ganeti

Ganeti Virtual Machine Cluster in Production

Hezron Mwangi
Senior Systems Administrator
hmwangi@kenet.or.ke

05th December 2018

kenet

Kenya Education Network

Acknowledgements

- Some Parts of this material originated from the Virtualization and Cloud Computing Workshop that was presented during the PacNOG 21 (Pacific Network Operators Group 21) held on December 4th - 8th, 2017 in Nuku'alofa, Tonga.
- The material was presented by Phil Regnauld from NSRC (https://www.nsrc.org/).

PacNOG

NSRC
Network Startup Resource Center

Introduction

- Ganeti is a virtual machine cluster management tool developed by Google.
- It uses:
  - either Xen, KVM or LXC as the virtualization platform,
  - LVM for disk management, and
  - Optionally DRBD for disk replication across physical hosts.
- Since 2007 Ganeti is developed and released as free and open-source software.
- Very actively developed.
  - preferred platform is Debian, others can be used.

Introduction Cont'd.

- Ganeti controls:
  - Disk creation management,
  - Operating system installation for instances (in co-operation with OS-specific install scripts), and
  - Startup, shutdown, and failover between physical systems.
- Ganeti is designed to facilitate cluster management of virtual servers and to provide fast and simple recovery after physical failures using commodity hardware.

Ganeti Features

- Ganeti provides the following features for managed instances:
- Support for Xen virtualization
    - Support for PVM and HVM instances
    - Support for live migration
    - Virtual console (on PVM) or VNC (on HVM) to control instances
    - Support for virtio or emulated devices

Ganeti Features Cont'd.

- Support for KVM virtualization
  - Support for live migration
  - Support for fully virtualized instances
  - Support for semi-virtualized instances (kernel residing on the host)
  - Support for VNC or serial access
  - Support for virtio or emulated devices

Ganeti Features Cont'd.

- Instance disk partitioning
- Export/import mechanism for backup purposes or migration between clusters
- Automated instance migration across clusters
- Remote API for integration with other systems

Ganeti Features Cont'd.

- Disk management
  - Plain LVM volumes
  - Files
  - Across-the-network RAID1 (using DRBD) for quick recovery in case of physical system failure
  - Support for third party storage solutions using External Storage Providers and shared filesystems (disk images on NFS, for example)

kenet
Kenya Education Network

Ganeti caveats

- Command-line based
    - Has a learning curve
    - Lots of features
    - Still fairly simpler to manage
- No web interface included
- Doesn't natively support the idea of different users and roles
    - but Ganeti Web Manager adds this

Terminology

- instance = a virtual machine (guest)
- node = a physical server (host)
- cluster = all nodes

Cluster management

- One of the nodes in the cluster is the master.
  - It maintains all the cluster state and copies it to other nodes.
  - It sends commands to the other nodes.
- It has an extra IP alias (the cluster address).

Cluster management Cont'd.

- All commands are executed on the master.
- If the master fails, you can promote another node to be master.
    - This is a manual event.
    - Log on to the node that will become master:
    - Run the following command:

        gnt-cluster master-failover

kenet
Kenya Education Network

Cluster Installation

- Installation:
  - *apt-get install ganeti/wheezy-backports*
  - *apt-get update && apt-get install snf-ganeti*

kenet
Kenya Education Network

Cluster Initialization

- Initialization:

*gnt-cluster init --secondary-ip 192.168.x.x --vg-name=ganeti *

*--master-netdev=<mgmt-bridge> *

*--enabled-hypervisors=kvm *

--hypervisor-parameters
*kvm:kernel_path=,initrd_path=,vnc_bind_address=0.0.0.0 *

*--nic-parameters link=<default-vm-bridge> *

*--disk-parameters=drbd:metavg=ganeti *

*--maintain-node-health=yes *

*--default-iallocator hail *

*--specs-nic-count min=0,max=16 *

*--enabled-disk-templates plain,drbd,ext *

ganeti.example.co.ke

kenet
Kenya Education Network

Other Cluster Commands

- Add node to cluster:

  gnt-node add -s 192.168.x.y ganeti-node-02.example.co.ke

- Verifying Configuration:

  gnt-cluster verify

- Redistribute Configuration:

  gnt-cluster redist-conf

“Plain” instance

- Instance disks are Logical Volumes, stored on a
- Volume Group
- Can be expanded as required

“Plain” instance Cont’d.

- *Command:*

```
gnt-instance add -t plain \
    -o noop -n ganeti-node \
    --disk 0:size=10G --disk 1:size=10G \
    -B vcpus=1,memory=1G \
    --no-install --no-start --no-ip-check --no-name-check \
    instance-name.example.co.ke

gnt-instance modify -H
  boot_order=cdrom,cdrom_image_path=/var/ganeti/iso/debian-7.8.0-
  amd64-netinst.iso instance-name.example.co.ke

gnt-instance modify -H boot_order=disk,cdrom_image_path=""
  instance-name.example.co.ke
```

kenet
Kenya Education Network

“Plain” instance Cont’d.

- Command:

```
gnt-instance add -t plain \
    --no-ip-check --no-name-check --no-wait-for-
      sync \
    --disk=0:size=100G,vg=ganeti,metavg=ganeti \
    -B vcpus=4,memory=4G \
    -o snf-image+freebsd-11-2 \
    --net 0:link=<default-vm-bridge> \
    -n ganeti-node \
    instance-name.example.co.ke
```

kenet
Kenya Education Network

“DRBD” instance

- Instance (Guest) does not use LV directly
- DRBD layer replicates disk to secondary node (B)
- Failover / migration becomes possible

“DRBD” instance Cont’d.

- Command

```
gnt-instance add -t drbd \
    --no-ip-check --no-name-check --no-wait-for-
      sync \
    --disk=0:size=100G,vg=ganeti,metavg=ganeti \
    -B vcpus=4,memory=4G \
    -o snf-image+freebsd-11-2 \
    --net 0:link=<default-vm-bridge> \
    -n primary-node:secondary-node \
    instance-name.example.co.ke
```

kenet
Kenya Education Network

DRBD: Migration

- Copy RAM + state of Instance 1 from A to B
- Pause Instance 1 on A
- Copy RAM again, reverse DRBD roles
- Resume Instance on B - B is now primary for Instance 1
- Command:

gnt-instance migrate instance-name.example.co.ke

kenet
Kenya Education Network

Unplanned node failure

- Instance 2 was running on node B
- It can be restarted on its secondary node (instance failover)
- This is a manual event
- Command:
- If the instance is off:

  gnt-instance failover instance-name.example.co.ke
- If the instance is on (when there is no node failure):

  gnt-instance migrate instance-name.example.co.ke

kenet
Kenya Education Network

Available storage templates

- plain: logical volume
- drbd: replicated logical volume
- file: local raw disk image file
- sharedfile: disk image file over NFS etc
- blockdev: any pre-existing block device
- rbd: Ceph (RADOS) distributed storage
- diskless: no disk (e.g. live CD)
- ext: pluggable storage API

Networking

- You need one management IP address per node, plus one cluster management IP
    - all on the same management subnet
- Optional: separate replication network for drbd traffic
- Optional: additional network(s) for guests to connect to
    - so they cannot see the cluster nodes
    - reduces the impact of a DoS attack on a guest VM

kenet
Kenya Education Network

Ganeti Scaling

- Start with just one or two nodes!
- Recommended you limit cluster to 40 nodes.
- Beyond this you can just build more clusters.

Available instance OS definitions

- ganeti-os-noop
  - empty script, does nothing
- ganeti-instance-debootstrap
  - install Debian/Ubuntu from .deb packages downloaded over the net
- ganeti-instance-image
  - unpacks prepared dump or tar images
- snf-image
  - image cloning part of the synnefo cloud solution
- modify or write your own

kenet
Kenya Education Network

Limitations of each approach

- Install from ISO? Have to learn to attach VGA console. Have to do manual install each time
- Install using debootstrap? Currently have to manually install grub in the guest; or have to boot from kernel on the host filesystem
- Install from filesystem dump? Have to prepare the dump. Have to script post-install tweaks
- Clone disk image? May need to resize partitions and filesystems to desired size
- snf-image does this for you

Note for Windows guests

- Attach a second CD-ROM with the RedHat Windows Virtio drivers (free)
    - cdrom2_image_path=...
    - cdrom_disk_type=ide
- Install Windows with virtio disk, network and balloon memory drivers
- Makes Windows work much better in the virtual environment

gnt-instance start windows-test.example.co.ke -H cdrom_image_path=/home/ganeti/iso/windows.iso,cdrom2_image_path="/home/ganeti/iso/virtio-win-0.1-100.iso"

kenet
Kenya Education Network

Web interfaces

- The Ganeti RAPI (Remote API) allows other systems to interface into Ganeti
- There are two standalone web front-ends available
  - Ganeti Web Manager (OSU OSL: Oregon State University Open Source Labs)
  - Ganeti Manager (GRNET NOC: Greek Academic, Research and Education Network NOC)
- Both are Python/Django applications
  - aka “non-trivial installations”
- Both projects are active.

Ganeti Web Manager

grnet

Ganetimgr

kenet
Kenya Education Network

Ganetimgr

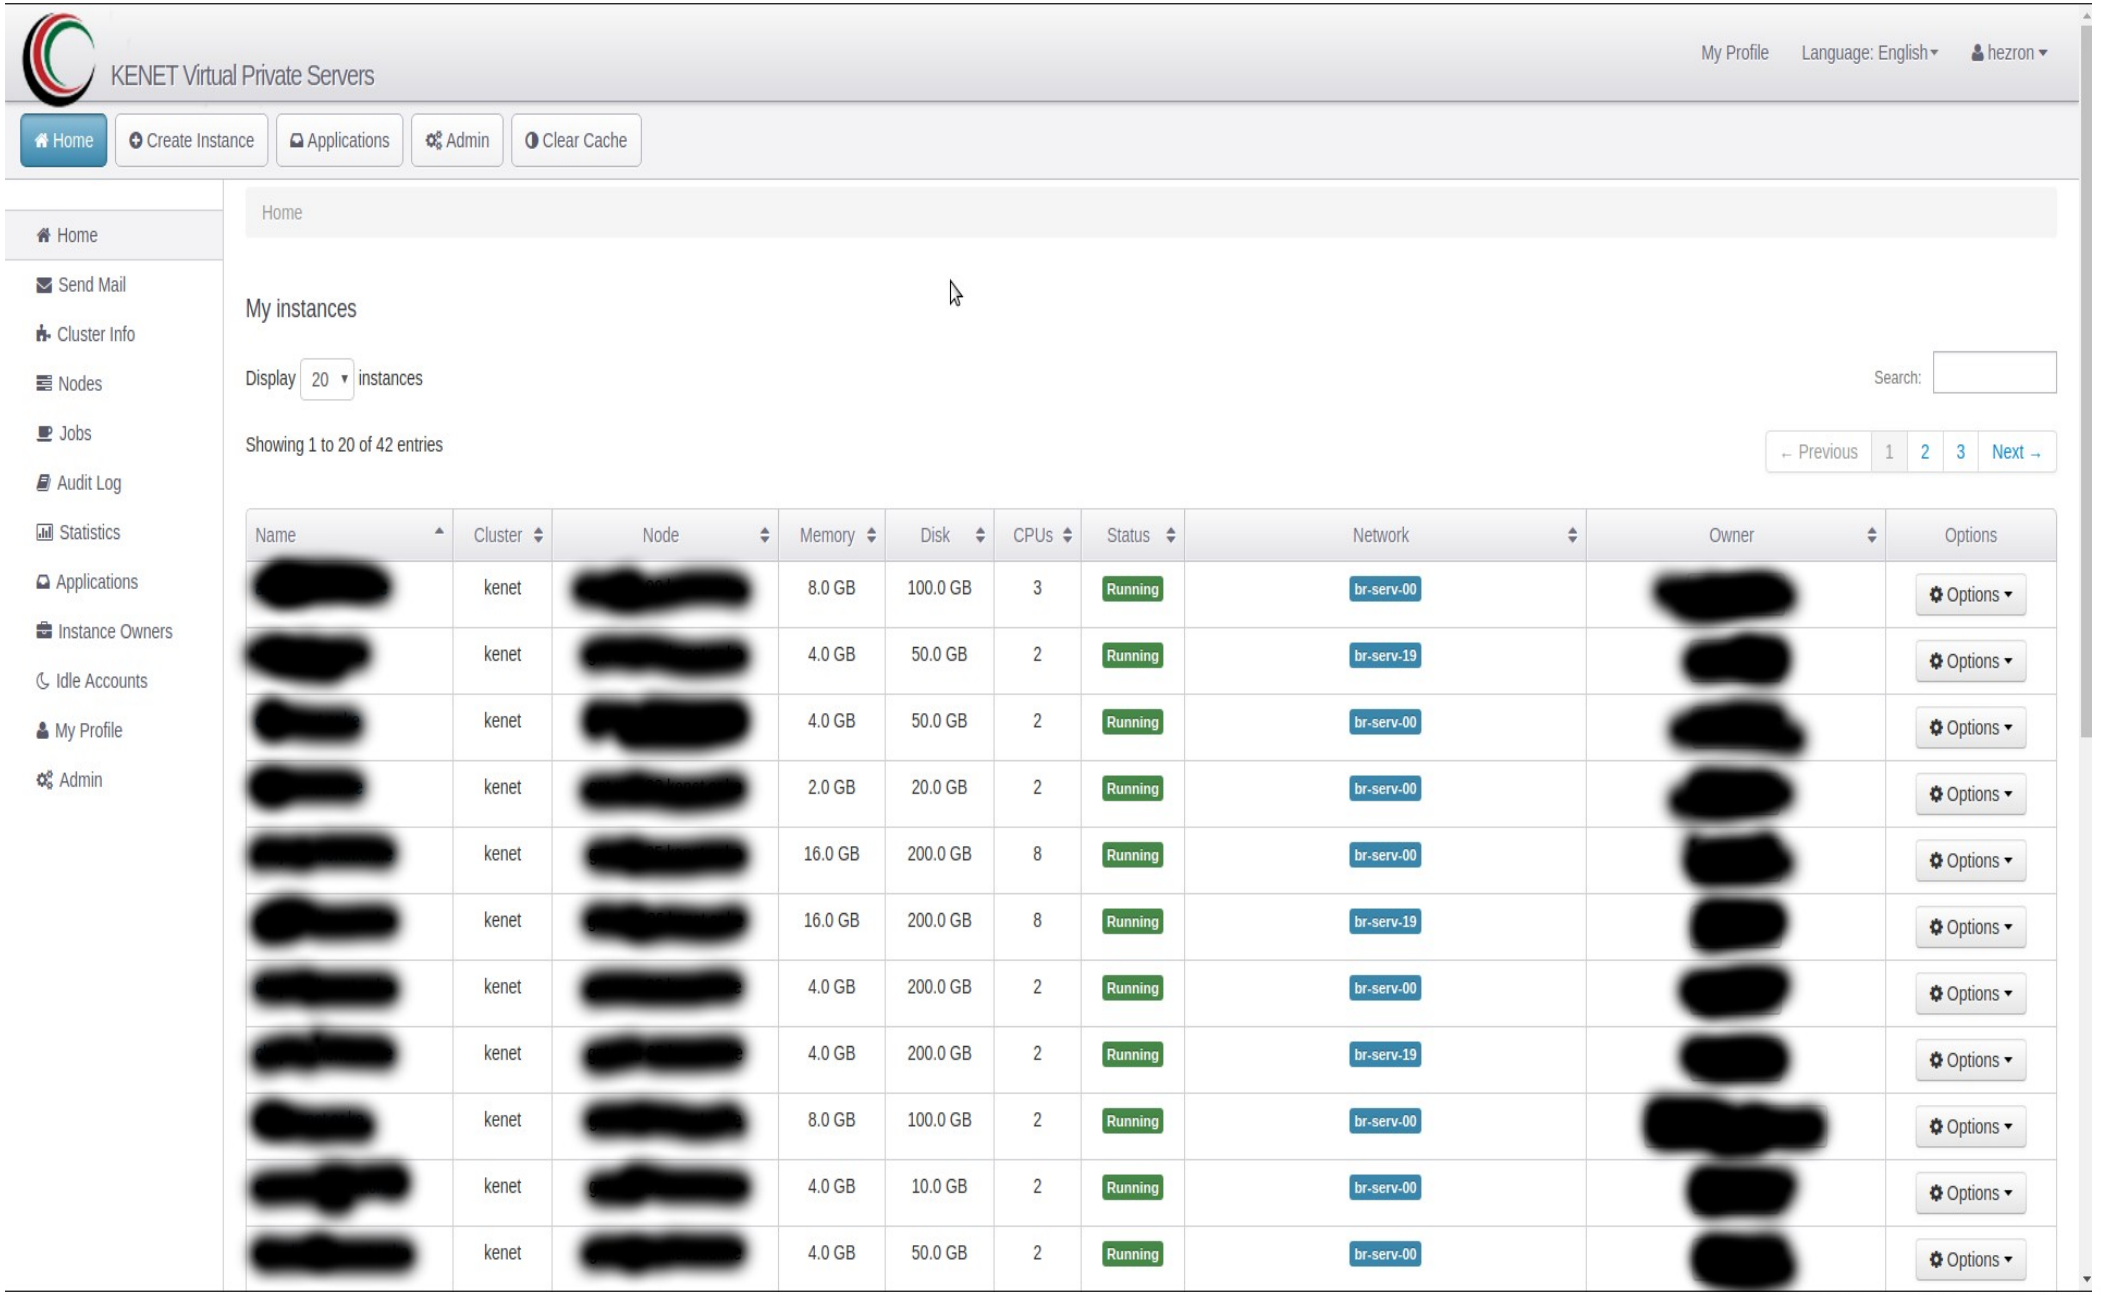

Ganeti Web Manager

grnet

Ganetimgr

kenet Kenya Education Network

Ganeti Web Manager

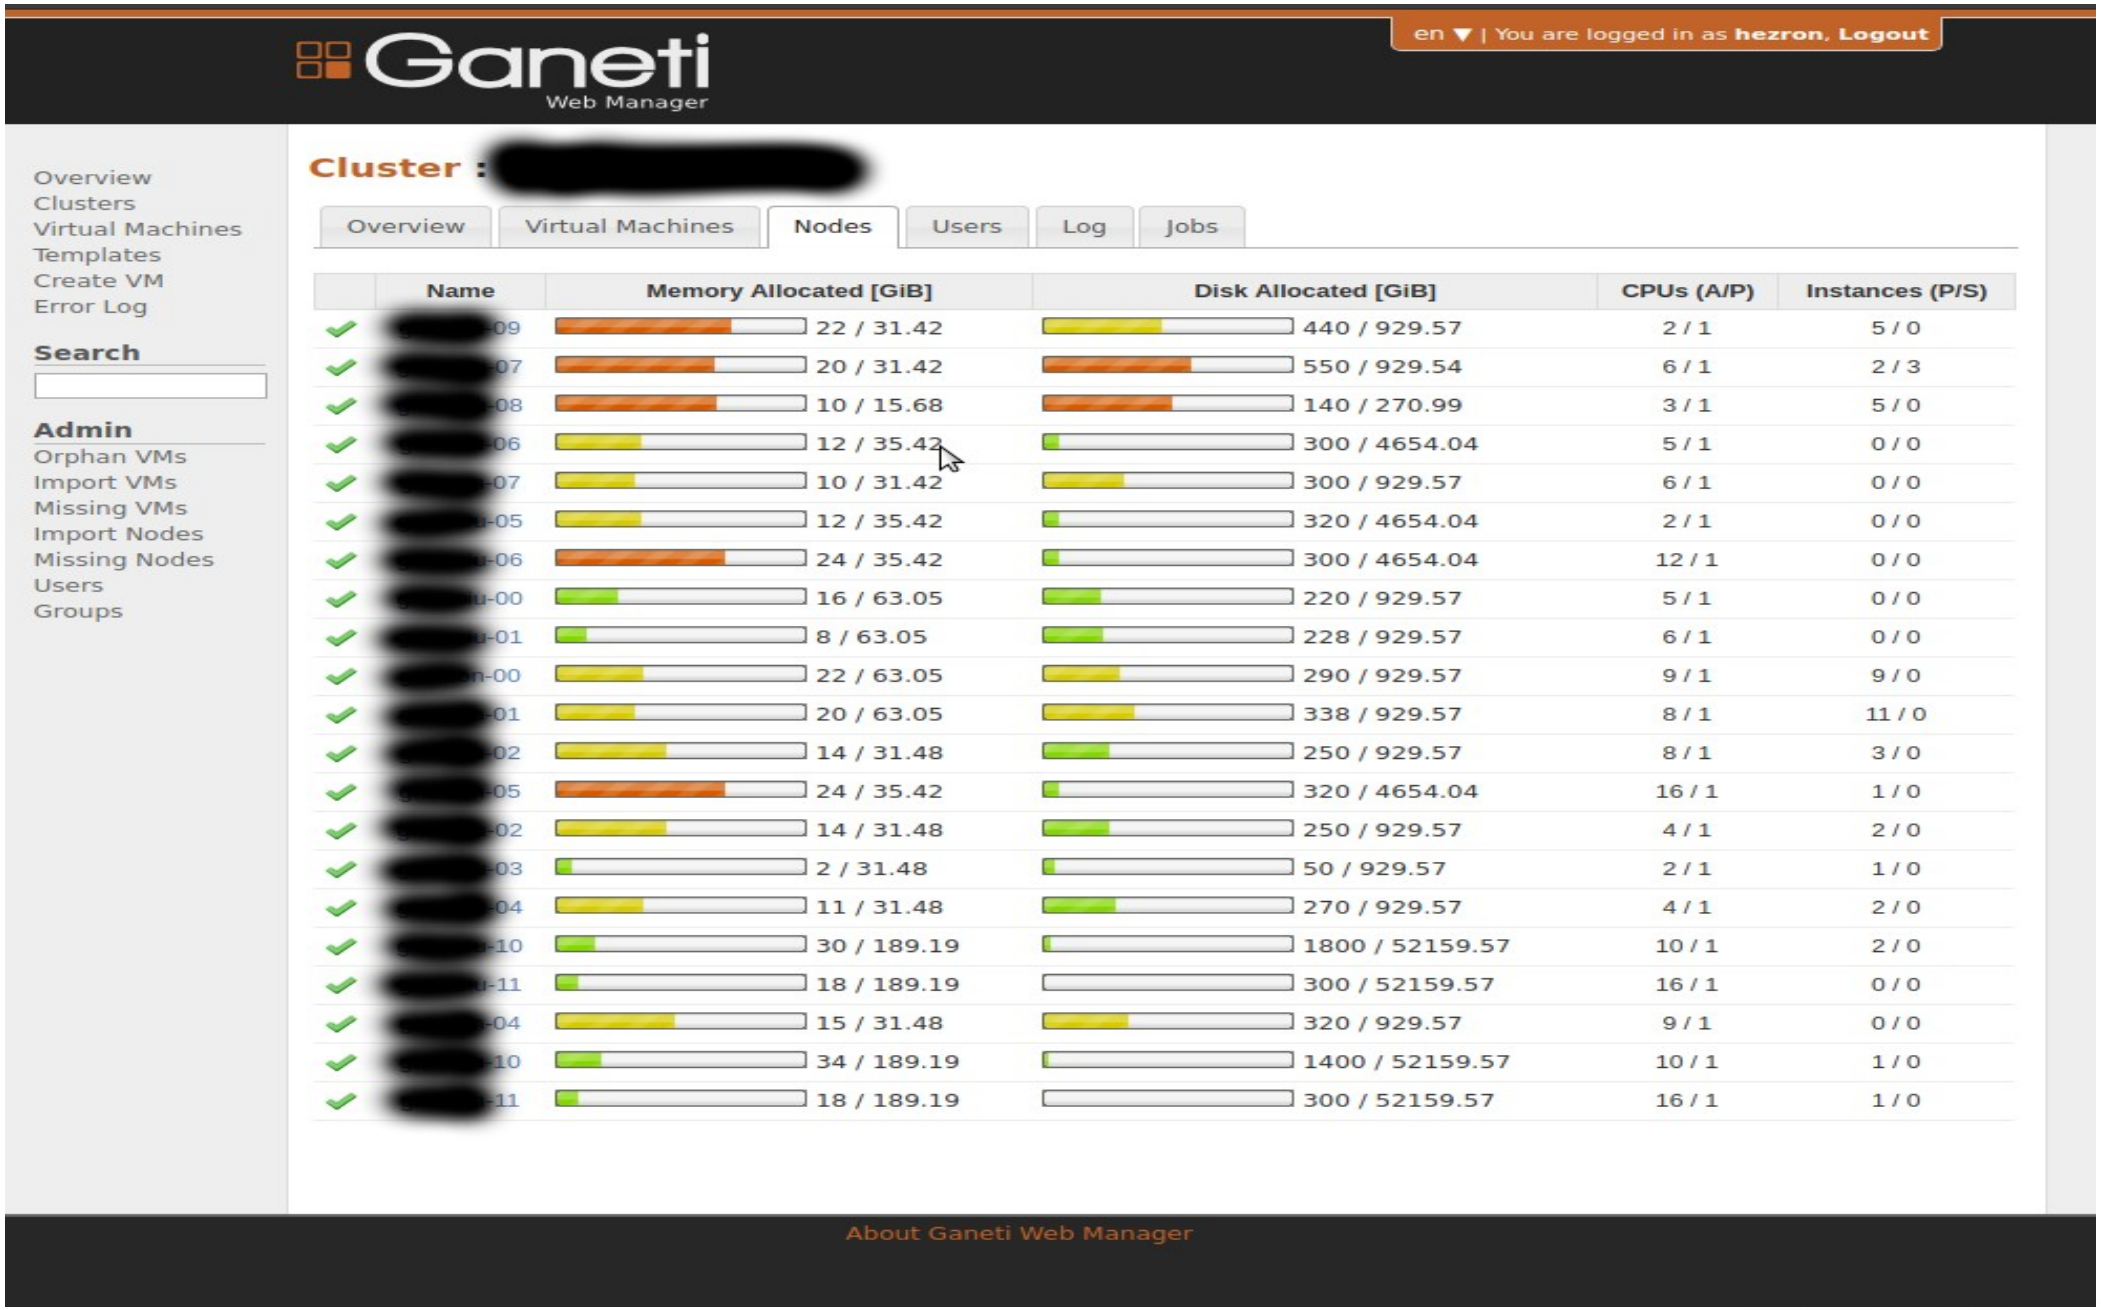

Ganeti Web Manager

Ganetimgr

Cloud Lightweight options

- Ganeti + ganetimgr
  - Basic provisioning of VMs via web tickets
  - Thin front-end to Ganeti cluster RAPI
  - Ideal for university-type environment
- Synnefo
  - A full cloud system on top of one or more Ganeti clusters

Synnefo components

- "Cyclades": compute, network, image, volume
  - Uses Ganeti as backend
- "Pithos": file and object storage
  - Uses Ceph (radosgw) or NFS storage
- "Astakos": identity and account services
- Provides web interface and OpenStack APIs

synnefo

kenet
Kenya Education Network

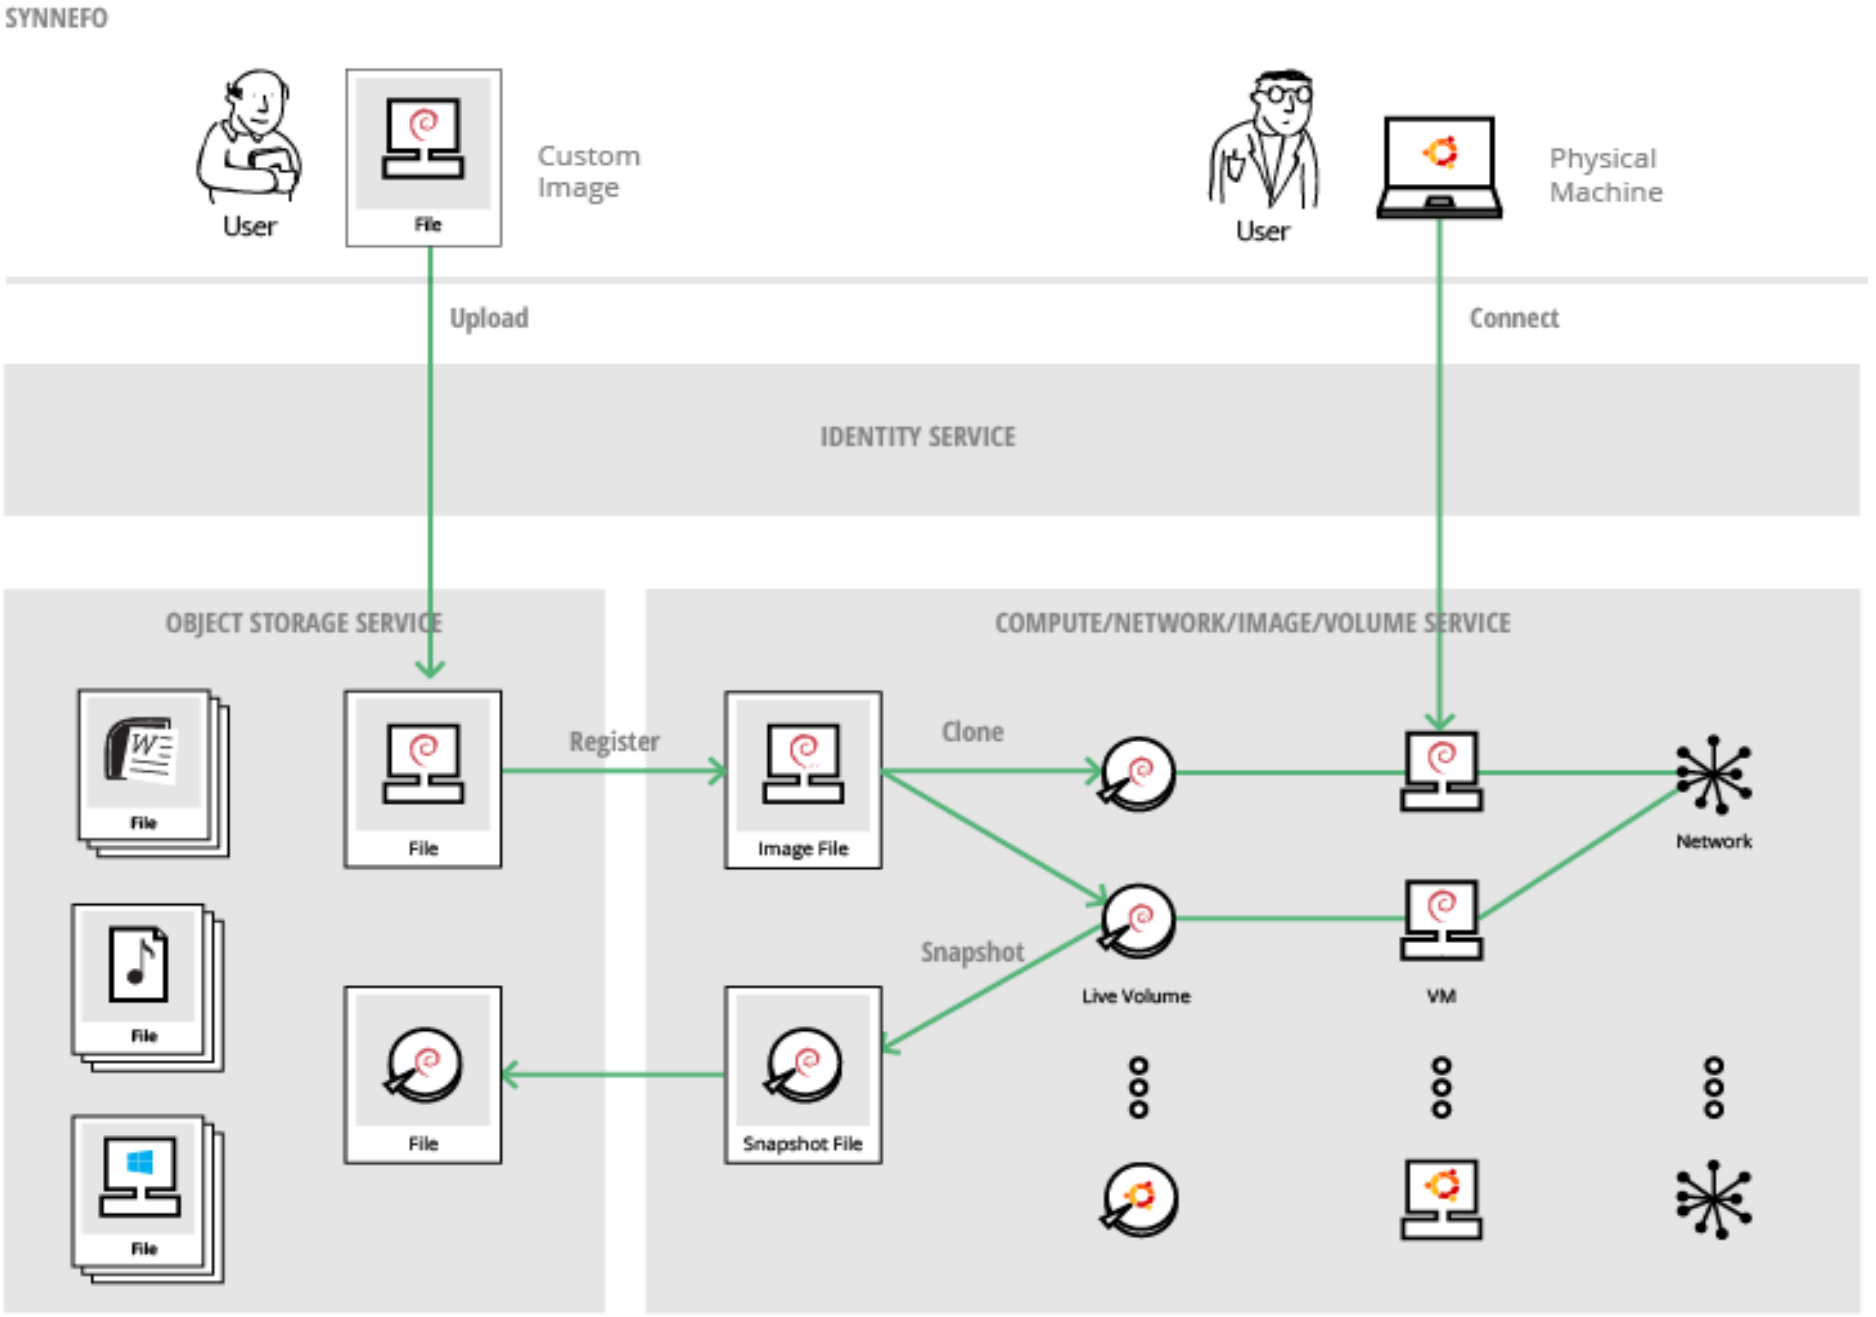
SYNNEFO
User
File
Custom Image
User
Physical Machine
Upload
Connect
IDENTITY SERVICE
OBJECT STORAGE SERVICE
COMPUTE/NETWORK/IMAGE/VOLUME SERVICE
File
File
File
File
Register
Image File
Clone
Network
Snapshot
Live Volume
VM
File
Snapshot File

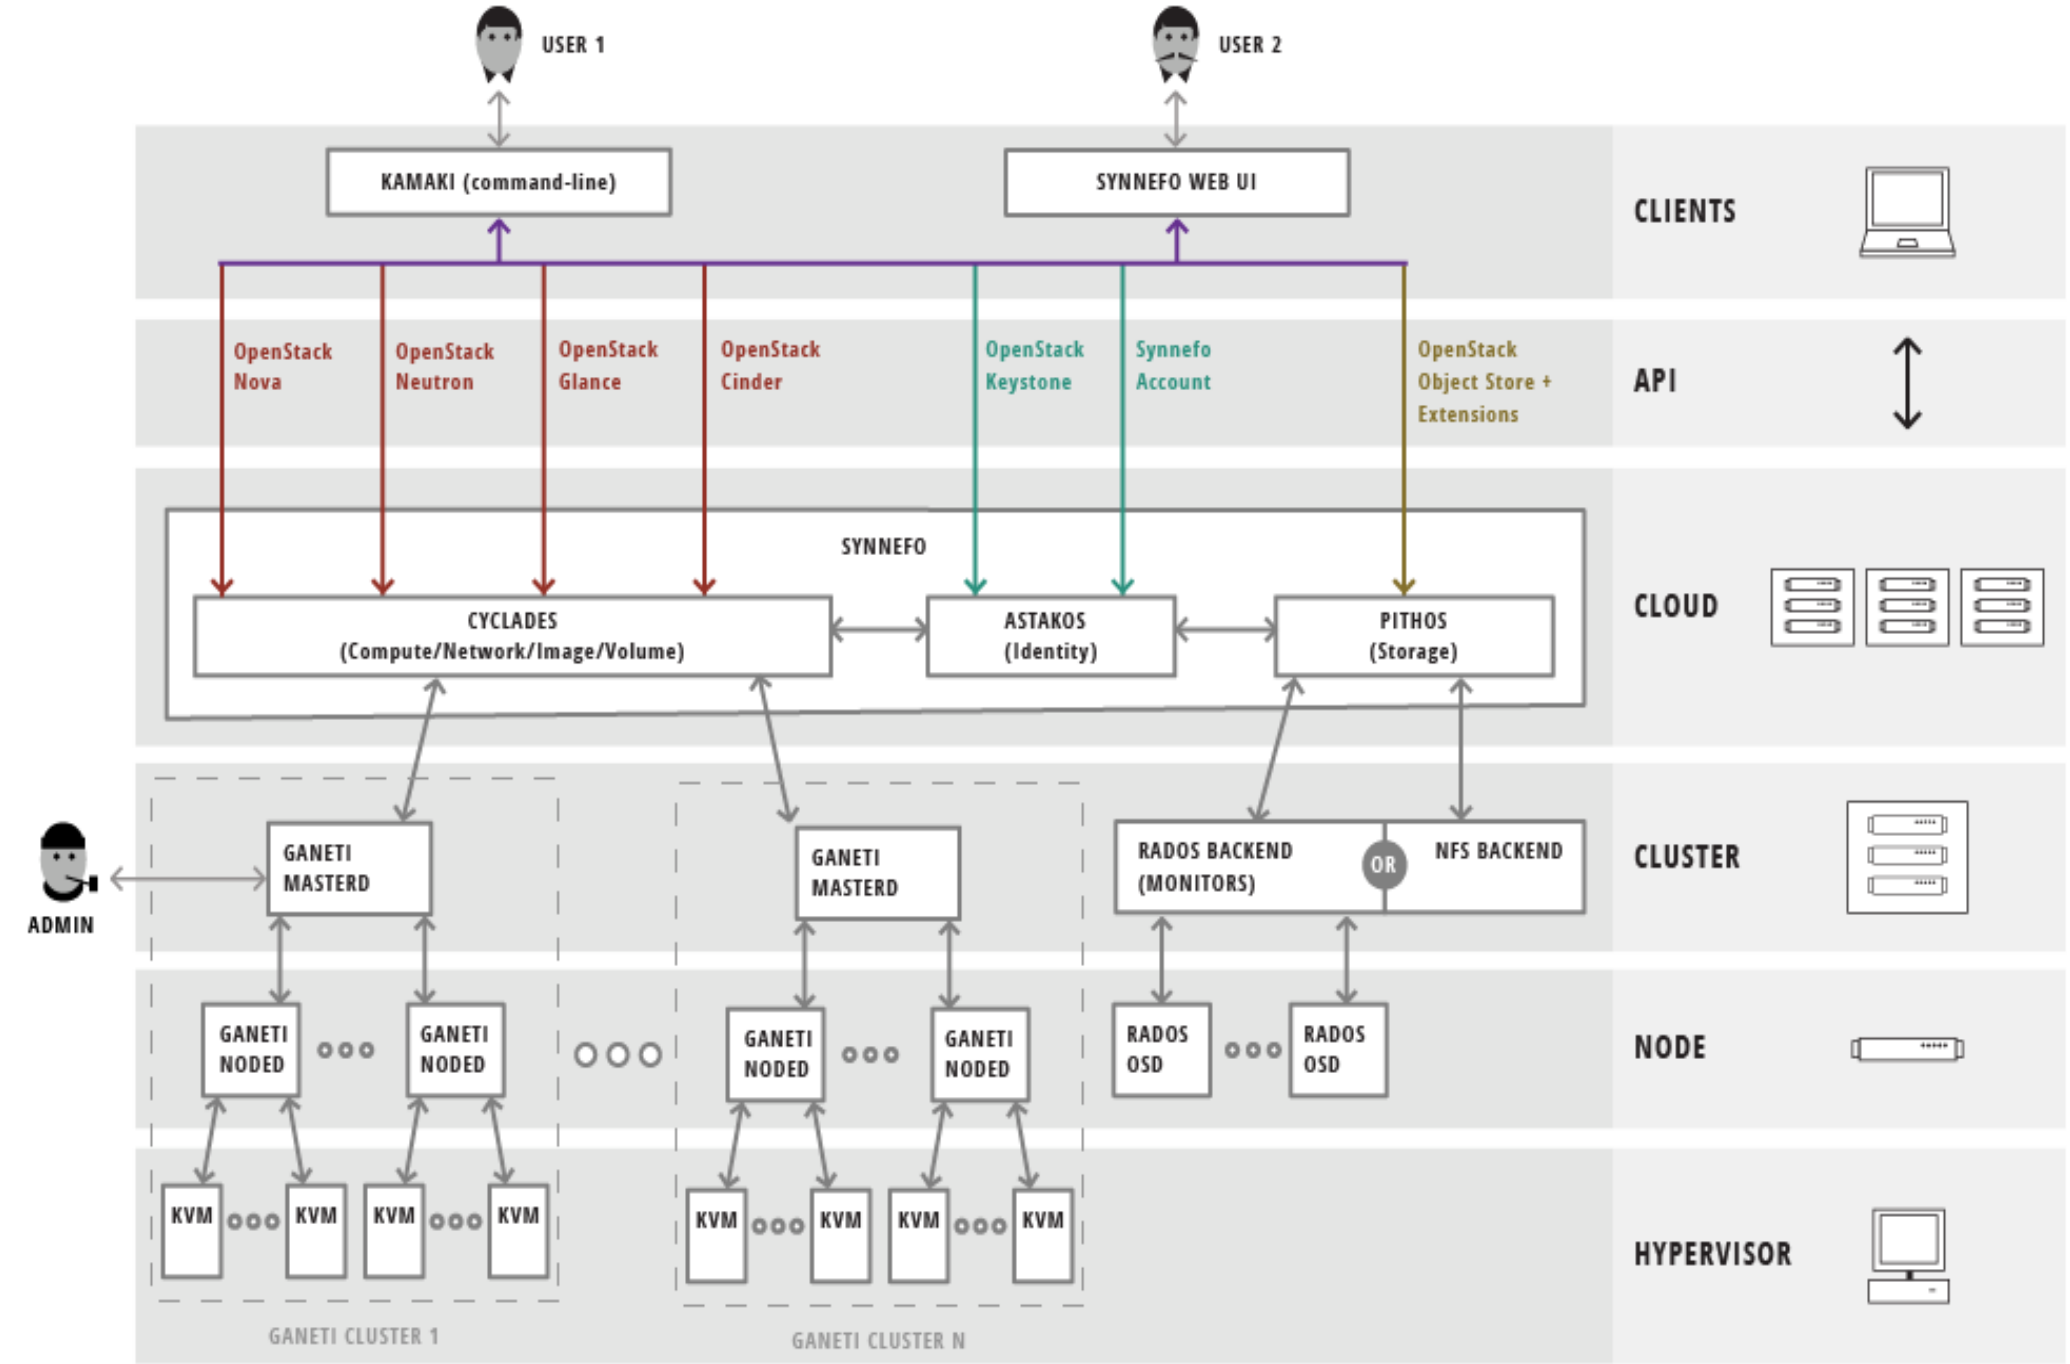

synnefo

kenet
Kenya Education Network

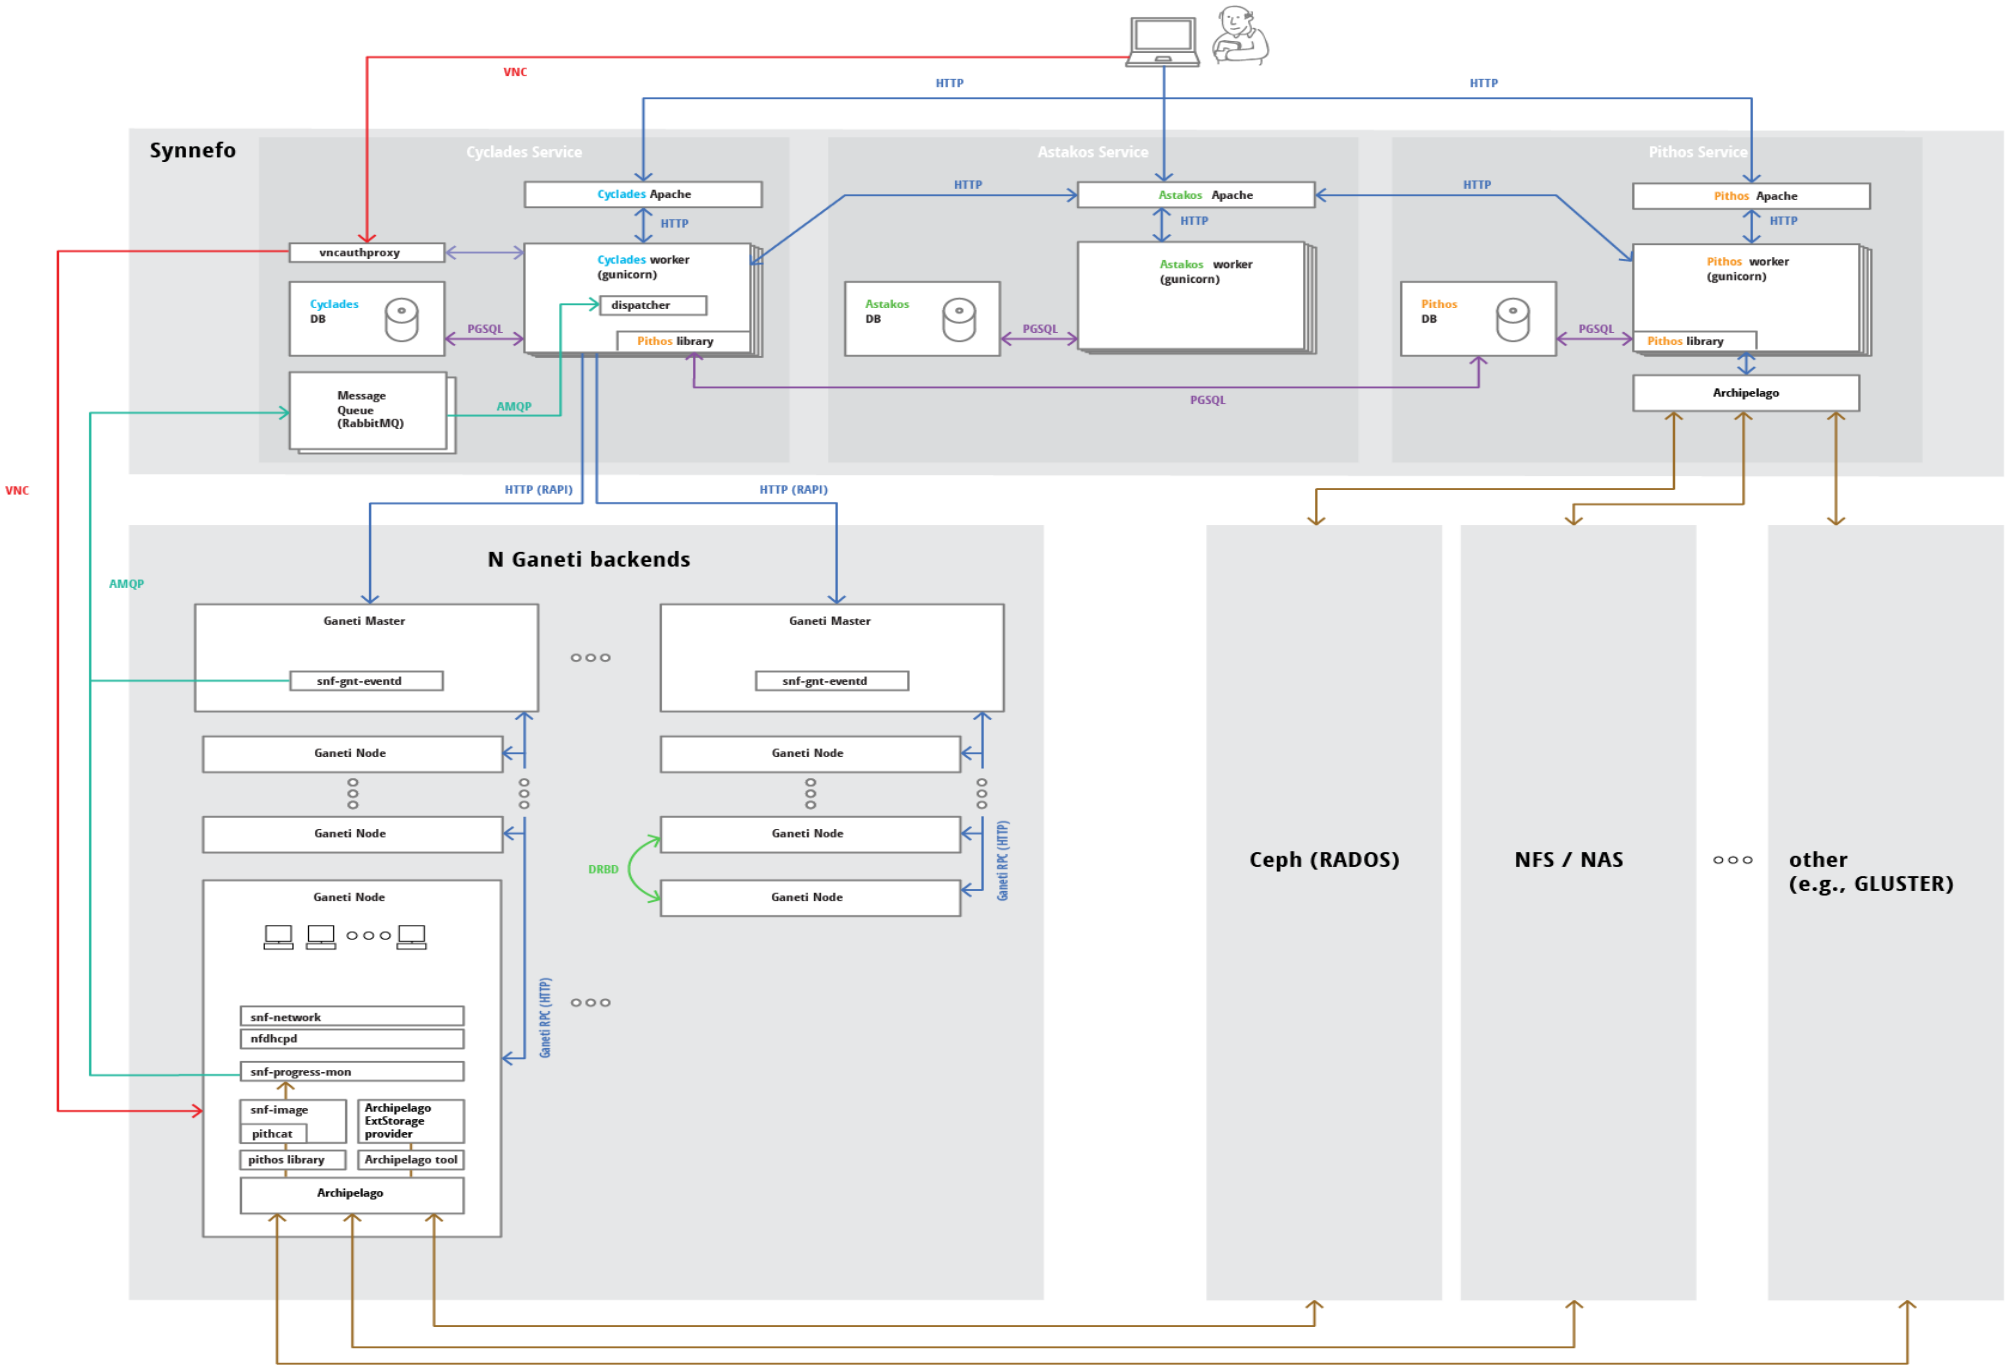

synnefo

kenet
Kenya Education Network

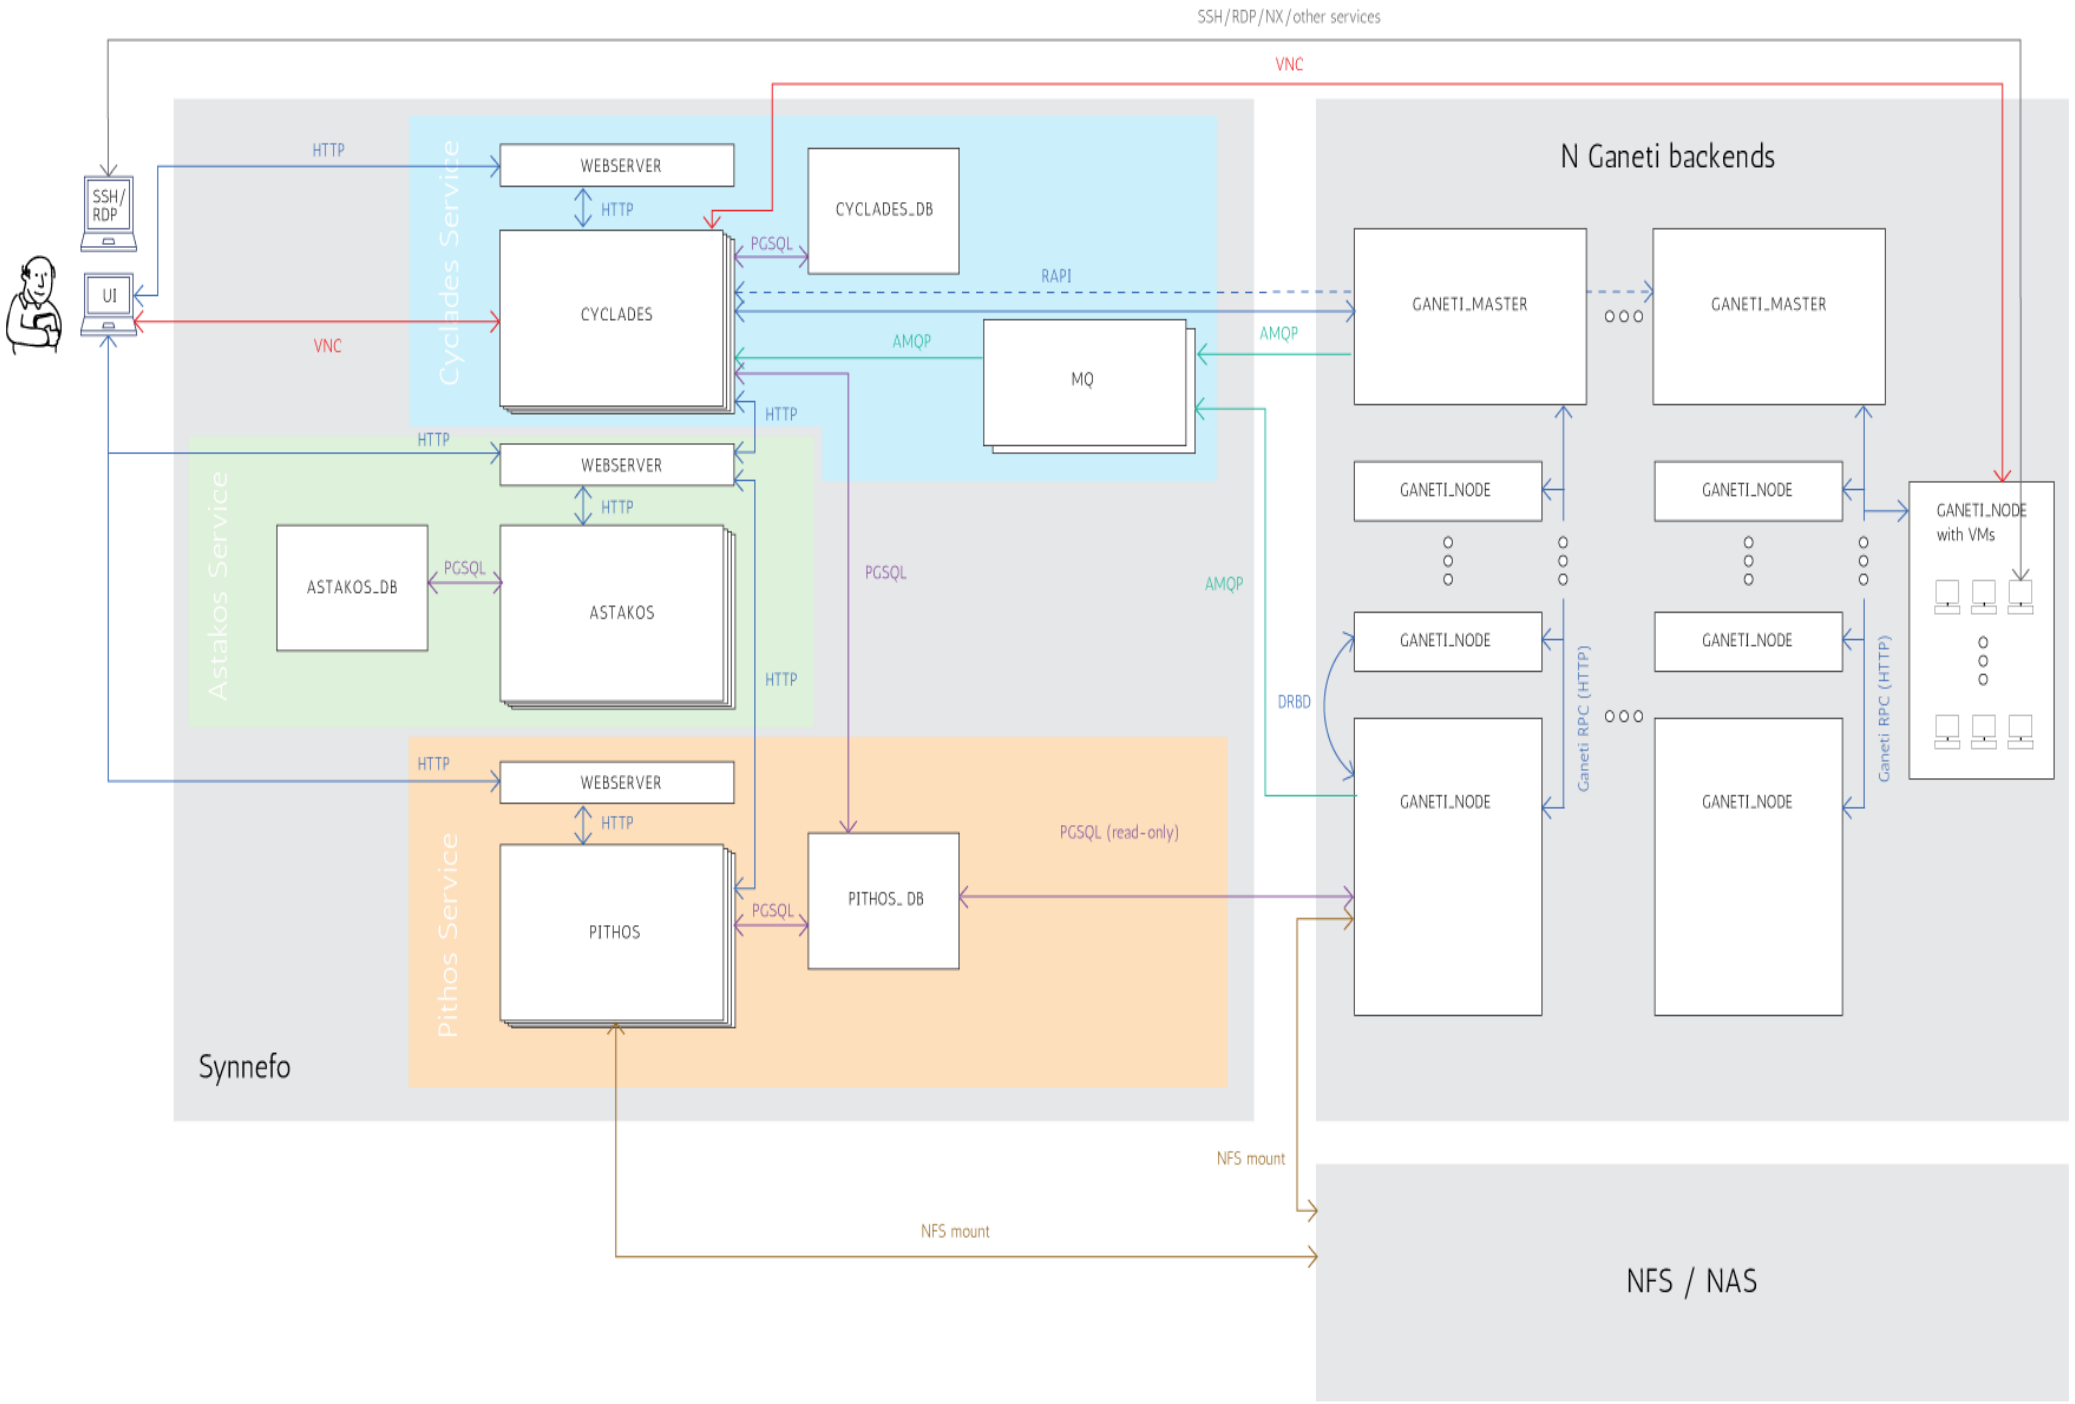

synnefo

kenet
Kenya Education Network

Synnefo

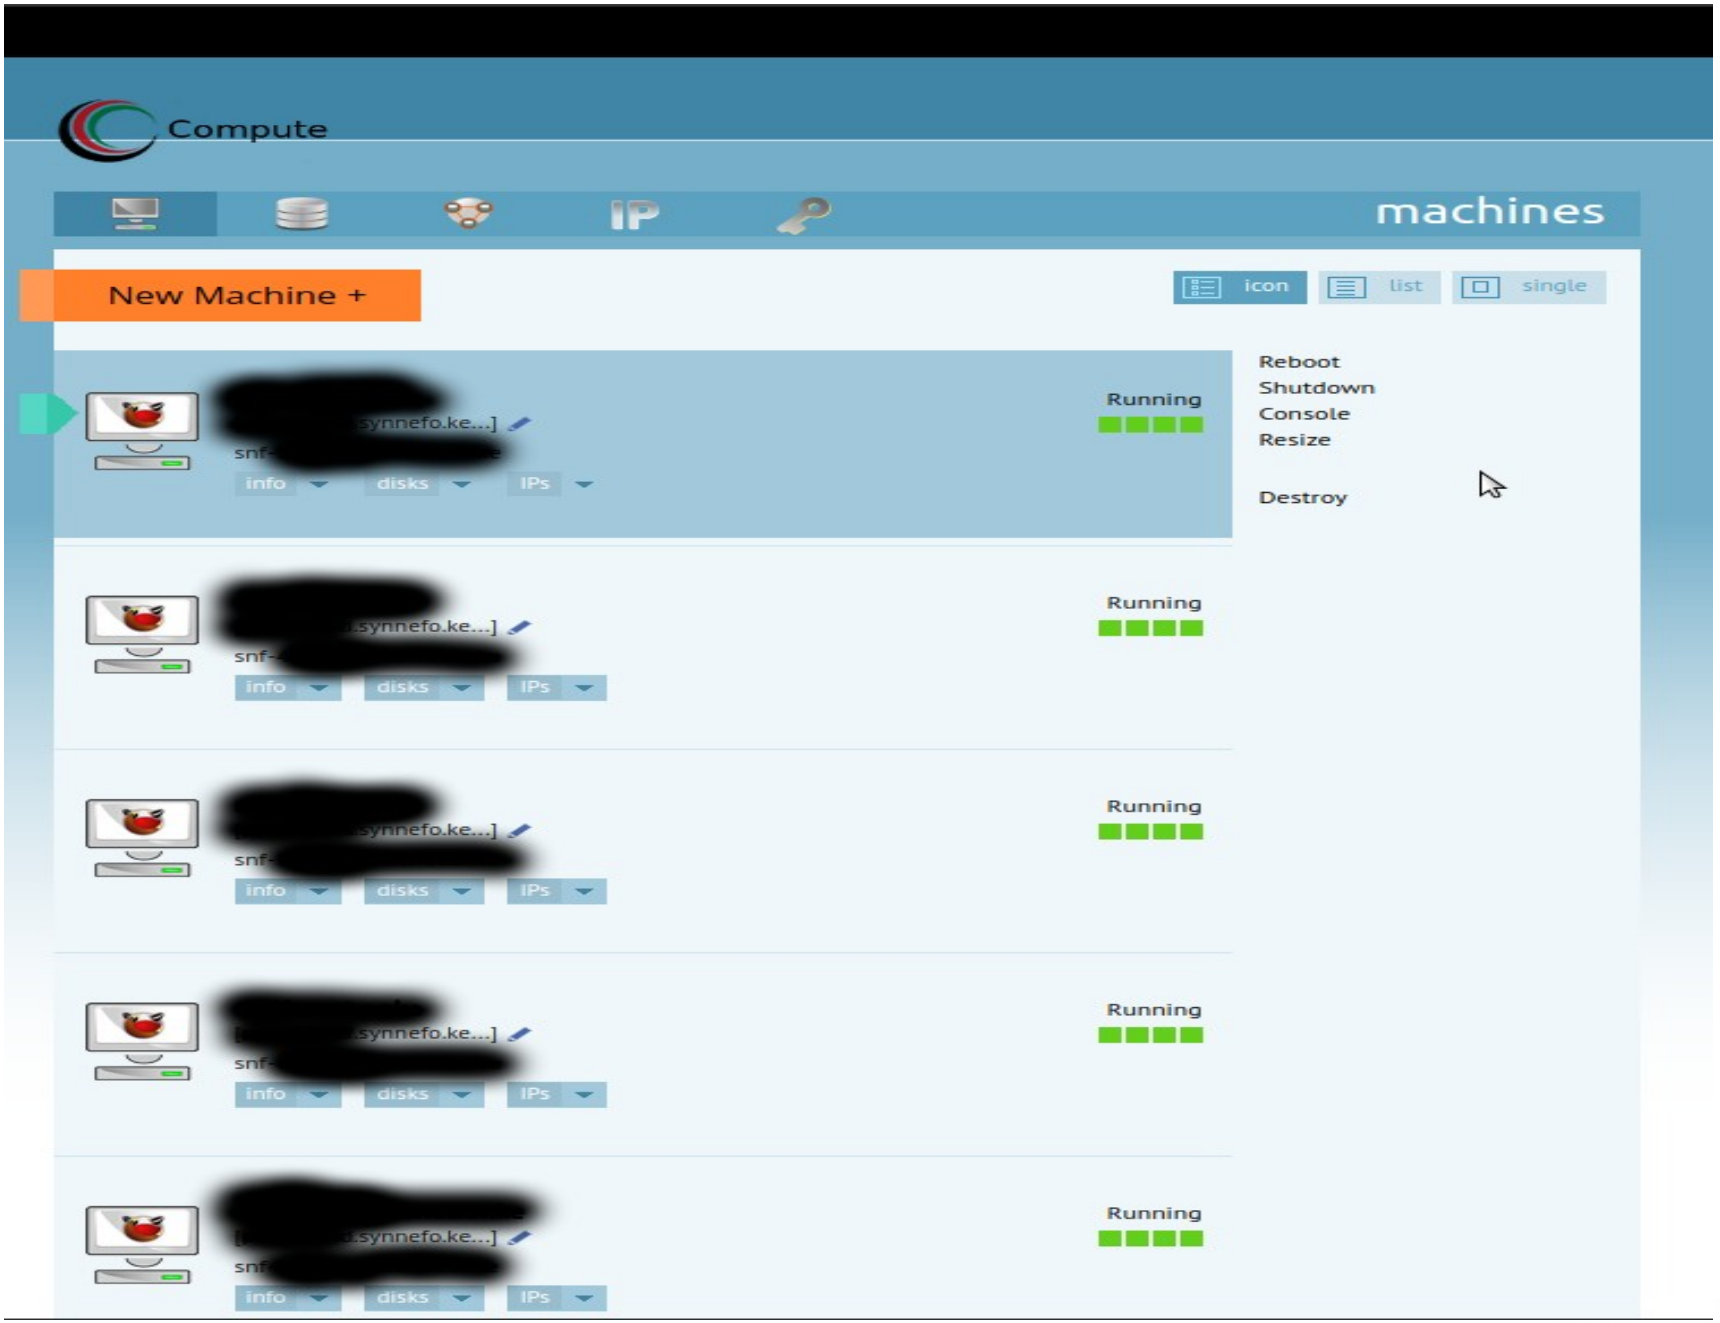
Compute
IP
machines
New Machine +
icon
list
single
Reboot
Shutdown
Console
Resize
Destroy
Running
info
disks
IPs

What distinguishes IaaS from just "virtualization"?

- Self-service
  - Users create, manage and destroy their own VMs and storage
- Multi-tenant
  - Multiple users on the same infrastructure, but isolated from each other
- Resource accounting
  - Feeds into billing

synnefo

kenet
Kenya Education Network

References

- http://www.ganeti.org/
- http://docs.ganeti.org/
- https://ganeti.googlesource.com/wiki/+/master/DocumentationByCategory.md
- https://en.wikipedia.org/wiki/Ganeti

- https://github.com/grnet/ganetimgr
- https://ganetimgr.readthedocs.io/en/master/install.html

- https://github.com/osuosl/ganeti_webmgr
- https://ganeti-webmgr.readthedocs.io/en/latest/
- https://en.wikipedia.org/wiki/Ganeti_Web_Manager

- https://www.synnefo.org/
- https://www.synnefo.org/resources/
- https://en.wikipedia.org/wiki/Synnefo

Q&A.

?

kenet
Kenya Education Network

kenet
Kenya Education Network

THANK YOU!