

**INTRODUCING THE SYNNEFO OPEN SOURCE
IAAS PLATFORM. VANGELIS KOUKIS
TECHNICAL LEAD, ~OKEANOS**

Motivation: ~okeanos

FOSDEM 2013

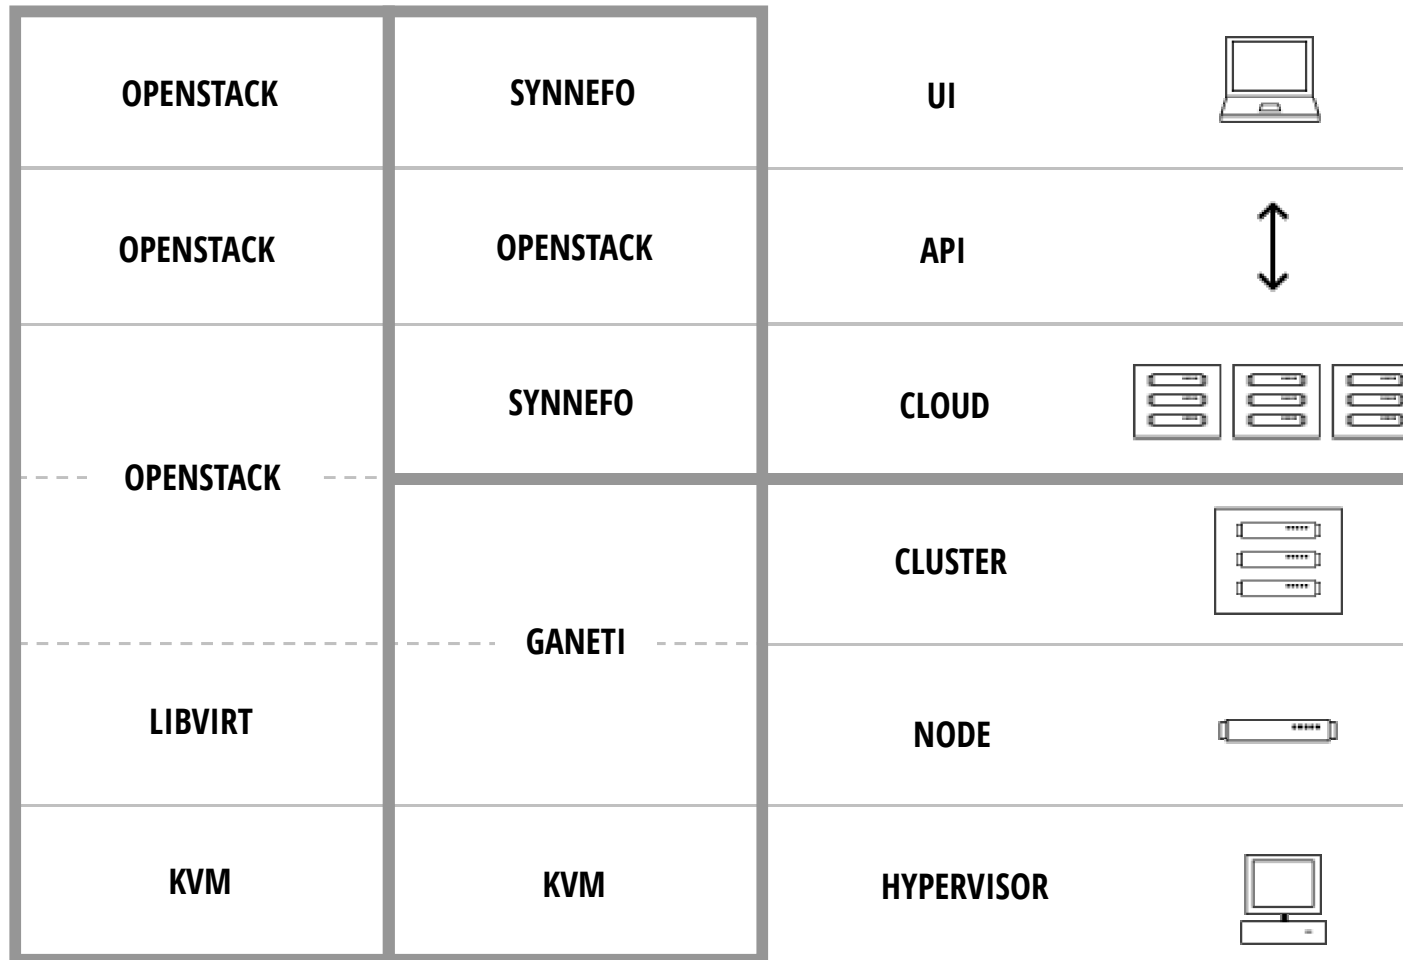
vkoukis@grnet.gr

Goals

- Everything open source
- Production-quality IaaS
- Persistent VMs
- Commodity hardware – No SAN
- Large-scale Installations
- Super-simple UI

How?

Cluster vs Cloud



Google Ganeti

FOSDEM 2013

vkoukis@grnet.gr

Mature, production-ready VM cluster management

- used in Google's corporate infrastructure

Multiple storage backends out of the box

- DRBD, LVM, Files on local or shared directory, RBD (Ceph/RADOS)

External Storage Interface for SAN/NAS support

Ganeti cluster = *masterd* on master, *noded* on nodes

Easy to integrate into existing infrastructure

- Remote API over HTTP, pre/post hooks for every action!

Compute – Network: Cyclades

FOSDEM 2013

vkoukis@grnet.gr

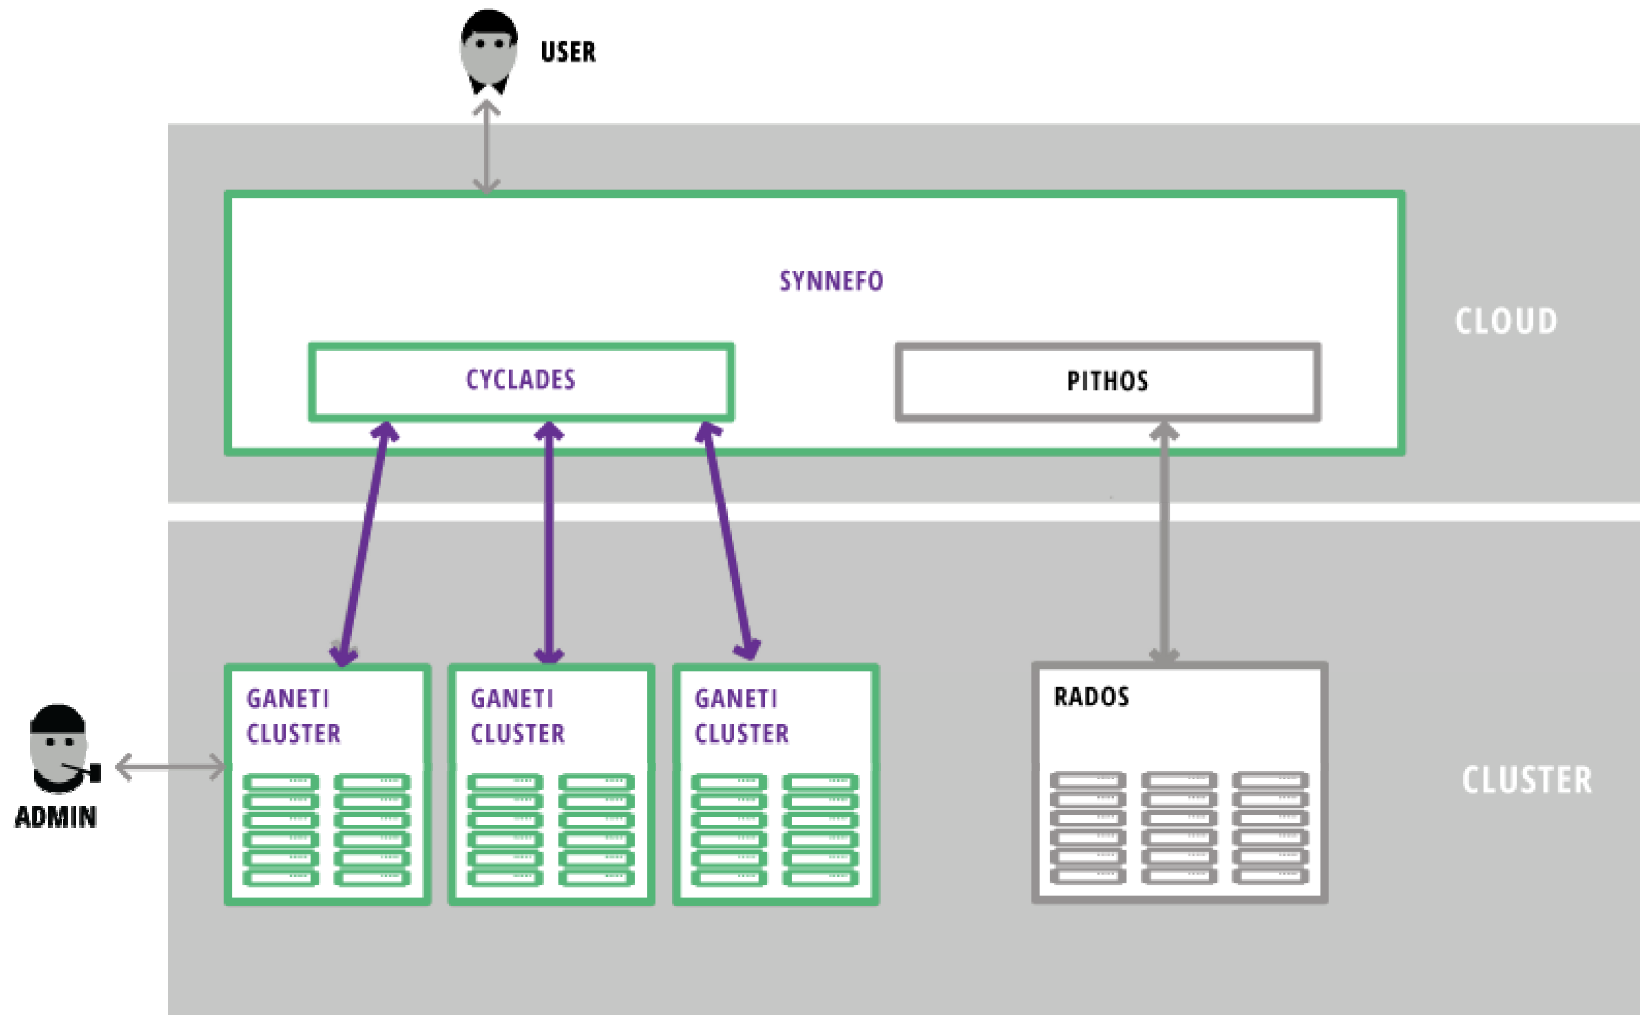
Thin Compute layer over Ganeti

- Python/Django
- Supports *multiple* Ganeti clusters, for scaling
- OpenStack Compute API

Networking

- No restrictions on deployment – it's the *Ganeti* side
- IPv4/IPv6 public networks, complete isolation among VMs
- Thousands of private networks, private L2 segments over single VLAN

Cyclades



Storage: Pithos

FOSDEM 2013

vkoukis@grnet.gr

Every file is a collection of blocks

Content-based addressing for blocks

Partial file transfers, deduplication, efficient syncing

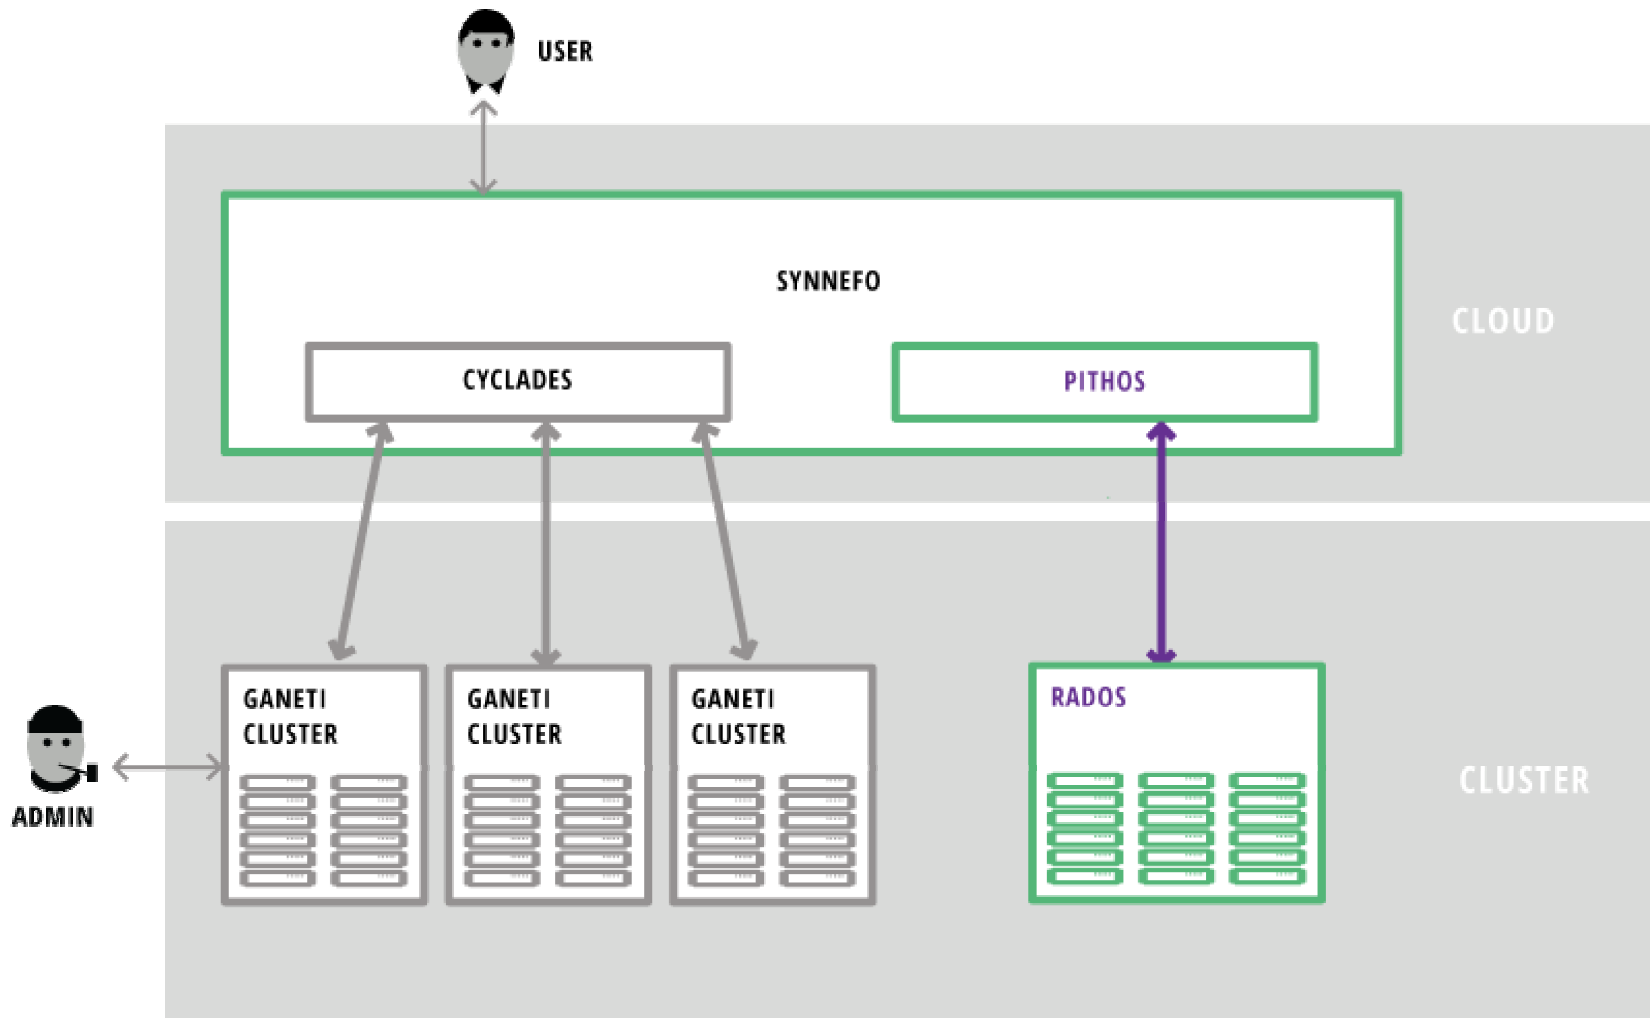
Independent of storage backend (NFS, RADOS, ...)


OpenStack Object Storage API plus extensions

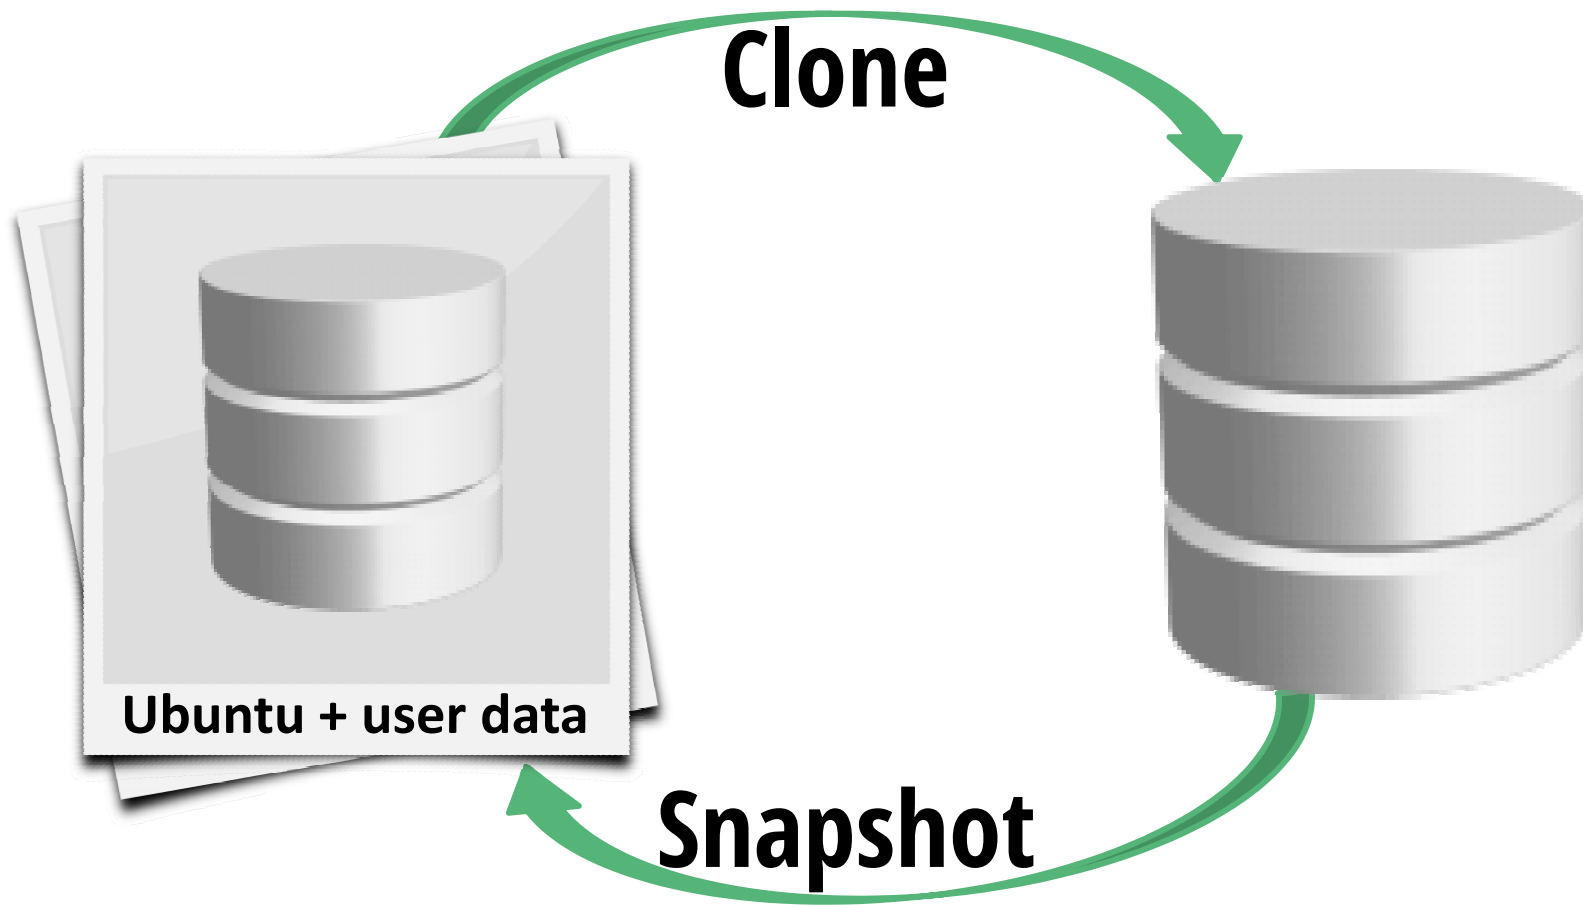
An integral part of Synnefo


- Single store for Files, VM Images
- Uses common backend with Archipelago


Pithos

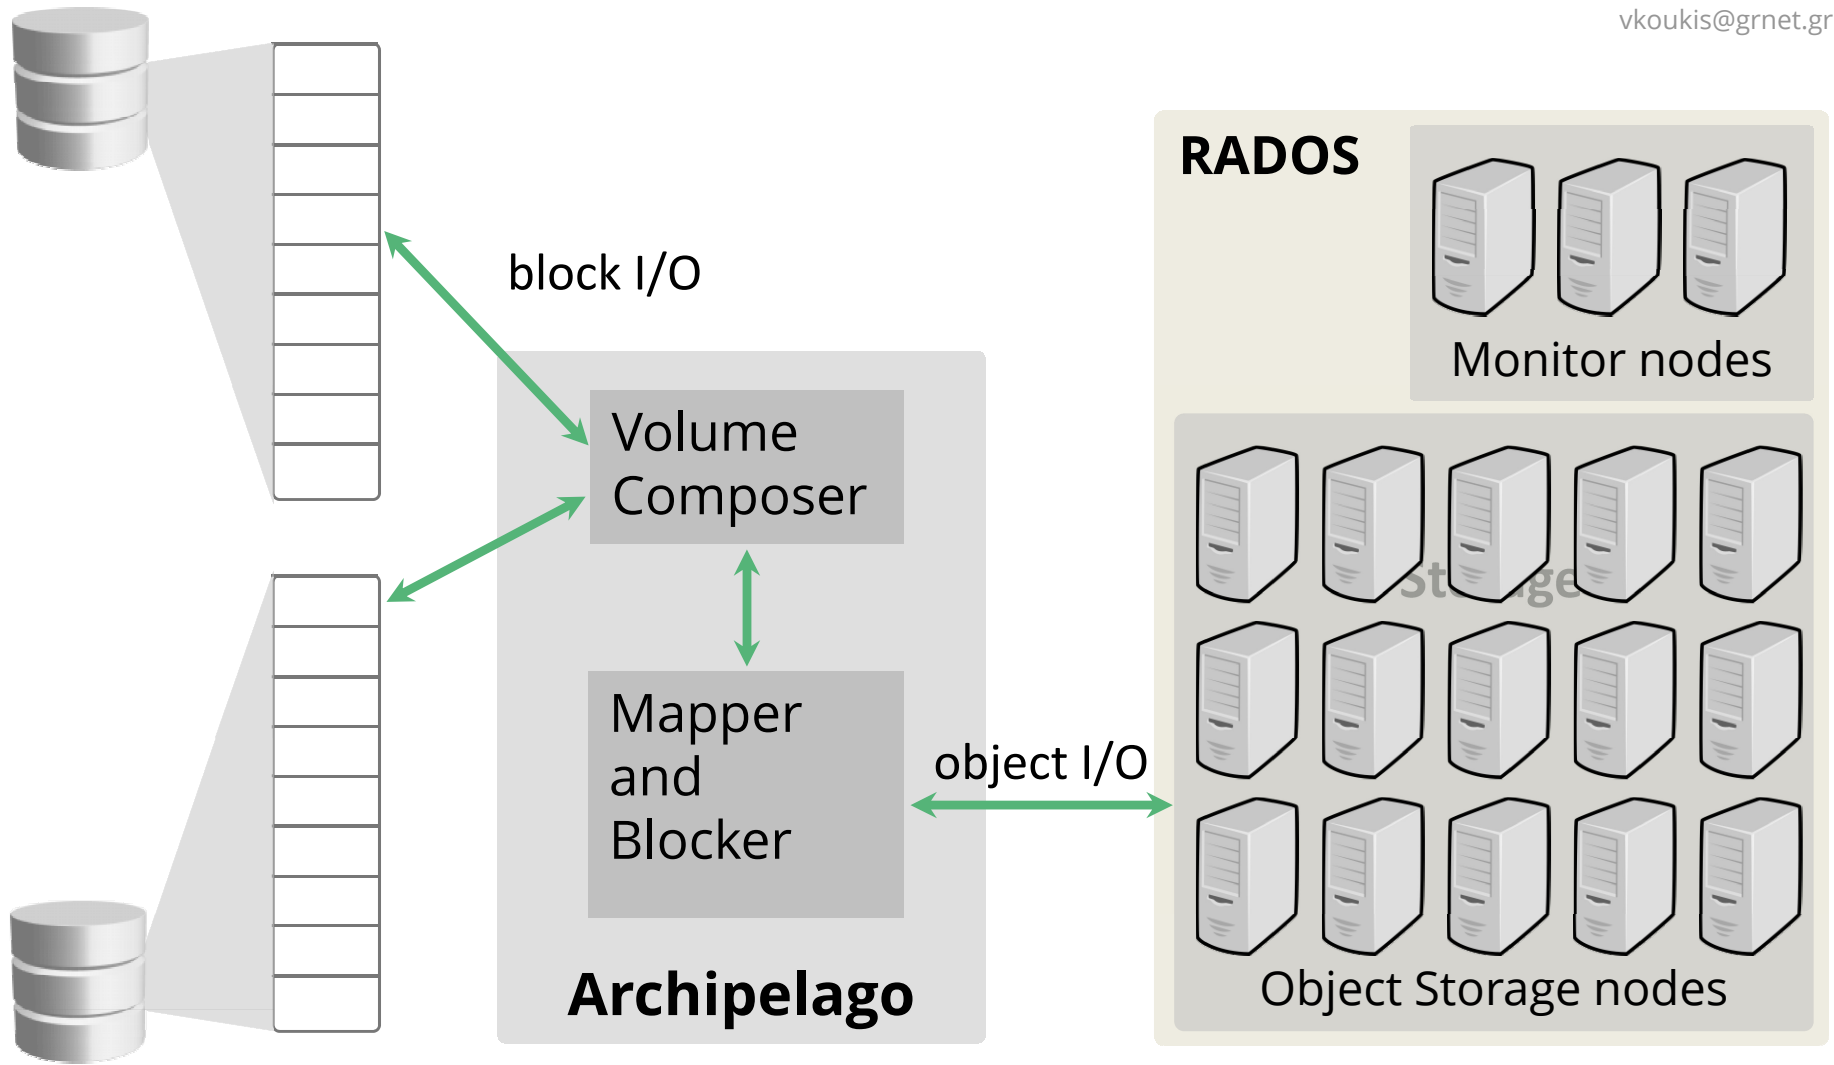












Volumes: Archipelago

FOSDEM 2013

vkoukis@grnet.gr

Unified storage for Files, Images ↔ Volumes

Thin layer over the actual storage cluster

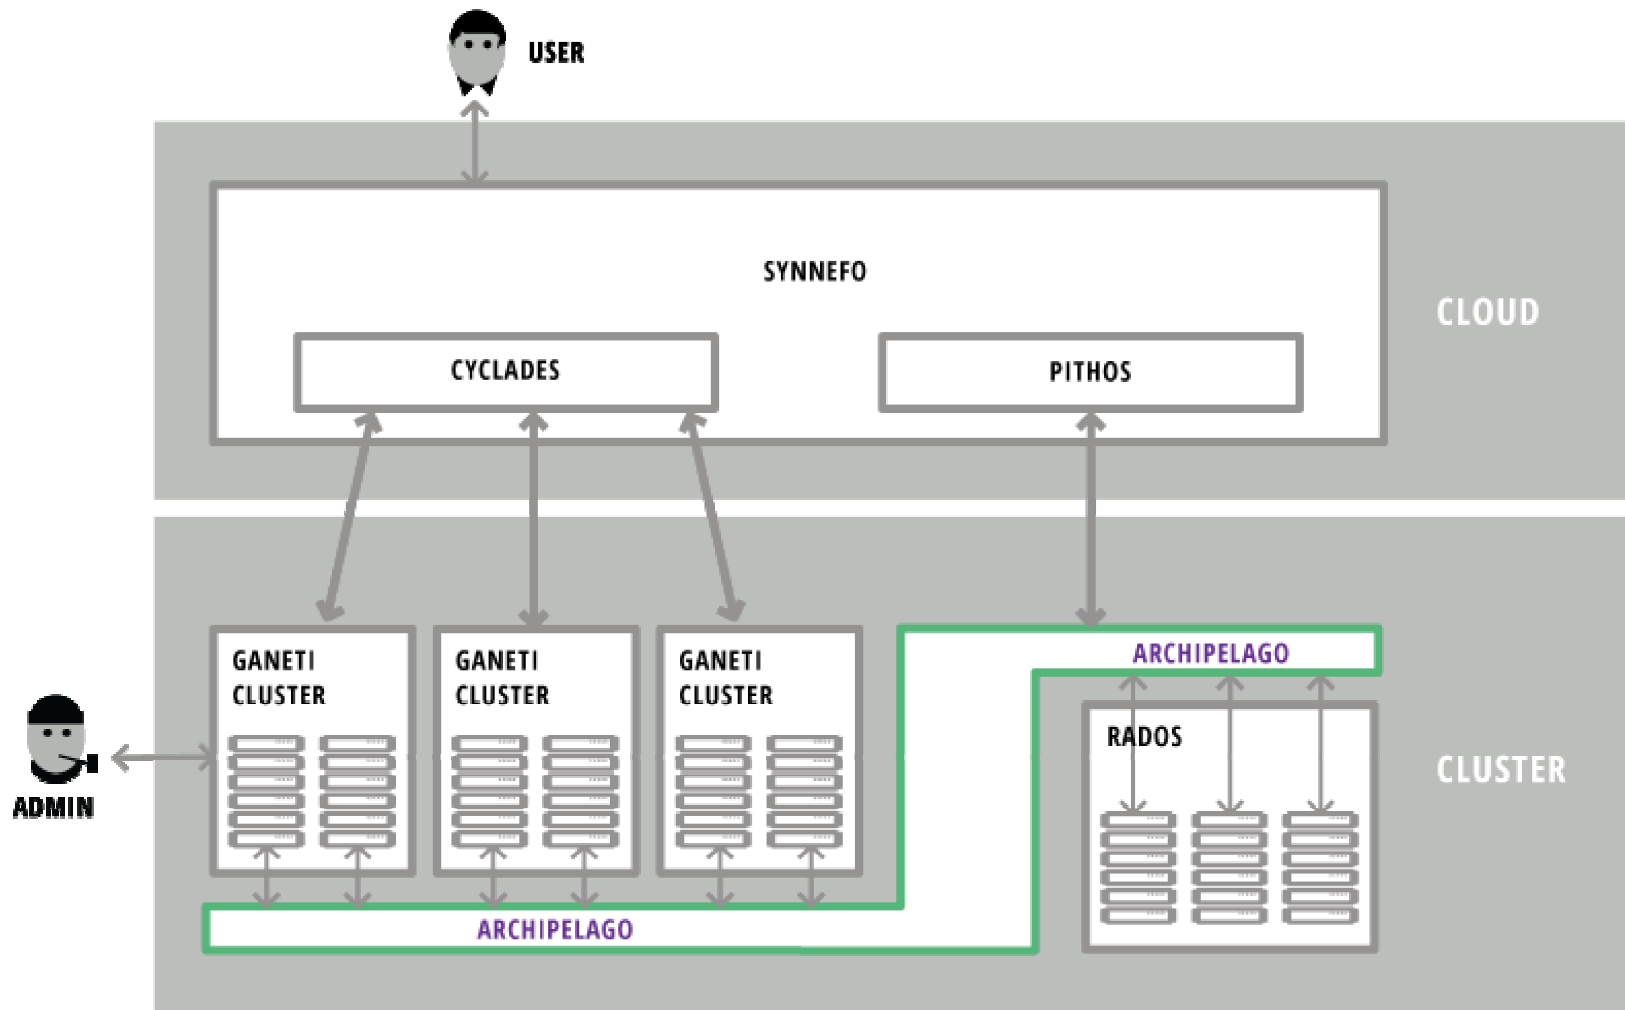
Storage backend agnostic

- NFS, RADOS, ...

Efficient syncing / sharing of Images as files on Pithos

Zero-copy cloning of volumes from Images

Archipelago



Current production ~okeanos

FOSDEM 2013

vkoukis@grnet.gr

Rolling hardware and software upgrades

- Node evacuations with live VM migrations
- Cross-datacenter move, Intel → AMD, no VM downtime
- IP renumbering of all VMs

> 2600 VMs, > 1800 users

Ready to scale in beta

- 7 Ganeti clusters currently, move to > 10.000 VMs

synnefo



Try it out!

FOSDEM 2013

vkoukis@grnet.gr

<http://www.synnefo.org>

<http://www.okeanos.io>