

# BUILDING A CLOUD INFRASTRUCTURE WITH OPEN SOURCE SOFTWARE

Nimal Ratnayake  
Senior Lecturer/Electrical & Electronic Eng.  
University of Peradeniya

# Building a cloud infrastructure

- Hardware
  - Low-cost, commodity hardware
  - Redundent resources to ensure high availability
  - Approach: Use low-cost hardware with software providing redundancy and avoid high-end hardware
- Software
  - Role of software is greater than that of hardware
  - A collection or “stack” of software is needed
  - Proprietary software stacks with vendor lock-in
  - Several open source software stacks available

# Software stack

- Because of the complexity, software required to run a cloud is organized as layers of a “stack”
  - Hypervisor: enables the creation of several virtual machines (Vms) on a single physical machine (node)
  - Clustering: use of multiple nodes / storage devices to serve as “clusters” of computing power (compute clusters) and storage capacity (storage clusters) with redundancy built in
  - Virtual Networks: need to create virtual networks that connect VMs to each other and then to the outside world
  - Resource Management: management of compute clusters, storage clusters, networks, software images etc.
  - User Management: User privileges, authentication etc.
  - Billing: Based on resource usage

# Open Source Cloud software stacks

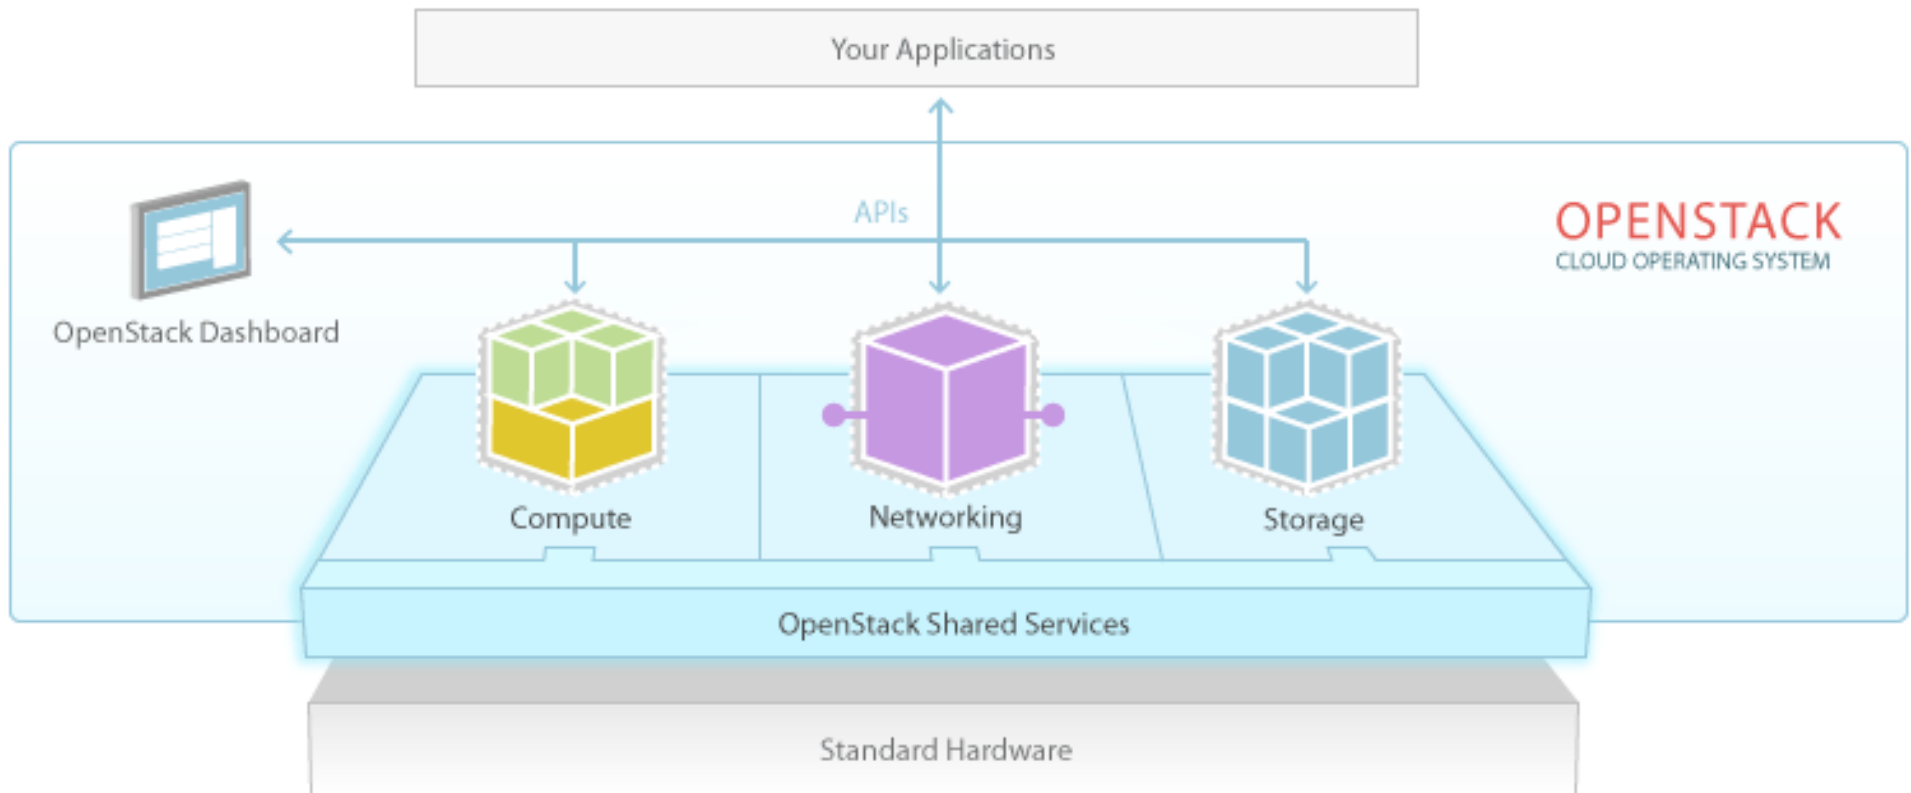
- OpenStack
  - Initiated by Rackspace / NASA
  - Governed by consortium of vendors
- Ganeti + Synnefo
  - Ganeti released to public by Google
  - Synnefo developed by GRNET runs on top of Ganeti
- OpenNebula
  - Focus on the requirements of “users”, not “providers”
  - VoneCloud - Open replacement for VMWare's vCloud
- Others

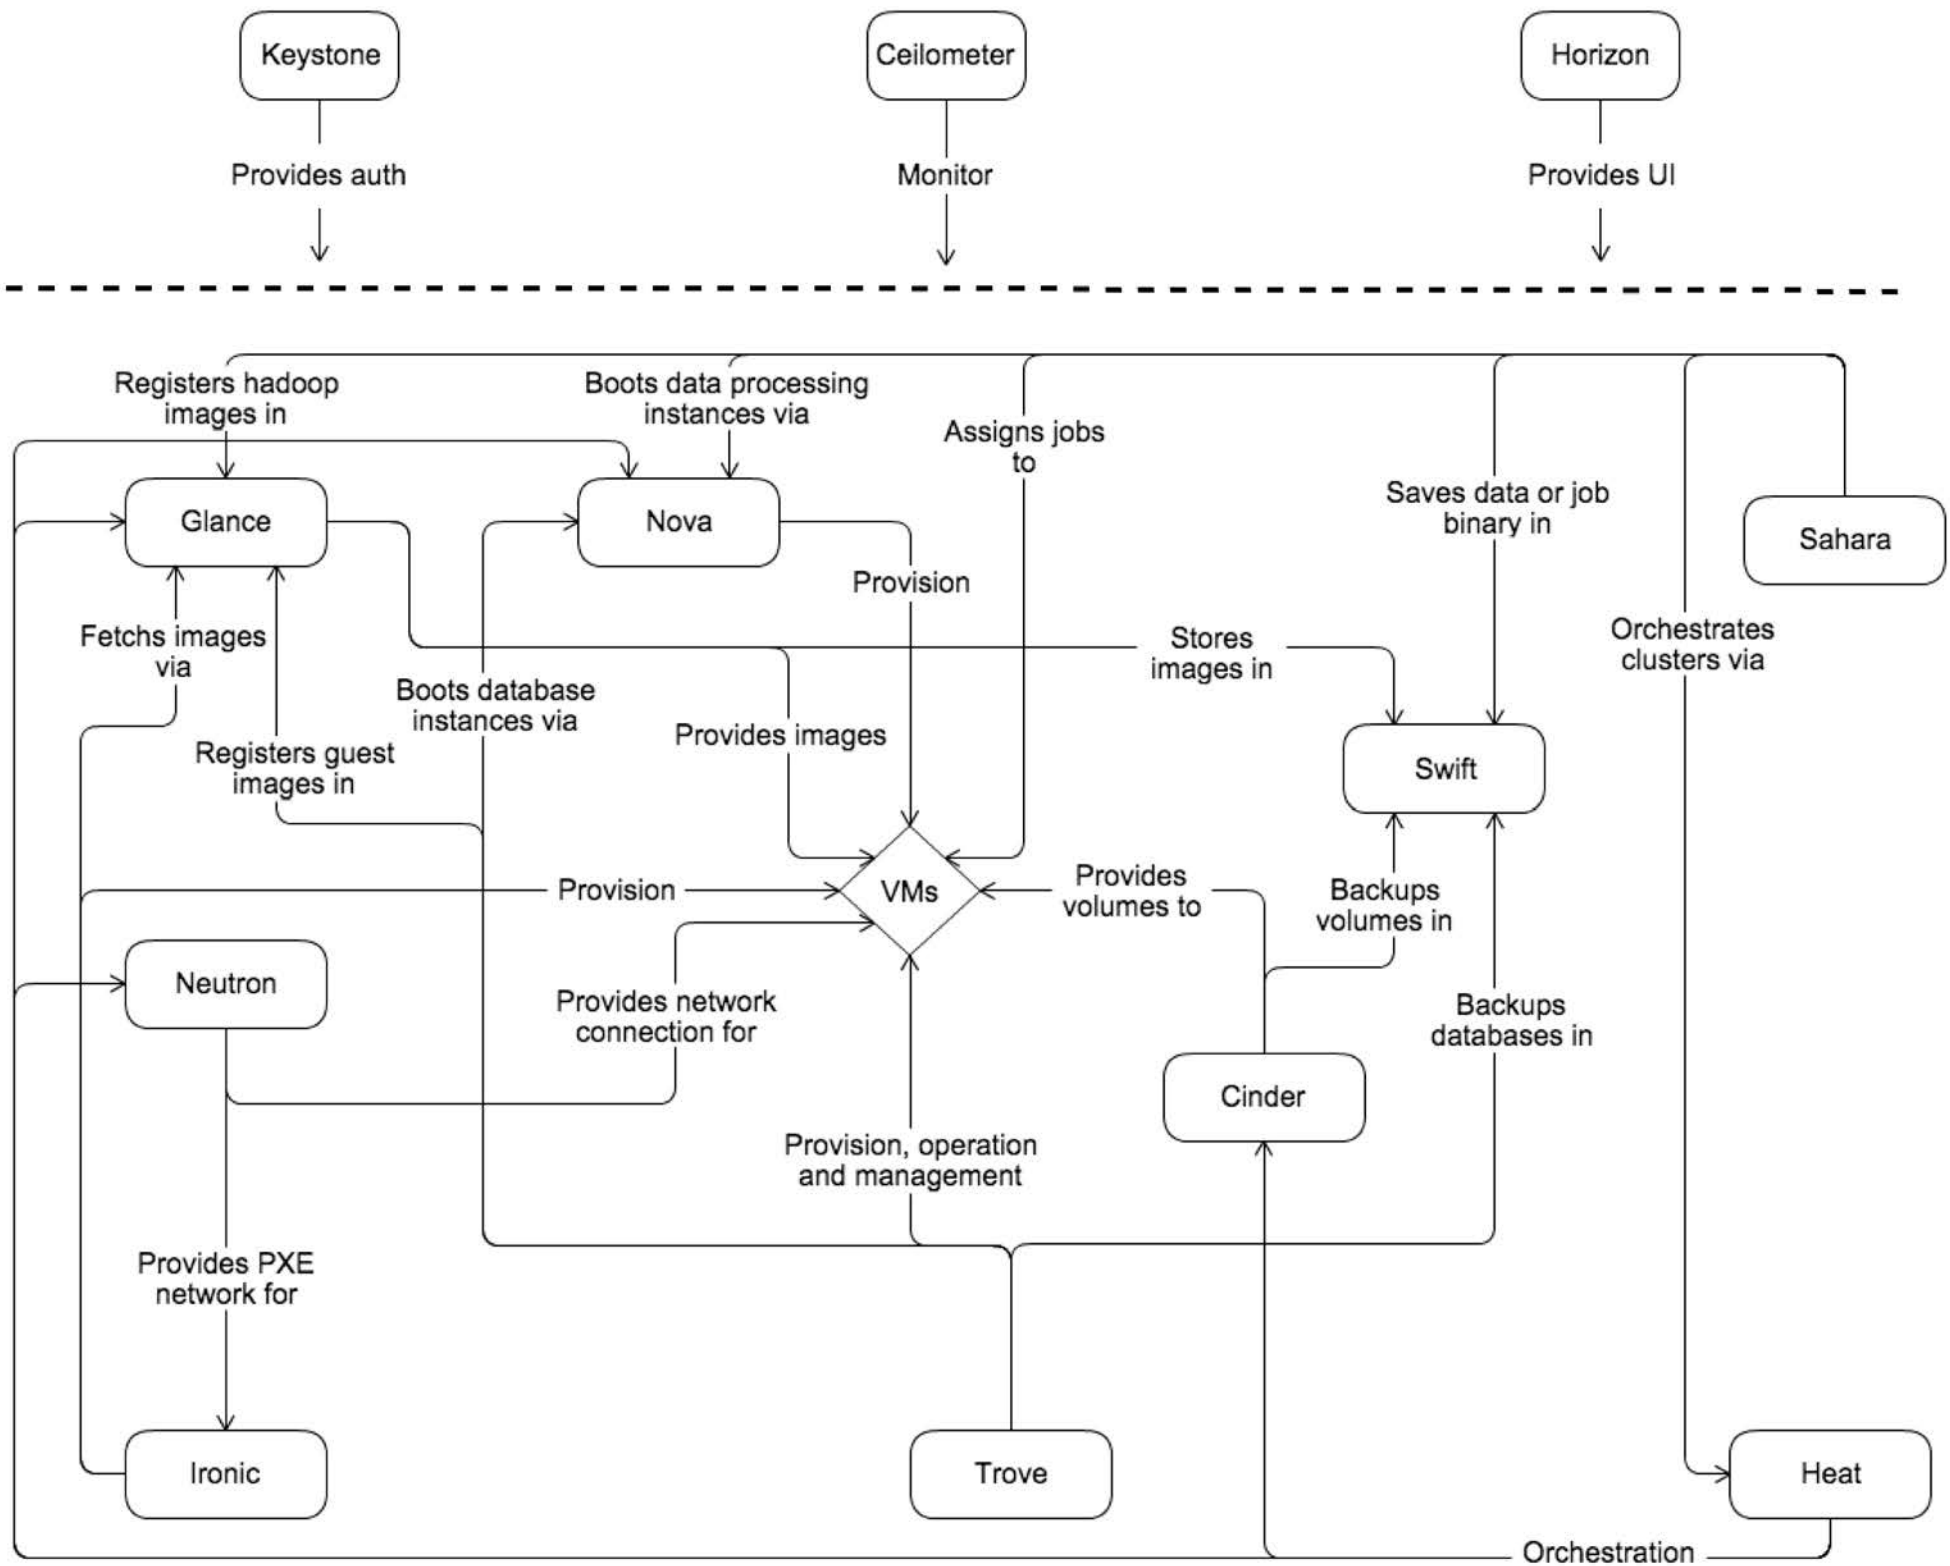
# OpenStack



- Development since 2010
- Managed by OpenStack Foundation
  - 500+ members including most cloud giants
- Supports many hypervisors
  - KVM/QEMU, Xen, VMWare vSphere, HyperV
- Current version is “Kilo”
  - Consists of 16 components
  - Installation from scratch is not easy
- Vendors sell “Distributions”
  - Usually not free, except for small scale use

# Openstack Architecture



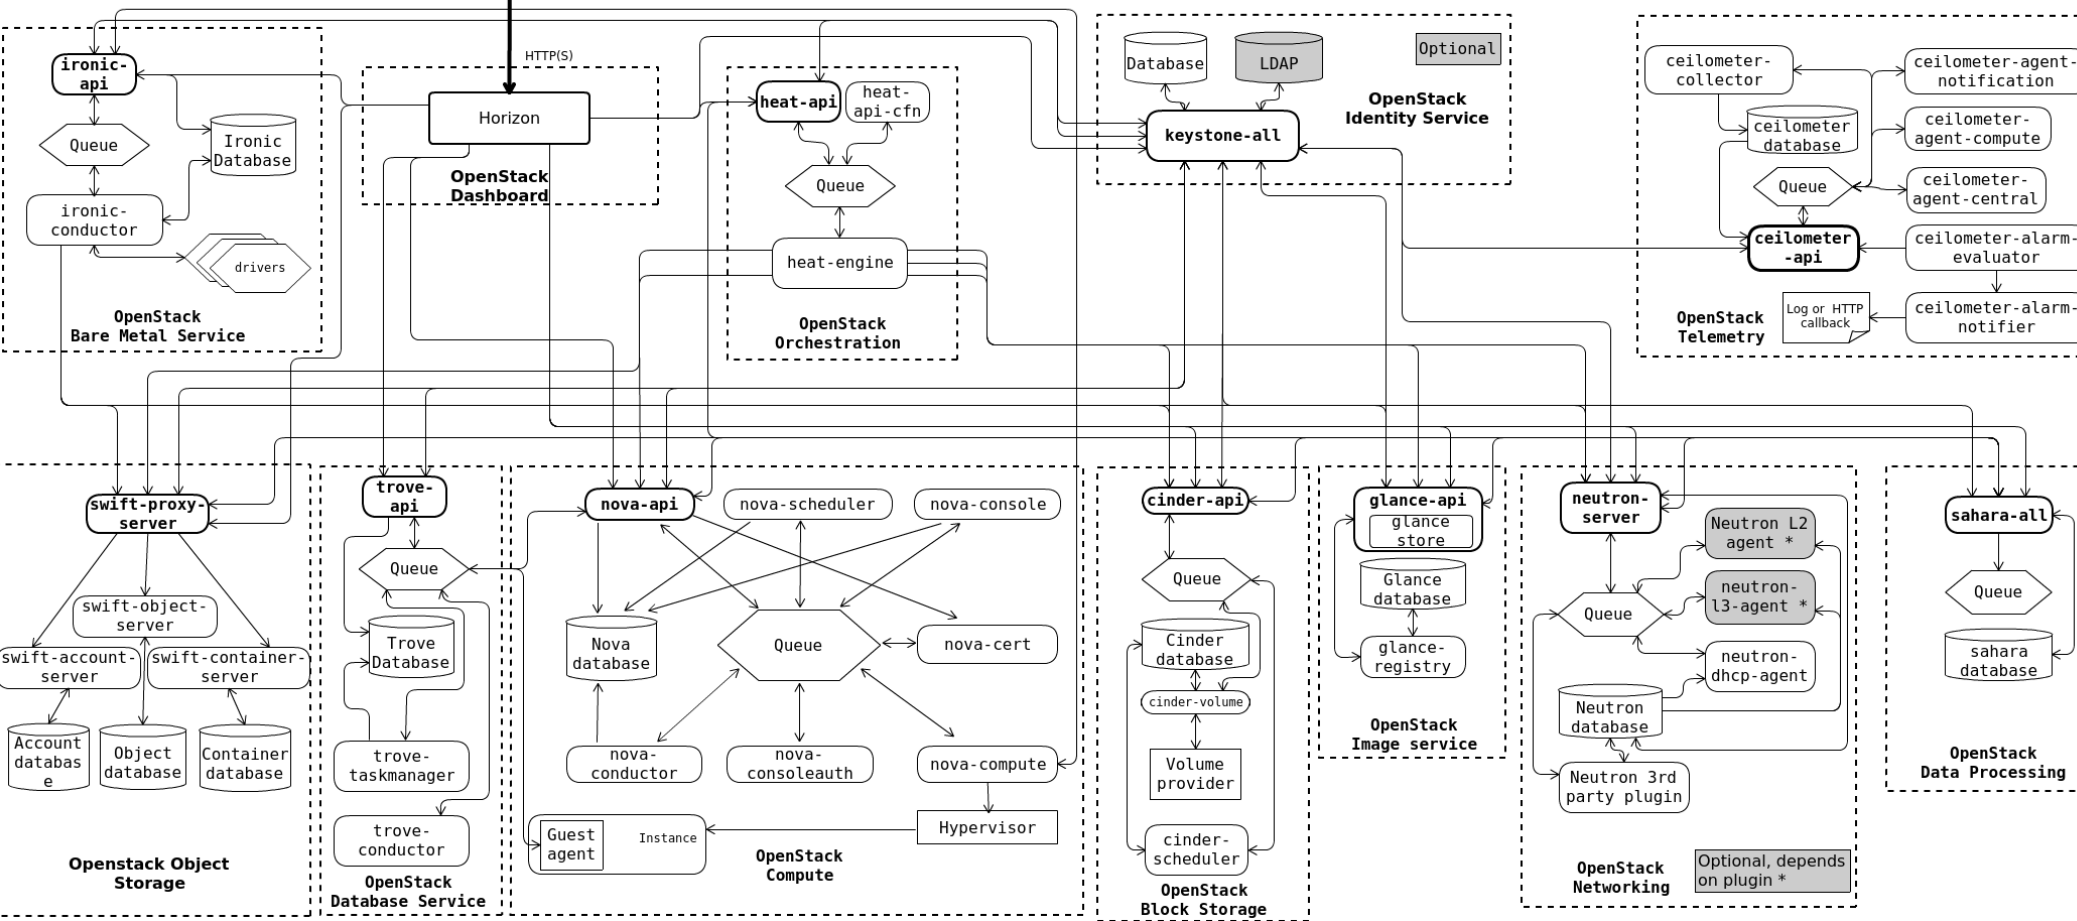


Internet



- CLI clients(nova, cinder, neutron and so on)
- Cloud management tools
- GUI tools

HTTP(S)



Optional

OpenStack Identity Service

OpenStack Telemetry

Log or HTTP callback

Openstack Object Storage

OpenStack Database Service

OpenStack Compute

OpenStack Block Storage

OpenStack Image service

OpenStack Networking

Optional, depends on plugin \*

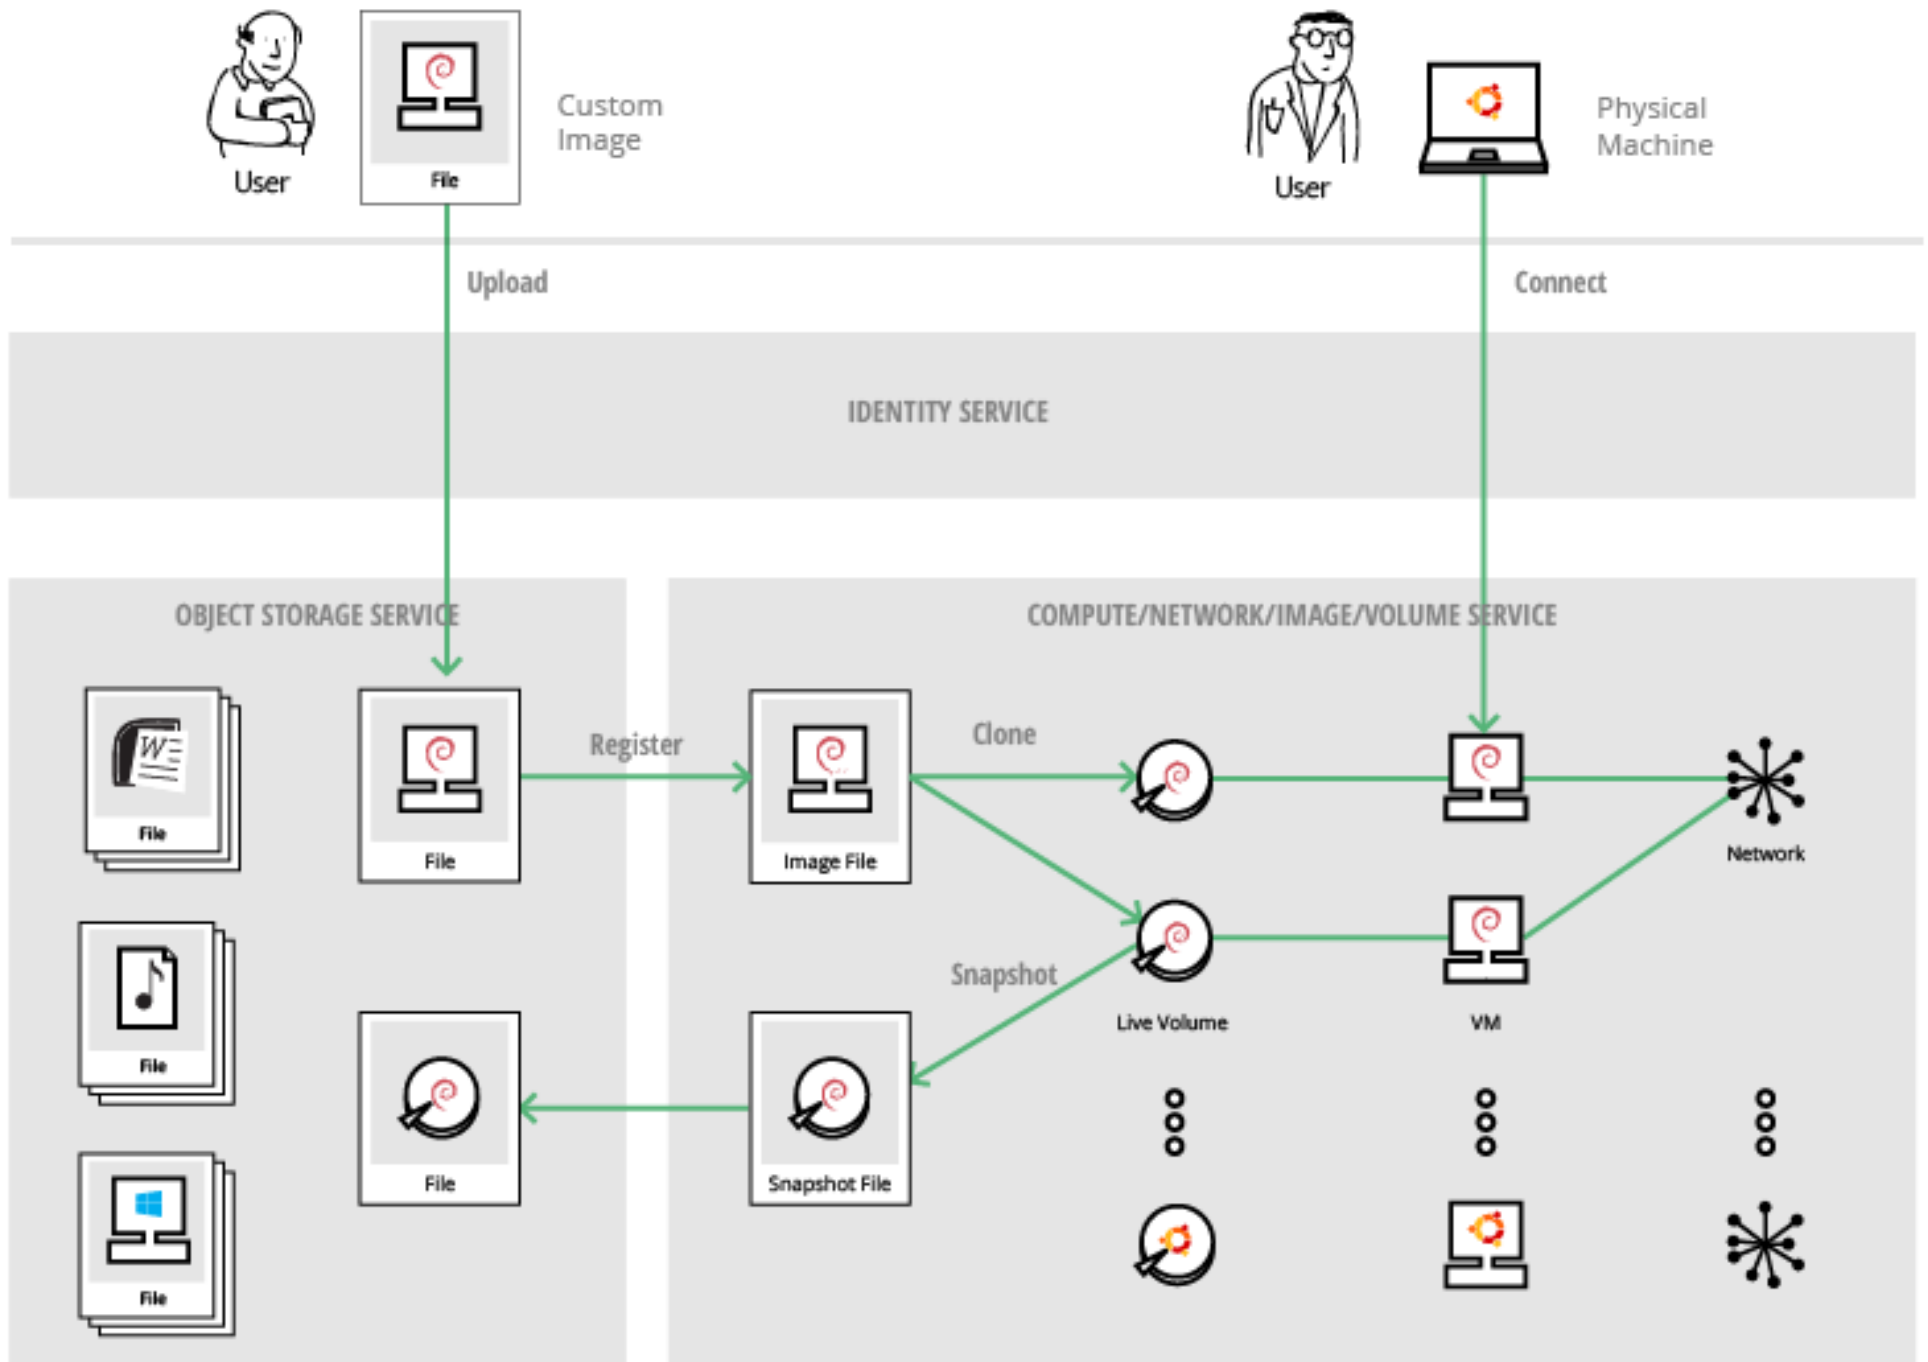
OpenStack Data Processing

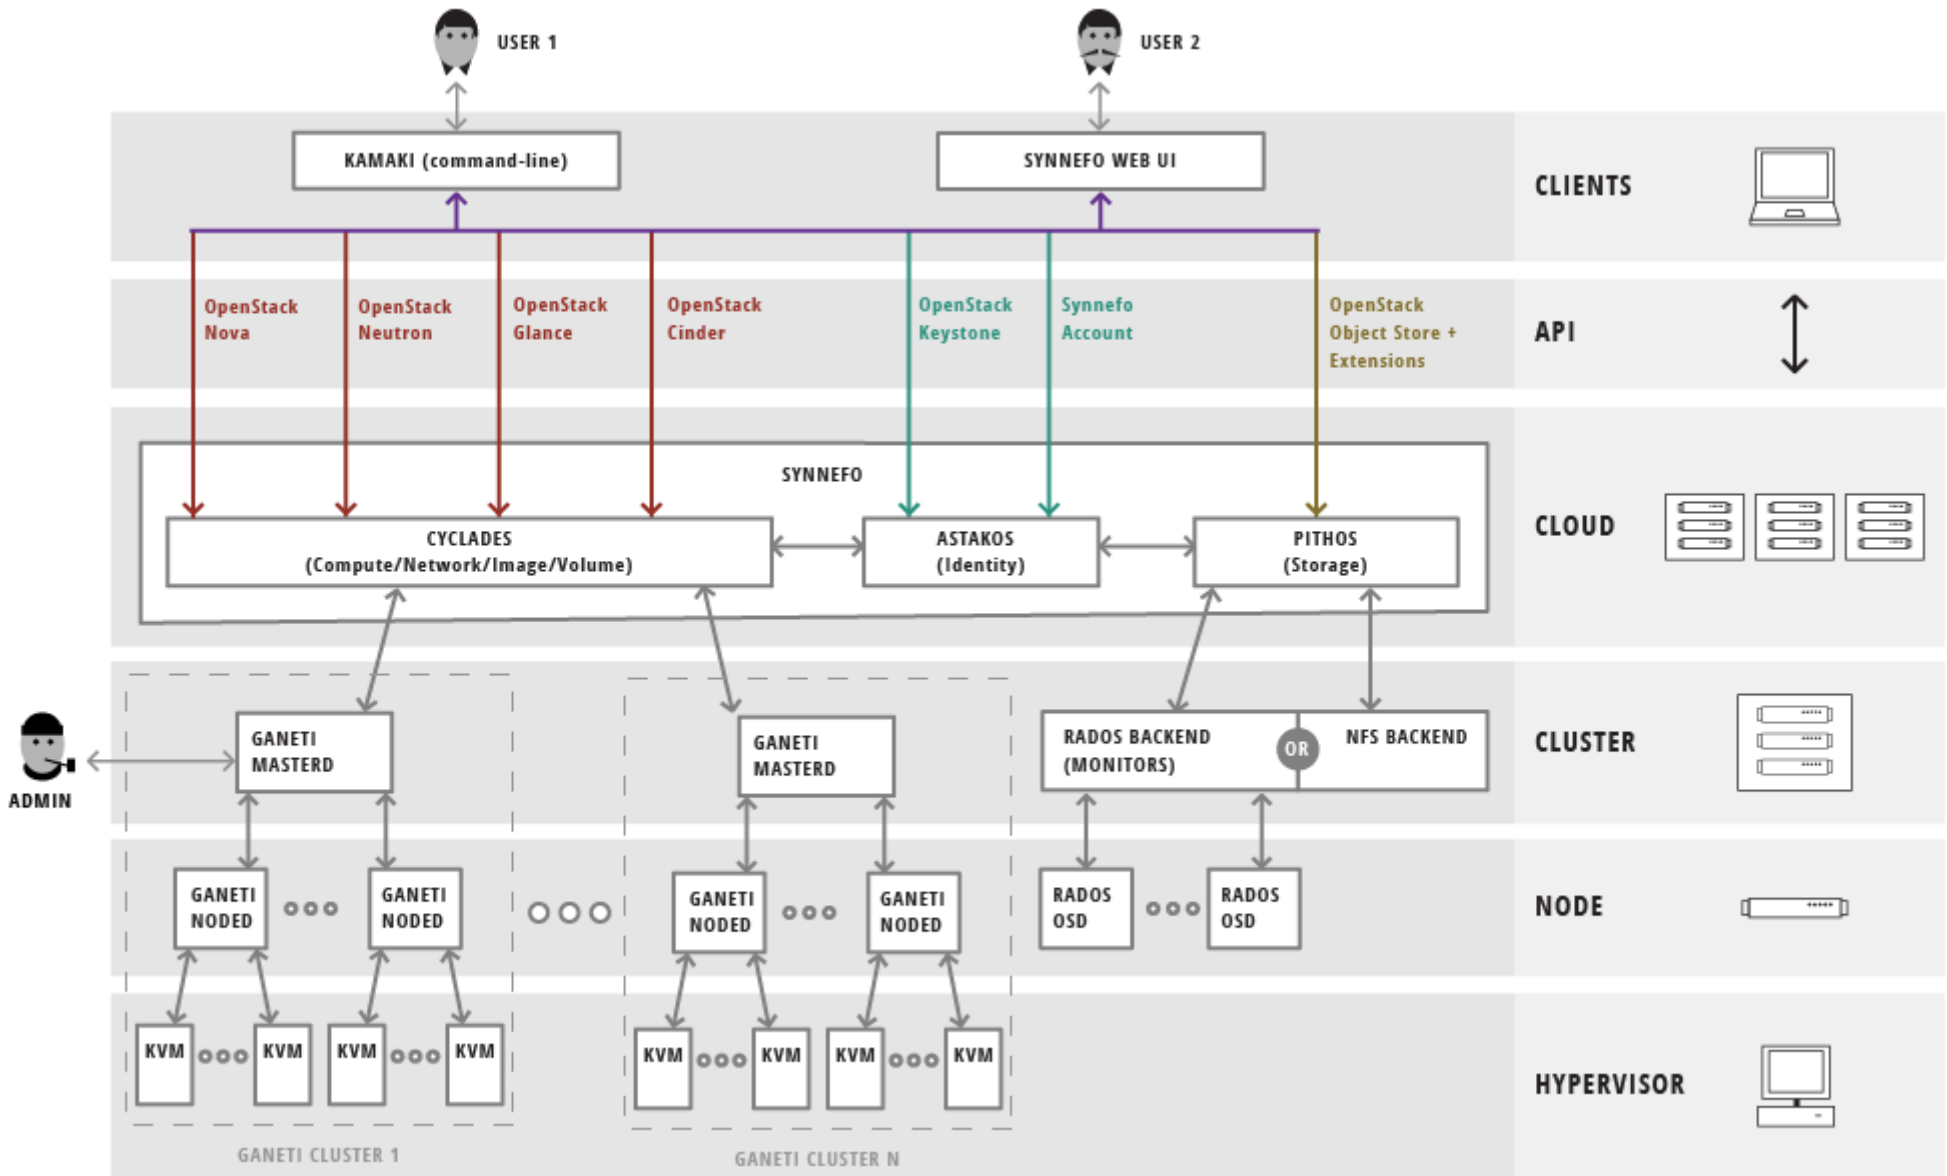


# Ganeti + Synnefo



- Ganeti developed by Google, released as open source – mainly provides clustering of KVM nodes
  - Current version 2.14
- Synnefo developed by GRNET – provides a cloud management interface to Ganeti clusters
  - Uses the OpenStack API
  - Current version 0.16
- Fewer components compared to OpenStack
- Installation and maintenance is simpler
- Less dependent on vendors





# OpenNebula

- Focus on Simplicity, Openness, Reliability and Flexibility
- Originated as a research project in 2005
  - Architecture is different from others
  - Current version is 4.12
- vOneCloud
  - Open replacement for VMWare's vCloud

## OpenNebula core

### Monitoring

- SSH-pull
- UDP-push

### Storage

- DFS like Lustre, GlusterFS, ZFS, GPFS, MooseFS...
- LVM
- VMware (VMFS)

### Network

- 802.1Q VLANS
- ebtabel
- Open vSwitch
- VMware network

### Virtualization

- Xen
- KVM
- VMware

### Hybrid

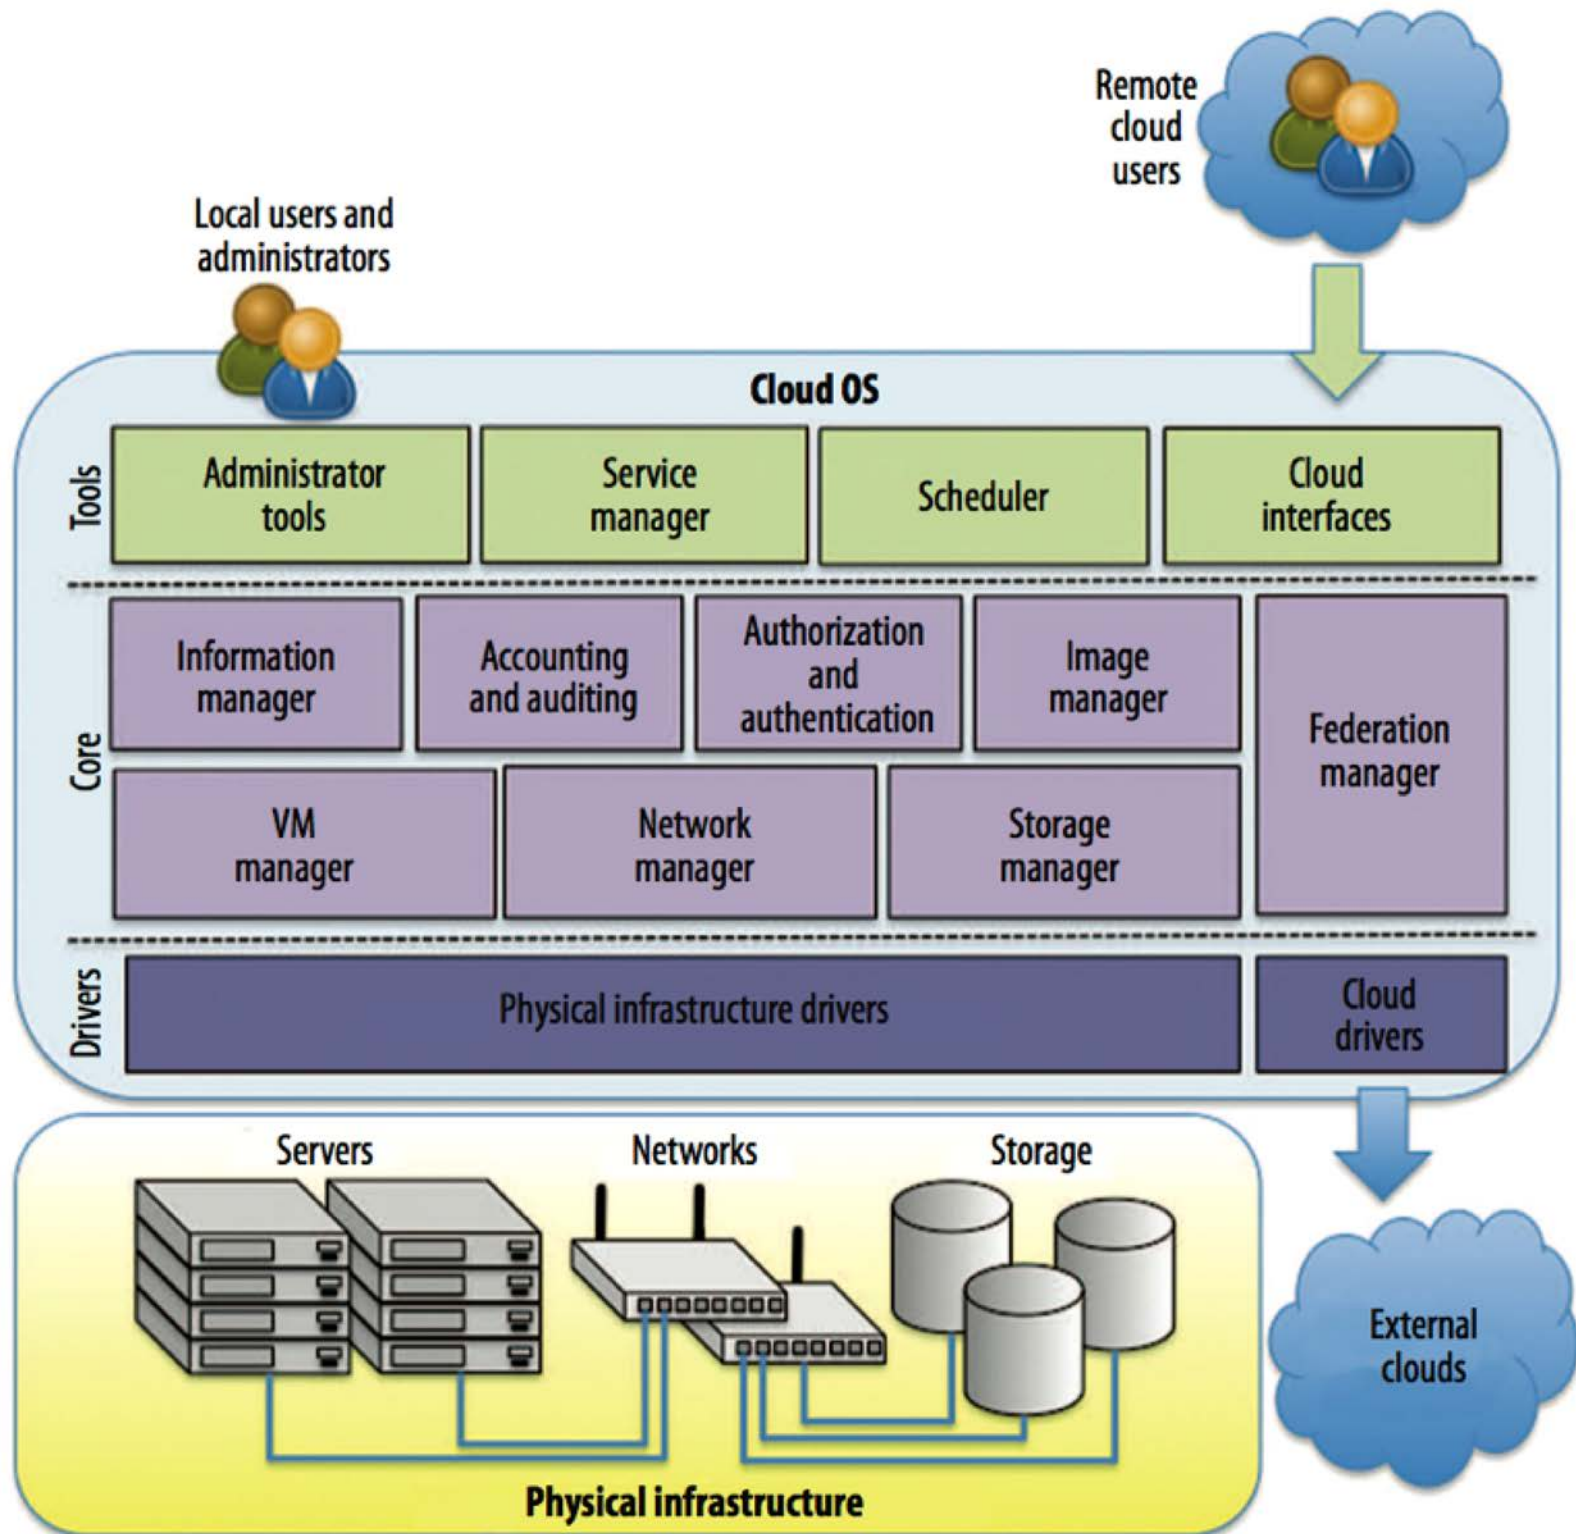
- Amazon EC2

### Auth

- password,
- ssh
- X509
- Ldap
- Active Directory

### Database

- sqlite
- mysql



# Containers

- Use of containers is changing the virtualization and cloud landscape significantly
- In conventional (hypervisor based) virtualization, each virtual machine (guest) runs an operating system on top of the host operating system
  - Inefficient
- Linux kernel features supporting virtualization has matured to a point where running a guest OS is no longer essential

# Containers

- Linux kernel features supporting virtualization
  - Namespaces: provide an isolated instance of a global resource
  - Control Groups (cgroups): allows allocation of resources (CPU time, system memory, network bandwidth, ...) among user-defined groups of tasks (processes) running on a system
- Container standards
  - Docker
  - OpenContainer

