nova-network: The Dirty Details

Ryan Richard, RHCA
OpenStack Architect - Private Cloud
ryan.richard@rackspace.com
@rackninja

April 2013

Why nova-network?

- Pre-existing installs
- Folsom Deployments
- Quantum:
  - http://docs.openstack.org/trunk/openstack-network/admin/content/ch_overview.html
  - https://wiki.openstack.org/wiki/Quantum

nova-network overview

- Provides networking for instances
- flat, flatDHCP,flatVLAN
- iptables, ebtables, linux bridge
- “behind the scenes” - no direct API
- http://docs.openstack.org/folsom/openstack-compute/admin/content/list-of-compute-config-options.html

nova-network overview

- Host Network - Physical server communication, management network
- Fixed Network - L3 network range for instances, instance to instance communication

nova-network overview

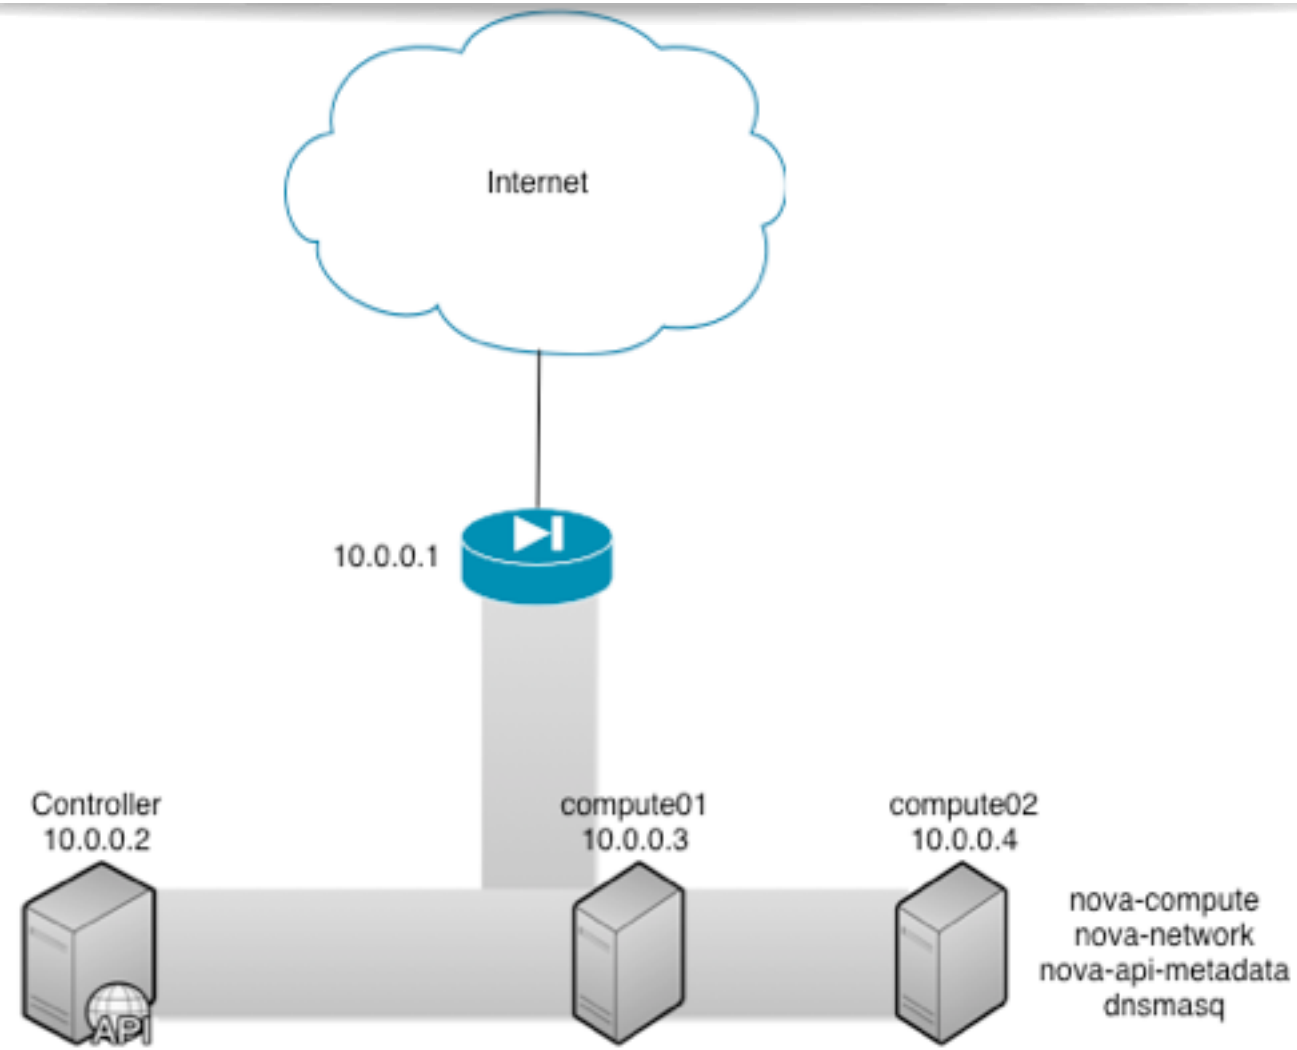

nova-network overview

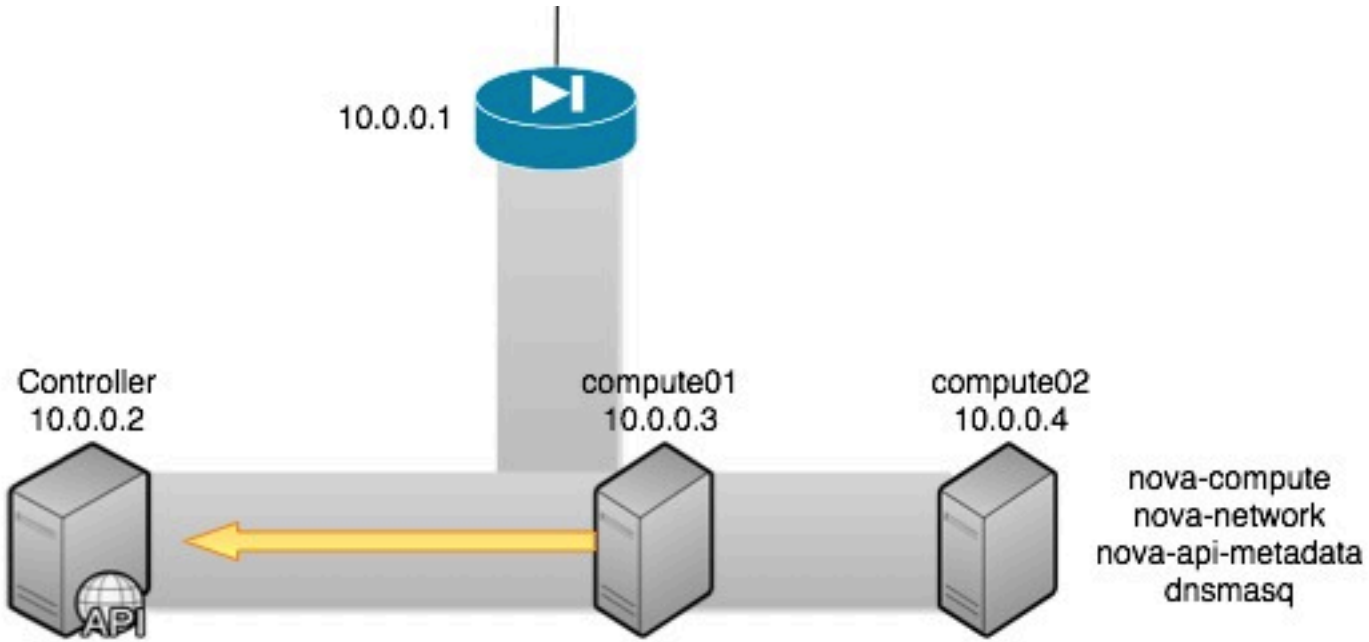

1. Compute requests IP from Fixed range

nova-network overview

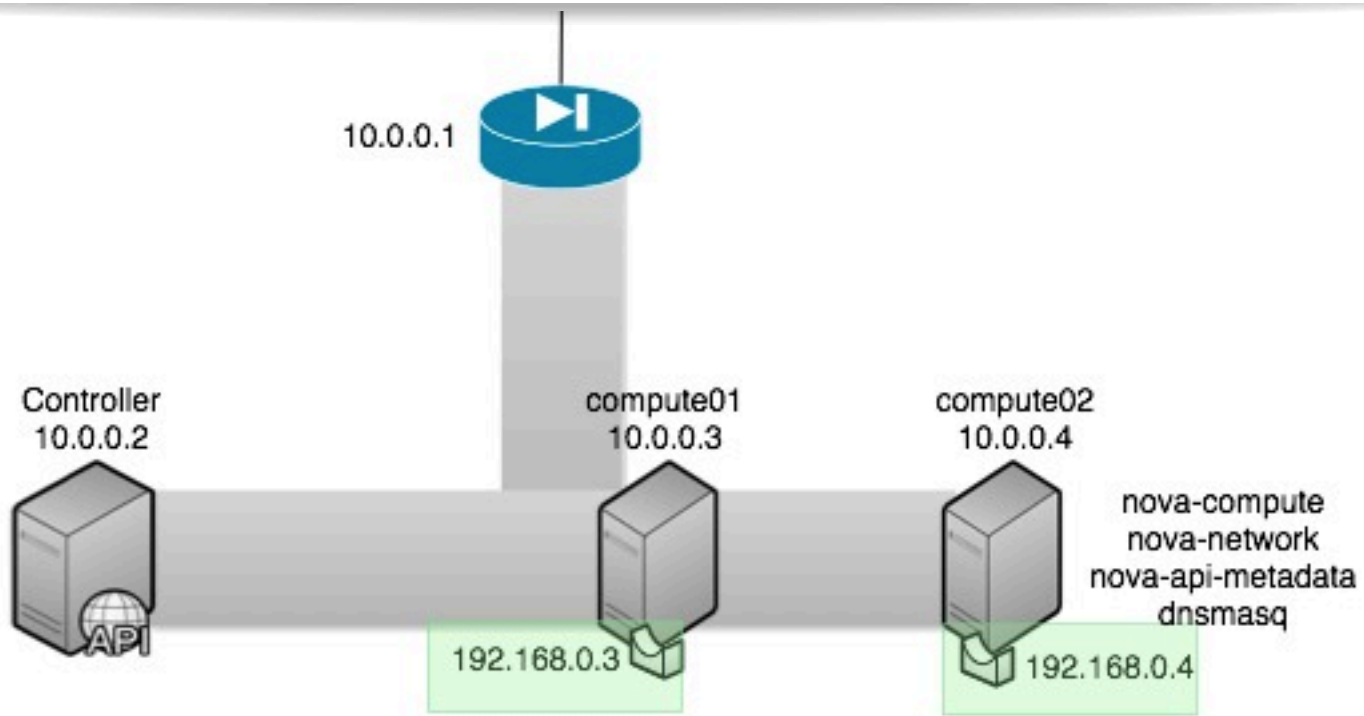

2. Creates Bridge, Place specified interface in bridge, Put IP on bridge

nova-network overview

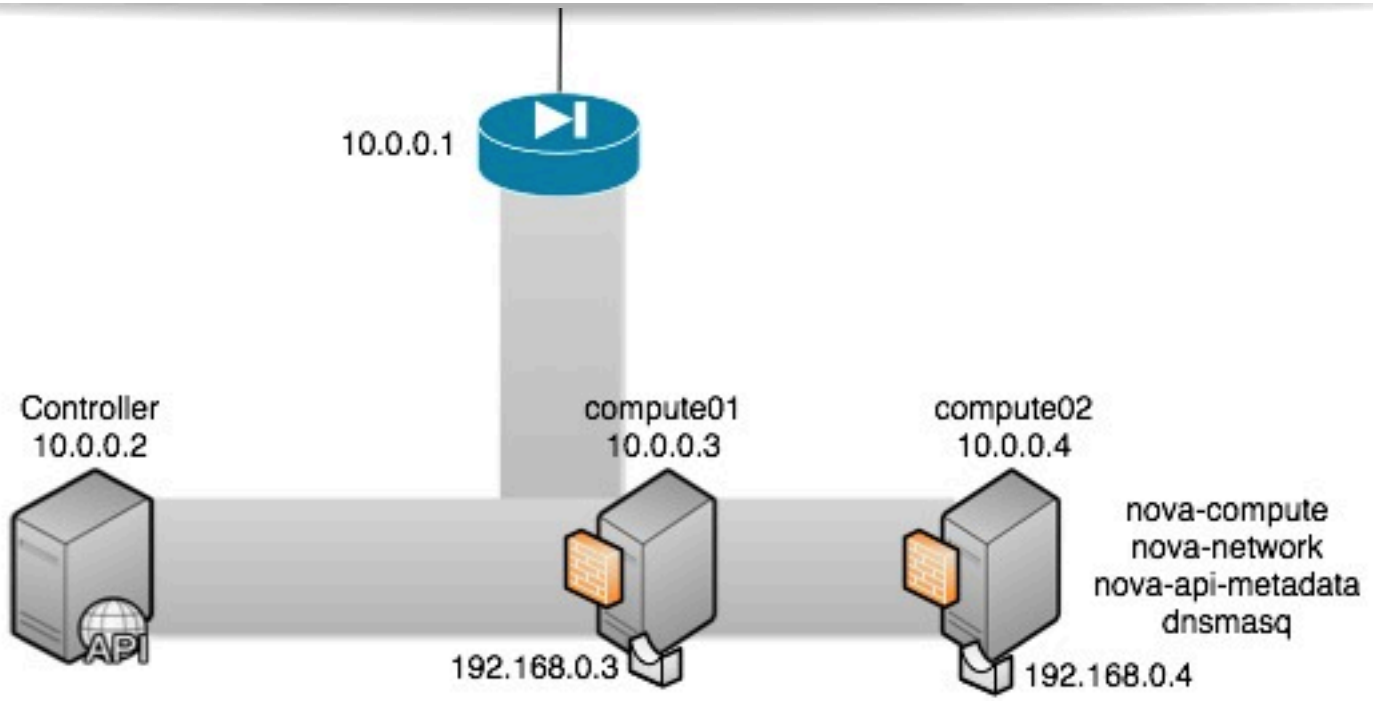

3. nova-network sets up initial iptables/ebtables

nova-network overview

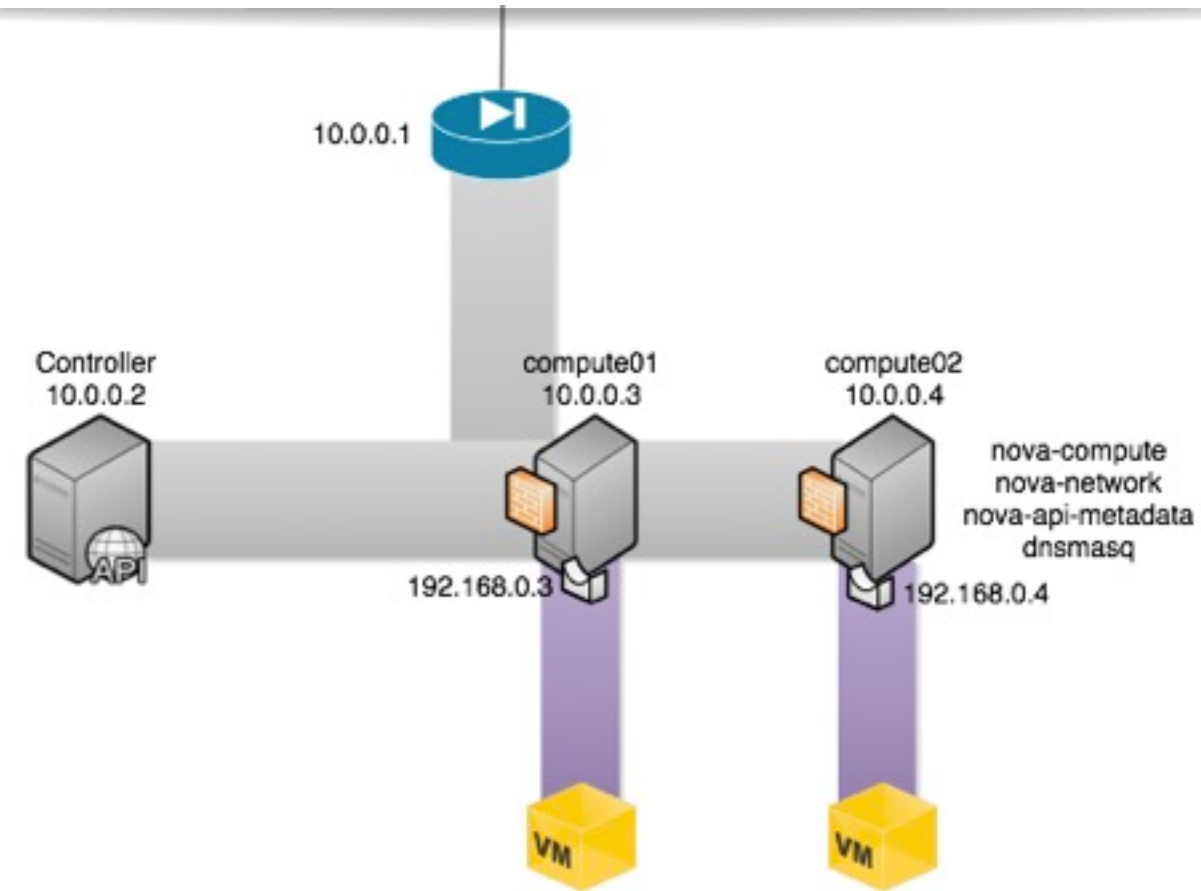

4. Instance is created, security groups set (iptables)

nova-network overview

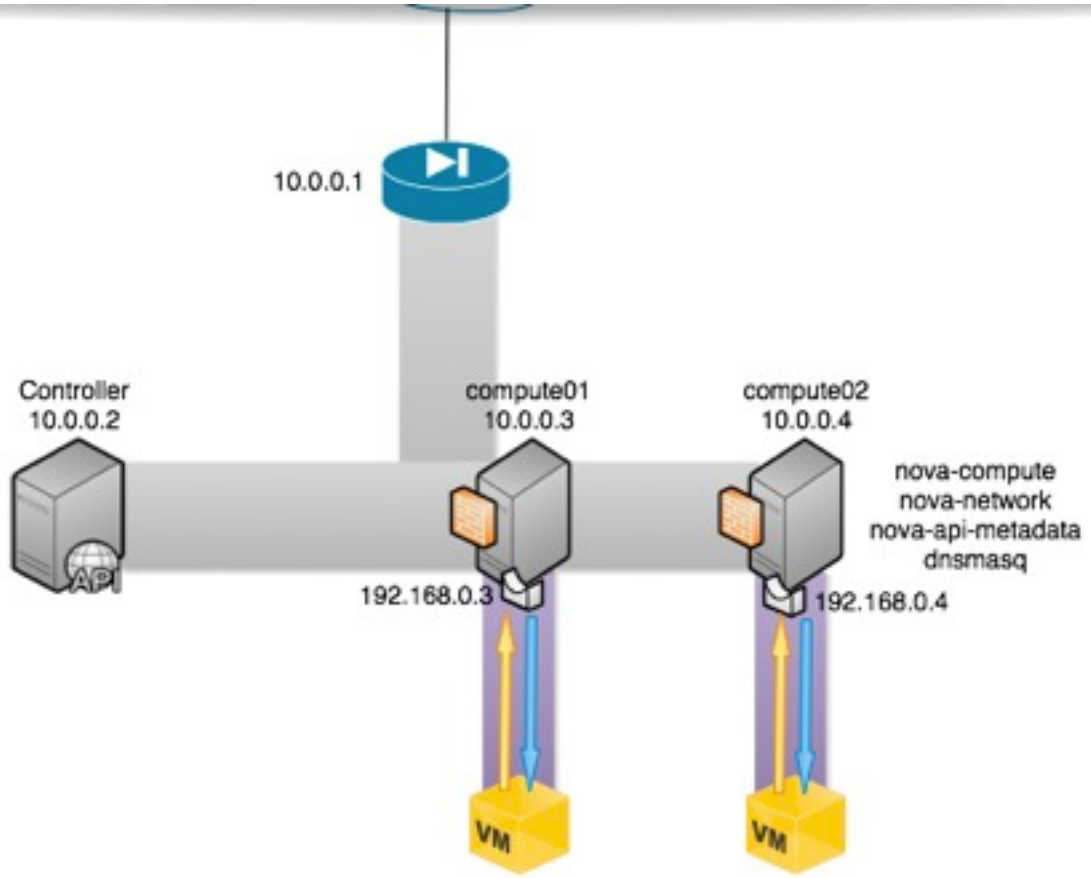

5. Instance sends DHCP request, response from dnsmasq

nova-network overview

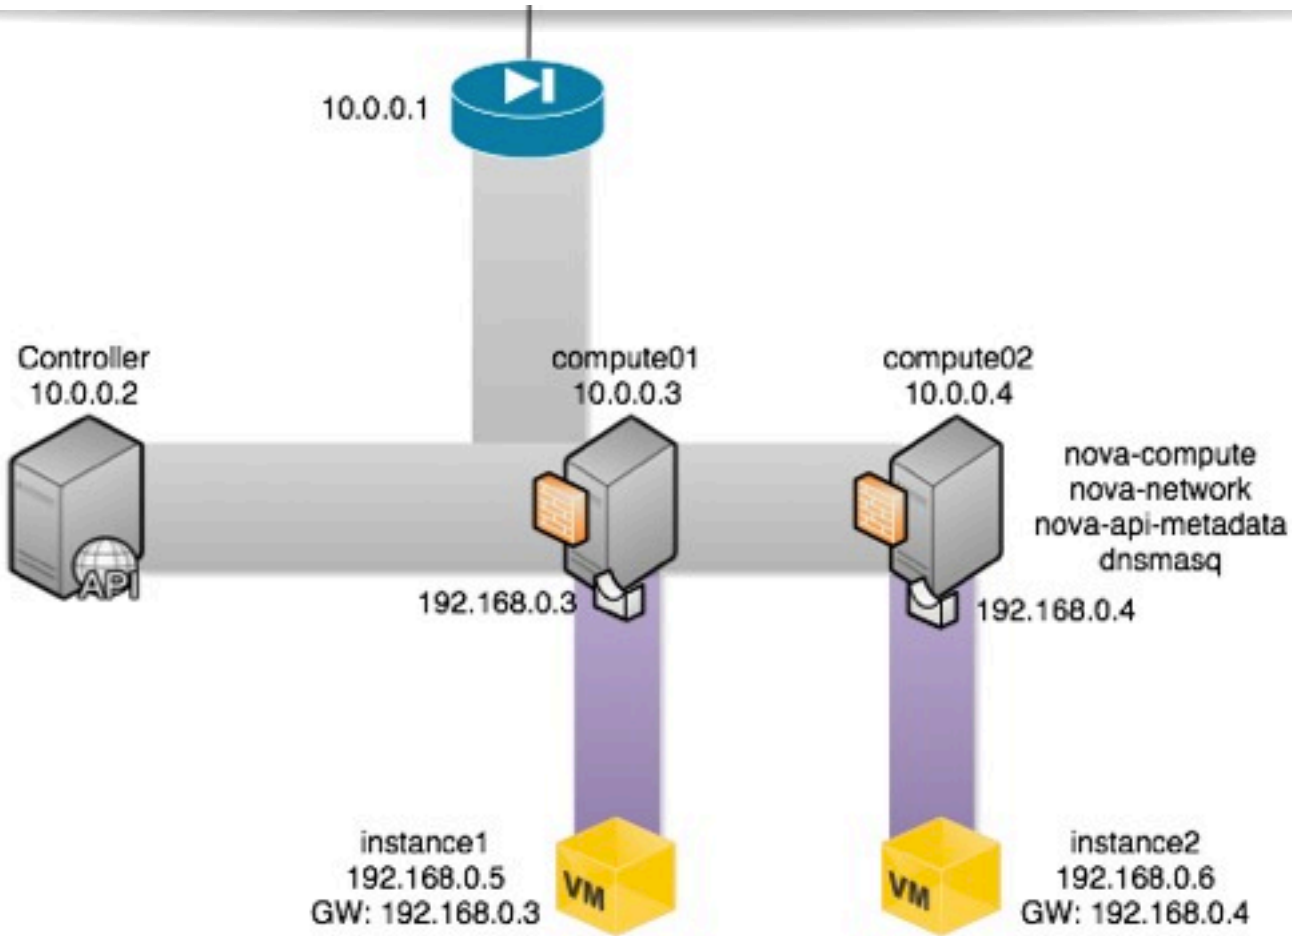

rackspace
the open cloud company

6. Instance sets IP/netmask/Gateway/etc

nova-network options

- 50+ options for networking config
- multi_host = multiple nova-network processes (1 per compute host)
- DNS, DHCP, public_interface, dmz_cidr

public interface

- Decides which interface the default SNAT rule applies
- # iptables -t nat -nvL nova-network-snat
- public internet access

nova-network options

- dnsmasq options
  - DHCP Lease
  - Hardware Gateway
  - DNS domain

nova-network options

- DMZ_CIDR
  - NAT exclusion list
  - ACCEPT rule in iptables NAT
  - # iptables -t nat -nvL nova-network-POSTROUTING

iptables & ebtables

- iptables
  - Security Groups implementation - 1 chain per instance
  - Default: Restrict all access
  - Responsible for NAT
  - Chain example: nova-compute-inst-771

iptables & ebtables

- ebtables
  - IP/MAC/ARP spoofing protections
  - Only 1 IP per instance
  - defined in /etc/libvirt/nwfilter/ (libvirt implementations)

floating IPs

- Easy to Add
- MUST be associated with the public_interface flag
- Don't get assigned inside the instance but instead rely on iptables (SNAT/DNAT)
- Dynamically assigned

floating IPs

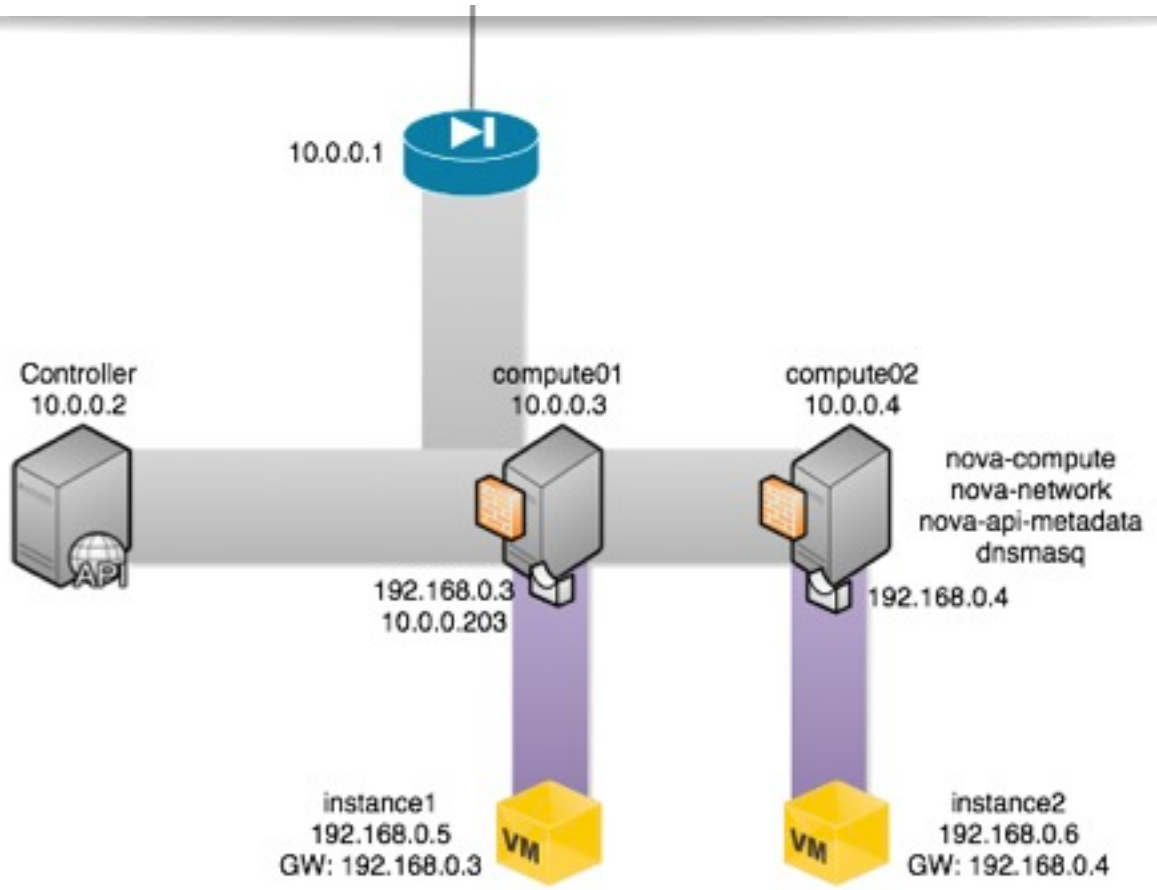

1. Floating IP gets assigned to bridge

floating IPs

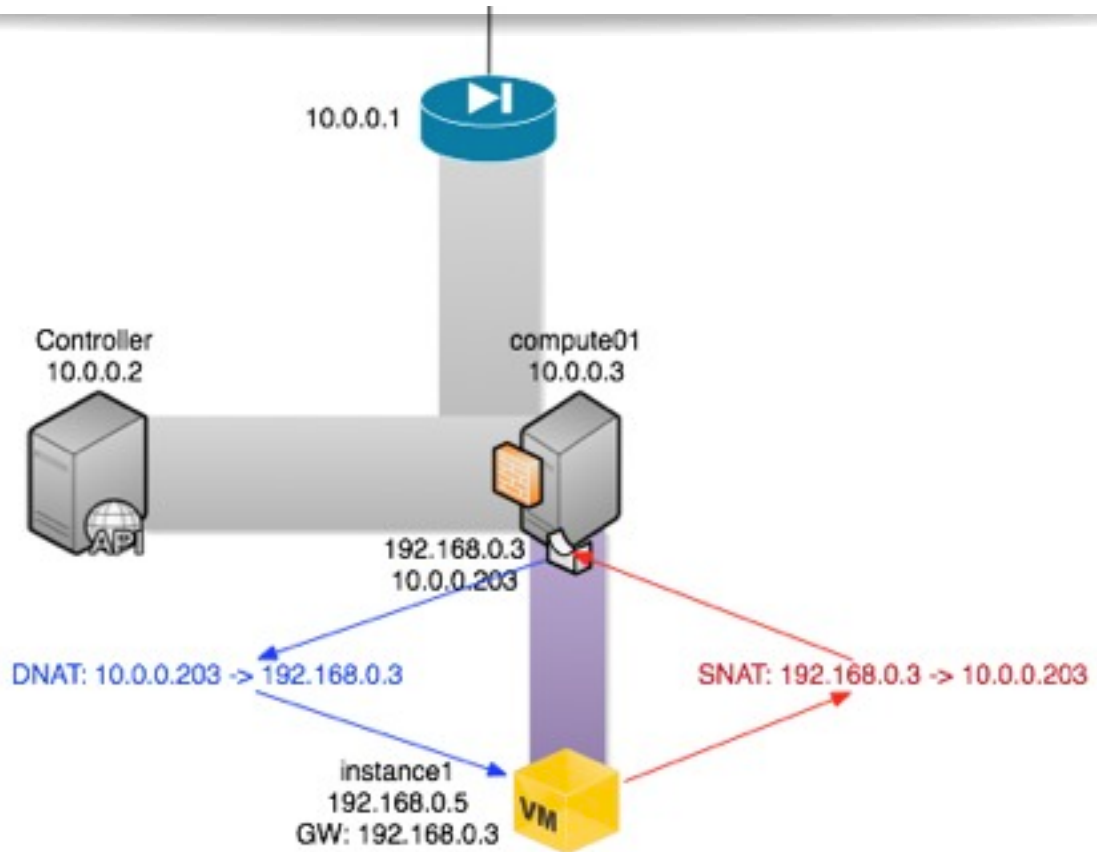

2. DNAT/SNAT rules get created via "iptables -t nat"

Integrating

- Difficult
- OpenStack is IPAM (partially)
- DNS integration is lacking

Example

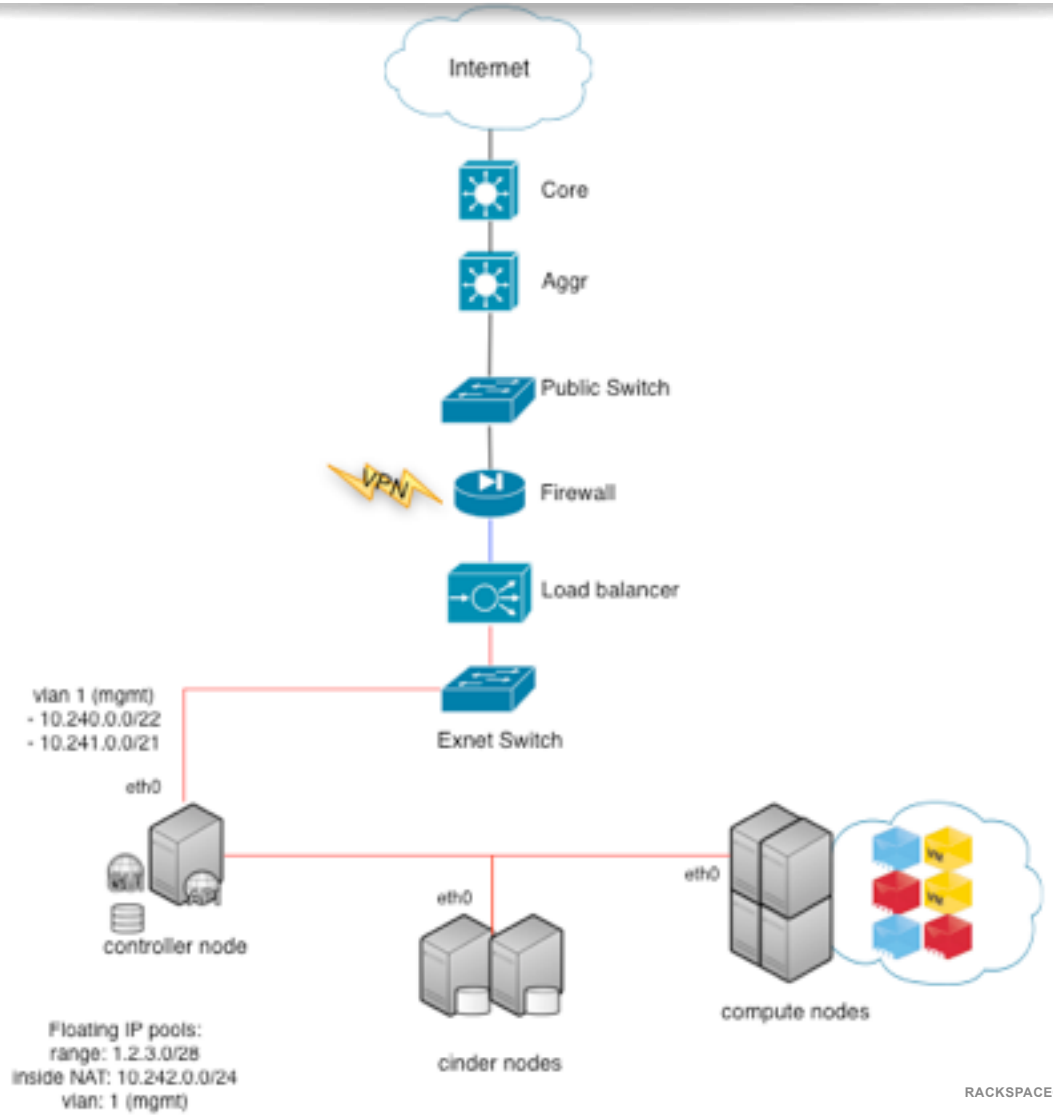
Internet
Core
Aggr
Public Switch
VPN
Firewall
Load balancer
vlan 1 (mgmt)
- 10.240.0.0/22
- 10.241.0.0/21
Exnet Switch
eth0
controller node
eth0
cinder nodes
eth0
compute nodes
Floating IP pools:
range: 1.2.3.0/28
inside NAT: 10.242.0.0/24
vlan: 1 (mgmt)

Example

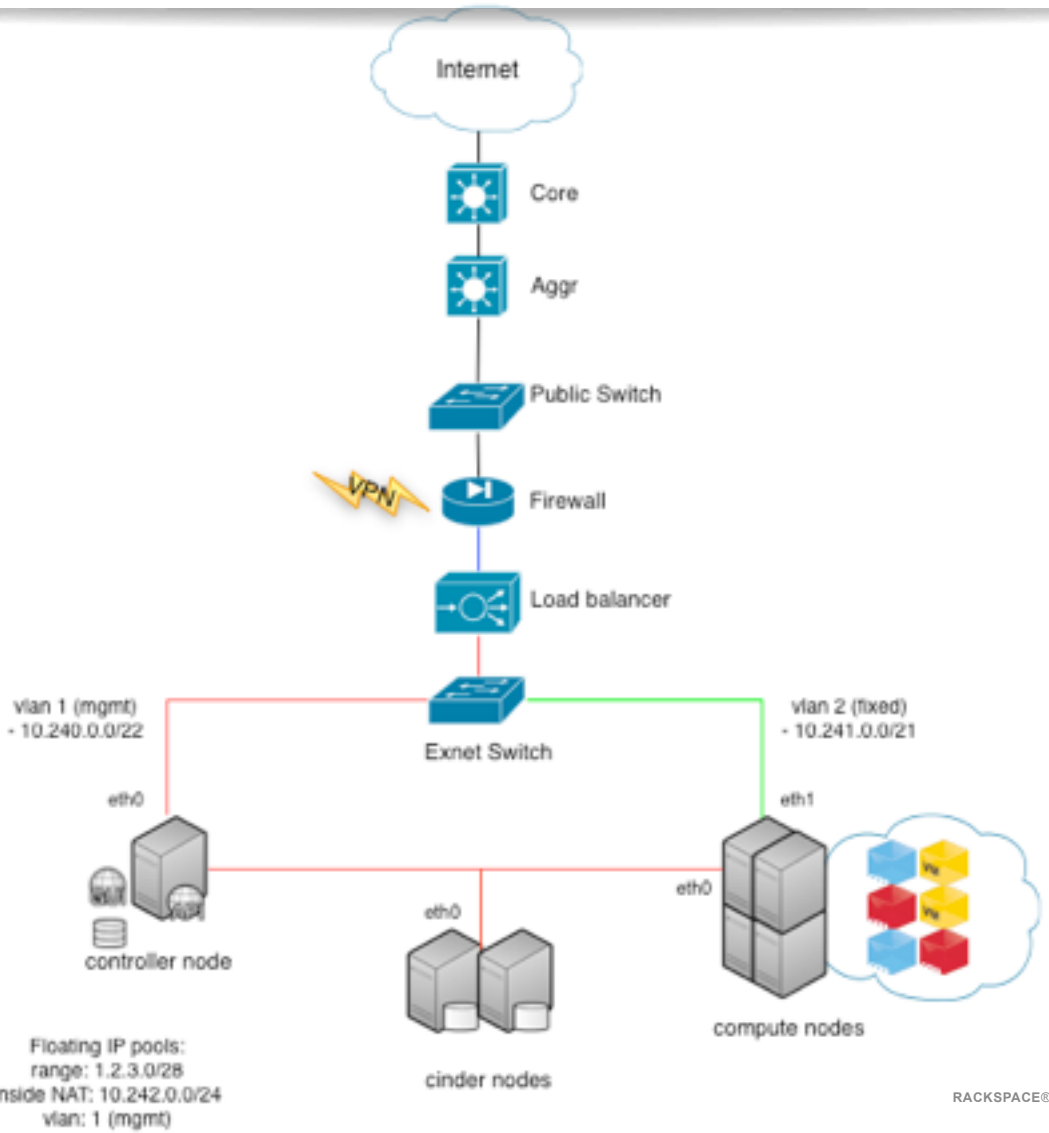

rackspace
the open cloud company

Open to discussions/thoughts/questions

THANK YOU

Rackspace is hiring

www.rackertalent.com

RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78218
US SALES: 1-800-961-2888 | **US SUPPORT:** 1-800-961-4454 | WWW.**RACKSPACE**.COM