



Cloud Networking (VITMMA02)

OpenStack

Markosz Maliosz PhD

Department of Telecommunications and Media Informatics
Faculty of Electrical Engineering and Informatics
Budapest University of Technology and Economics

Spring 2020



OpenStack

- » Free and open source IaaS cloud platform
 - » manages processing, storage, network resources
 - » on commodity hardware
 - » flexible configuration options
- » Based on collection of open source software
- » started as a joint project of Rackspace and NASA in 2010
 - » AT&T, Ericsson, Huawei, Intel, IBM, HP, RedHat, Cisco, Dell, etc.
 - » <http://www.openstack.org/foundation/companies/>
- » written in python
- » well documented
- » modular architecture
- » ApacheLicense 2.0



Can

Commercial Use
 Modify
 Distribute
 Sublicense
 Private Use
 Use Patent Claims
 Place Warranty

Cannot

Hold Liabile
 Use Trademark

Must

Include Copyright
 Include License
 State Changes
 Include Notice

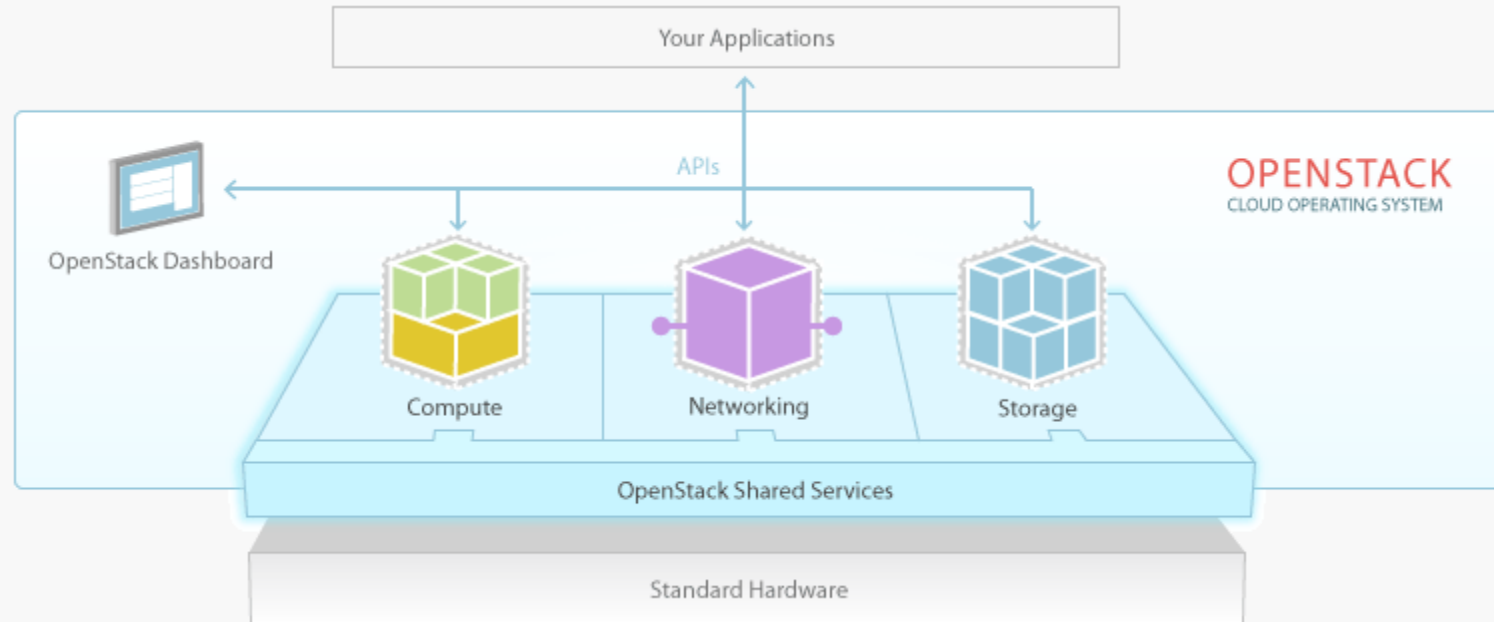


Releases

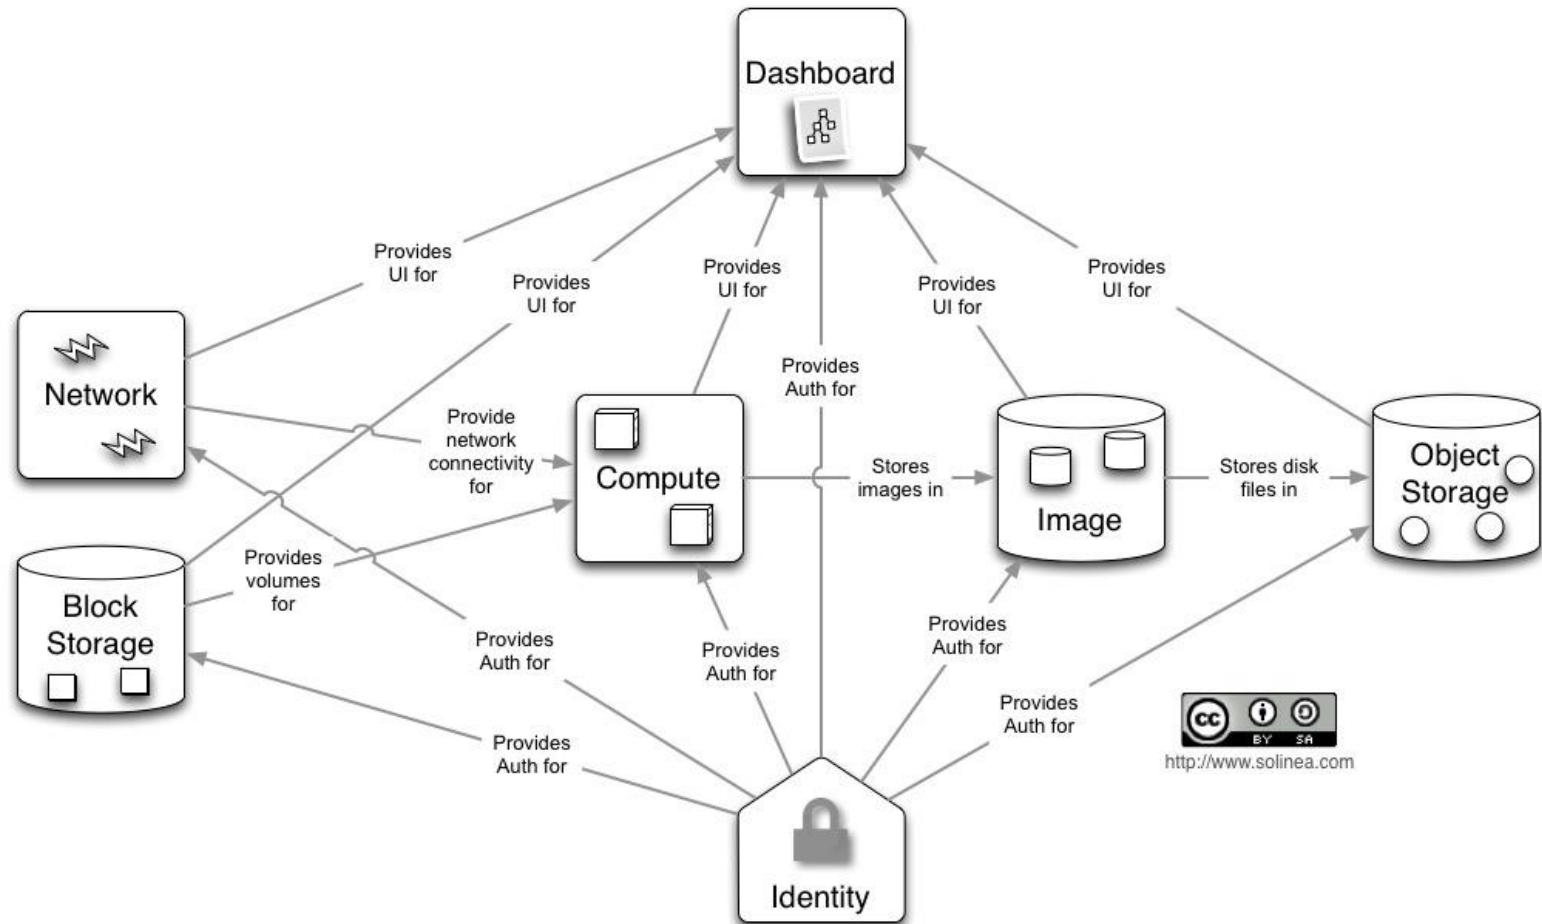
Release Name	Release Date	Included Components
Austin	21 October 2010	Nova, Swift
Bexar	3 February 2011	Nova, Glance, Swift
Cactus	15 April 2011	Nova, Glance, Swift
Diablo	22 September 2011	Nova, Glance, Swift
Essex	5 April 2012	Nova, Glance, Swift, Horizon, Keystone
Folsom	27 September 2012	Nova, Glance, Swift, Horizon, Keystone, Quantum, Cinder
Grizzly	4 April 2013	Nova, Glance, Swift, Horizon, Keystone, Quantum, Cinder
Havana	17 October 2013	Quantum-> Neutron, +: Ceilometer, Heat
Icehouse	17 April 2014	+: Trove
Juno	October 2014	+: (DBaaS), Sahara (data processing)
Kilo	April 2015	+: Sahara, Ironic (bare metal)
Liberty	October 2015	+: Searchlight, Designate (DNS), Zaqr (messaging), Barbican (key manager), Manila (shared file system)
Mitaka	April 2016	+: cloudkitty (billing and charging), freezer (backup and recovery), magnum (container orchestration), monasca (monitoring), senlin (clustering), solum (app. lifecycle framework), tacker (NFV)
Newton	October 2016	+: panko (telemetry), virtage (Root Cause Analysis), watcher (resource optimization)
Ocata	February 2017	focusing on resolving scalability and performance issues , Congress (Governance Service), Cells (partition compute nodes into smaller groups)
Pike	October 2017	Zun (container mgmnt.), Kolla (lifecycle mgmnt.)
Queens	February 2018	high availability features, GPU support, multi-attach storage, edge computing, etc.
Rocky	August 2018	focus on bare metal clouds, fast forward upgrades, and hardware accelerators
Stein	April 2019	enhanced bare metal and network management, and containers functionality
Train	October 2019	Extends Security and Data Protection, Adds New AI and Machine Learning Support



Architectural Overview



Looking into a little bit closer

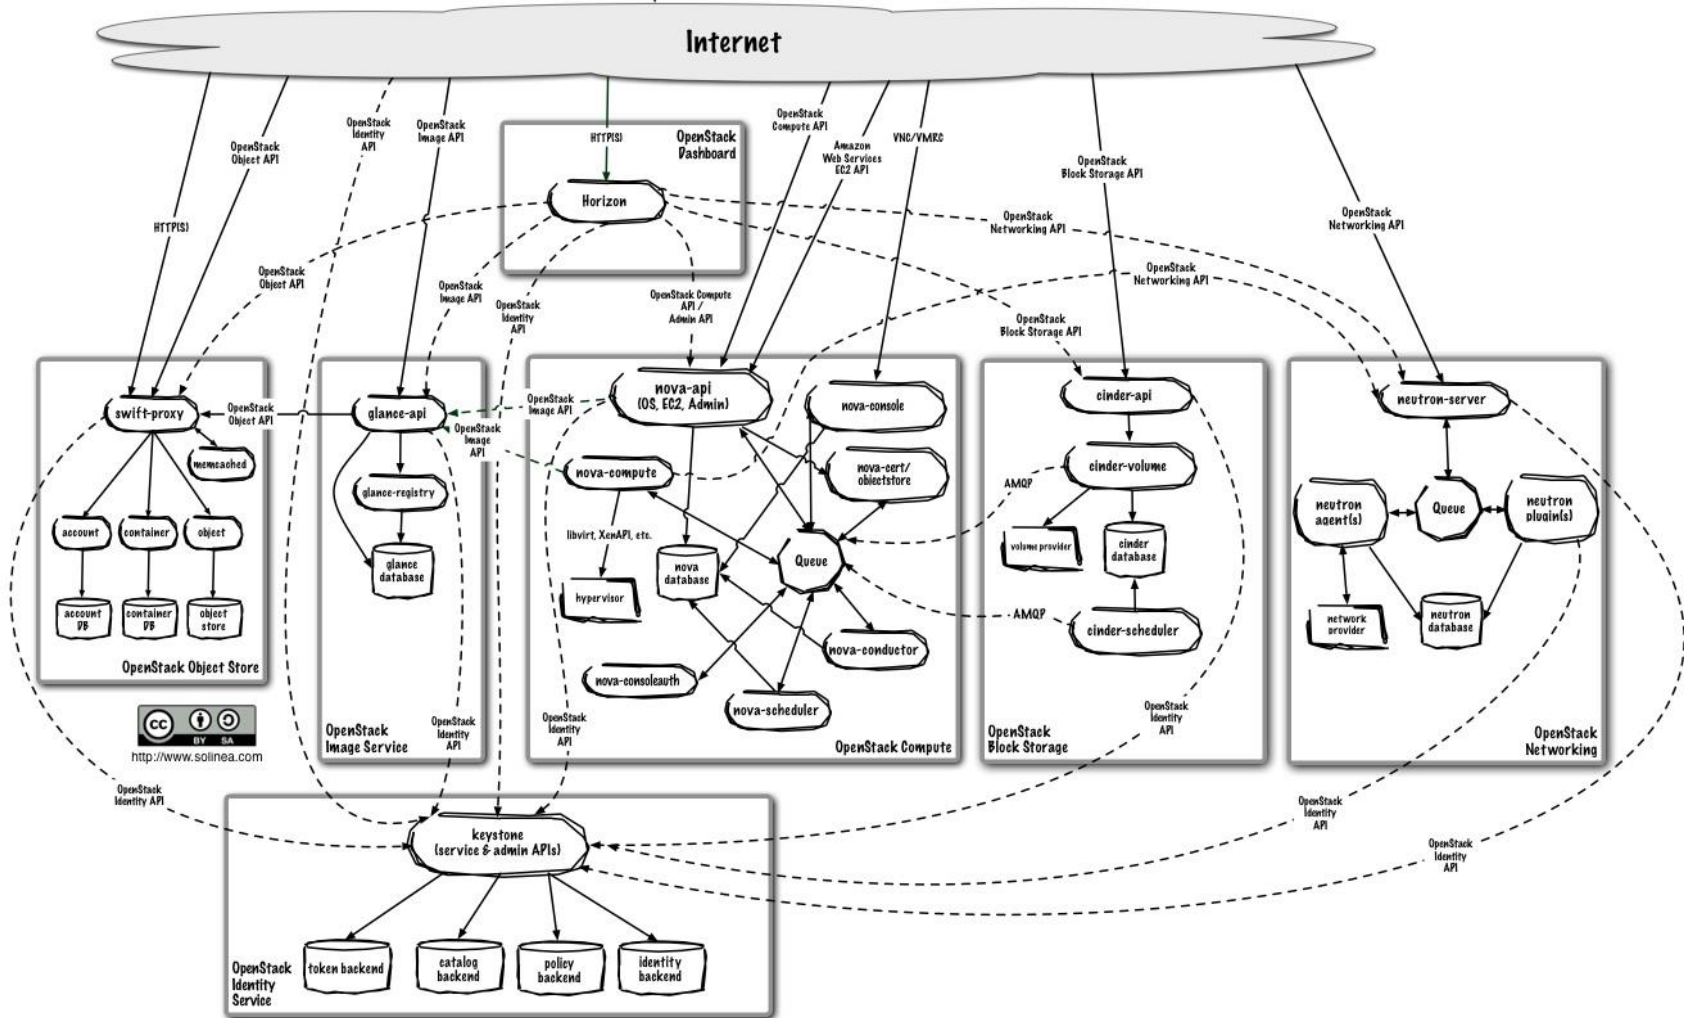




OpenStack is not that simple



- Command-line interfaces (nova, neutron, swift, and so on)
- Cloud Management Tools (Rightscale, Enstratus, and so on.)
- GUI tools (Dashboard, Cyberduck, iPhone client, and so on.)





OpenStack components

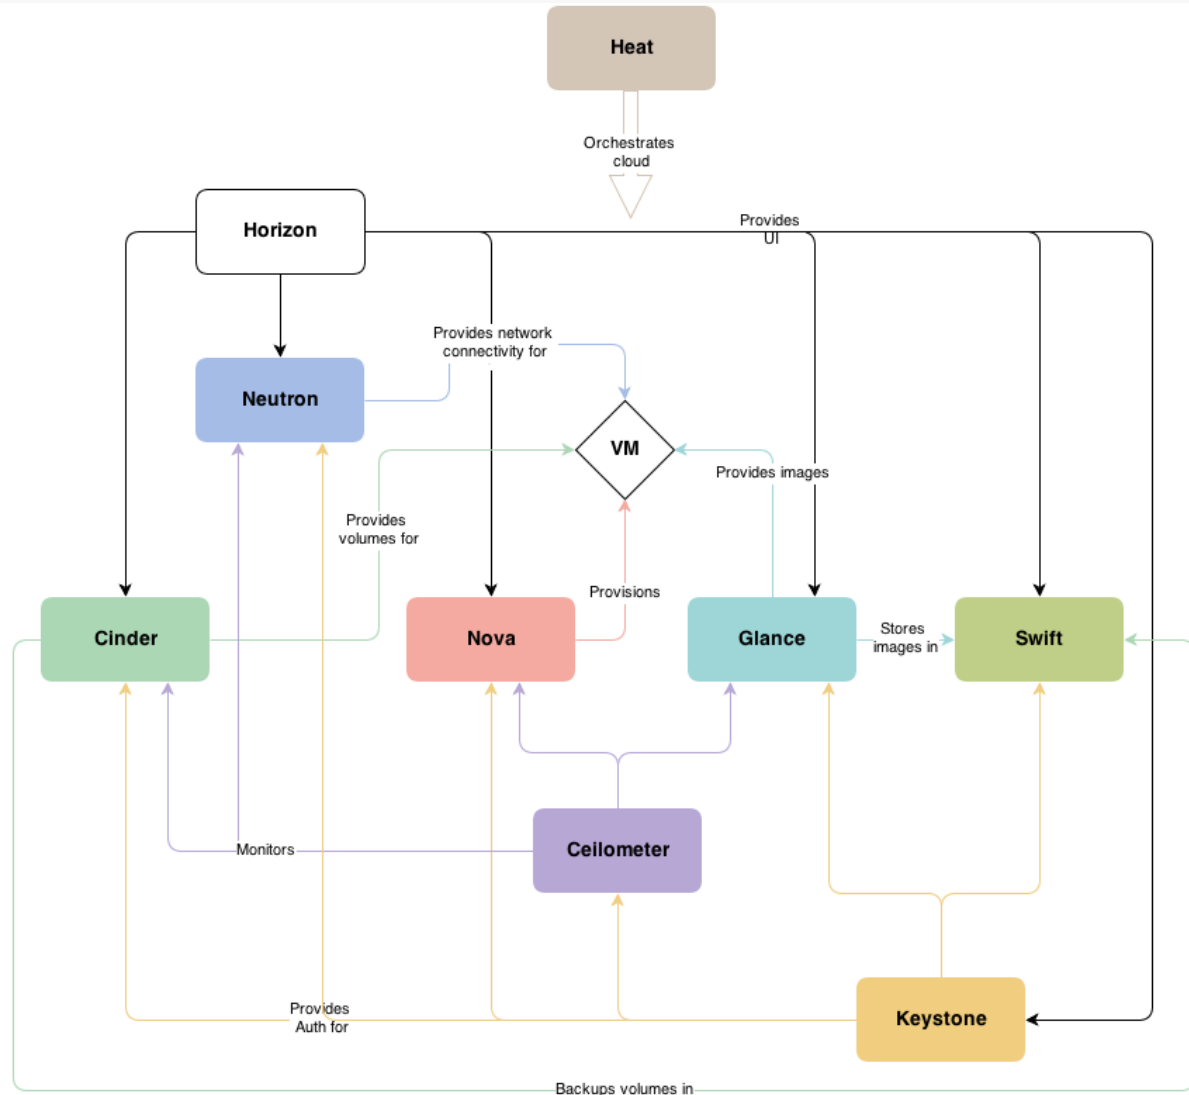
Stats about the project:

<https://www.stackalytics.com/>

- » Dashboard ("Horizon"): web interface
- » Compute ("Nova"): running VMs, reading VM images, storing VM images with the help of Image service ("Glance")
- » Network ("Neutron"): provides virtual networking for a Compute nodes
- » Block Storage ("Cinder"): virtualizes the management of block storage for Compute nodes
- » Object Storage ("Swift"): store and retrieve data objects
- » Image ("Glance"): VM image management, storage with e.g. Object Storage ("Swift")
- » Identity ("Keystone"): central authentication
- » Telemetry (Ceilometer): usage monitoring
- » Orchestration (Heat): automated VM management
- » Database as a Service (Trove)
- » etc.



Interactions among components



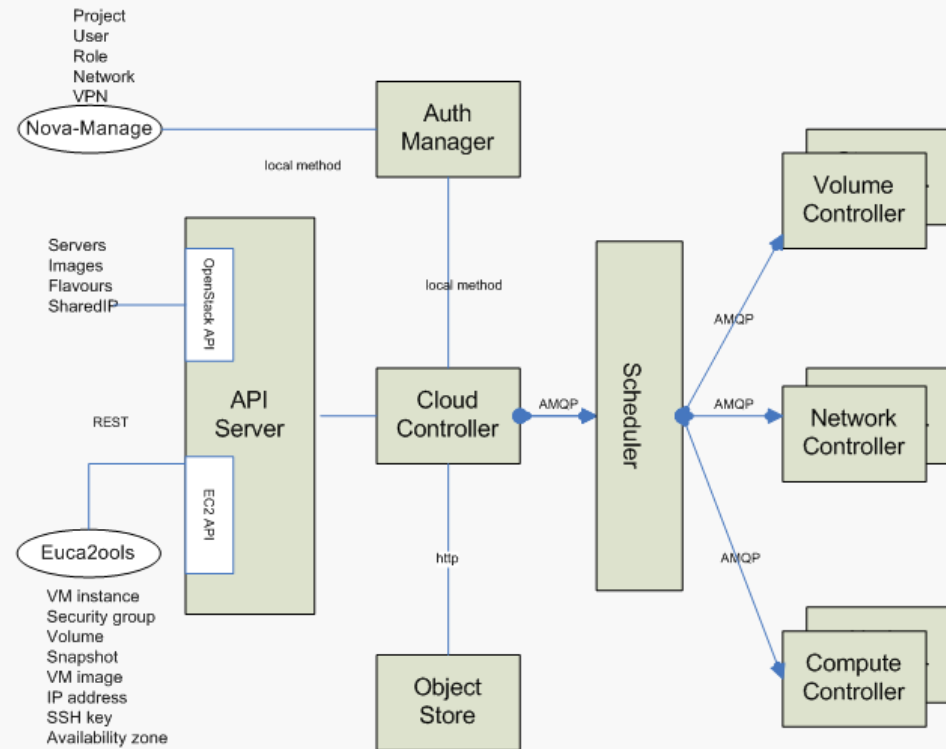



General modules

- » Message Queue
 - » interaction and information exchange between services
- » Storage for metadata, configuration data, etc.
 - » databases
- » Scheduler
 - » serving a new VM request



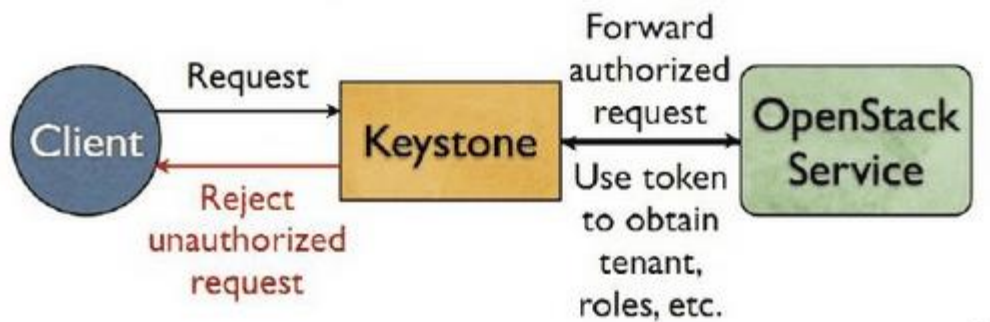
Message Queue



Identity Service: Keystone

» Main Services

- » Identity: auth credential validation and data about *users* and *groups*
- » Resource: provides data about *projects* and *domains*
- » Token: validates and manages tokens used for authenticating requests once a user's credentials have already been verified
- » Service catalog: list of registered services
- » Policy: rule-based authorization engine

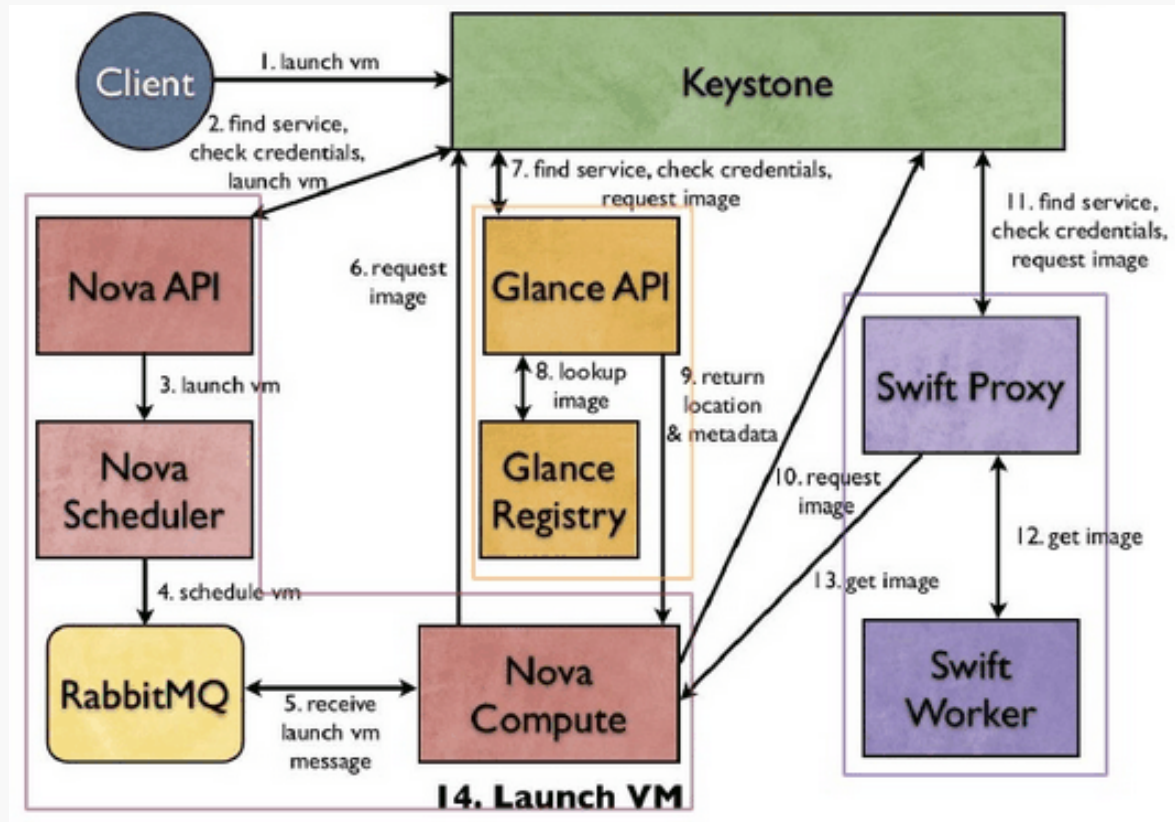




Compute (Nova)

- » Processes
 - » nova-api: public interface
 - » nova-compute: executing VM instances and managing their lifecycle
 - » wide hypervisor support
 - » KVM, Xen, XenServer, Hyper-V, etc.
 - » nova-volume: managing permanent storage
 - » nova-network: networking for VMs
 - » nova-schedule: schedule the VM to compute nodes
- » Horizontal scaling
 - » commodity hardware without special requirements

Starting a VM





Swift (Object storage)

- » similar to Amazon S3 (Simple Storage Service)
- » scalable, redundant, highly available
- » ideal for storing unstructured data that can grow without bound
- » replication on multiple hard drives
- » storage software for
 - » any binary object (data)
 - » e.g. VM image, backup, files, etc.
 - » can have user *metadata* associated with them
- » an object is handled as a unit
 - » ideal for data that is mostly read



Cinder (Block level storage)

- » persistent storage
- » typically for file systems: partition, volume
- » accessible via API
 - » create, delete, attach
 - » resize, snapshot
- » multiple backend implementation: local server, Ceph, GlusterFS, external storage systems from third-party vendors, etc.
- » simpler than Swift, but replication is hard to achieve with multiple vendor backend
- » ideal for
 - » VM file system
 - » database with frequent write



Glance (Image service)

- » storage, catalogue and retrieval for disk and container images
 - » VM/container templates and associated metadata
- » formats: raw, QCOW, VMDK, VHD, ISO, OVF, etc.
- » backend service
 - » file system
 - » Swift
 - » Amazon S3



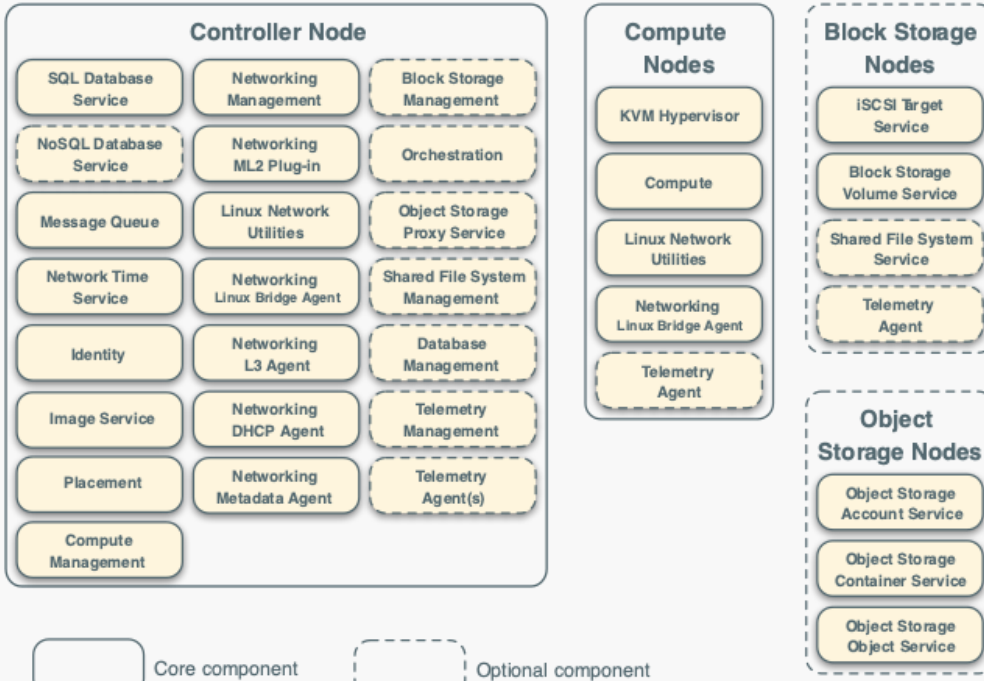
Network architecture

- » „Networking in OpenStack is a complex, multifaceted challenge.“ /OpenStack Operations Guide/
- » Network as a Service
- » Functions
 - » Managing network objects
 - » network
 - » subnetwork
 - » network port
 - » IP address management
 - » static, DHCP
 - » floating IP
 - » virtual networks
 - » flat, VLAN
 - » self-service operation
- » Neutron components
 - » neutron-server
 - » database: storing configuration data and states
 - » plug-in agents: interface towards the native Linux network mechanisms, external devices or SDN controllers



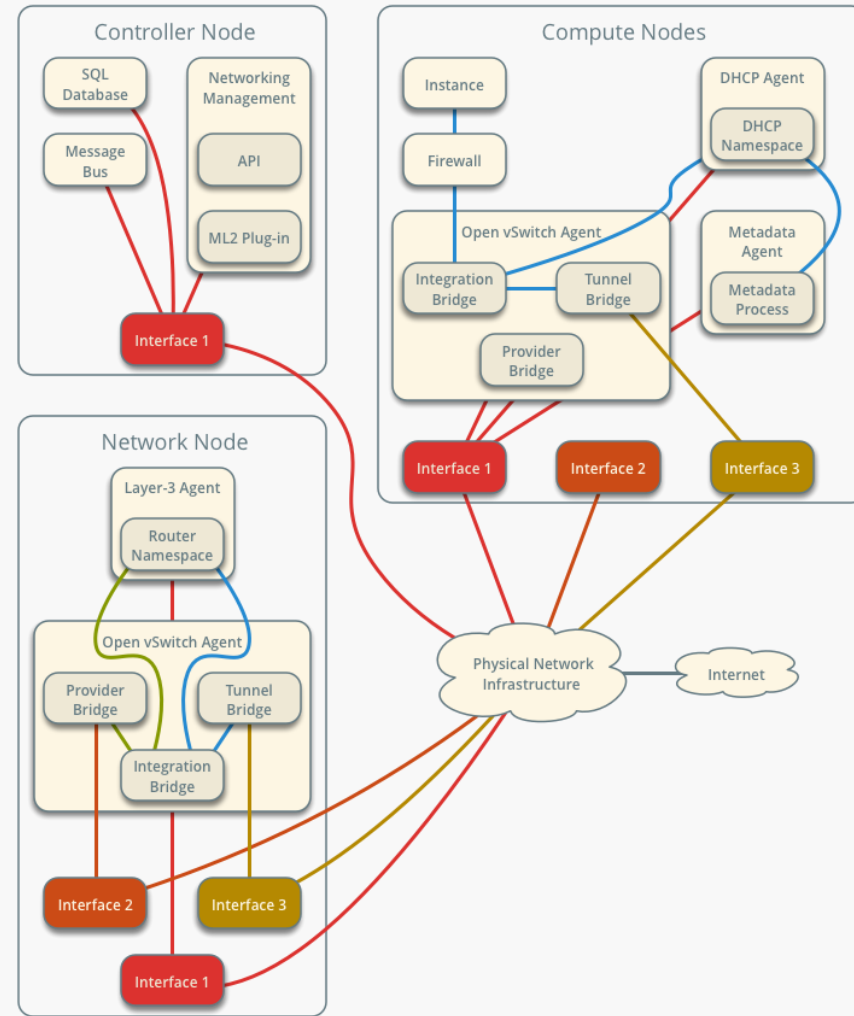
Neutron network

Service Layout



Open vSwitch - Self-service Networks

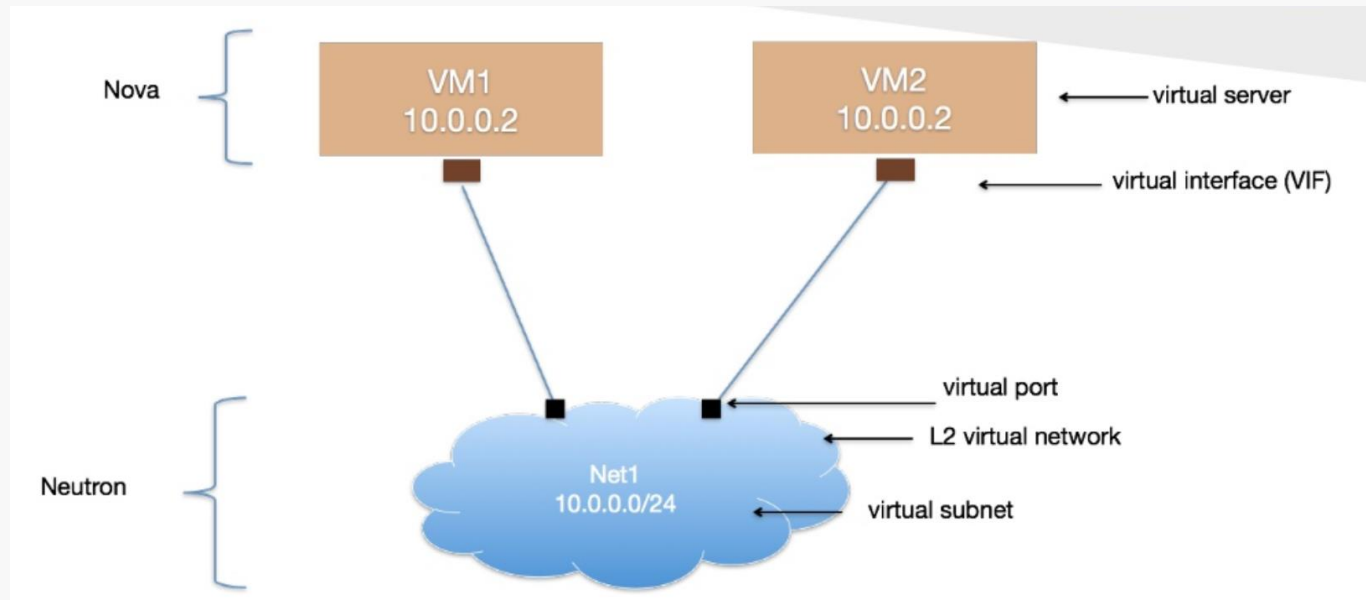
Overview



- Management network 10.0.0.0/24
- Provider network Aggregate
- Self-service network
- Overlay network 10.0.1.0/24
- Provider network

Networks terminology

- » Internal or management network
 - » connects physical nodes
 - » for communication between internal components of OpenStack
- » External or public network
 - » controller external IP address
 - » public IP addresses for VMs (floating IP)
 - » assigned dynamically to instances
- » Elements of a virtual network





Neutron Provider Networks

- » Layer-2 connectivity for VMs
- » optionally: DHCP and metadata service
- » Mapping to an existing physical Layer-2 network inside the datacenter (Flat or separation with VLANs – for a small number of tenants)
- » Simple, but not flexible
 - » Only admins can configure, because the physical network is affected by the configuration
 - » Only Layer-2, thus no support for routing or floating IP
 - » Layer-3 functions provided by the physical network



Neutron Self-service Networks

- » Virtual networks
- » Self-service operation, without administrator
- » Connectivity to the provider or external (e.g. Internet) network via a virtual router
- » Isolation
 - » Flat: no isolation
 - » VLAN: mapped to physical VLANs
 - » overlay tunneling (VXLAN, GRE)



Neutron services

- » Virtual Router / Distributed Virtual Router (DVR)
- » Load Balancing as a Service
- » VNP as a Service
- » Firewall as a Service



OpenStack installation alternatives

- » OpenStack install guide
 - » step-by-step: installing and configuring Linux packages
- » Deployment tools
 - » General tools
 - » Chef, Puppet, Juju, Ansible
 - » OpenStack specific automation tools / scripted
 - » deploying, testing and maintaining
 - » hardware discovery
 - » native Openstack: <https://www.openstack.org/software/project-navigator/deployment-tools>
 - » OpenStack-Ansible
 - » Kolla
 - » TripleO - OpenStack on OpenStack
 - » ...
 - » provisioning server (with GUI)
 - » Ubuntu: MaaS + conjure-up
 - » console: RedHat packstack – ssh
- » Developer / tester version
 - » DevStack
 - » setting up a configuration file



OpenStack vs. DevStack

- » OpenStack
 - » components run as Linux daemons
- » DevStack
 - » for development and testing
 - » minimal configuration
 - » startup/shutdown by script
- » Deployment options
 - » all-in-one physical server / VM
 - » multi-node physical servers / VMs



Sketch of Practice 2 and 3

- » DevStack Multi-Node Lab
 - » <https://docs.openstack.org/devstack/latest/guides/multi-node-lab.html>
- » 2 VirtualBox virtual servers: DevStack nodes
 - » controller + compute
 - » compute
- » Network model
 - » Self-service neutron network
- » Practice
 - » starting VMs on DevStack nodes
 - » investigate network architecture



Sources

- » <http://www.openstack.org>
- » <http://docs.openstack.org>