an application.

Policy RuleSet: Set of Classifier / Actions describing how Policy Groups communicate.

Policy Classifier: Traffic filter including protocol, port and direction.

Policy Action: Behaviour to take as a result of a match. Supported actions include “allow” and “redirect”

Service Chains: Set of ordered network services between Groups.

L2 Policy: Specifies the boundaries of a switching domain. Broadcast is an optional parameter

L3 Policy: An isolated address space containing L2 Policies / Subnets

Neutron Cisco Application Policy Infrastructure Controller (APIC) Driver and Plugin

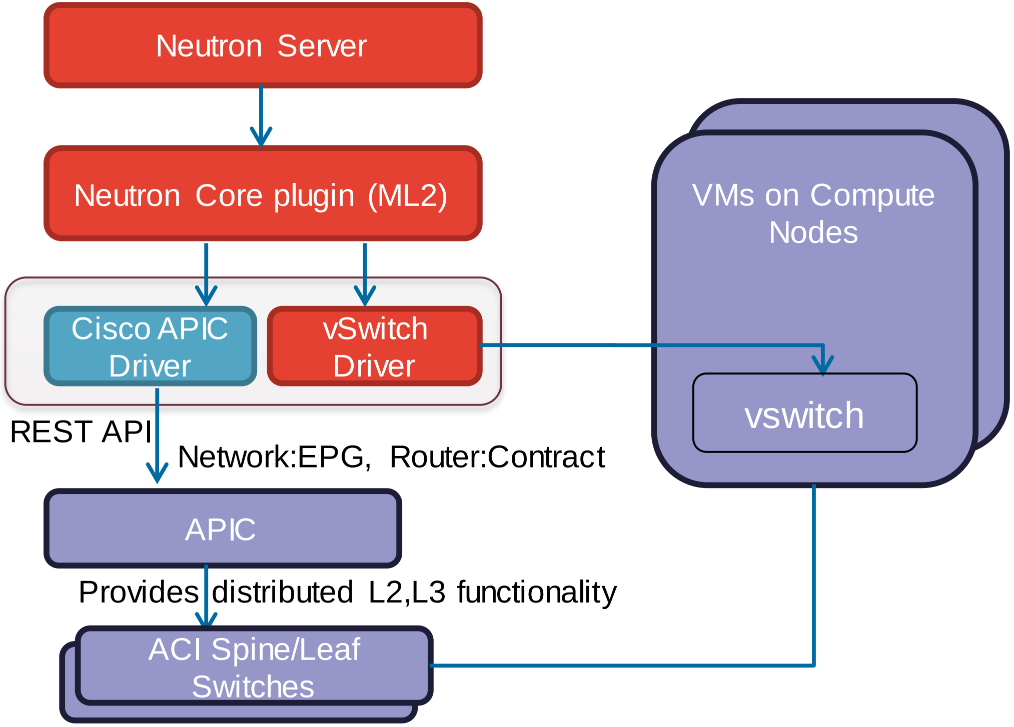

Developing Integration with APIC Using OpenStack Neutron Group Policy API

OpenStack Deployment

Controller Node

Supporting Services
- Database MySQL
- MessageBroker RabbitMQ

Basic Services
- Identity Keystone
- Image Glance
- Compute Nova Mgmt
- Dashboard Horizon

Optional Services
- Block store Cinder Mgmt
- Object store Swift Proxy
- Orchestration Heat
- Telemetry Ceilometer

Network Basic Services
- Networking ML2 Plugin OVS Driver APIC Driver DHCP

Network Interfaces
- Management
- Fabric

Generally combine network and controller node for ACI

Compute Node

Basic Services
- Compute Nova Hypervisor
- Networking ML2 Plugin APIC Host Agent OVS Agent

Network Interfaces
- Management
- Fabric

APIC

APIC manages all tenant and physical networking policy. At least 3 APICs are deployed for redundancy

Network Interfaces
- Management
- Fabric

Nexus Standalone Integration

- 1 Controller, 2 Compute nodes
- Separate Management (eth0) and Data networks (eth1)
- ToR switch connections and config
- Separate Nova availability zones
- Neutron Cisco plugin -> ML2 Driver
- Not running Neutron L3 agent on controller server
- VLAN range managed by vswitch plugin (ovs, n1kv)
- Supported on Grizzly or later releases
- Requires ncclient on control-server for NETCONF

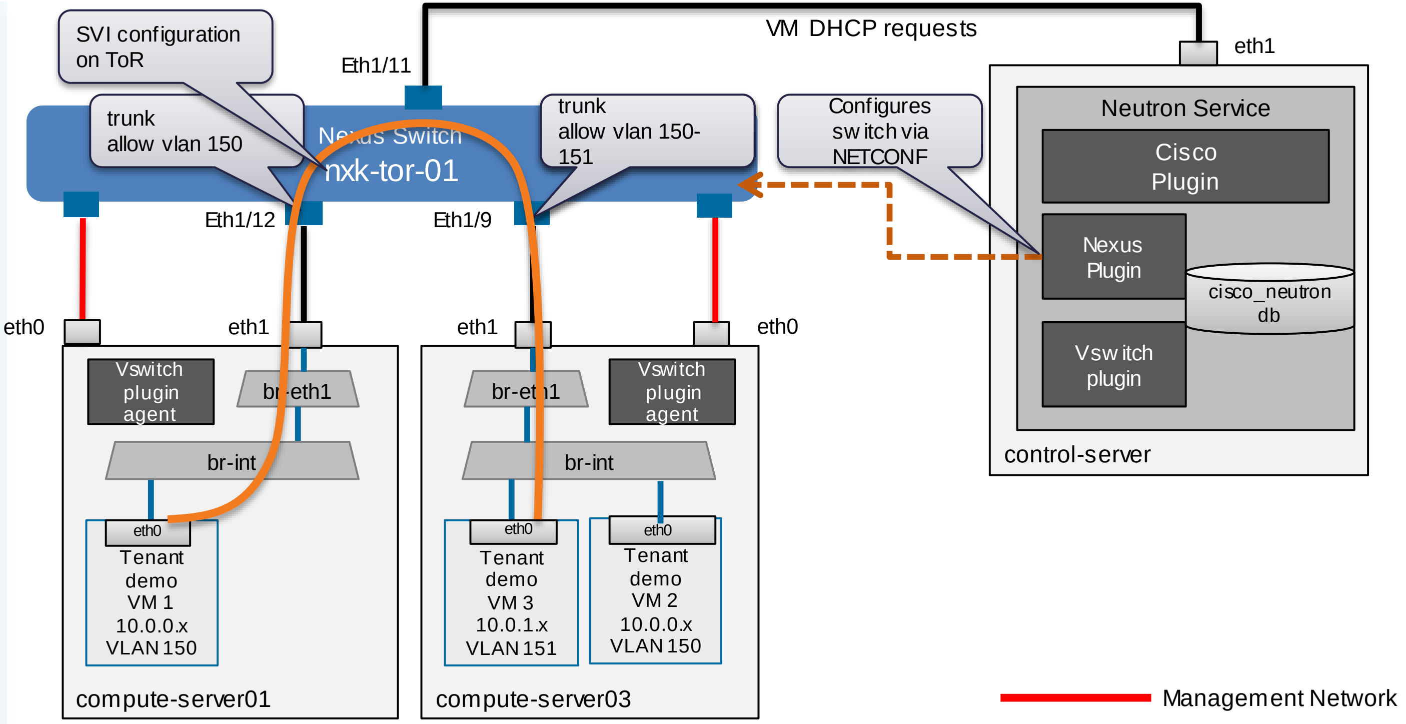

*There is an additional linux bridge on the host which has not be shown for simplicity

Why Cisco ACI and OpenStack

1 GROUP-BASED POLICY SUPPPORT

- Automation
- Intent-driven

2 PHYSICAL + VIRTUAL

- Zero-touch Performance
- Physical server
- Multi-hypervisor

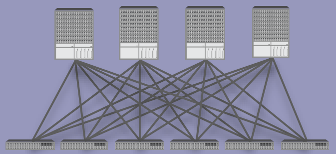

3 FABRIC TUNNELS

- Automatic VXLAN
- Distributed L2
- Distributed L3

4 SERVICE CHAINING

- Service chaining and redirection

5 TELEMETRY AND OPERATIONS

- Health Metrics
- Visibility
- Troubleshooting

Cisco live!

Advanced Services Integrations

- ACI supports a rich device package and service chaining feature for managing network services.
    - Automatic service deployment
    - Automatic service chaining

- However, it is also possible to use the LBaaS / FWaaS APIs in Neutron as well
    - Neutron handles device configuration via API
    - No direct support for service chaining.
        - Physical devices as external next-hop routers
        - Service VM per tenant set as gateway for Neutron network

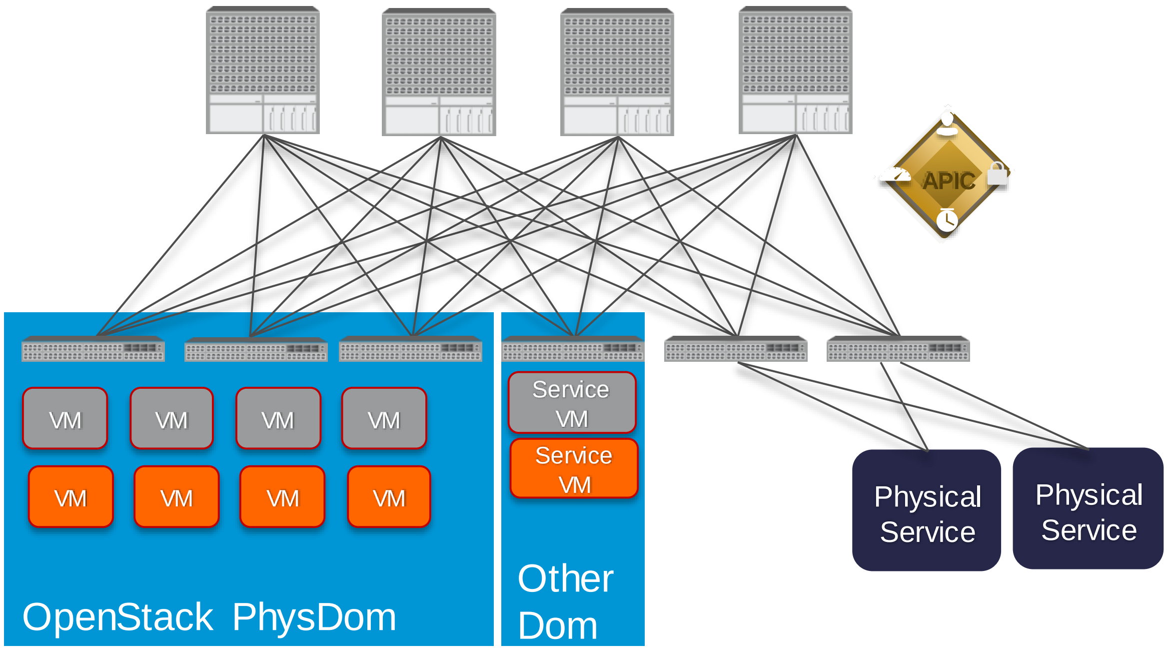

Two Options for ACI

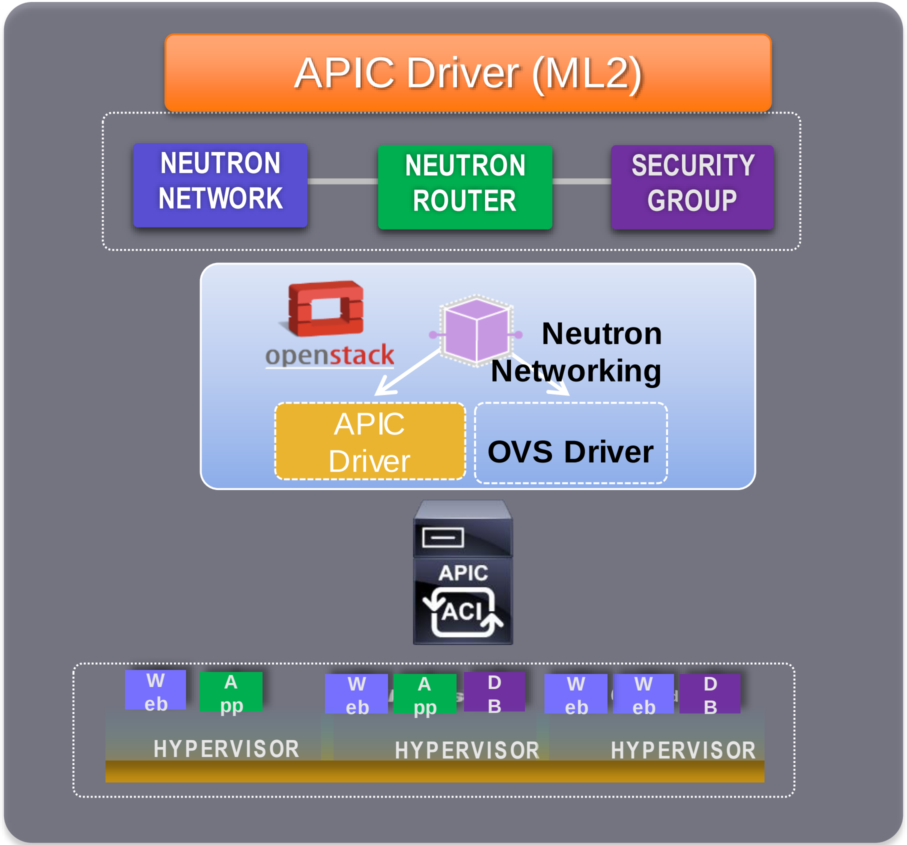

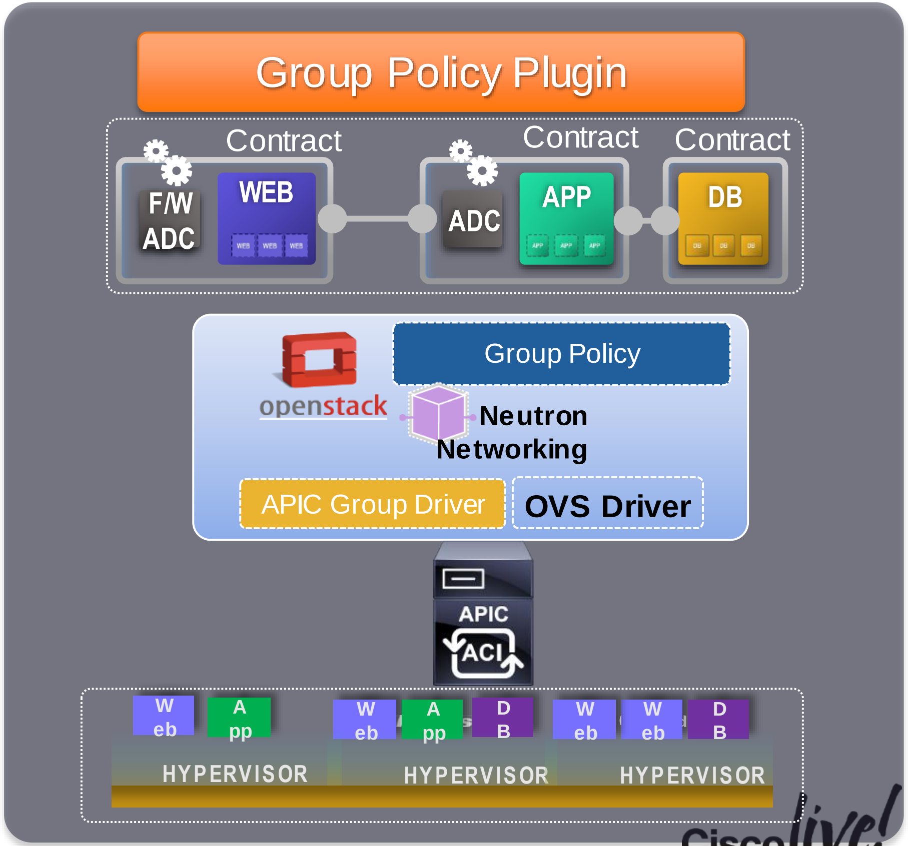

APIC Driver for OpenStack

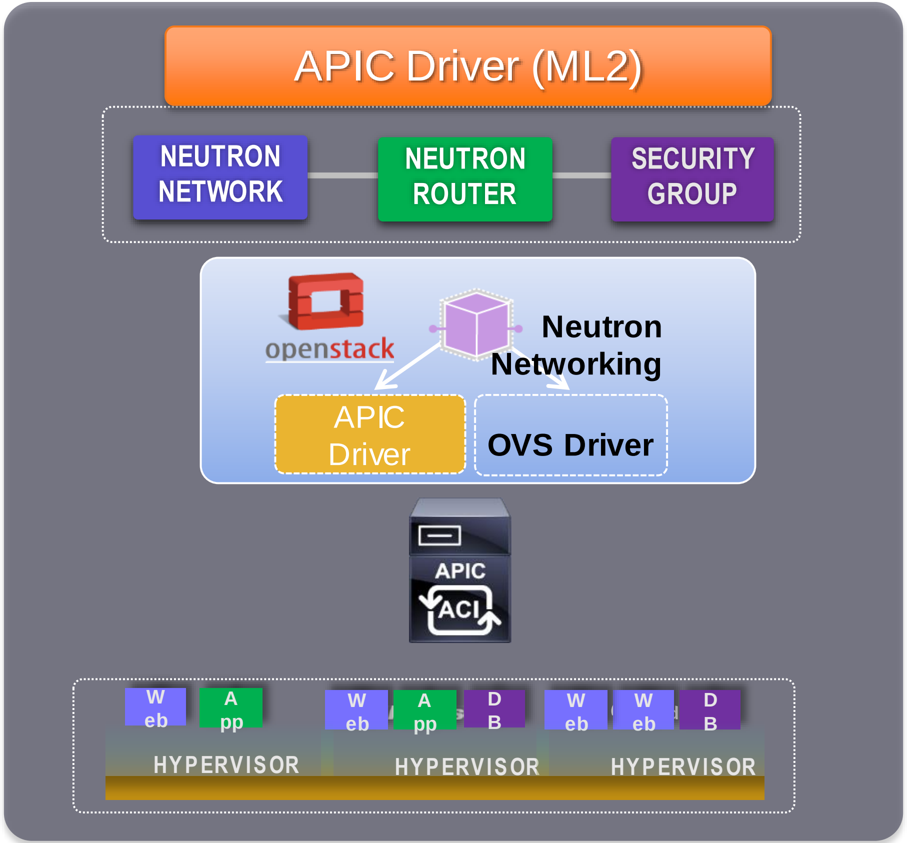

- ML2 (modular level 2) driver supporting existing Neutron APIs: network, router, security group, LBaaS, etc.
- Automation of neutron ports for virtual machines
- Relies on OVS in hypervisor
- Shipping today from Cisco
- Available on Openstack IceHouse, Juno, etc.

APIC Driver Details

ACI Fabric Offers:
- VXLAN tunnels
- Distributed L2
- Distributed default gateway

Hypervisor:
- Enforces security groups

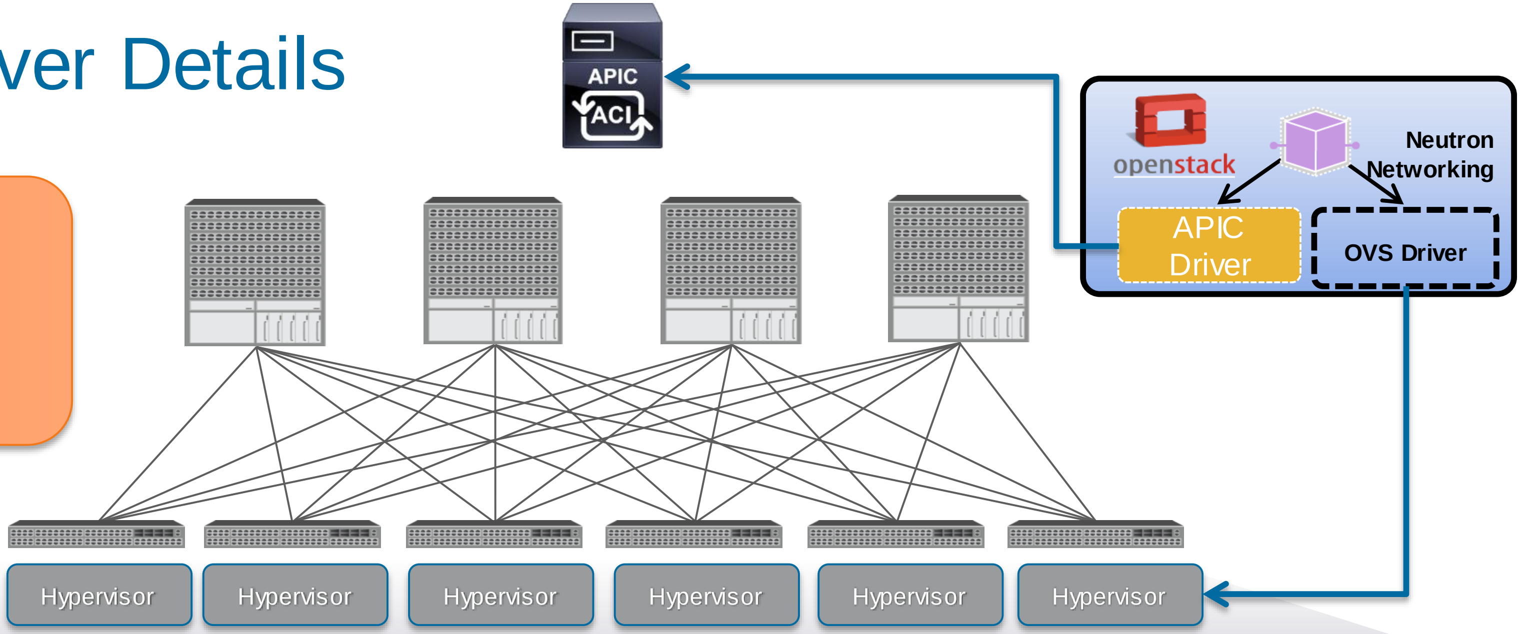

Neutron Workflow

1. User creates a network / router / etc. through Neutron CLI / Horizon / Heat
2. OVS Driver selects VLAN from VLAN pool. VLAN is configured in Open vSwitch
3. APIC Driver maps neutron object to APIC policy model
4. IP Tables in Linux Hypervisor provides host-based security group enforcement
5. Open vSwitch tags each Neutron network with VLAN
6. ACI ToR translates VLAN into VXLAN, providing distributed L2 and distributed default gateway support.

Group-Based Policy

- OpenStack extensions on top of Neutron exposing a policy API
- Supports policy API to APIC
- Backwards compatible with existing neutron plug-ins (works with Nexus 9000 standalone)
- Available for Openstack Juno (Q1 CY 15)
- Open approach
- Enables Openstack customers to deploy, scale and modify policy across teams fast

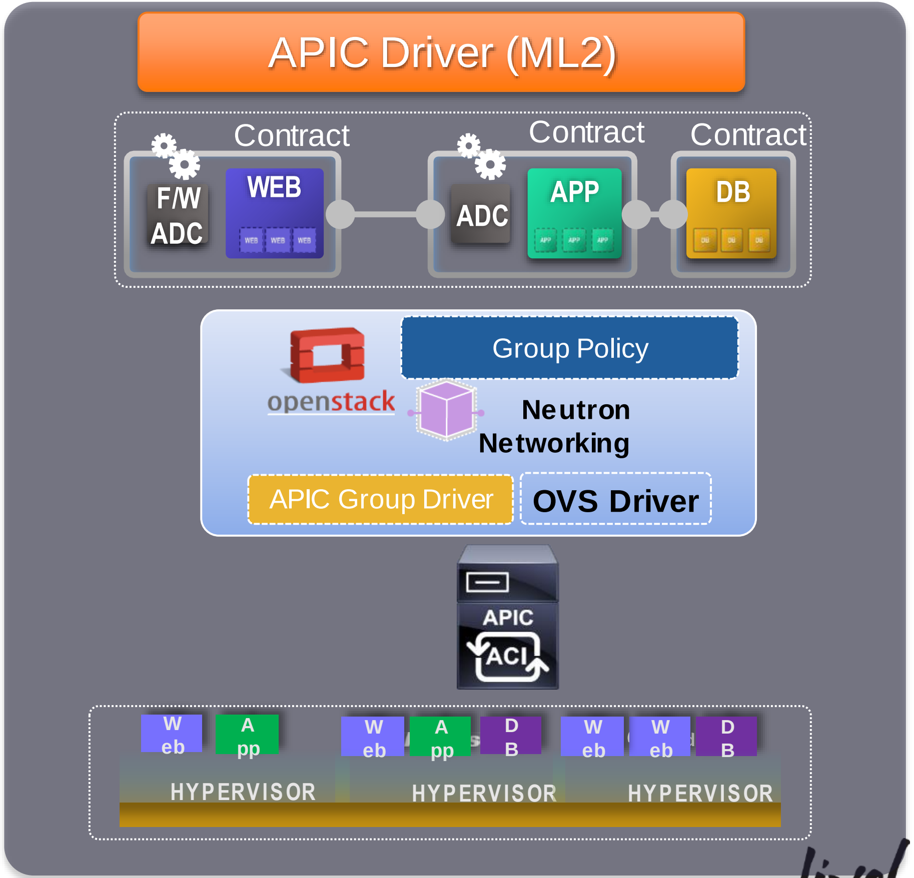

Group Policy Plugin

ACI Fabric Offers:

- VXLAN tunnels
- Distributed L2
- Distributed default gateway
- Security enforcement

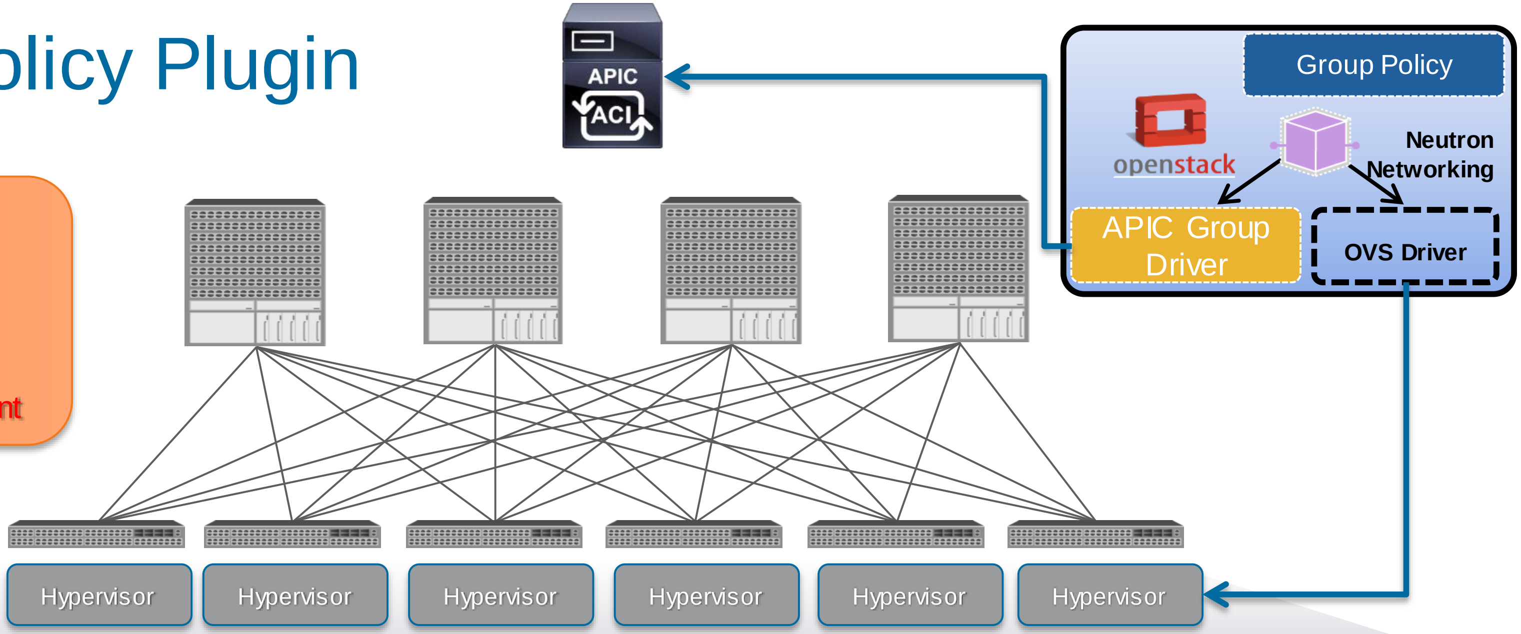

Neutron Workflow

1. User creates Group-Based Policy through CLI / Horizon / Heat.
2. OVS Driver selects VLAN from VLAN pool. VLAN is configured in Open vSwitch
3. APIC Driver maps GBP to APIC policy
4. Non-OpFlex: All inter-EPG traffic sent to ToR for enforcement (note, with OpFlex switching and enforcement may occur in OVS).
5. Open vSwitch tags each group with VLAN
6. ACI ToR translates VLAN into VXLAN, providing distributed L2, security policy, and distributed default gateway support.

- Automated configuration of Layer 2 and 3 and Layer 4-7 services

Scaling Open Stack Deployments

What’s Missing in Today’s Public Clouds?

- The network, a critical component of application performance, is abstracted away from the developer.
  - High-performance applications need to be able to understand how the network is performing to optimise their behaviour, and optimally could report their performance indicators to influence network behaviour
- SPs are migrating network services away from managed CPE toward Network Function Virtualisation.
  - This places performance demands on the network that cannot easily be achieved in public clouds
- Public clouds have spotty support for advanced features like IPv6, dynamic routing, and multicast

Tenant Resource View (L2 – L3)

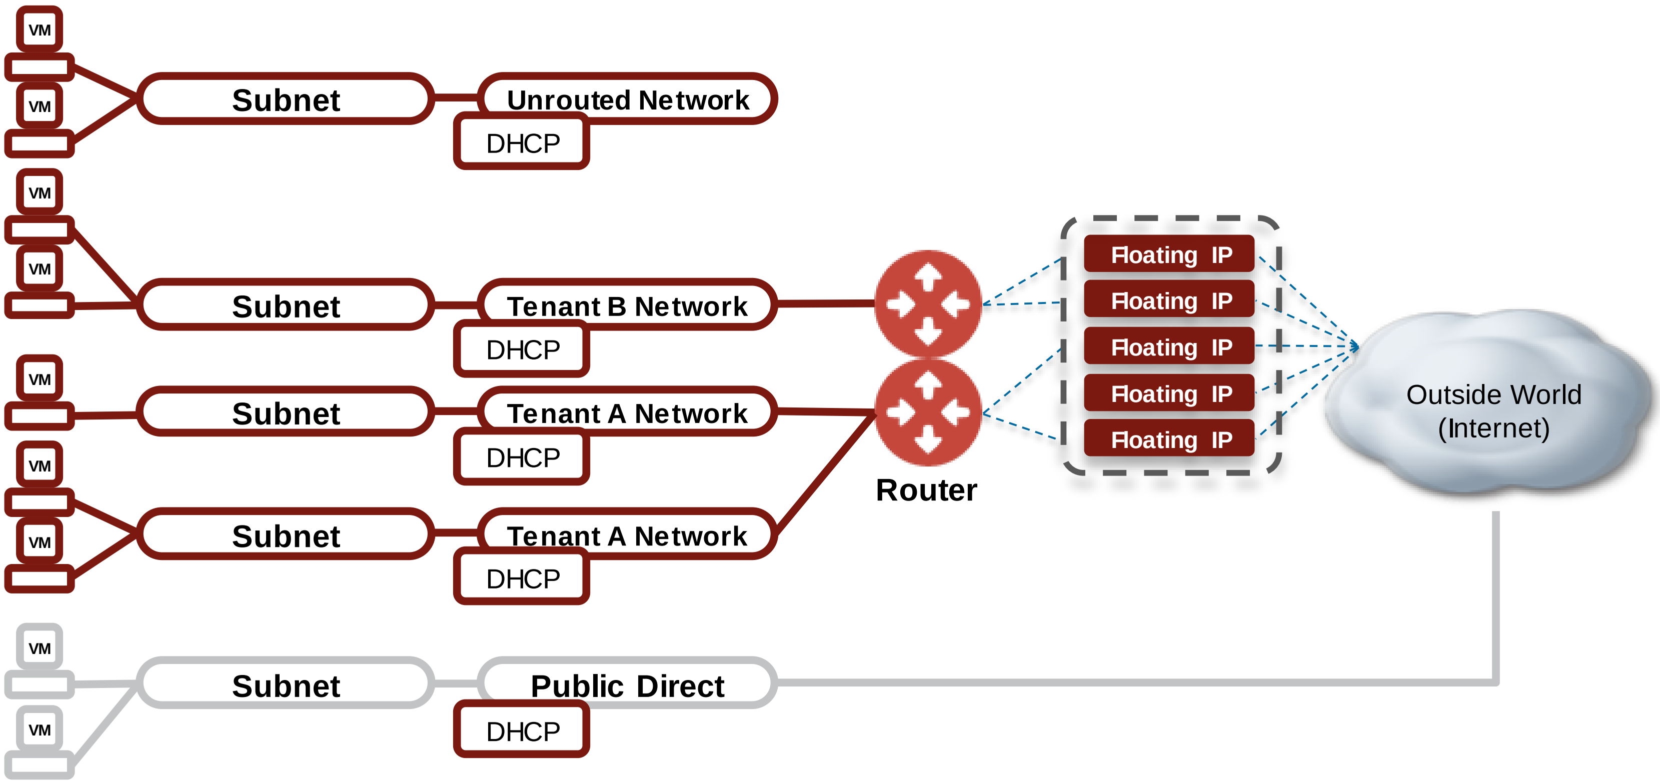

DC LAN with Cloud Mgmt. (Classic)

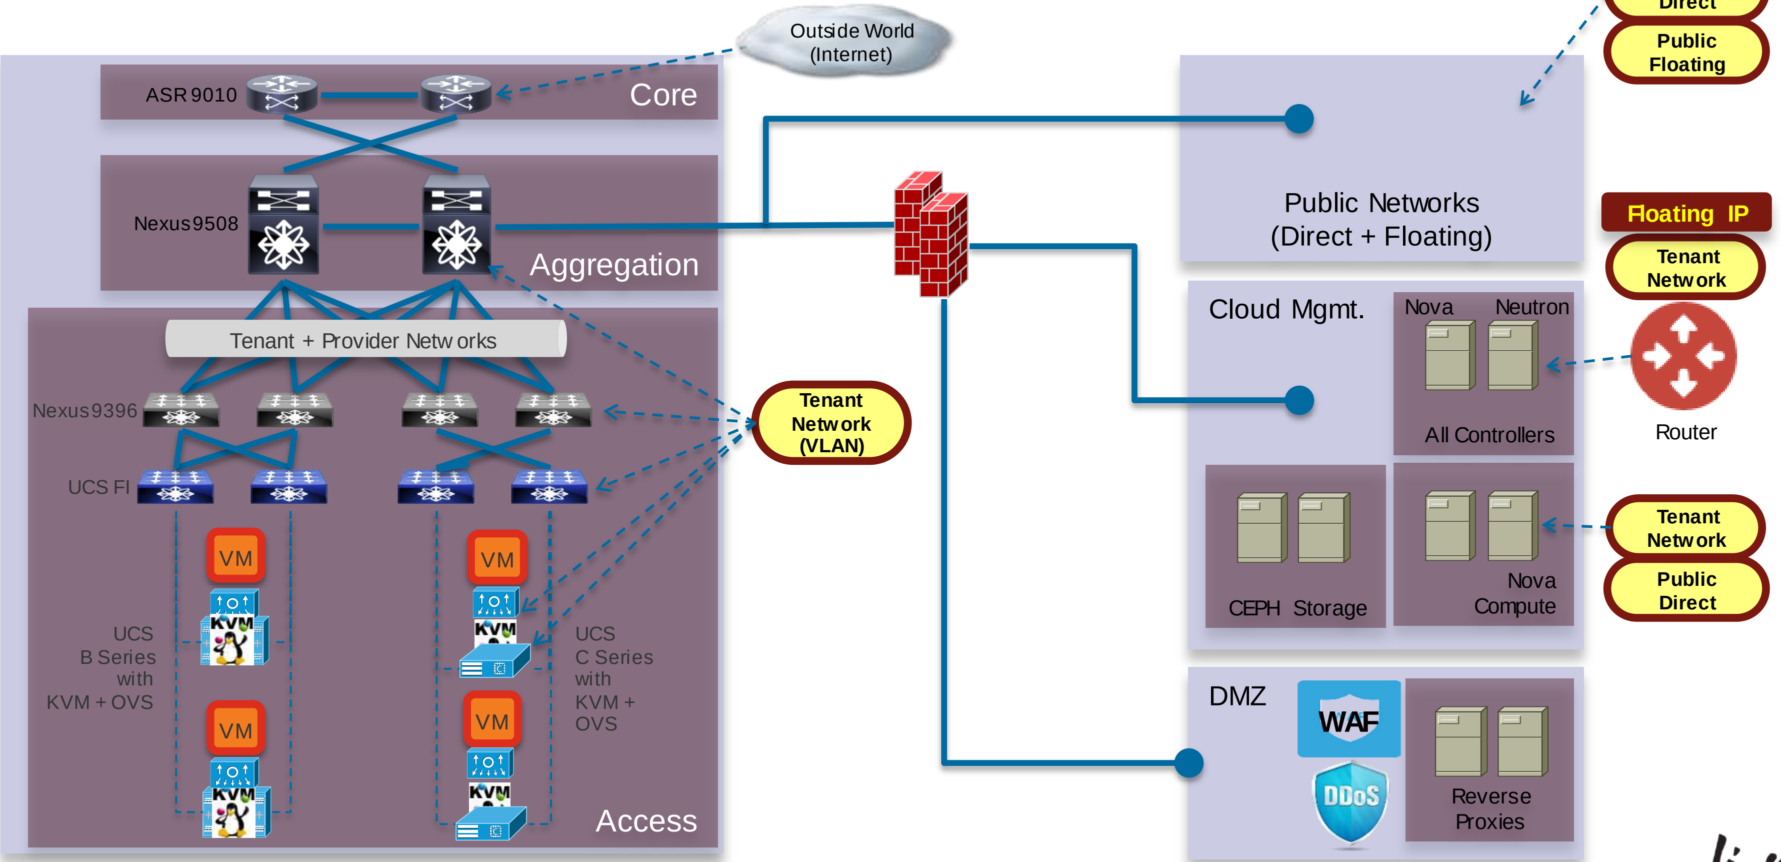

DC Network - Product Requirements and Challenges

Openstack

OpenStack out-of-box networking capabilities on top of traditional DC network

Capabilities	Current State
DC Network	9k VLAN Based DC network
Network Scale	1200 VLANs system wide
Tenant Scale	500 tenants
High availability	Poor – OpenStack provided
Performance	Medium – OpenStack in data path
Stability	Limited – OpenStack provided
Control (policy)	Basic (no policy)

Challenges

Infrastructure Scale Demands

Stable, highly available, performing & secure network topology for cloud scale

Capabilities	Cloud Scale
Network Scale	100s-1000s s networks / rack
DC scale	10s of racks
Tenant Scale	1000s of tenants
High availability	Highly available network providing 99.9 + % SLA
Performance	High performance Cisco DC networking down to the VM (near 10G throughput from hypervisor)
Stability	Stable control and data path components using Cisco technologies
Control (policy)	Centralised control w/Policy

Interim Step – Smart Plugin (VLAN based)

Address the Tenant & Tenant NW scale with Smart OpenStack Plugin –

- VLAN is dynamically provisioned on TOR if/when needed
- 1200 Tenants (assuming 3 NW on average/tenant)
- 3600 tenant networks
- 12K Floating IP (10/tenant)
- Addresses Data Plane HA Problem

DC LAN with Cloud Mgmt. (Classic with Tenants GW + Floating IP HA)

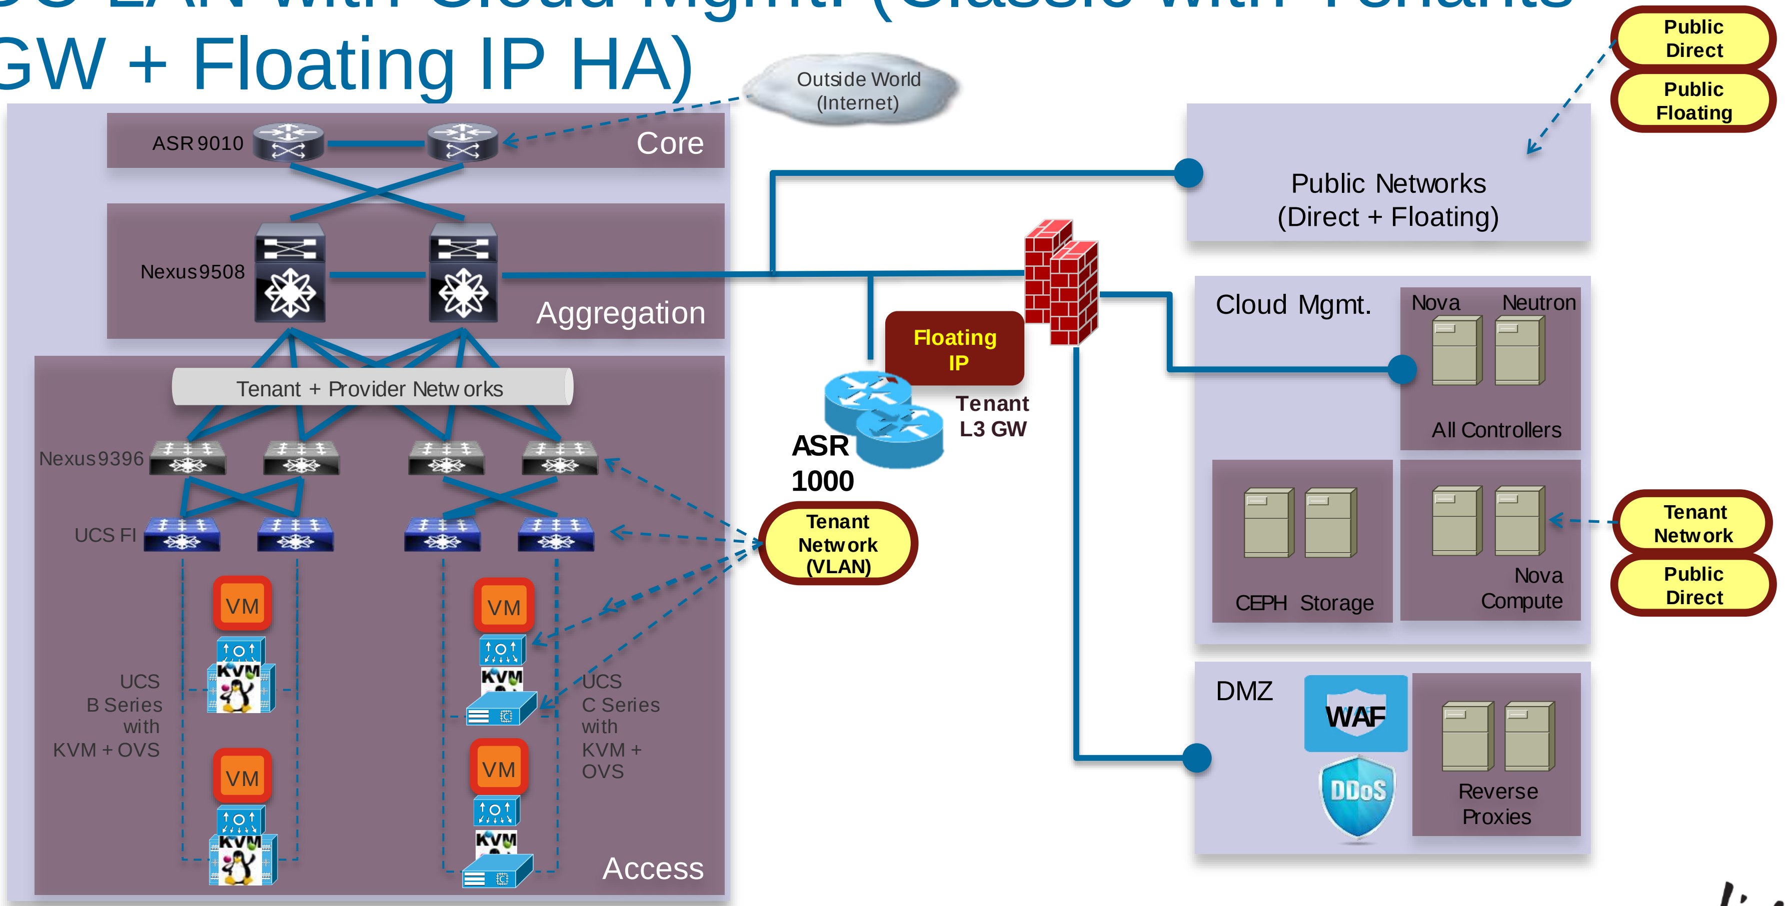

Sample Scale Out Architecture for OpenStack

Standalone Network with SDN Controller and OpenStack Plugin

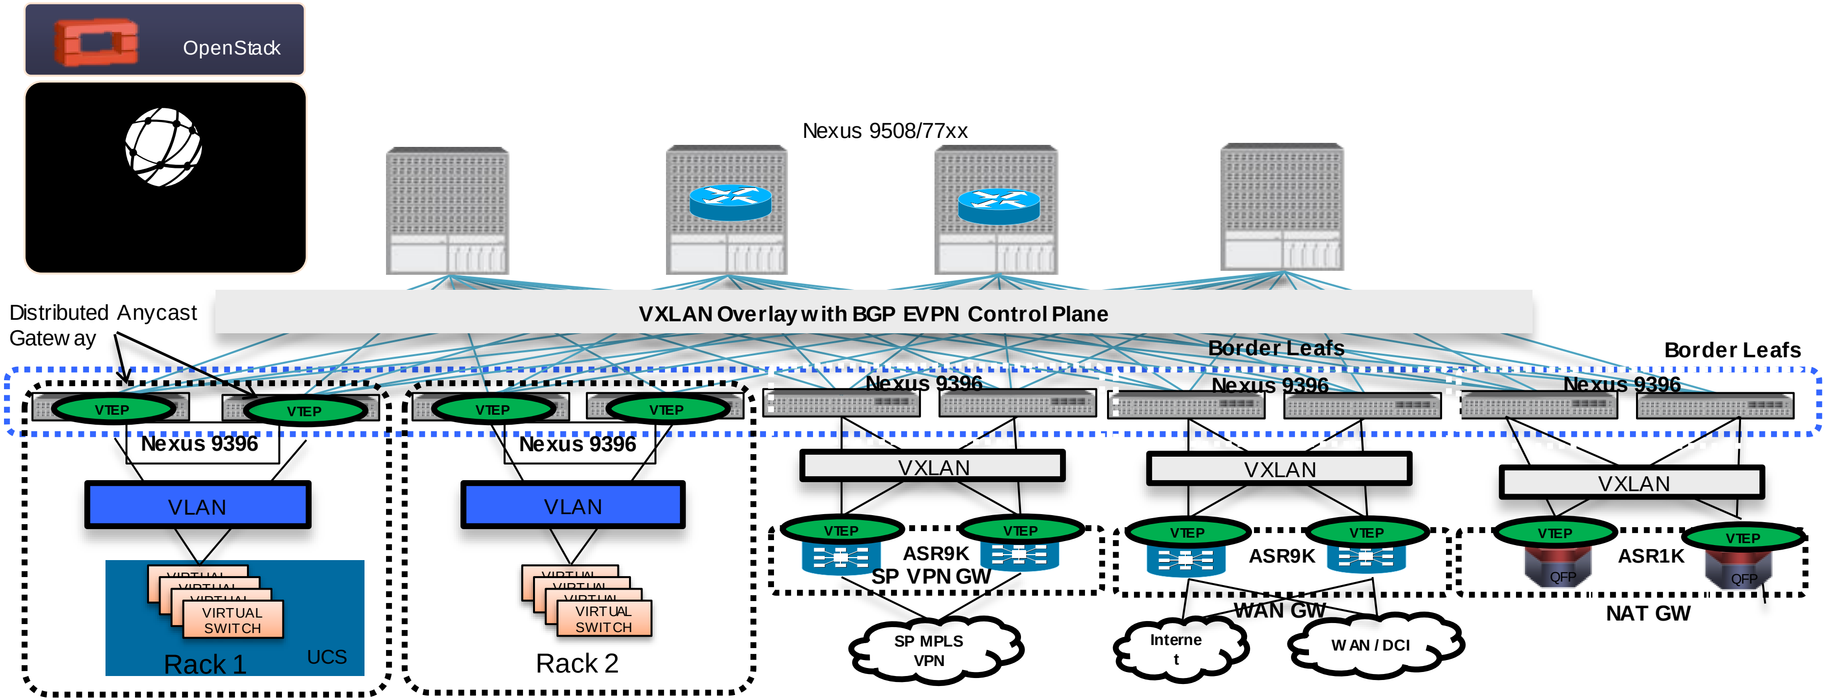

Scale Out Architecture w/ACI

- Scalable NAT/PAT on OVS
- Distributed GW on the ToR (L3 FW) – East/West Traffic
- Flexible Group based Policy Plugin
- OpFlex For Tighter Integration & Local Switching
- ARP suppression

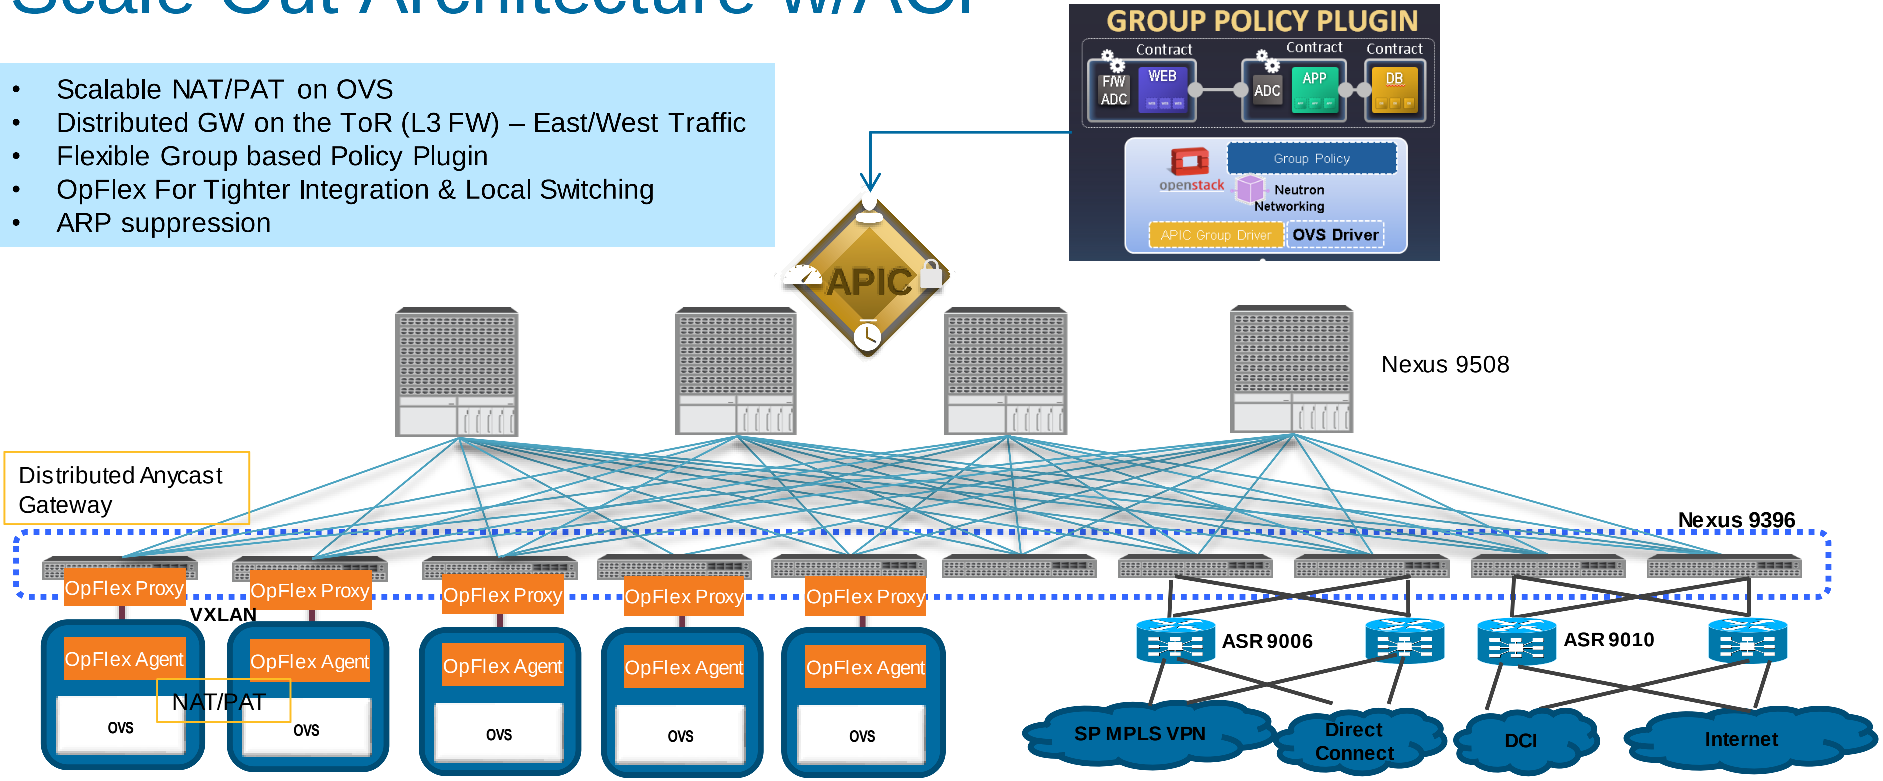

Conclusion

Agenda

- Trends
- Introduction to OpenStack
- Infrastructure Consideration
- GBP and OpenStack
- Scaling OpenStack Deployments
- Conclusion

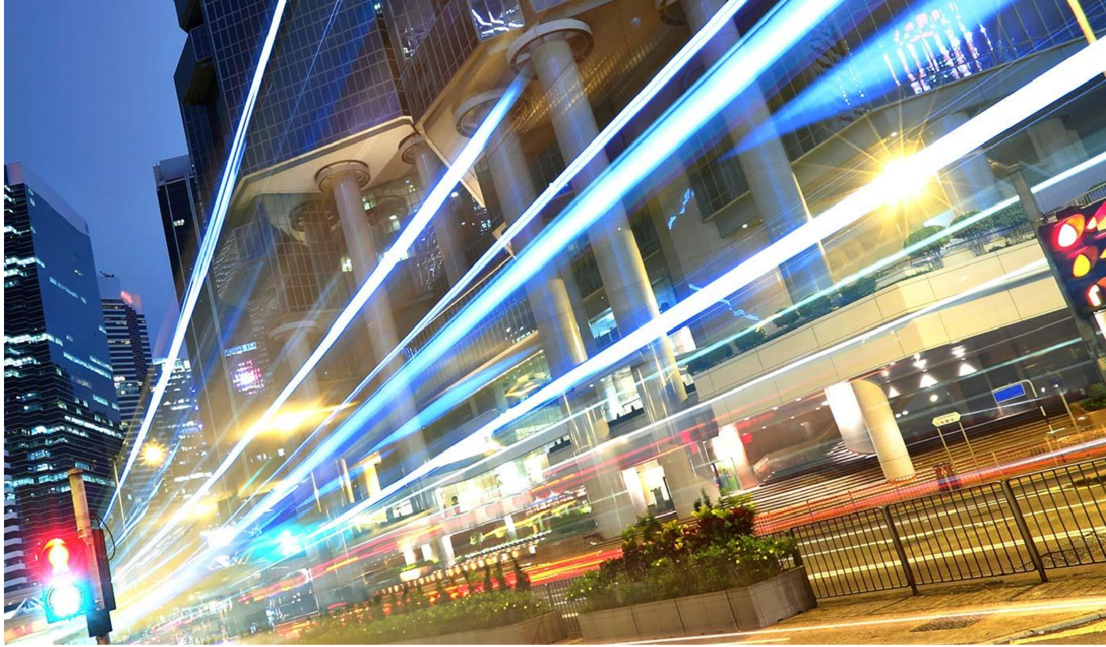

Cisco's Investment and Expertise in OpenStack

openstack

Community Participation

- OpenStack Foundation board member
- Code contributions across core services
- Prolific reviewer of completed blueprints
- One of the leading contributors of code to the Neutron project
- Expanding beyond Neutron on bare metal and group policy code

Engineering Investment

- Neutron plug-ins for Cisco Nexus®
  - Cisco® ACI and APIC plug-ins
  - VLAN programming
  - Cisco Nexus 1000V portfolio for KVM
- Cisco UCS plug-ins for Neutron and Ironic (incubation)
- Cisco UCS OpenStack Cisco Validated Design

Customers and Partners

- Cisco OpenStack Advanced Services
- Driving innovation through real-world use cases
- Comcast, Photobucket, Cisco WebEx®, large service providers

Cisco live!

Continue Your Education

- Demos in the Cisco Campus
- Walk-in Self-Paced Labs
- Meet the Expert 1:1 meetings

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2015 T-Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site http://showcase.genie-connect.com/clmelbourne2015
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm

Learn online with Cisco Live!
Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

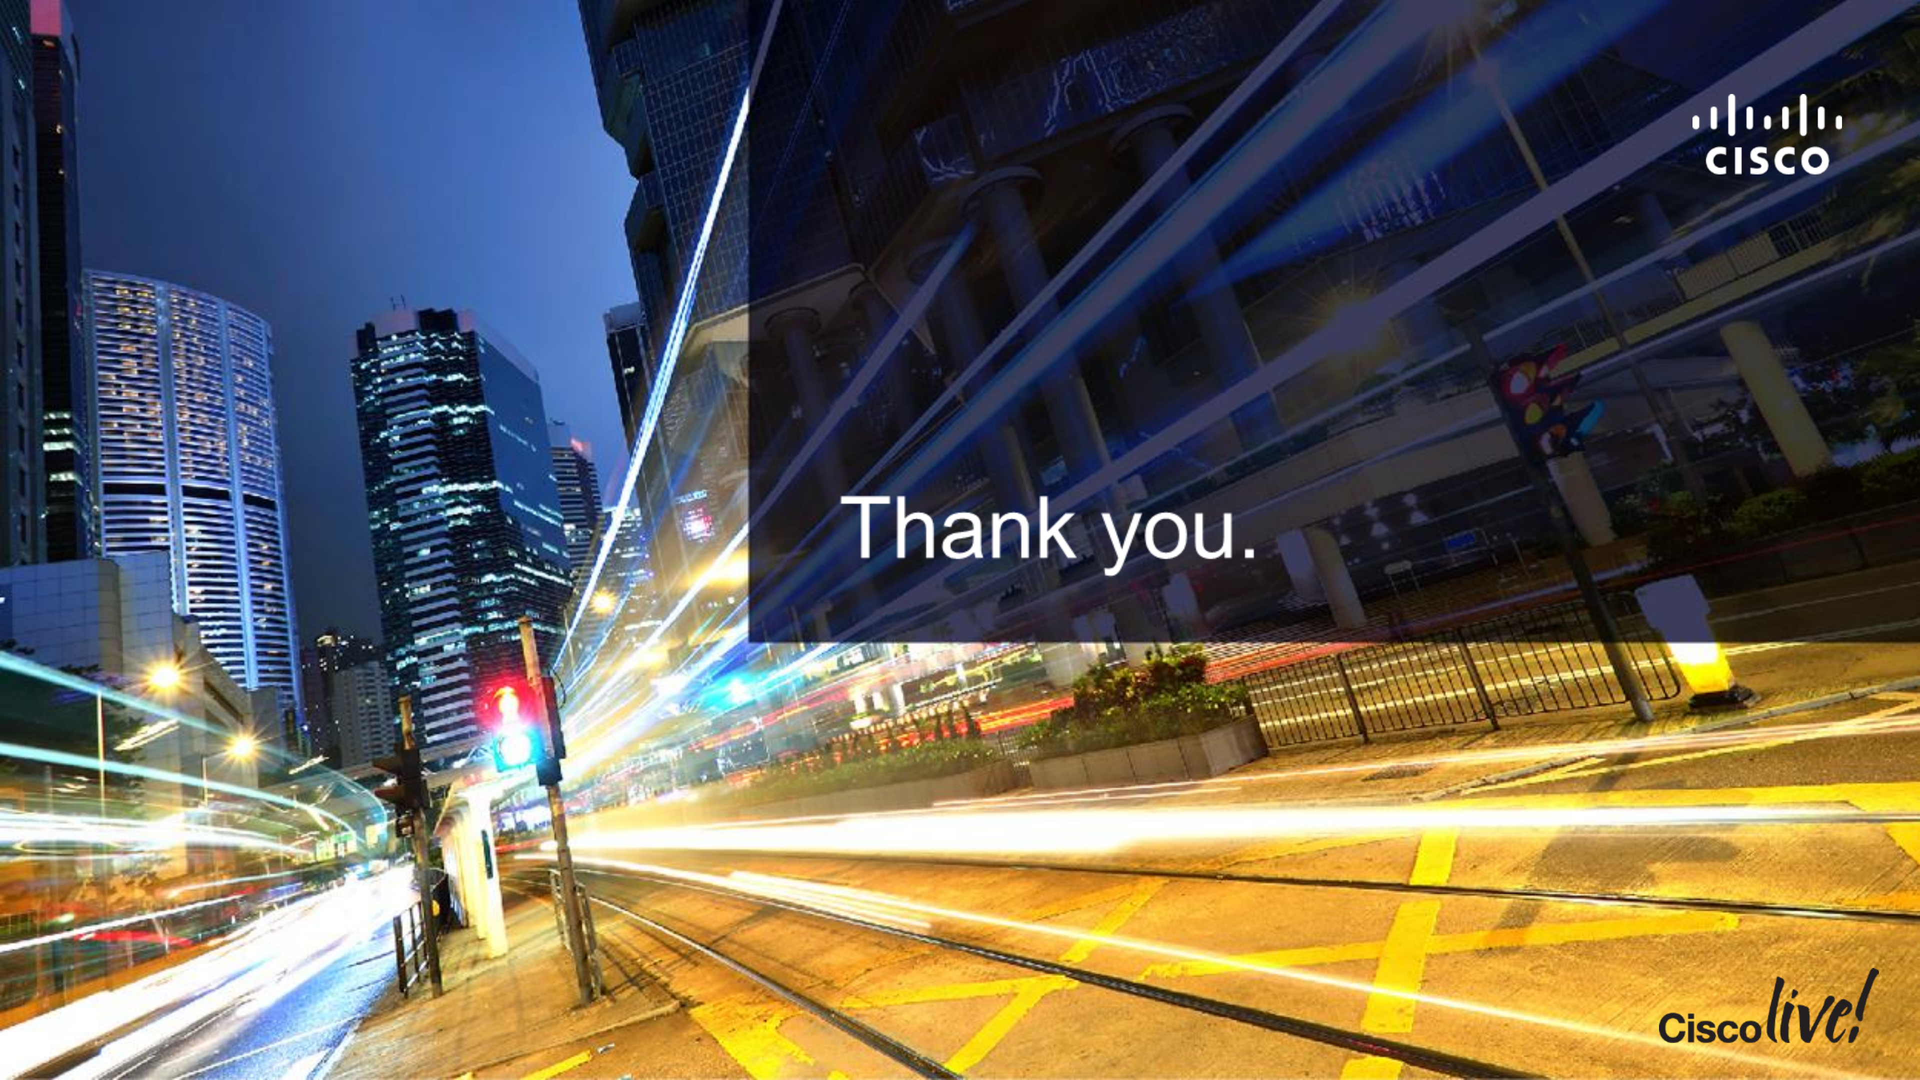
Thank you.

CISCO