Towards XaaS and Facility-wide Flexible IT

John Hover <jhover@bnl.gov>

Xin Zhao <xzhao@bnl.gov>.

Big Picture

Grid= BSaaS (Batch Shell as a Service):

Site is responsible for everything up to the shell login -> heterogeneity.

Cloud=IaaS (Infrastructure as a Service):

Site is responsible for providing x86 CPU(s), Memory, Disk, and Network.

Everything else specified by end user.

Facility+Cloud = Virtual data centers

Configure all physical resources as cloud compute hosts. All users, internal and external, invoke resources dynamically when needed, for as long as they are needed.

- Purely internal usage simplifies site admin.
- Adding external usage allows centrally administered distributed facilities.

BNL Openstack Instances

First:

- 47 nodes x 16 cores (752 core) x 24GB x 500GB
- Grizzly
- Controller + DB + Neutron
- Flat Networking
- Currently running ~260 jobs (prod analy combined)

Second:

- 47+? x 16core x 32GB X 1TB
- Havana
- Controller + DB + 2xQpid + Neutron w/ 3-way Load balancing agents
- Will begin offering Swift (S3 block storage service)

Instances will alternate between prod and testbed.

Openstack at CERN

Heavily involved and deploying the new data center entirely on Openstack.

- But, heavily customized architecture,
- Grizzly based.
- 2 Cells w/ front end load balancing controller.
- External Oracle DB (1per cell)
- Multi-node rabbitmq cluster (2x3)
- Custom-written network driver that interacts w/ CERN network systems
- CERN IT (not native Openstack) manages IPs and default routes (NAT)
- Glance w/ Ceph back end used for VM image storage.

Heavy personnel requirements, but leveraging existing CERN IT capabilities (networking, databases).

Openstack Project Roadmap

FYI Openstack version sequence:

- Essex (Jan 2012)
- Fulsom (Sep 2012)
- Grizzly (Apr 2013) BNL Instance One
- Havana (Oct 2013)
- Icehouse (April 2014)
- Juno (October 2014)

Issues/ Roadblocks

Key problem is networking (public IPs vs. NAT)

- Network module provides routing for virtual subnets, and NAT for outbound connectivity.
  - 1 public IP per Openstack instance: NAT host (network daemon) is bandwidth bottleneck.
  - 1 public IP per compute host?: possible w/ Essex, not w/Grizzly/Havana
  - 1 public IP per VM: avoids problem, but uses a lot of IPs

NOTE: “Outbound” means outside of OS virtual subnets.

How are others avoiding this bottleneck?

CERN and ATLAS Sim@P1 bypass this by using non-Openstack or hybrid OS/local networking mechanisms, i.e. they tie into physical network rather than using purely virtual networking.

Networking capabilities

Essex had *multi-host* mode.

- Each compute host runs an instance of nova-network daemon (virtual router)
- Permits each compute host to NAT outbound traffic
- Grizzly removed this mode, with no replacement.
- All outbound traffic for all VMs for all tenants go through single network node.

Havana provides *network-scheduler* plugin

- Still one network daemon, but outbound routing can be load balanced across multiple network agent hosts (HAProxy or F50

Essex Network Diagram
1 Head Node per Rack: Public IP
Compute nodes: Private IP
VMs: 1 Private VLAN per project:

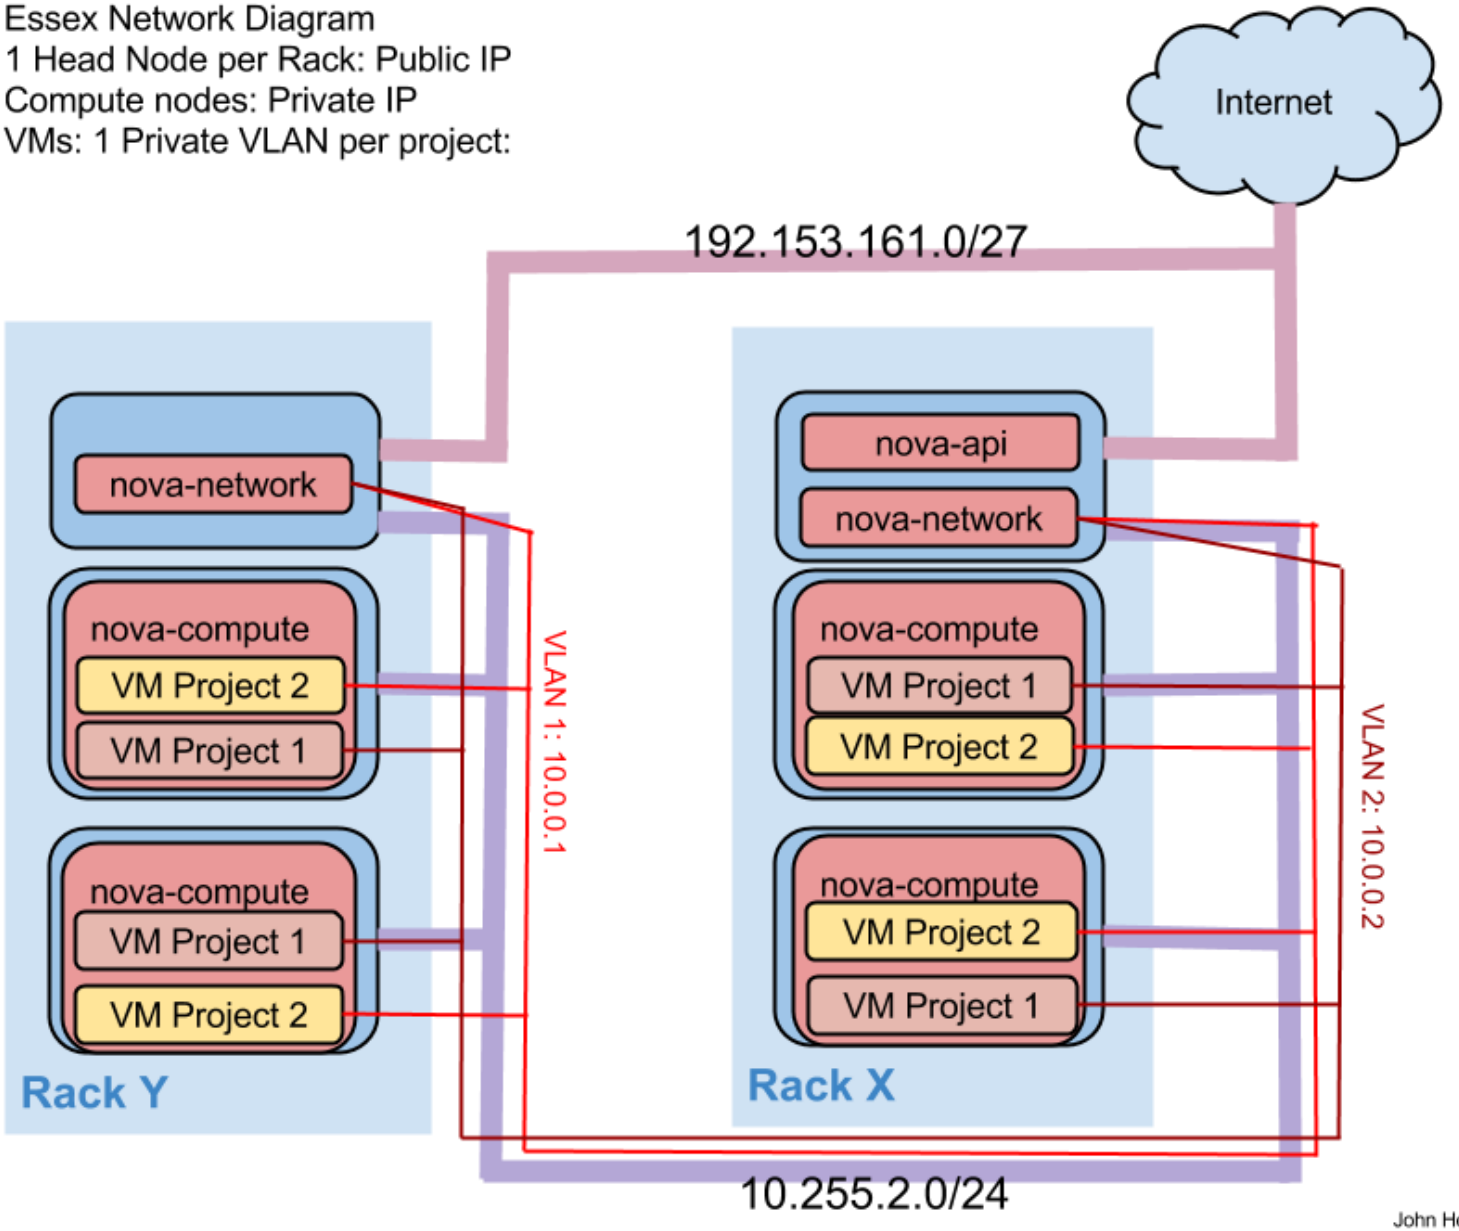

John Hover, BNL

Grizzly Network Diagram
1 Head Node for Instance: Public IP
Compute nodes: Private IP
VMs: 1 Private VLAN per project:

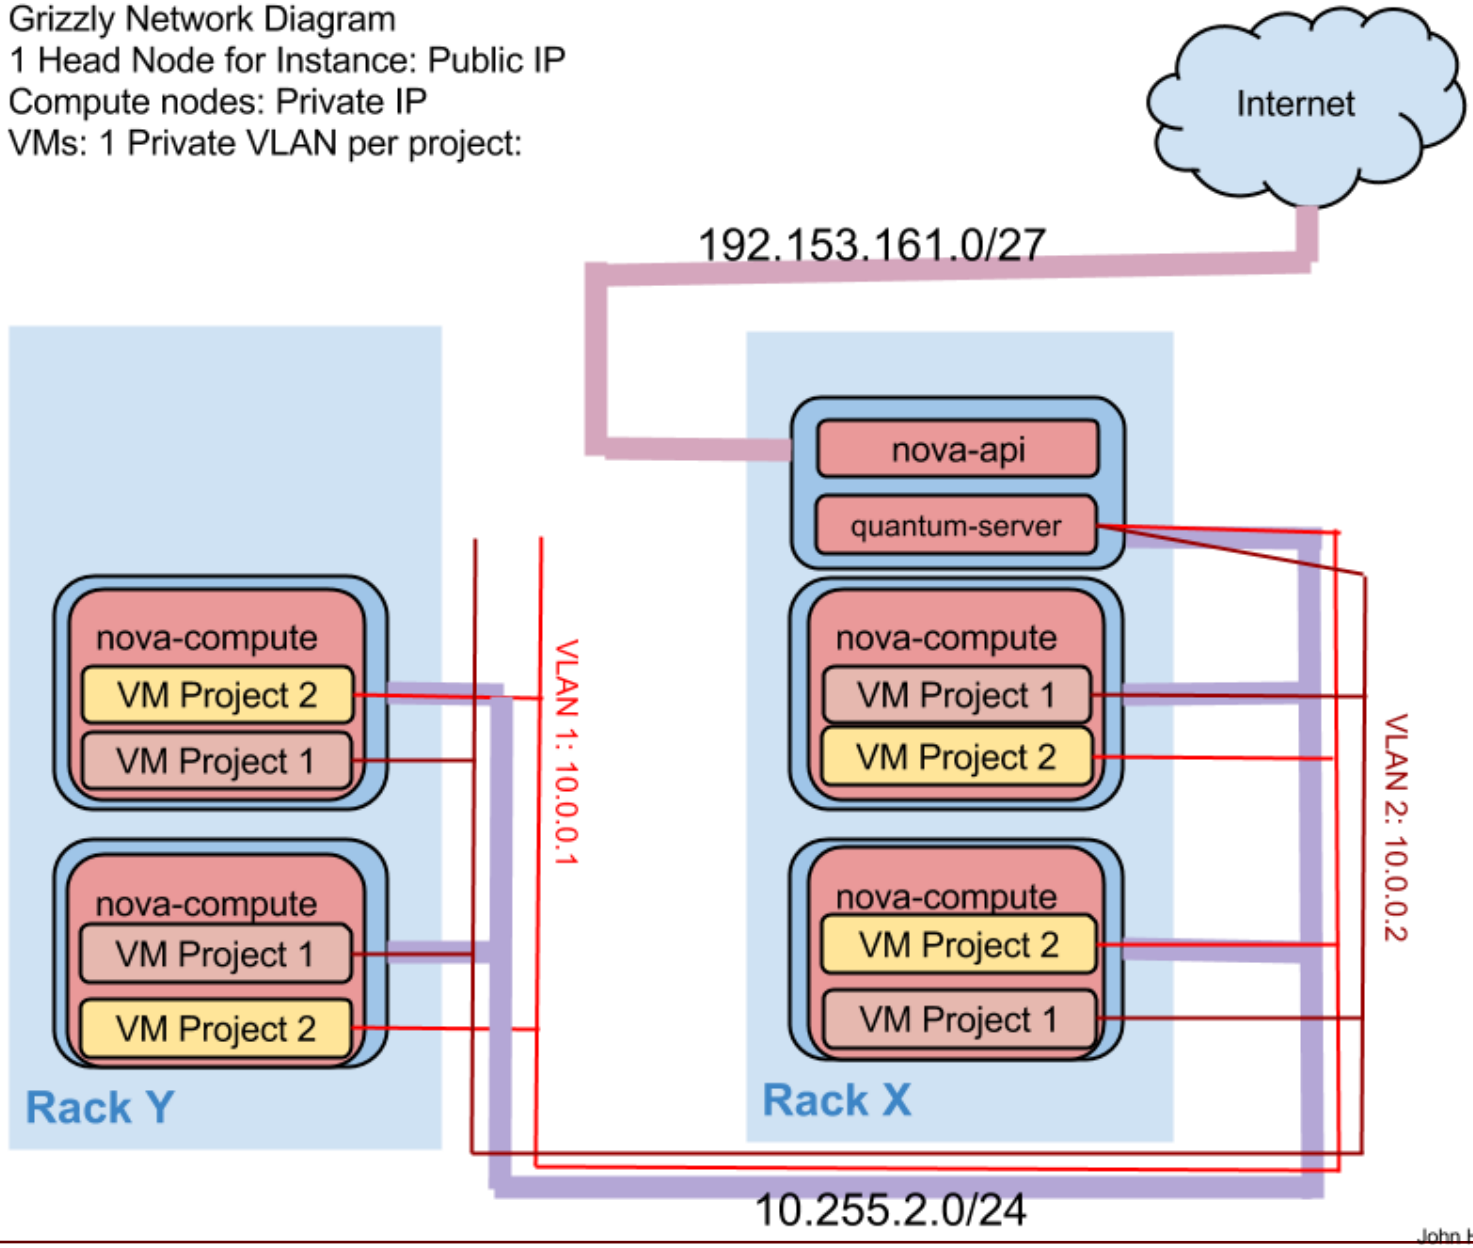

Other Issues

Security

If you want to insulate tenants from one another, you need VLANs (pre-allocated, or dynamic SDN). This is orthogonal to the NAT question, but is lost if you tie into native site networking.

Scaling

Messaging service (Qpid or RabbitMQ) clearly requires multi-node architecture.

Database must be run on a separate host.

We have had problems with a fraction of newly-started nodes failing to get Ips. Likely due to message queue congestion.

Hardware/Facility Requirements

Physical NICs (2 per compute node, 3 for controller). This should be achievable with a single NIC by using virtual interfaces (e.g. eth0:1, eth0:2) rather than requiring physical NIC per subnet.

XaaS Rationale

Why do this?

A lot of what we do involves node-oriented services

- E.g. efficient CVMFS config is done at the node level, not per-job.

Once you learn to dynamically invoke and manage an instance of something (difficult) you can:

- Invoke the same thing at another site (private cloud).
- Invoke the same thing on an academic or commercial cloud.
- Create wide-area cross-site systems.

IaaS Cost/Benefit

Each site needs to ask:

To what degree are we spending effort on worker node /grid environment administration? Cloud outsources this.

How many institutional tenants use the site?

How uniform is the network bandwidth/connectivity within cluster? For IaaS network needs to be flat and performant.

How does data work?

Simplest case is use IaaS to create site SEs. Otherwise model stays the same.

What SE technology is most cleanly able to be scaled by invoking new instances?

More refined case is to move toward global wide-area queues, with central stage-out. (for Sim?)

Questions

How long until ***stock*** Openstack has ability to scale to 1000+ nodes?

- Unclear. Ideal would be the return of multi-host network mode allowing router/NAT specified per OS compute host.

How many sites are entirely handled by configuration management framework?

- Is Puppet the prevailing tool? If BNL provided a standard architectural recipe and Puppet manifests, would that tip the scales?

How much custom work/appliances/procedures would be acceptable for USATLAS sites to move forward now?

Sites which can issue public IPs to all VMs (infra + batch) could use Openstack very soon--just hook into local networking, i.e. let VMs get normal DHCP network config.

Question/Thought experiment

What would a fully IaaS-enabled ATLAS look like?

All grid sites provide IaaS interface? Alongside standard CEs?

- What are the issues with hybrid site access?

OSG provides tools and workflows to generate, label, distribute, and register virtual machine images in multiple formats.

- Do images get deployed to standard SEs? Or do they get uploaded via EC2/Openstack-style image registration process?

OSG provides several base images pre-configured for standard work, e.g. wn-client, OSG_APP in OASIS (CVMFS), and standard contextualization mechanisms for locating site Squid, services, etc.

- Usable on EC2, other public clouds out of the box?

What changes about monitoring? Do sites need to be able to observe on-VM processes?

Other Steps

Eliminate two-NIC compute host requirement.

Fully usable Imagefactory-based, templated VM authoring tool. (Done).

- Talk to me if you're interested.

Cloud-oriented LSM:

- FAX for input.
- T1 for output, but could use T2

Fully dynamic in-cloud cluster

- Working in concert w/ Doug B. T3 analysis strategies (e.g. POD) all piggy-back effectively on Condor w/o shared filesystem.

Discuss...

Extra Slides

Example Scenarios

Sites issue 3 tenant credentials: LocalAdmin, ATLAS, and OSG.

- For ATLAS, all of US cores (20k) could be in a single distributed cluster. FAX, CVMFS/OASIS, and the ATLAS event server are enabling technologies for this.
- MCORE resource allocation becomes easy.
- Permits central OSG provisioning services to function independently. No CEs, GUMS, tickets, or VO support issues.
- For local static infrastructural services (e.g. Squid, doors, gateways, etc.), new instances trivial to deploy, upgrade (minutes vs. day).

For storage, SEs could be set up on VMs directly, or backed by S3 block storage at the site.

- These service could be managed by the site, or centrally.

Could reduce the effort required by site-specific high-level

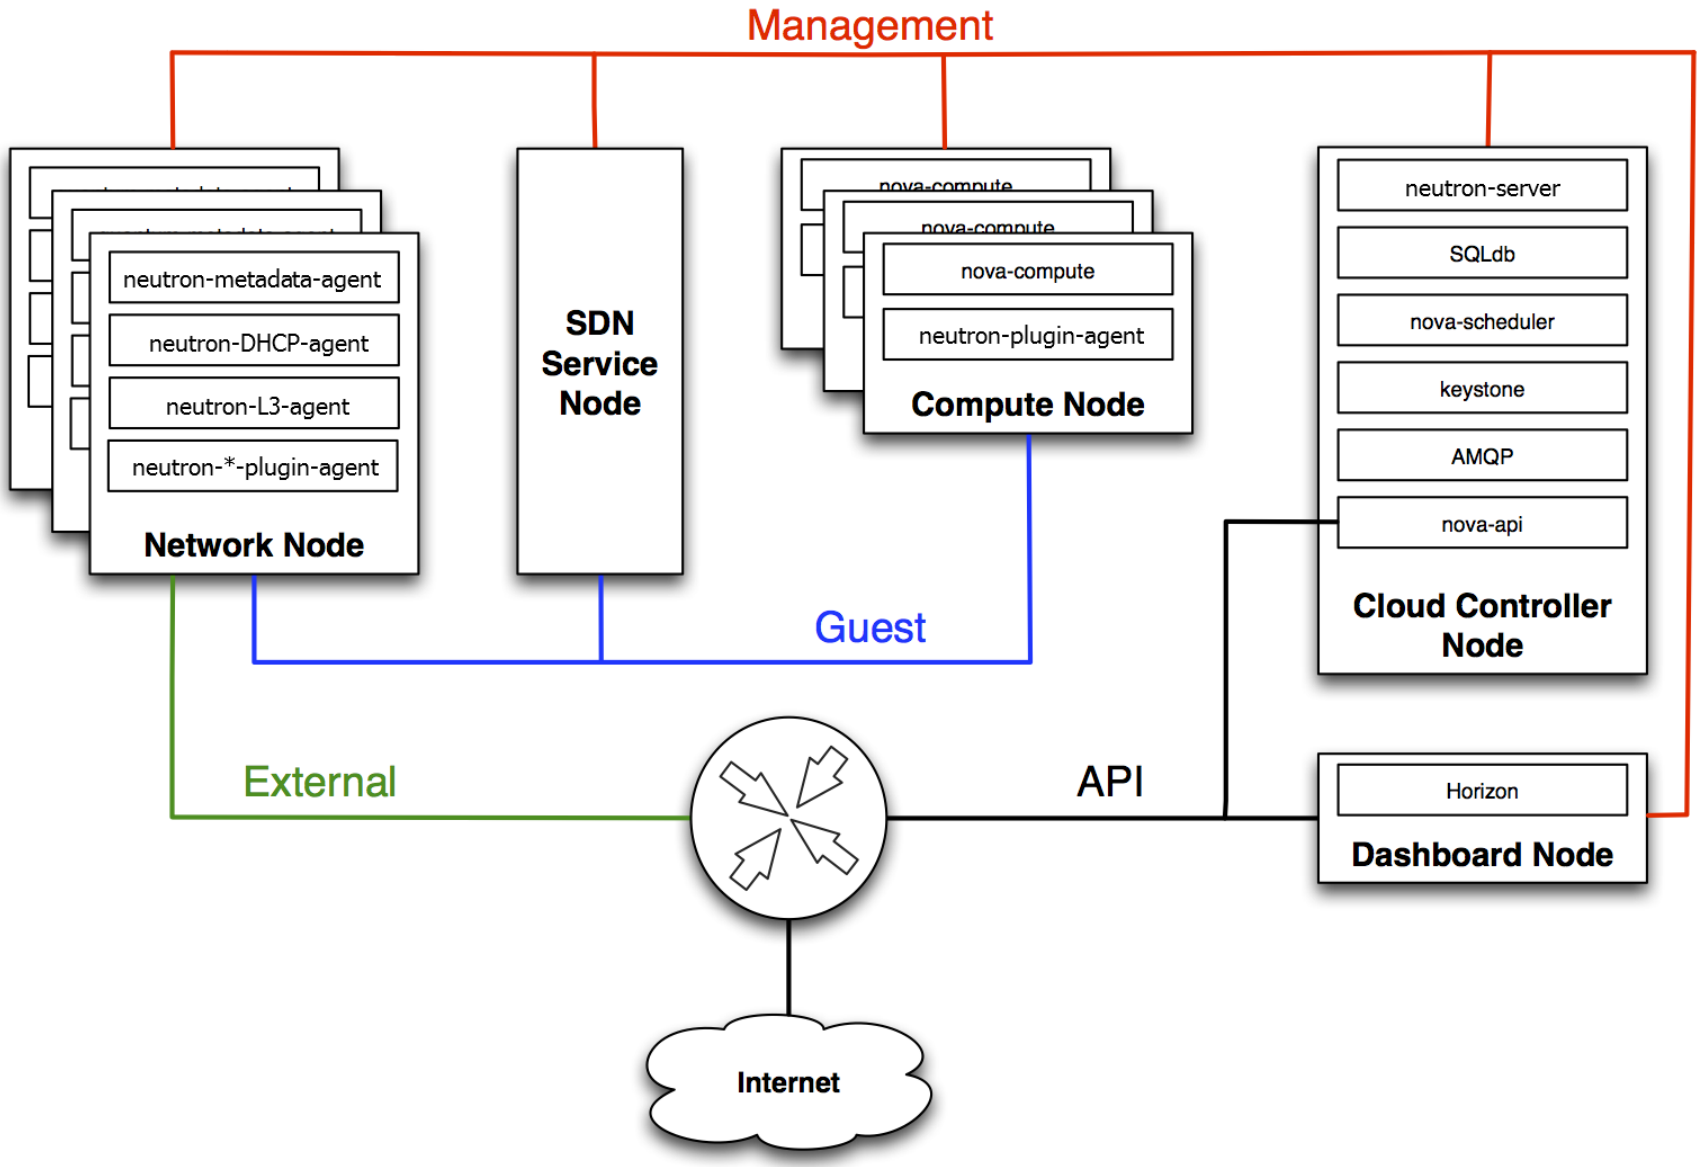
Management
neutron-metadata-agent
neutron-DHCP-agent
neutron-L3-agent
neutron-*-plugin-agent
Network Node
SDN Service Node
nova-compute
nova-compute
nova-compute
neutron-plugin-agent
Compute Node
neutron-server
SQLdb
nova-scheduler
keystone
AMQP
nova-api
Cloud Controller Node
Guest
External
API
Horizon
Dashboard Node
Internet

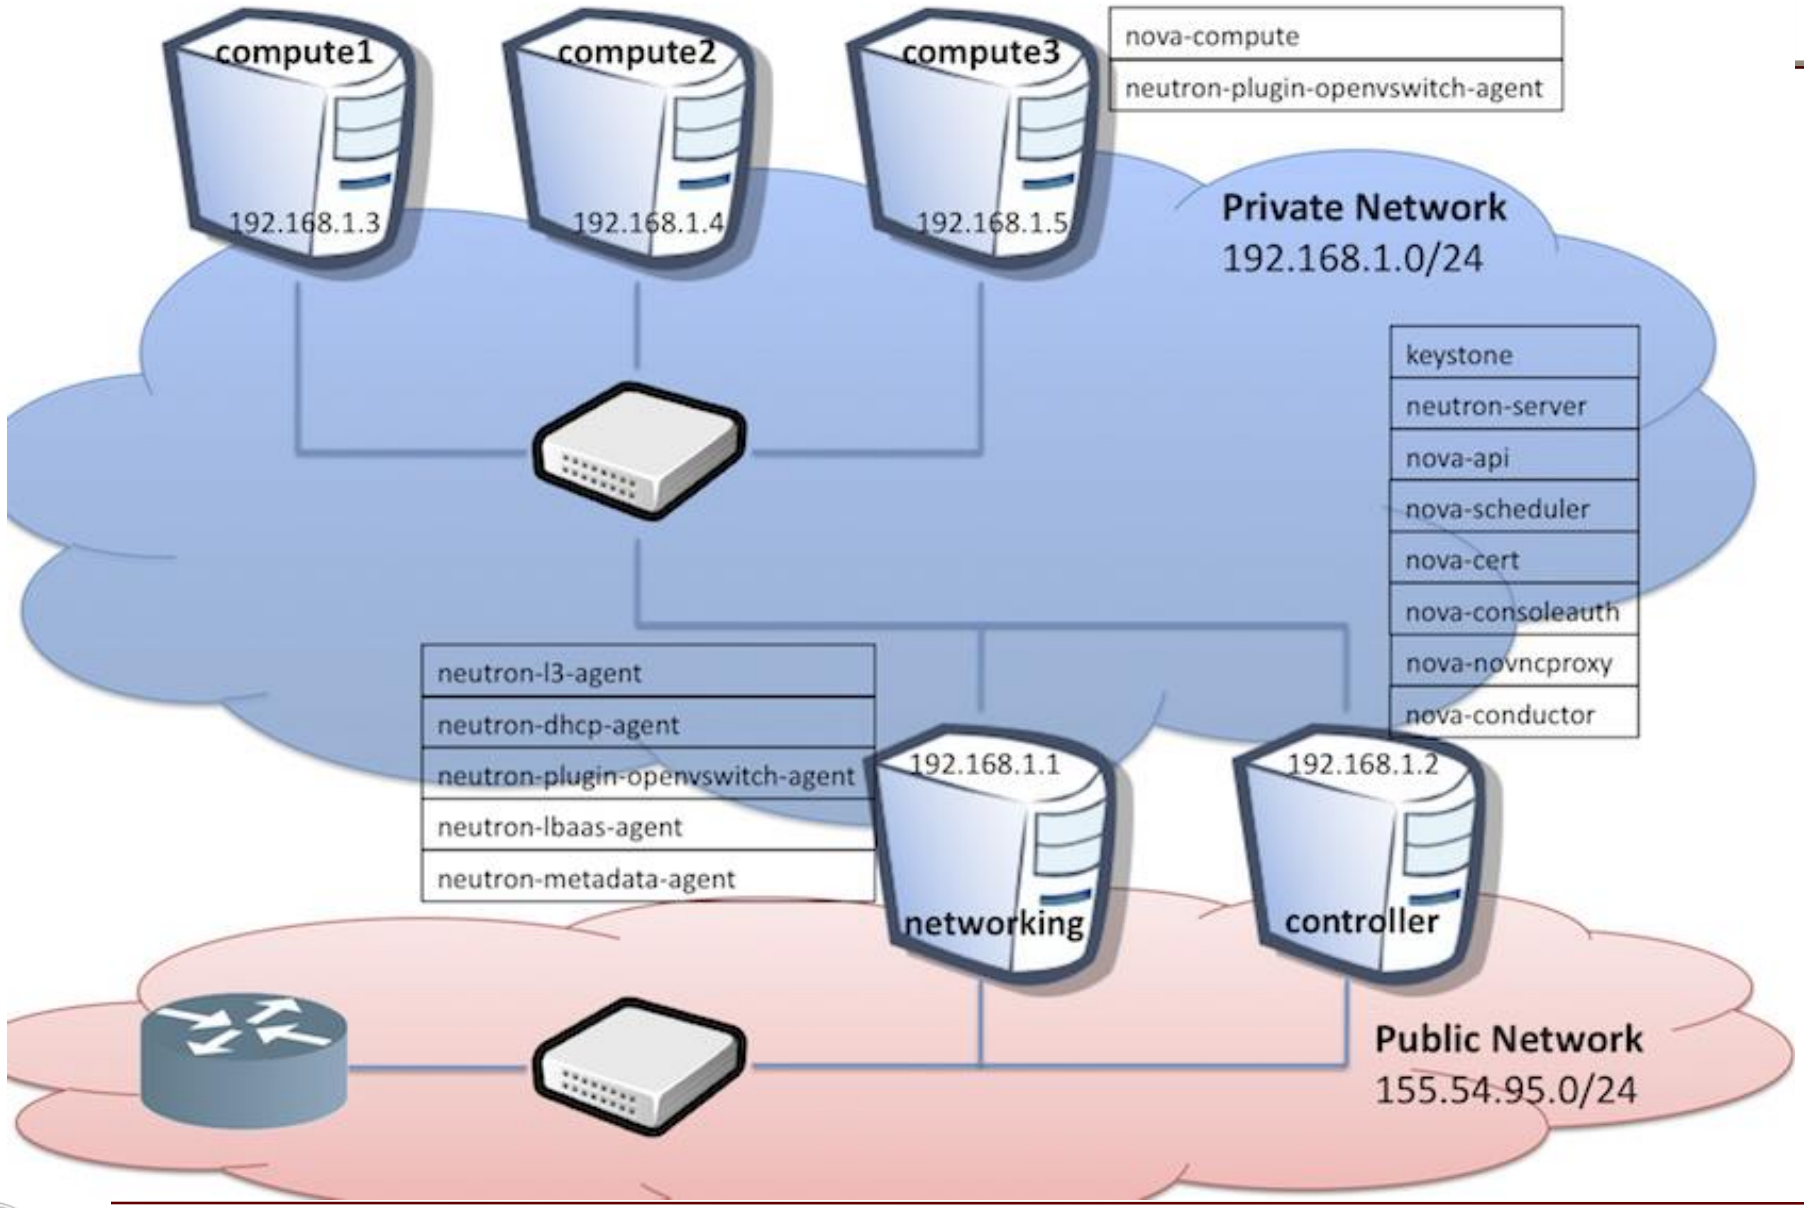
compute1
compute2
compute3
192.168.1.3
192.168.1.4
192.168.1.5
nova-compute
neutron-plugin-openvswitch-agent
Private Network
192.168.1.0/24
keystone
neutron-server
nova-api
nova-scheduler
nova-cert
nova-consoleauth
nova-novncproxy
nova-conductor
neutron-l3-agent
neutron-dhcp-agent
neutron-plugin-openvswitch-agent
neutron-lbaas-agent
neutron-metadata-agent
192.168.1.1
networking
192.168.1.2
controller
Public Network
155.54.95.0/24