



Quantum, network services for Openstack

Salvatore Orlando

sorlando@nicira.com

Openstack Quantum core developer

Twitter- [@taturiello](https://twitter.com/taturiello)

Caveats

- Quantum is in its teenage years: there are lots of things that it WILL do, but doesn't do yet;
 - nevertheless it has a great potential and energy
 - however it is a bit unruly and sometimes needs extra patience
- The main focus of this talk is on users wishing to deploy Quantum in their cloud

Outline

- What is OpenStack?
- Why Quantum?
- What is Quantum?
 - API Overview
 - Plugin Architecture
 - Quantum in Openstack Folsom
- Use Case: Nicira NVP Plugin
 - Network virtualization model
- Questions



What is OpenStack?

What is OpenStack?



“To produce the **ubiquitous** open source cloud computing platform that will meet the needs of public and private cloud providers regardless of size, by being **simple to implement** and **massively scalable.**”

Allow anyone to automate IT provisioning in a public or private IaaS cloud that meets (and exceeds) the capabilities of Amazon EC2.

Why is Openstack Compelling?

Incredible cross-industry mindshare and momentum



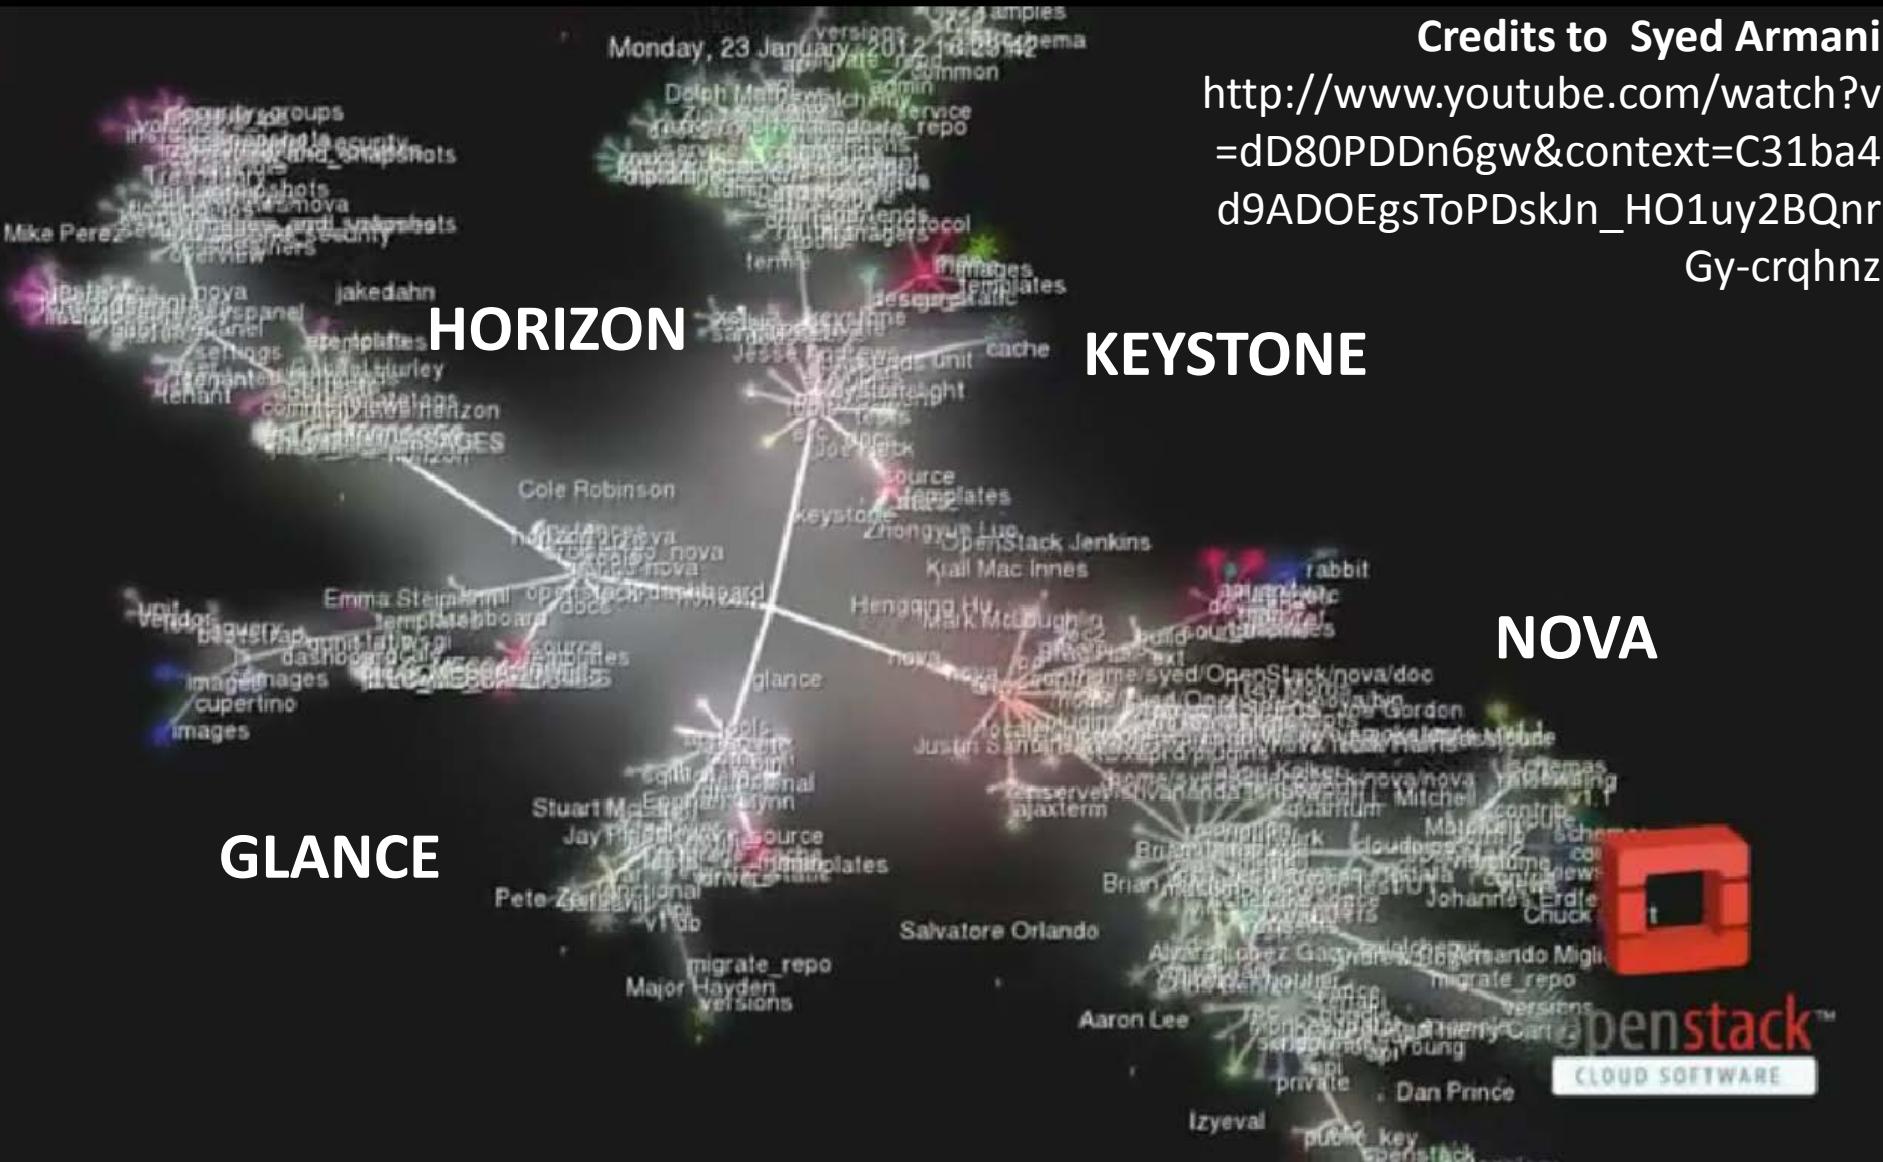
Real resource commitments from many companies:

200+ unique developers from 55+ companies in last 6-month release

Openstack core

Credits to Syed Armani

http://www.youtube.com/watch?v=dD80PDDn6gw&context=C31ba4d9AD0EgsToPDskJn_HO1uy2BQnrGy-crqhnz



Who is Deploying OpenStack?

Hosting Providers



Carrier / Telco



Large Enterprise



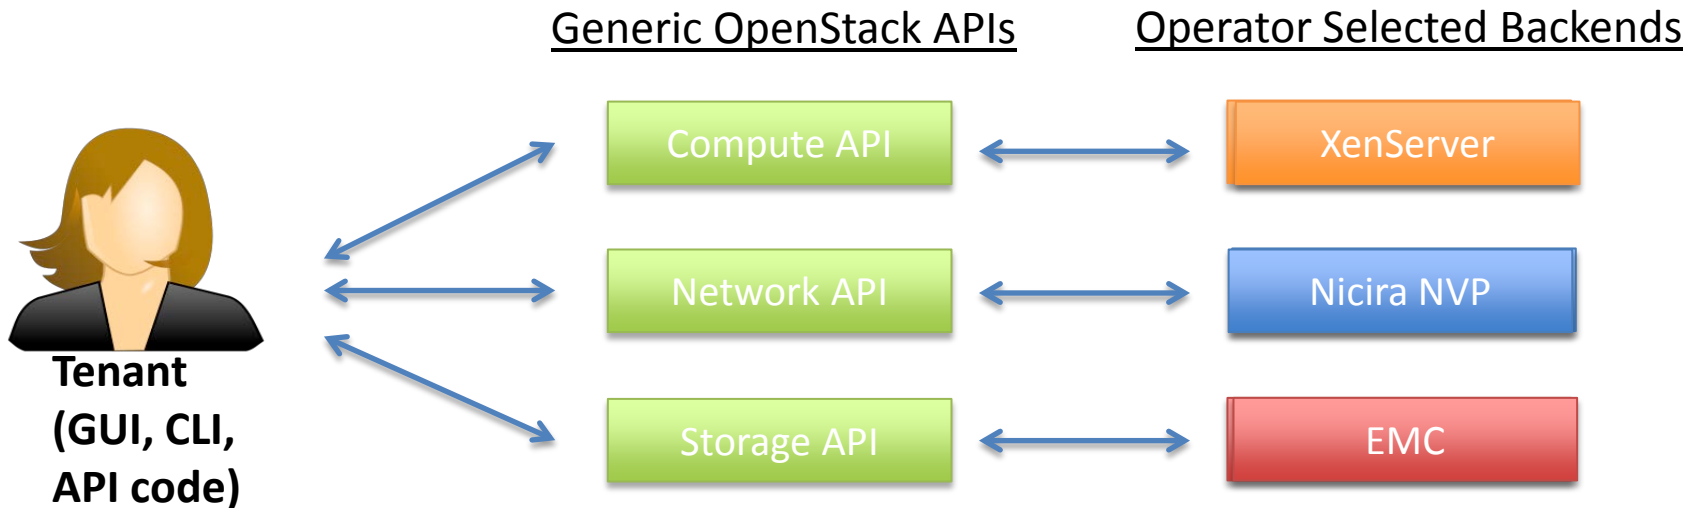
Government



More information: <http://www.openstack.org>

~~10,000 ft~~ (3 km) Architecture View

- Each service is a separate piece of software, includes:
 - A tenant-facing API that exposes logical abstractions for consuming/monitoring the service.
 - Pluggable backend to choose “best of breed” implementations (open source or vendor proprietary).



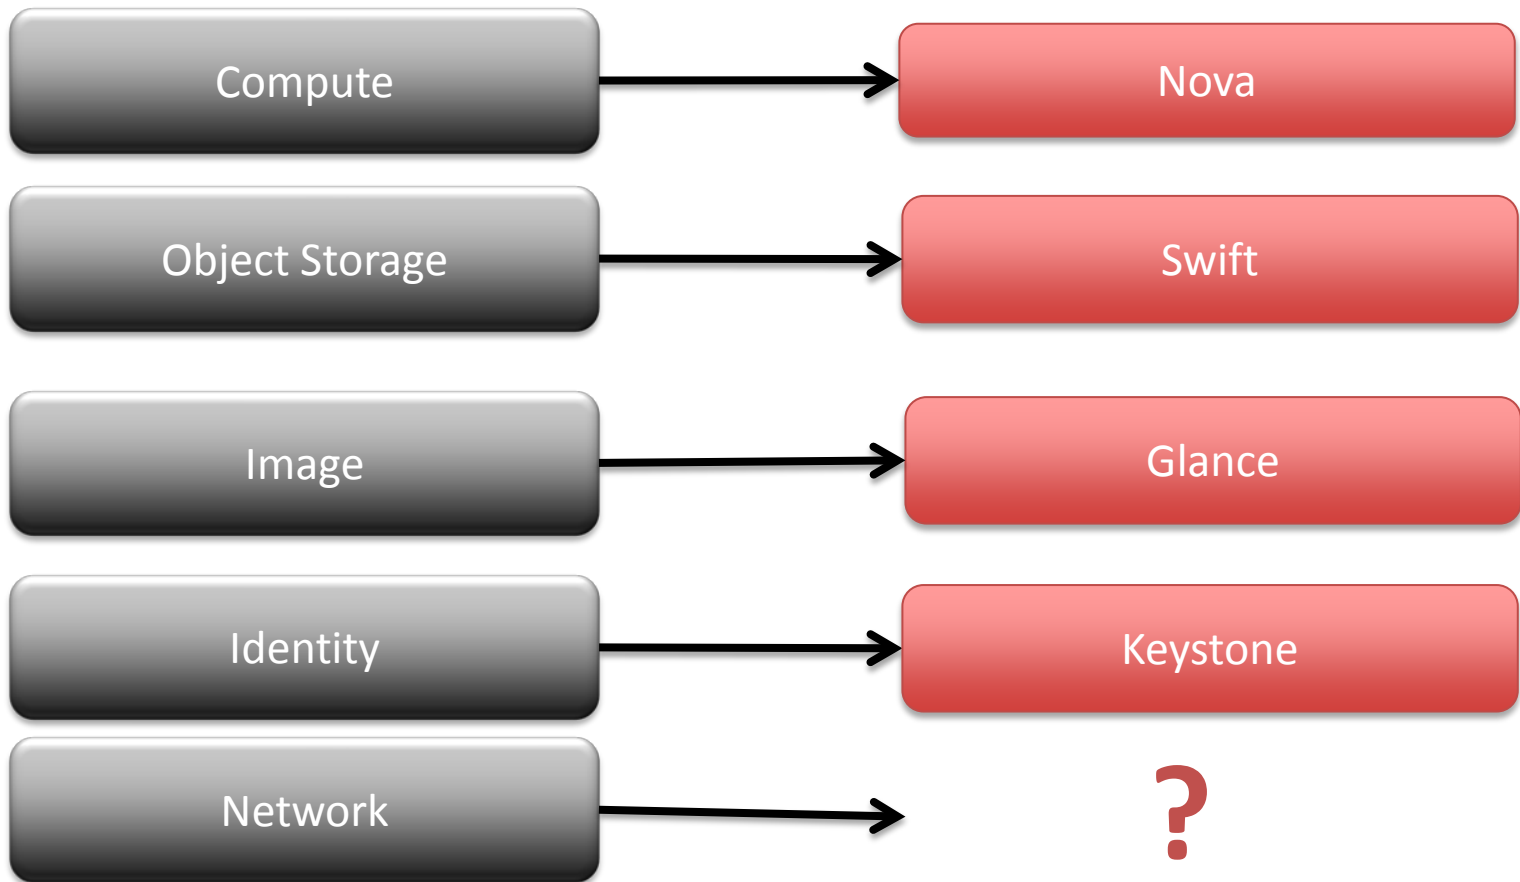


Why Quantum?

In the beginning..

*-as-a-Service Capability

OpenStack Service



Why Quantum?

- Networking was sub-component of OpenStack Compute layer (Nova).
- Two Key Problems:
 - #1: No tenant control of networking.
 - #2: Limited technology “baked in” to design.

Problem #1: No Tenant Control

To move enterprise apps to the cloud, tenants want to “copy and paste” their existing data center network topologies:

- Ability to create “multi-tier” networks (e.g., web tier, app tier, db tier)
- Control over IP addressing.
- Ability to insert and configure your own services (e.g., firewall, IPS)
- VPN/Bridge to remote physical hosting or customer premises (“cloudbursting”).



“You can have any color as long as its black.”

- Henry Ford about the Model-T

Why Quantum? Reason #1

On-demand Enterprise-Class Networking

- Tenants can:
 - create multiple private networks
 - control IP addressing
 - monitor basic network status.
- Quantum API extensions provide:
 - Advanced control + visibility:
Security policies, Quality-of-Service,
Monitoring + Troubleshooting.
 - Advanced Network Services:
routers, Firewalls, VPN, IDS, etc.



**Build rich networks,
customized to tenant
needs.**

Problem #2: Technology Limitations

- Cloud puts new stresses on networks:
 - High-density multi-tenancy, massive scale
 - Strict uptime requirements.
 - Integrate with legacy hosting environments / remote data centers.
 - VM mobility
 - On-demand service insertion
- But Nova was limited to basic VLAN model + Linux IPTables.



Who needs private networks?

Trunking all VLANs is a great idea!

- Stone Age Man

Why Quantum?

#2: Leveraging Advanced Technologies

- New networking technologies are emerging to try and tackle these challenges.
 - Overlay tunneling: VXLAN, NVGRE, STT
 - Software-defined Networking (SDN) / OpenFlow
 - VPN-based solutions (e.g., E-VPN).
 - L2 Fabric solutions: FabricPath, Qfabric, etc.
 - **Your Own Idea for virtual networking!**
- Quantum provides a “plugin” mechanism to enable different technologies (more later).
- Choice is a good thing!




Use advanced technologies to reach new heights.

...and finally

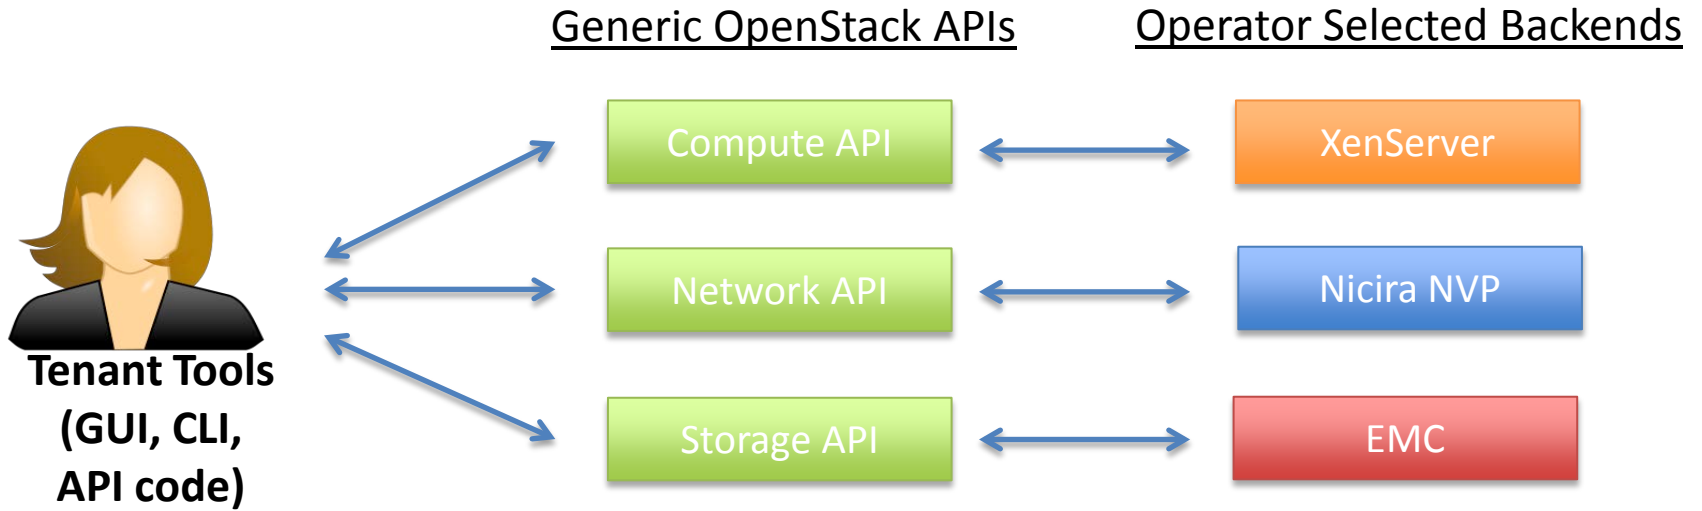
*-as-a-Service Capability

OpenStack Service



What is Quantum?

Quantum Architecture

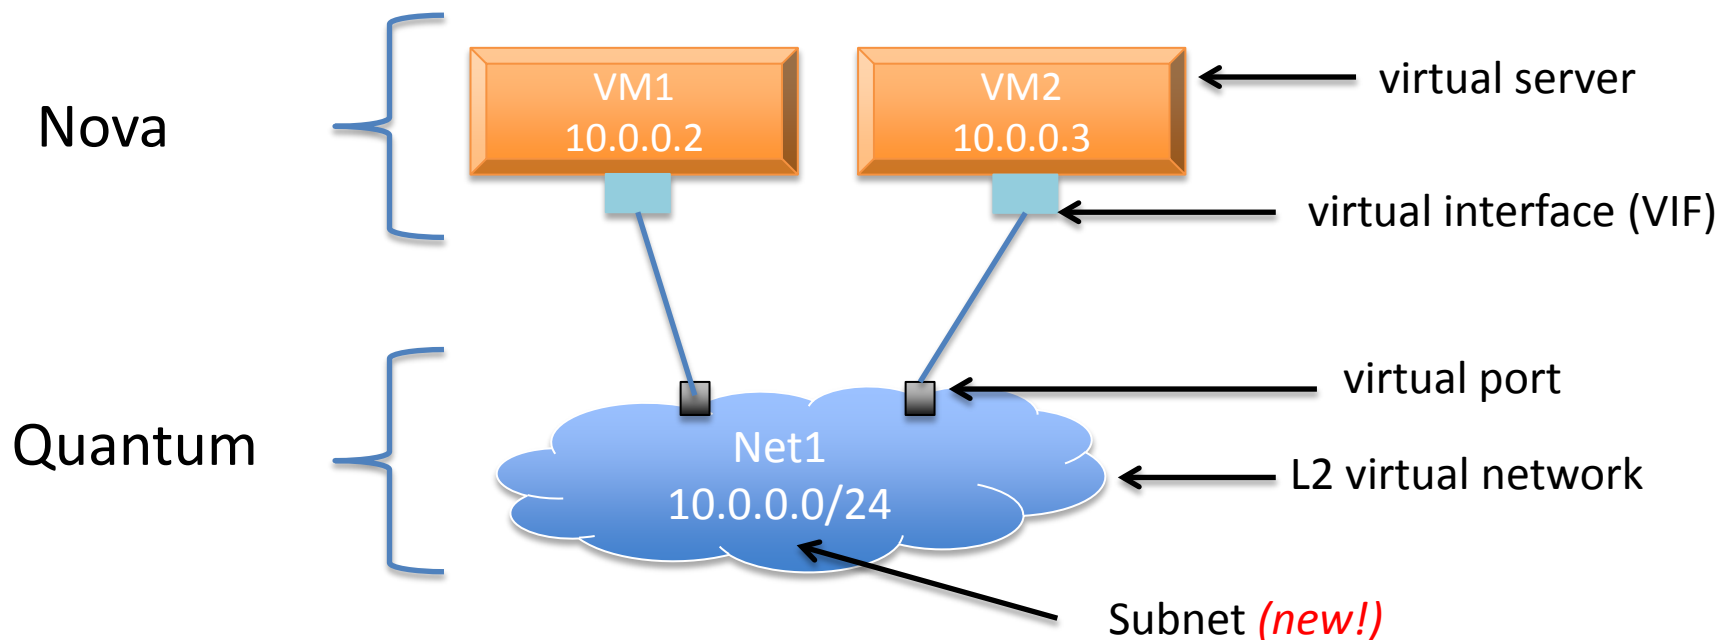


An eco-system of tools that leverage the Quantum API.

A generic tenant API to create and configure “virtual networks”

A “plugin” architecture with different back-end “engines”


Basic API Abstractions




“virtual networks” are fundamentally multi-tenant, just like virtual servers.

Tenant view vs Provider view

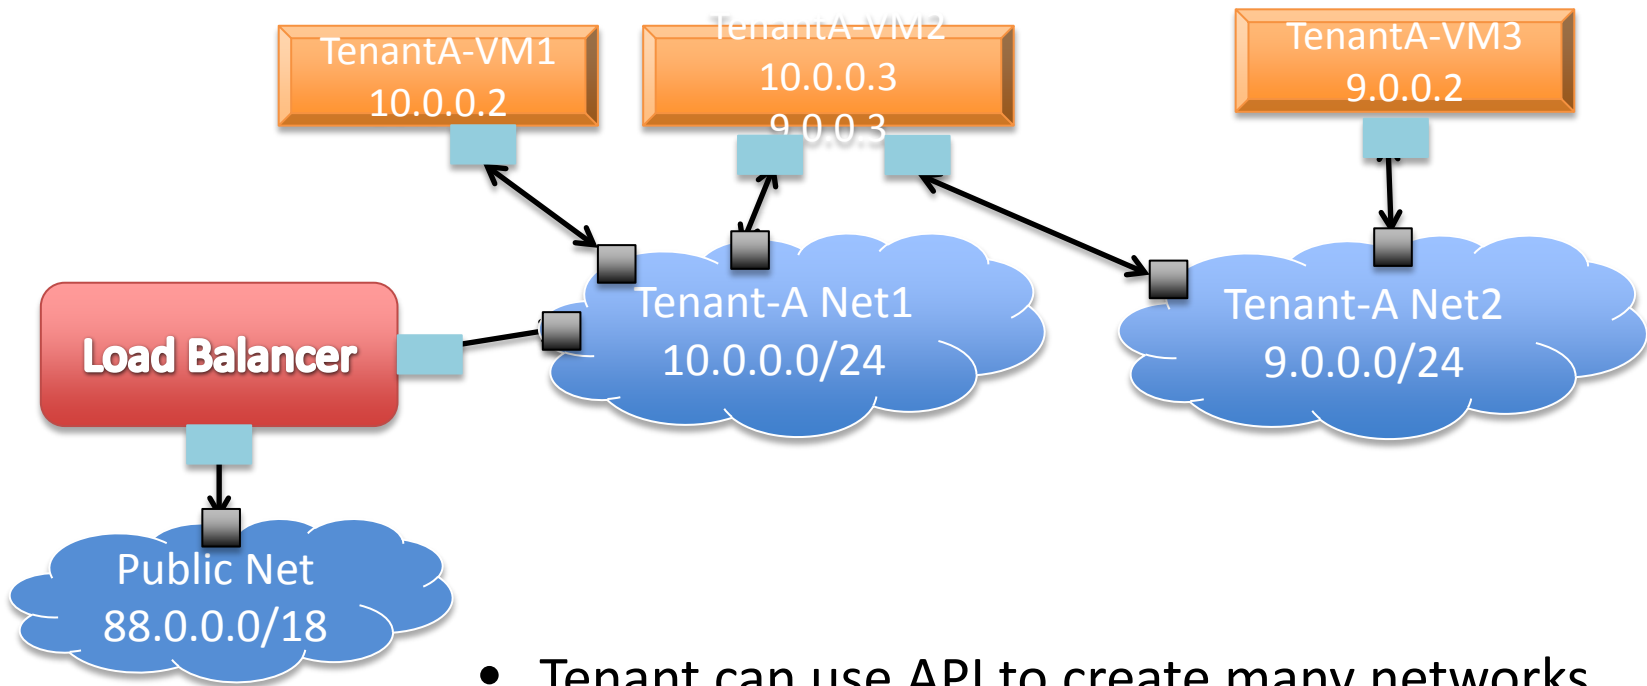
Tenant View



Provider View






Quantum Model: Dynamic Network Creation + Association

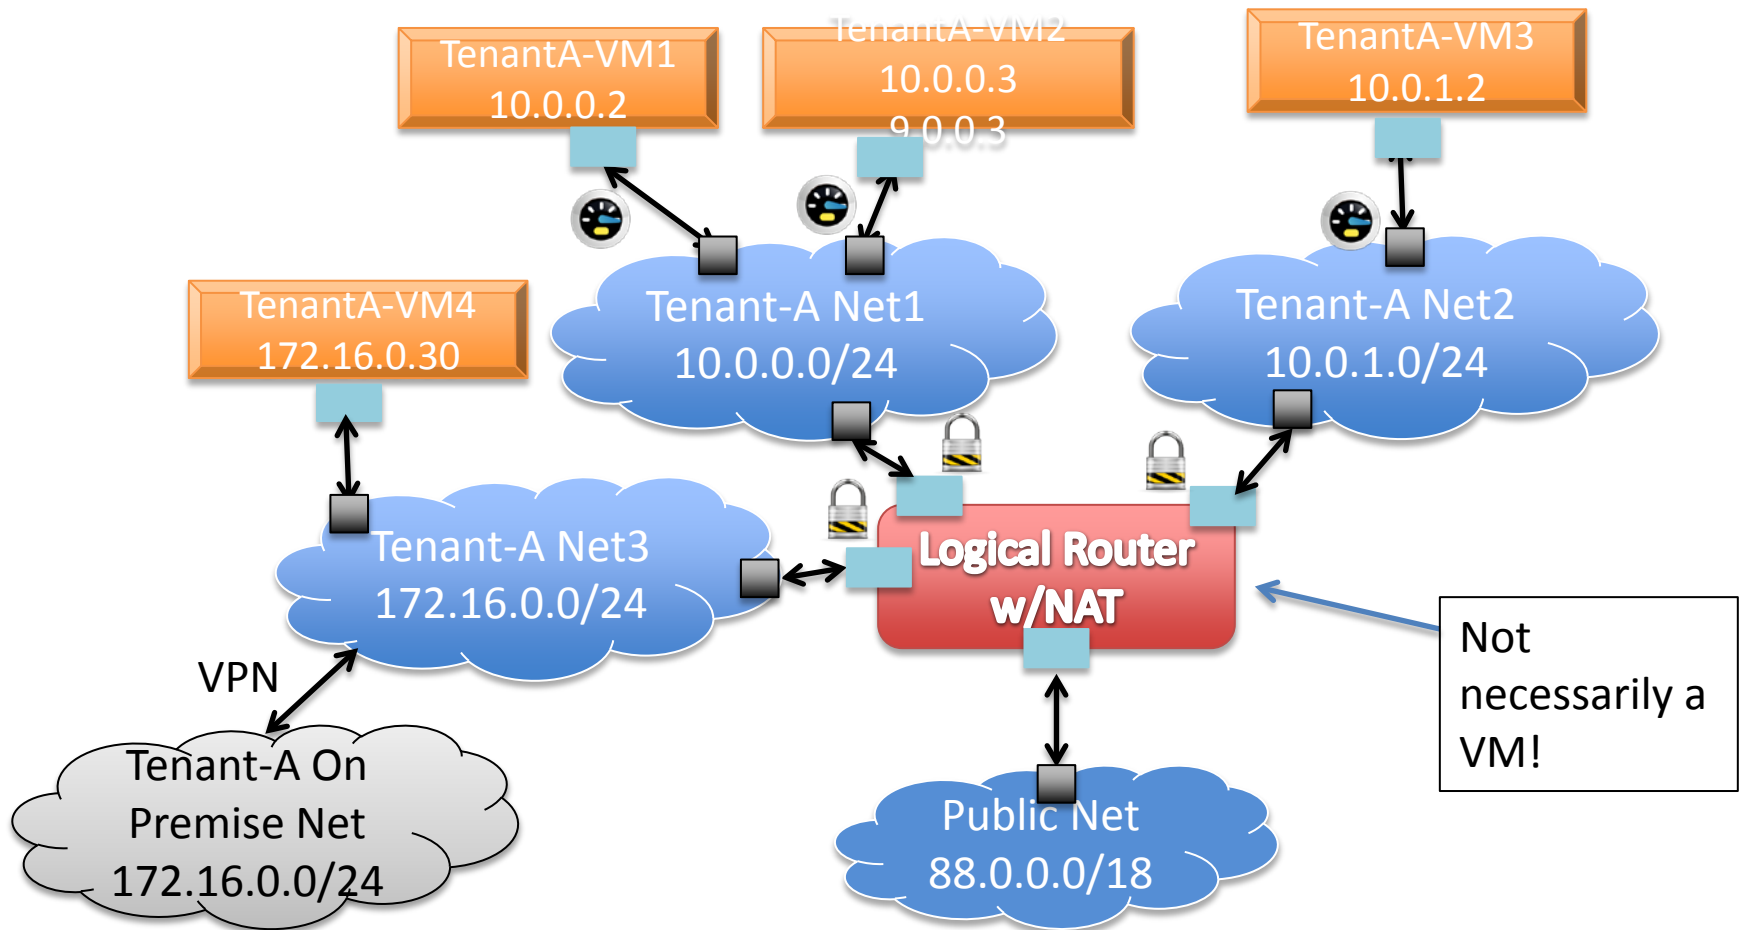


- Tenant can use API to create many networks.
- When booting a VM, define which network(s) it should connect to.
- Can even plug-in instances from other services (e.g., a load-balancing service).


Quantum API Extensions

- Enables innovation in virtual networking.
 - Tenants can query API to programmatically discover supported extensions.
 - Overtime, extensions implemented by many plugins can become “core”.
- Add properties on top of existing network/port abstractions:
 - QoS/SLA guarantees / limits 
 - Security Filter Policies 
 - port statistics / netflow 
- New Services
 - L3 forwarding, ACLs + NAT (“elastic” or “floating” IPs)
 - VPN connectivity between cloud and customer site, or another cloud datacenter.

Example: Quantum + Extensions



Quantum Architecture



An eco-system of tools that leverage the Quantum API.

A generic tenant API to create and configure “virtual networks”

A “plugin” architecture with different back-end “engines”

Quantum “Plugins”

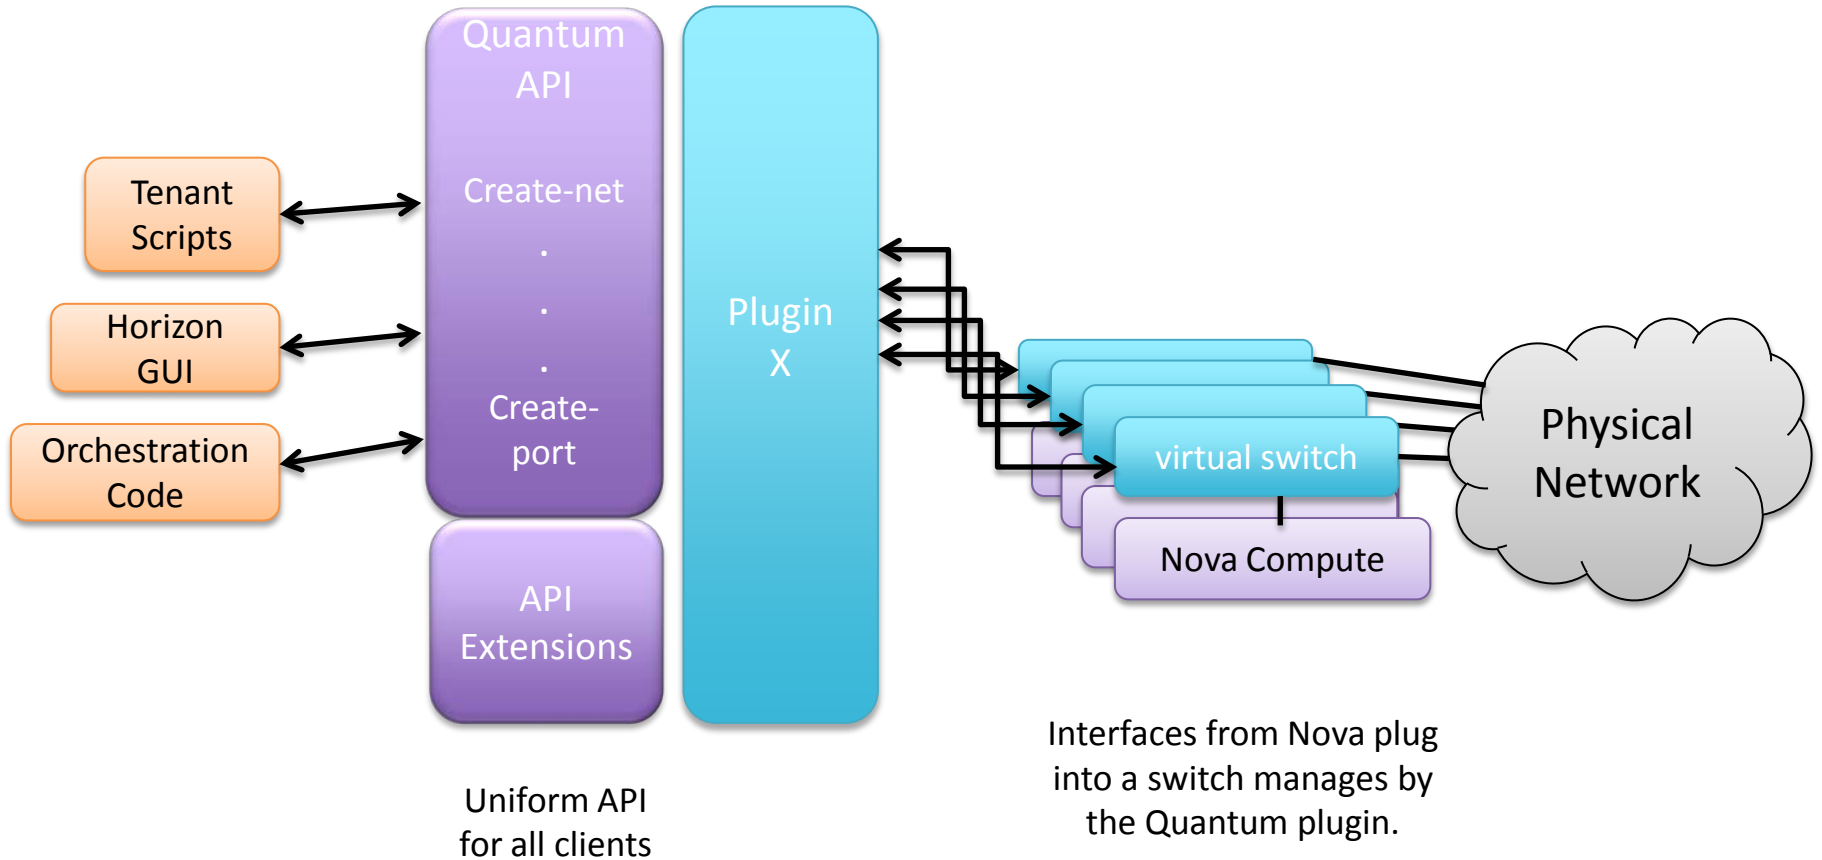
- Different plugin “engines” present different trade-offs:
 - Free vs. Commercially Supported
 - Advanced Features (exposed as API extensions)
 - Scalability
 - High Availability (control & data plane)
 - Hypervisor Compatibility
 - Network HW Compat (vendor specific? Allow L3 scale-out?)
 - Manageability / troubleshooting
- Cloud Operators weigh trade-offs, choose a plugin.
- Note: Back-end technology hidden behind core Quantum API
 - Example: VLANs vs. tunneling

Quantum Architecture (generic)

API Clients

Quantum Service

Backend X



Open Source and Commercial Quantum Plugins

Basic open source plugins based on Open vSwitch, Linux Bridge, and Ryu network operating system exist

The following vendors have publicly stated that they already have or are developing a Quantum plugin (others exist as well)



Quantum Project Releases

- Incubation release (OpenStack Essex, April '12)
 - v1 API, base L2 API abstractions
 - 5 plugins available.
 - In production at early adopters.
- First “core” release (OpenStack Folsom, Oct. '12)
 - v2 API, with L2 + L3 API abstractions
 - Additional plugins available.
 - Quantum becomes “default” networking option for OpenStack.

New in Folsom – Quantum API v2

- “merge” between Quantum v1 and Melange
 - API specification draft:
<http://wiki.openstack.org/Quantum/APIv2-specification>
- New resource: Subnet
 - IP address blocks for instances, IP gateway, dns nameservers, host routes and IP allocation pools
- IP and MAC allocation for ports
- Shared networks
- “Provider” networks (API extension)

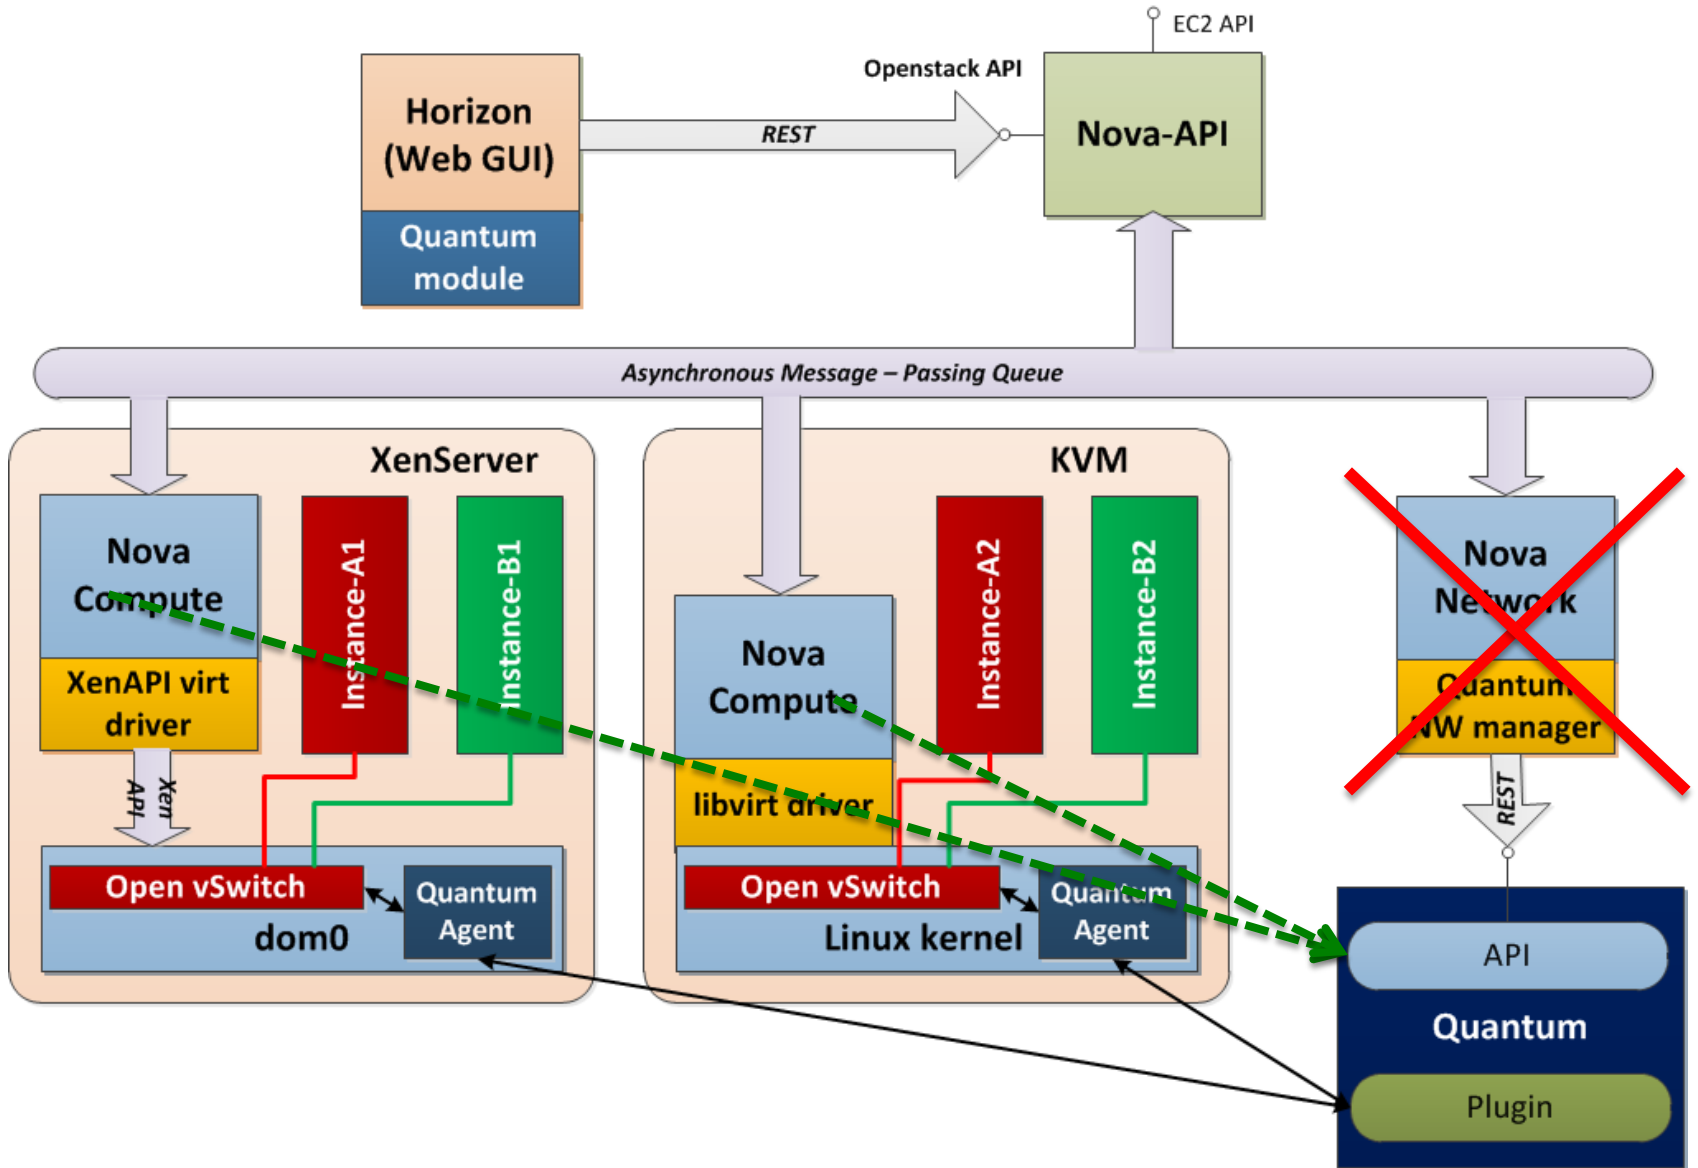


New in Folsom – Interaction with nova

- Nova-network not invoked anymore
- Compute manager calls Quantum directly
- Capabilities exposed by the “Quantum Manager” for nova network now provided by Quantum



Changes in Nova/Quantum interaction



New in Folsom – DHCP agent

- Sends IP configuration to instances
 - IP address and default gateway
 - DNS nameservers
 - Host Routes
- Configuration info fetched from Quantum database
- Attaches to Quantum networks and listens for DHCP requests





Quantum Case Study: Quantum Plugin for Nicira NVP

Nicira NVP Network Virt Model

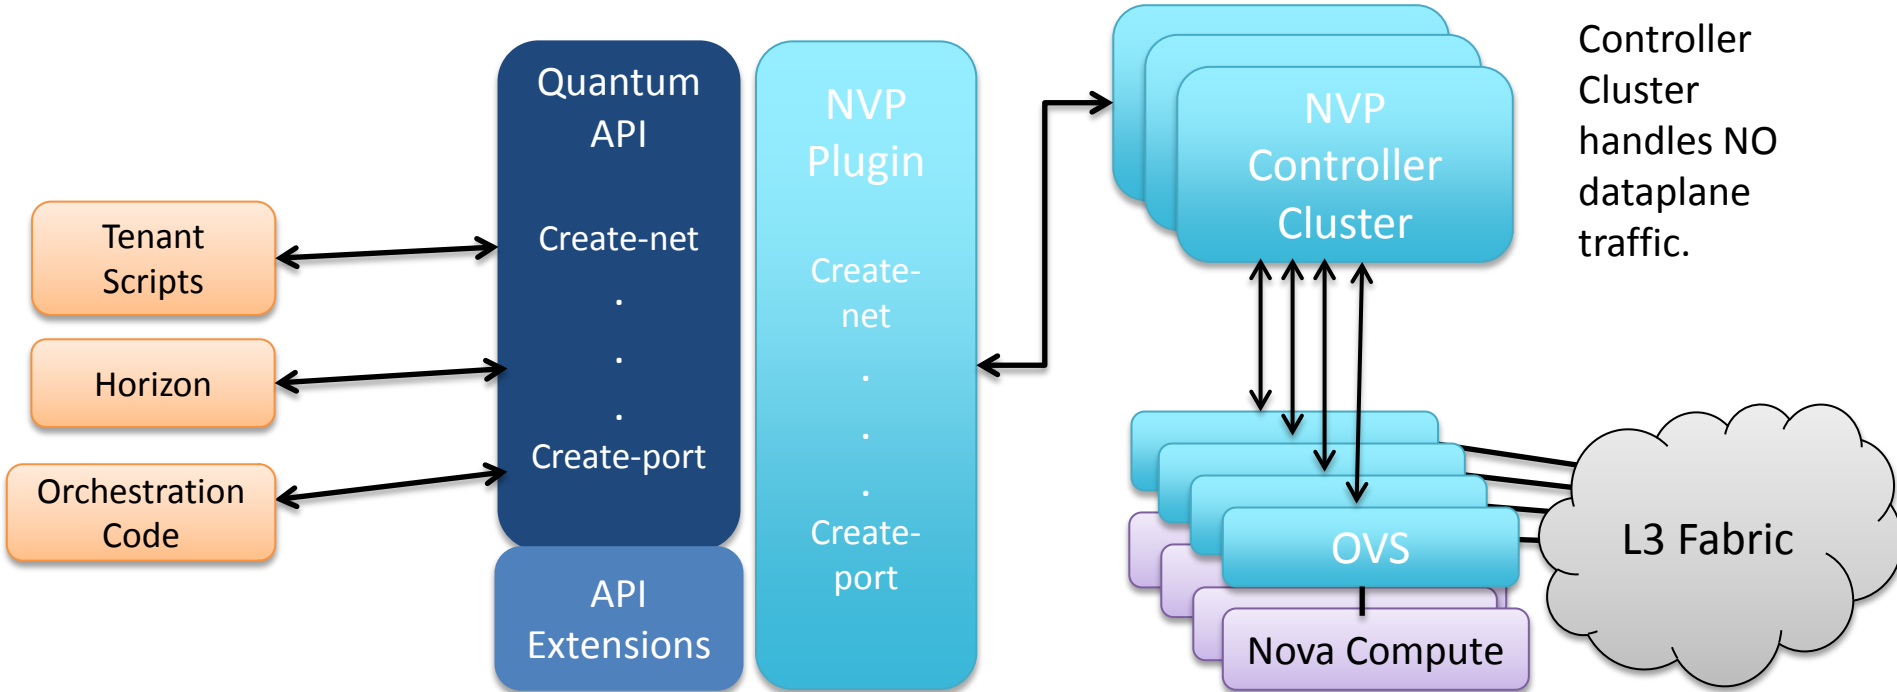
- Edge virtualization in hypervisor (open vswitch) with overlay tunneling decouples logical + physical topology.
 - Flexibility designing Fabric (requires only IP unicast)
 - Can use traditional design, or Fat-tree/Clos
 - No requirement for L2 adjacency, large MAC/ACL tables in HW
 - Place/move any workload anywhere in the DC.
- Control Plane work is Distributed across nodes within the NVP control cluster to provide scalability & fault tolerance.
- Quantum Tenants can dynamically create/modify/monitor rich networks abstractions via Quantum API.

Quantum w/NVP Architecture

API Clients

Quantum Service

Nicira NVP





Thanks!
Questions?

Salvatore Orlando

sorlando@nicira.com

OpenStack Quantum core developer

twitter - @taturiello