synnefo

grnet

SYNNEFO: AN INTRODUCTION AND UPDATE

CONSTANTINOS VENETSANOPOULOS,
PRINCIPAL CLOUD ENGINEER, GRNET

synnefo

grnet

Synnefo cloud platform

GanetiCon 2014
cven@grnet.gr

An all-in-one cloud solution

- Written from scratch in Python
- Manages multiple Google Ganeti clusters of VMs
- Uses Archipelago to unify all cloud storage resources
- Exposes the OpenStack APIs to end users

Live since 2011

- Came out of the ~okeanos public cloud service

synnefo

grnet

Synnefo cloud platform

GanetiCon 2014
cven@grnet.gr

A complete cloud platform

- Identity Service (Keystone API)
- Object Storage Service (Swift API)
- Compute Service (Nova API)
- Network Service (Neutron API)
- Image Service (Glance API)
- Volume Service (Cinder API)

synnefo

grnet

Unified view of storage resources

GanetiCon 2014
cven@grnet.gr

Files

- User files, with Dropbox-like syncing

Images

- Templates for VM creation

Volumes

- Live disks, as seen from VMs

Snapshots

- Point-in-time snapshots of Volumes

synnefo

grnet

Services Overview

GanetiCon 2014
cven@grnet.gr

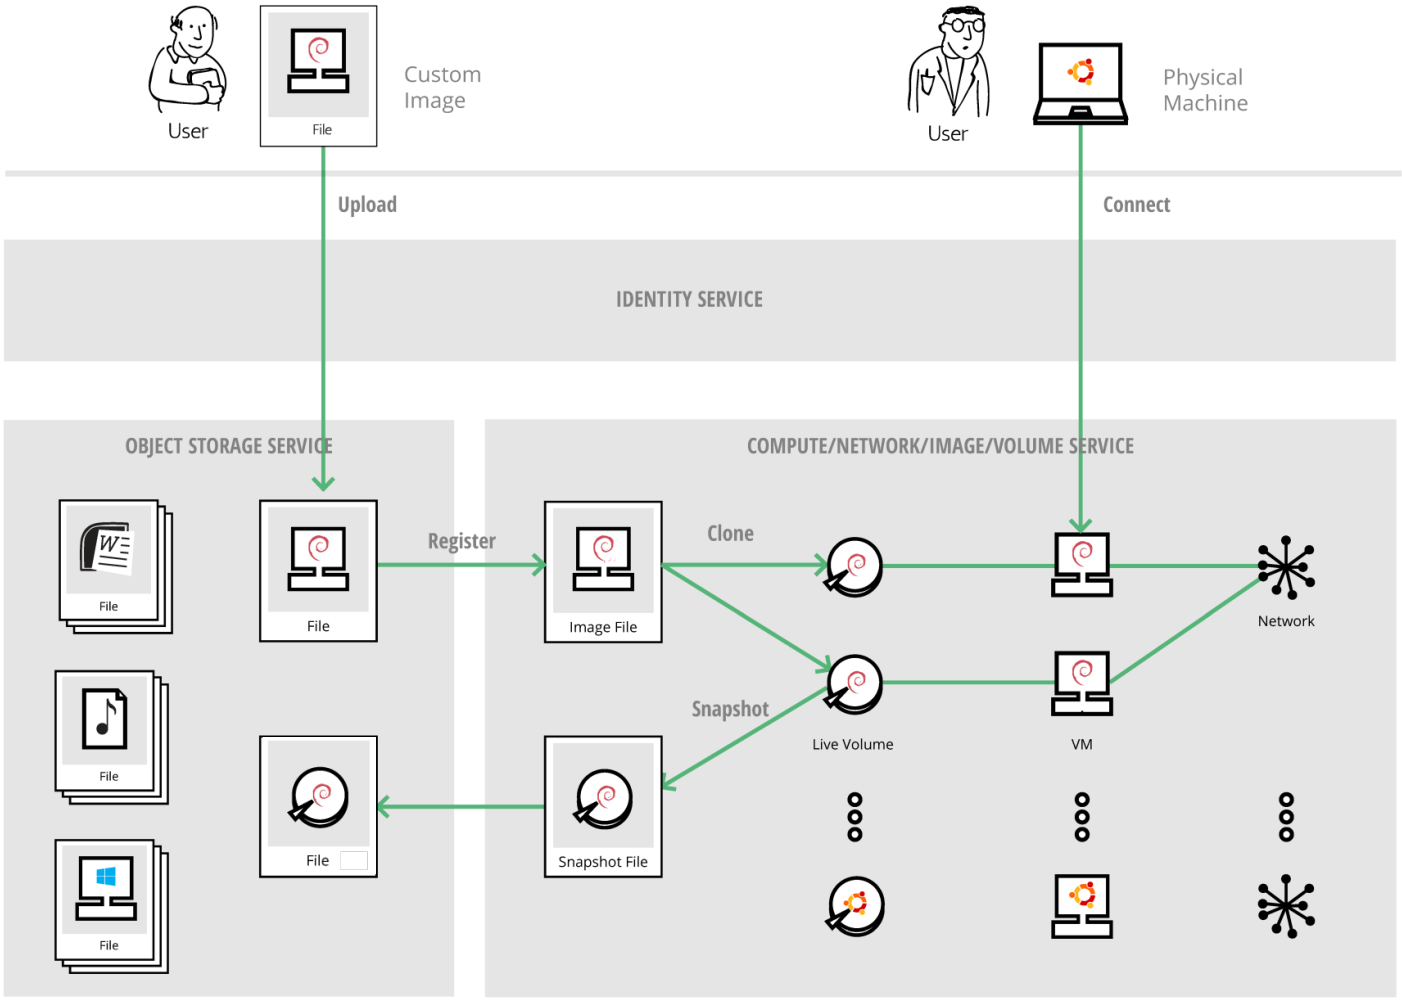

synnefo

grnet

Layered design

GanetiCon 2014
cven@grnet.gr

Client	OpenStack	Synnefo	UI	
vCloud	OpenStack	OpenStack	API	
vCloud	OpenStack	Synnefo	CLOUD	
vCenter	OpenStack	Ganeti	CLUSTER	
vSphere	libvirt	Ganeti	NODE	
ESXi	KVM / XEN	KVM / XEN	HYPERVISOR	

Architecture

GanetiCon 2014
cven@grnet.gr

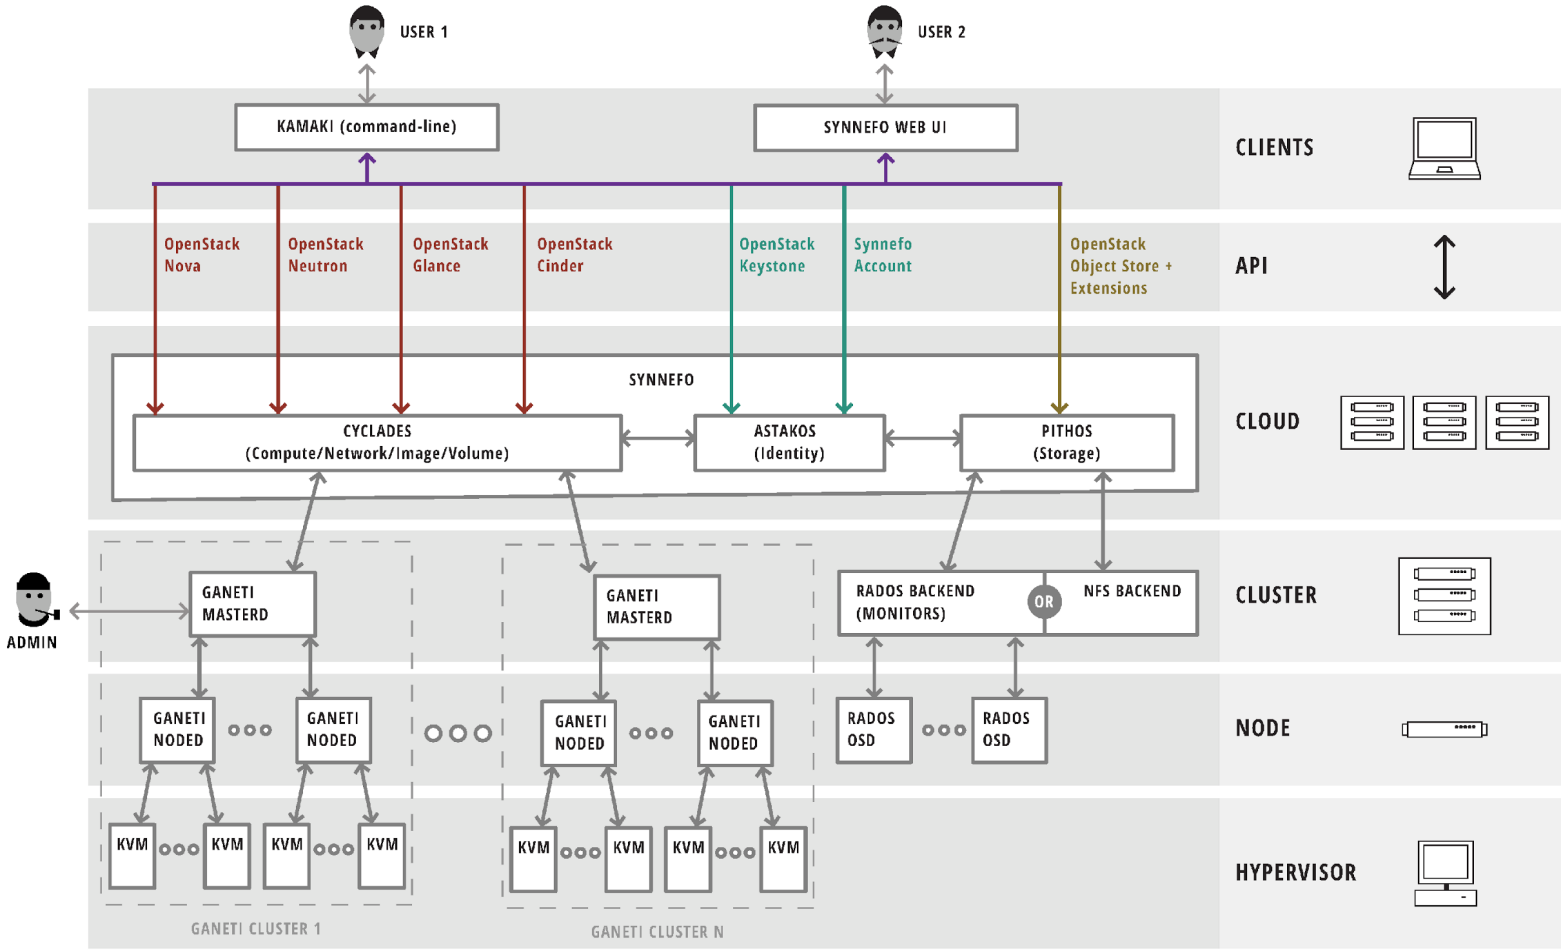

synnefo

grnet

Interaction with Ganeti

cven@grnet.gr

Support for all Ganeti storage templates including ExtStorage

OS Definition = snf-image

Networking = gnt-network +

snf-network (KVM ifup scripts) +

nfdhcpd (custom NFQUEUE-based DHCP server)

Asynchronous operation

- Effect path: Receive API requests from user, enqueue requests over RAPI to Ganeti
- Update path: Receive asynchronous notifications from Ganeti, update Synnefo DB, so the user can poll

synnefo

grnet

Archipelago

GanetiCon 2014
cven@grnet.gr

Storage Virtualization System

Decouples storage resources from storage backends

- Files / Images / Volumes / Snapshots

Unified way to provision, handle, and present resources

Decouples logic from actual physical storage

- Software-Defined Storage

synnefo

grnet

Interaction with Archipelago

GanetiCon 2014
cven@grnet.gr

A common storage substrate for Synnefo

Everything is a resource on Archipelago

The *same* resource is exposed as

- A File through the API of the Storage Service
- An Image through the API of the Image Service
- A live disk / VM Volume through the API of the Volume Service
- A Snapshot through the API of the Volume Service

All data remain in one place

No copying of data around

synnefo

grnet

Cloud Storage with Archipelago

GanetiCon 2014
cven@grnet.gr

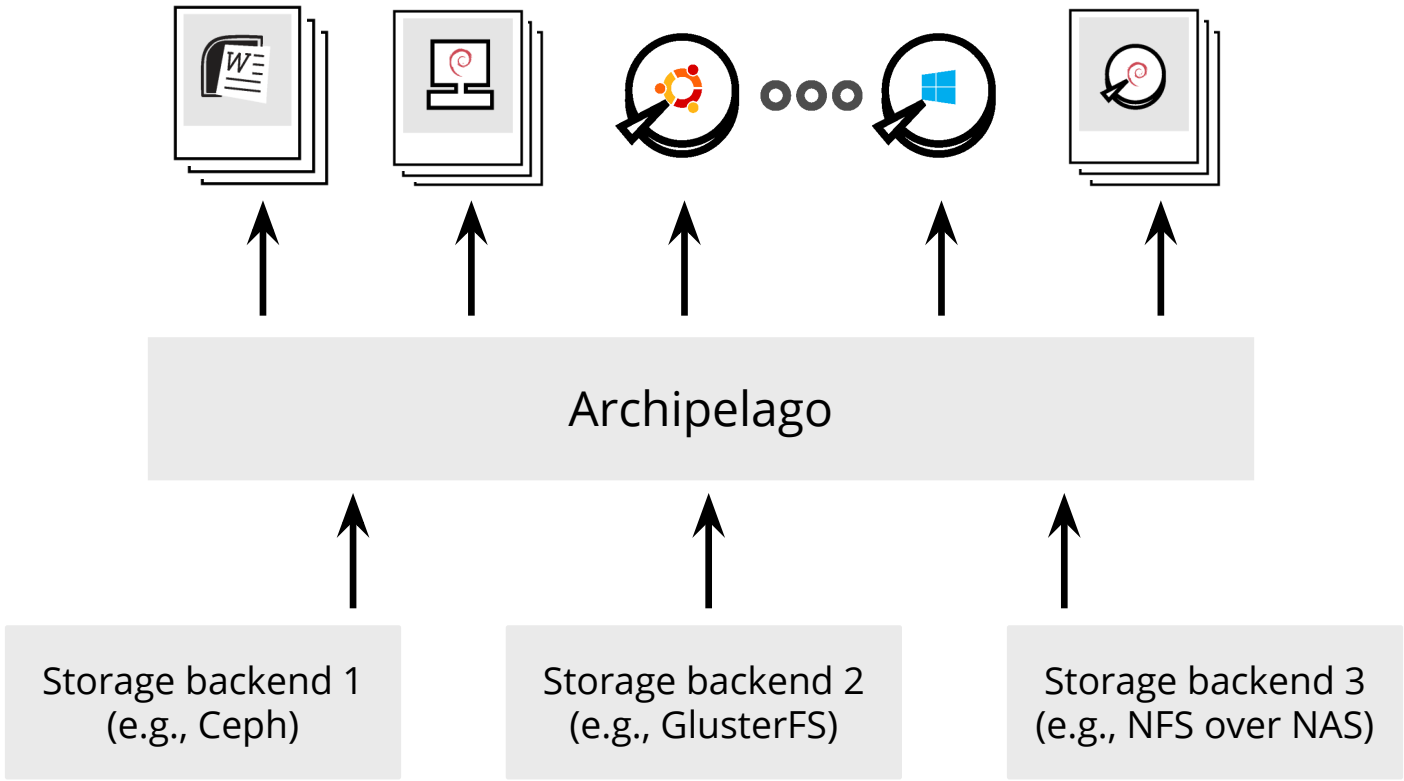

synnefo

grnet

Archipelago logic

GanetiCon 2014
cven@grnet.gr

Thin provisioning, with clones and snapshots

- Independent from the underlying storage technology

Hash-based data deduplication

Pluggable architecture

- Multiple endpoint (northbound) drivers
- Multiple backend (southbound) drivers

Multiple storage backends

- Unified management
- with storage migrations

synnefo

grnet

Archipelago interfaces

GanetiCon 2014
cven@grnet.gr

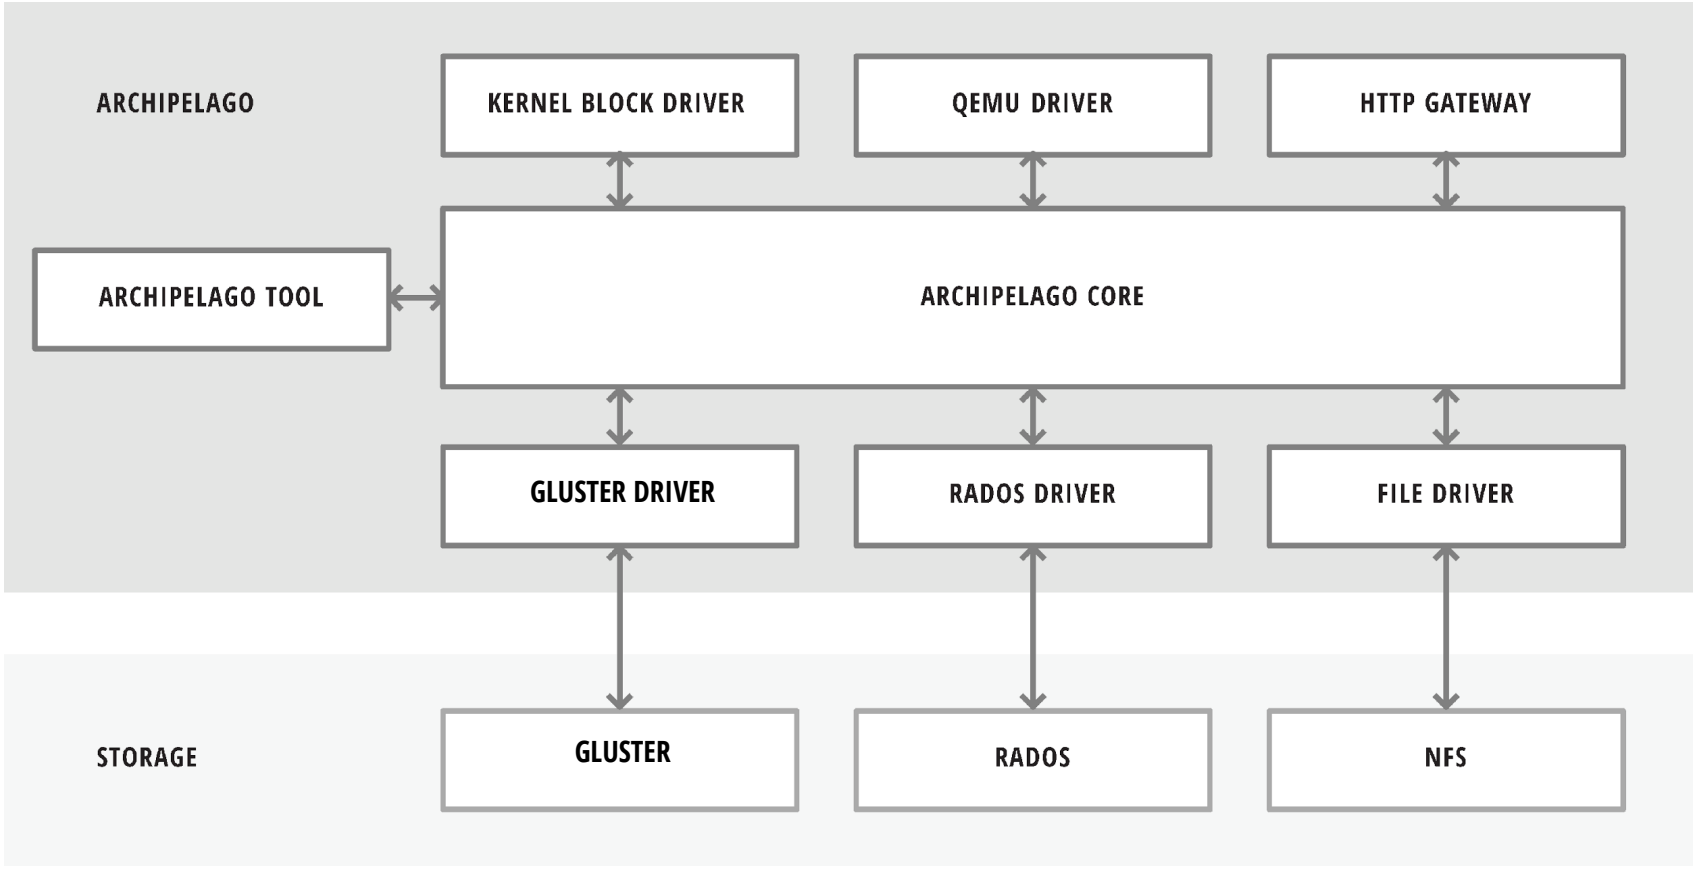

synnefo

grnet

Running Archipelago

GanetiCon 2014
cven@grnet.gr

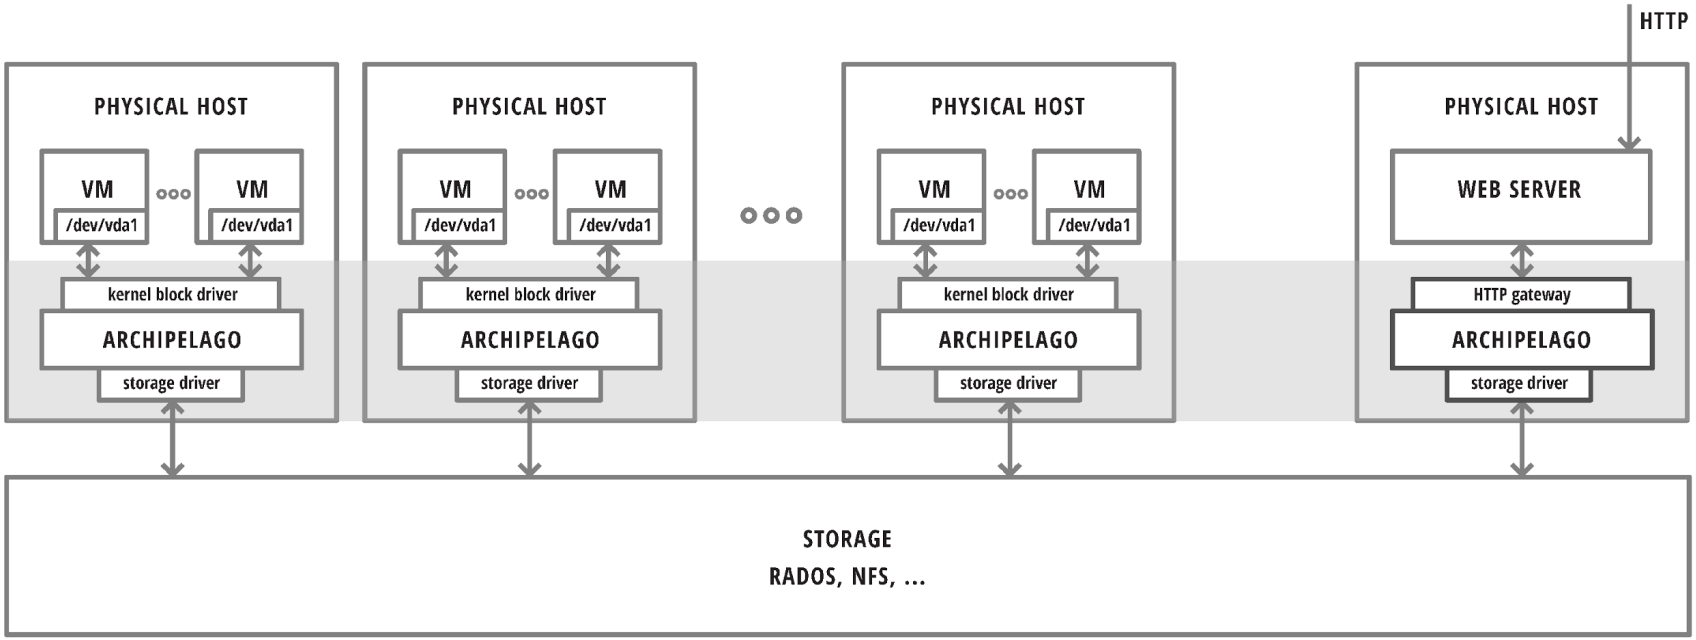

synnefo

grnet

Comparison to OpenStack?

GanetiCon 2014
cven@grnet.gr

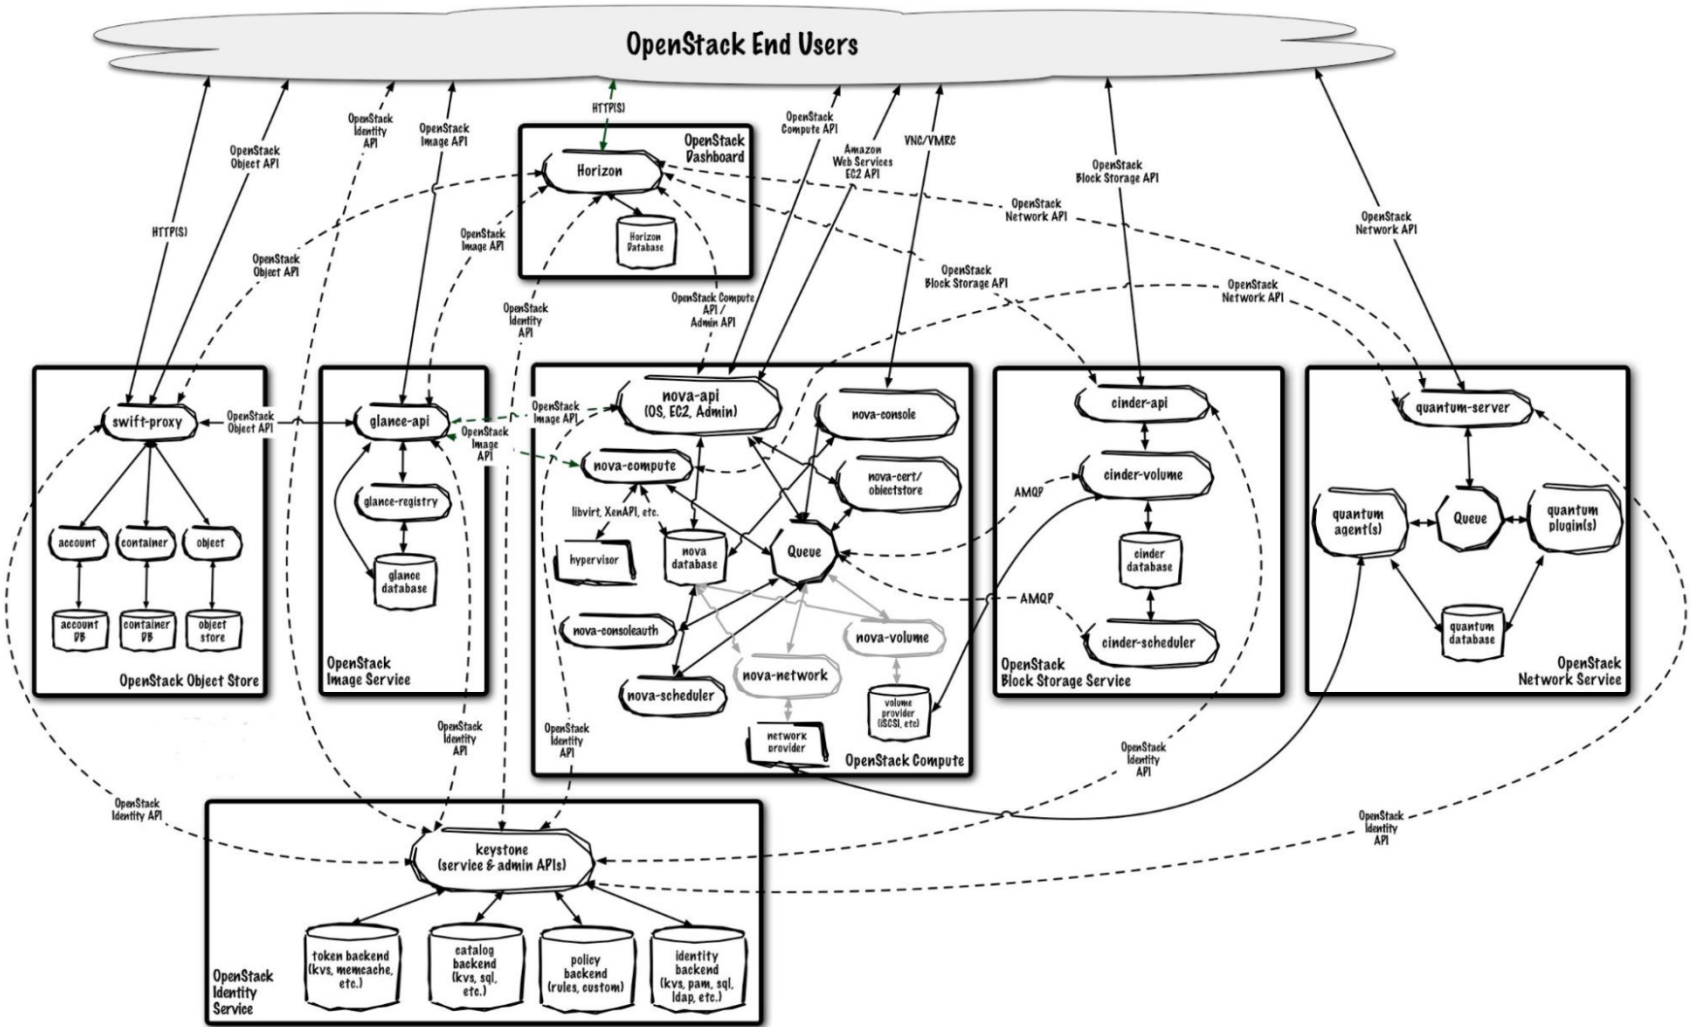

Synnefo

GanetiCon 2014
cven@grnet.gr

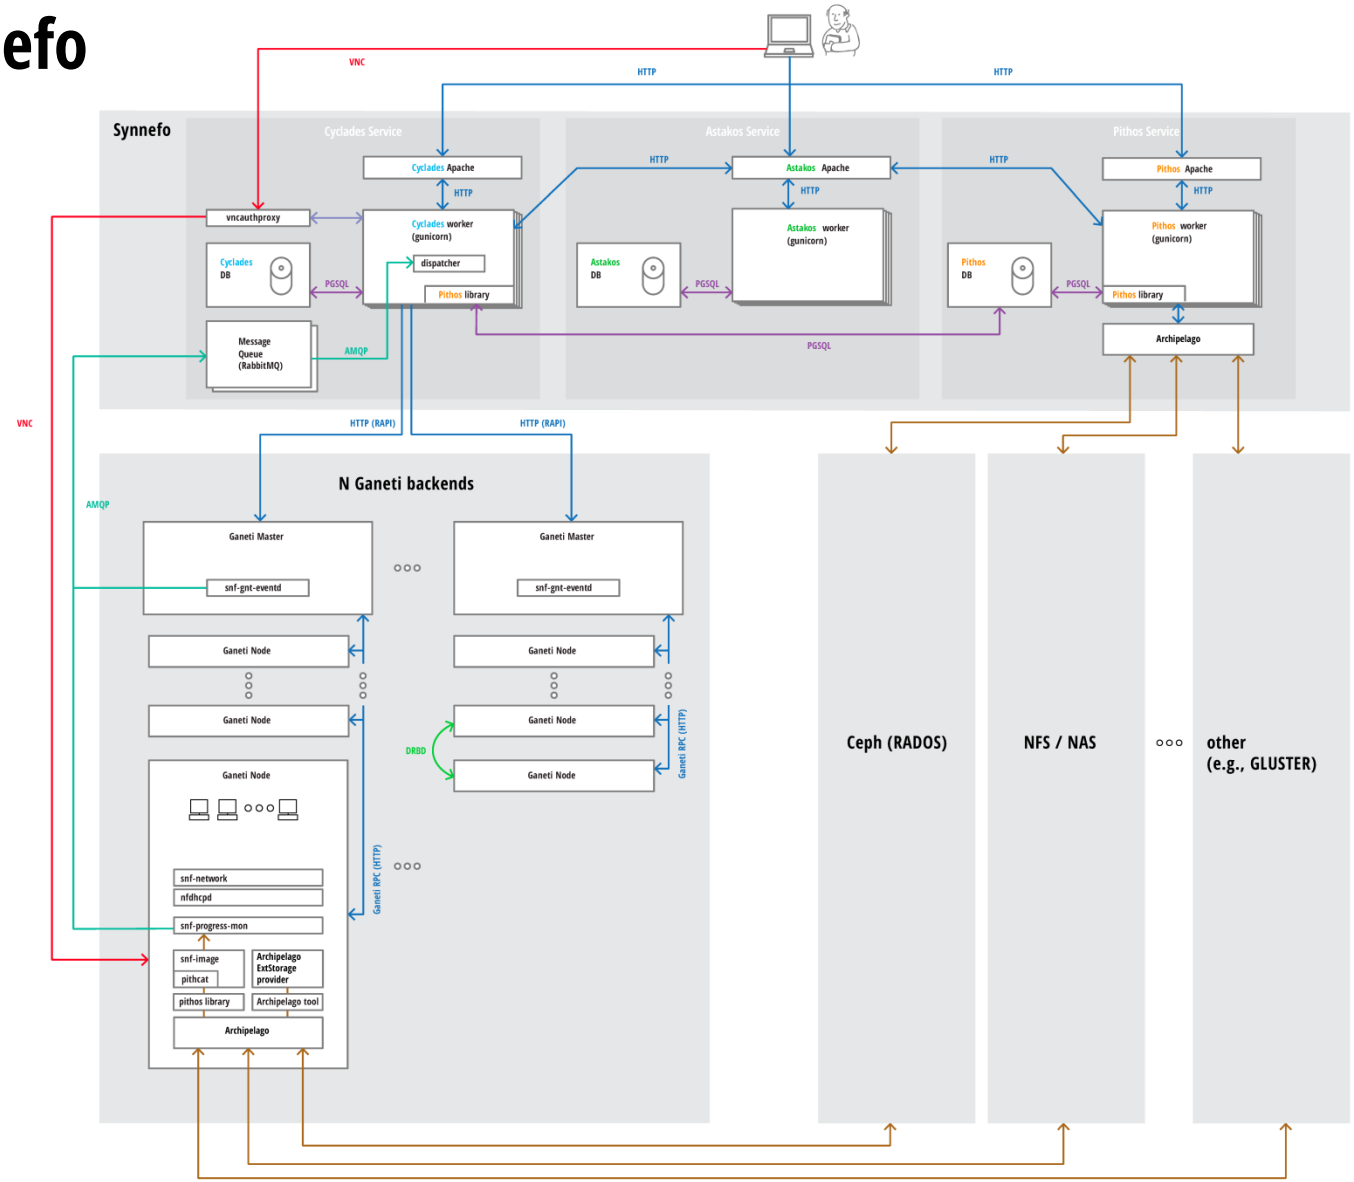

synnefo

grnet

Why Synnefo? A: Enterprise VMs at Cloud scale.

GanetiCon 2014
cven@grnet.gr

The best of both worlds

- Enterprise, persistent, stable VMs, live migrations (VMware-like)
  - Key technologies: Ganeti
- Over commodity hardware, no SAN needed
  - Key technologies: DRBD, Archipelago, Ceph
- at Cloud scale, accessible over Cloud APIs (OpenStack-like)
  - Key technologies: Synnefo

synnefo

grnet

Why Synnefo? B: Unified Cloud Storage.

GanetiCon 2014
cven@grnet.gr

Storage virtualization with Archipelago

- Common storage pool for everything
  - User files, Images (VM templates), live VM volumes, Snapshots
- Zero-copy thin cloning / snapshotting for super-fast provisioning
  - Over commodity hardware, no SAN needed
  - Less than 30 sec for a VM to be fully up and running
- Independent of the actual data store
- Pluggable storage: NFS/NAS, Ceph, Gluster, even SAN all at once
  - With inter-backend data moves

synnefo

grnet

Why Synnefo? C: Easier to run at scale.

GanetiCon 2014
cven@grnet.gr

Distinct management domains: Synnefo and Ganeti

- Management of self-contained Ganeti clusters
- Distinct Synnefo and Ganeti upgrade cycles
- Independent upgrades with no VM downtime

Limited access to DBs, decentralized VM state

- Only Synnefo workers need access to DBs
- No access from Ganeti nodes
  - Reduces impact of possible VM breakout
  - Boosts scalability to thousands of nodes
- Easier to firewall, easier to handle security-wise

synnefo

grnet

Why Synnefo? D: Survives failure.

GanetiCon 2014
cven@grnet.gr

Physical node management

- Dynamically add/remove/drain/set offline physical nodes
- Dynamically add/remove/drain/rebalance whole Ganeti clusters
- Evacuate failing nodes with live VM migrations, no VM downtime

Recovery from failure

- Built-in reconciliation mechanisms
- Able to recover from Synnefo/Ganeti downtime
  - Ensures in-sync state across components

Easier to contain failure

- Outages contained inside smaller domains
  - inside a node, or inside a Ganeti cluster

synnefo

grnet

What's new in upcoming Synnefo v0.16

GanetiCon 2014
cven@grnet.gr

Admin Dashboard

Implementation of Volumes + Cinder API

Add/Remove Volumes in running VMs (hotplug)

Revamped Projects

Snapshots

Archipelago becomes the Pithos backend

Coming in v0.17:

- Generic ACL mechanism for all Synnefo objects
- New settings mechanism

synnefo

grnet

The ~okeanos use case @ GRNET

GanetiCon 2014
cven@grnet.gr

Live since July 2011

Numbers

- Users: > 10000
- VMs: > 10000 currently active
- More than 380k VMs spawned so far, more than 110k networks

Physical Infrastructure

- 13 Ganeti Clusters, spanning a whole DC
- 1PB of raw storage capacity

synnefo

grnet

Try it out!

GanetiCon 2014
cven@grnet.gr

http://www.synnefo.org

synnefo

grnet

Thank you!

GanetiCon 2014
cven@grnet.gr

synnefo

grnet

Screenshots.

GanetiCon 2014
cven@grnet.gr

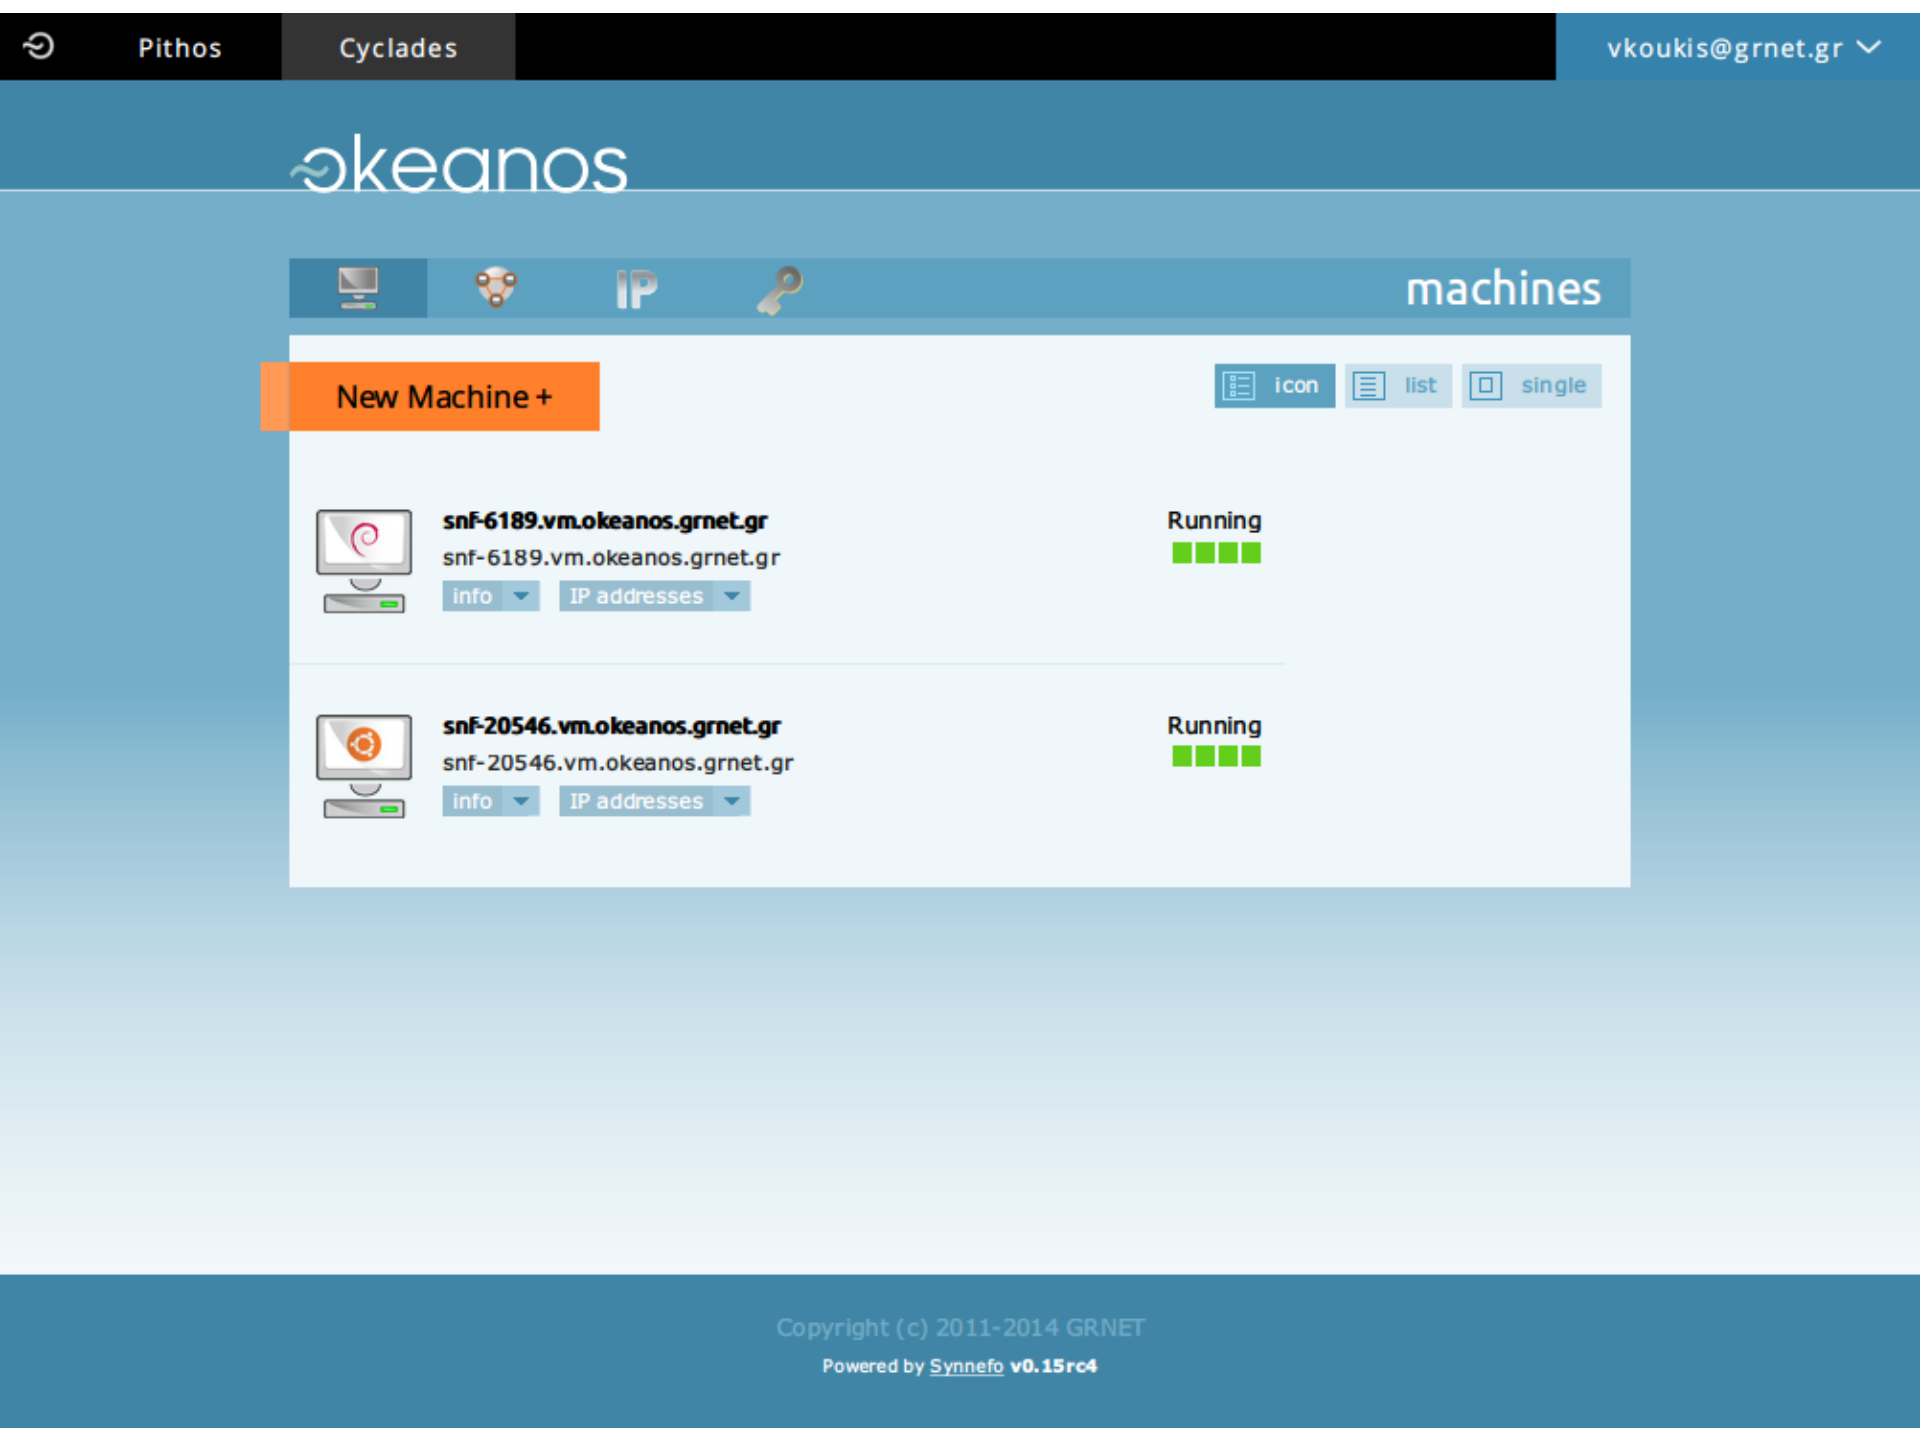
Pithos
Cyclades
vkoukis@grnet.gr
okeanos
IP
machines
New Machine +
icon
list
single
snf-6189.vm.okeanos.grnet.gr
snf-6189.vm.okeanos.grnet.gr
info
IP addresses
Running
snf-20546.vm.okeanos.grnet.gr
snf-20546.vm.okeanos.grnet.gr
info
IP addresses
Running
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

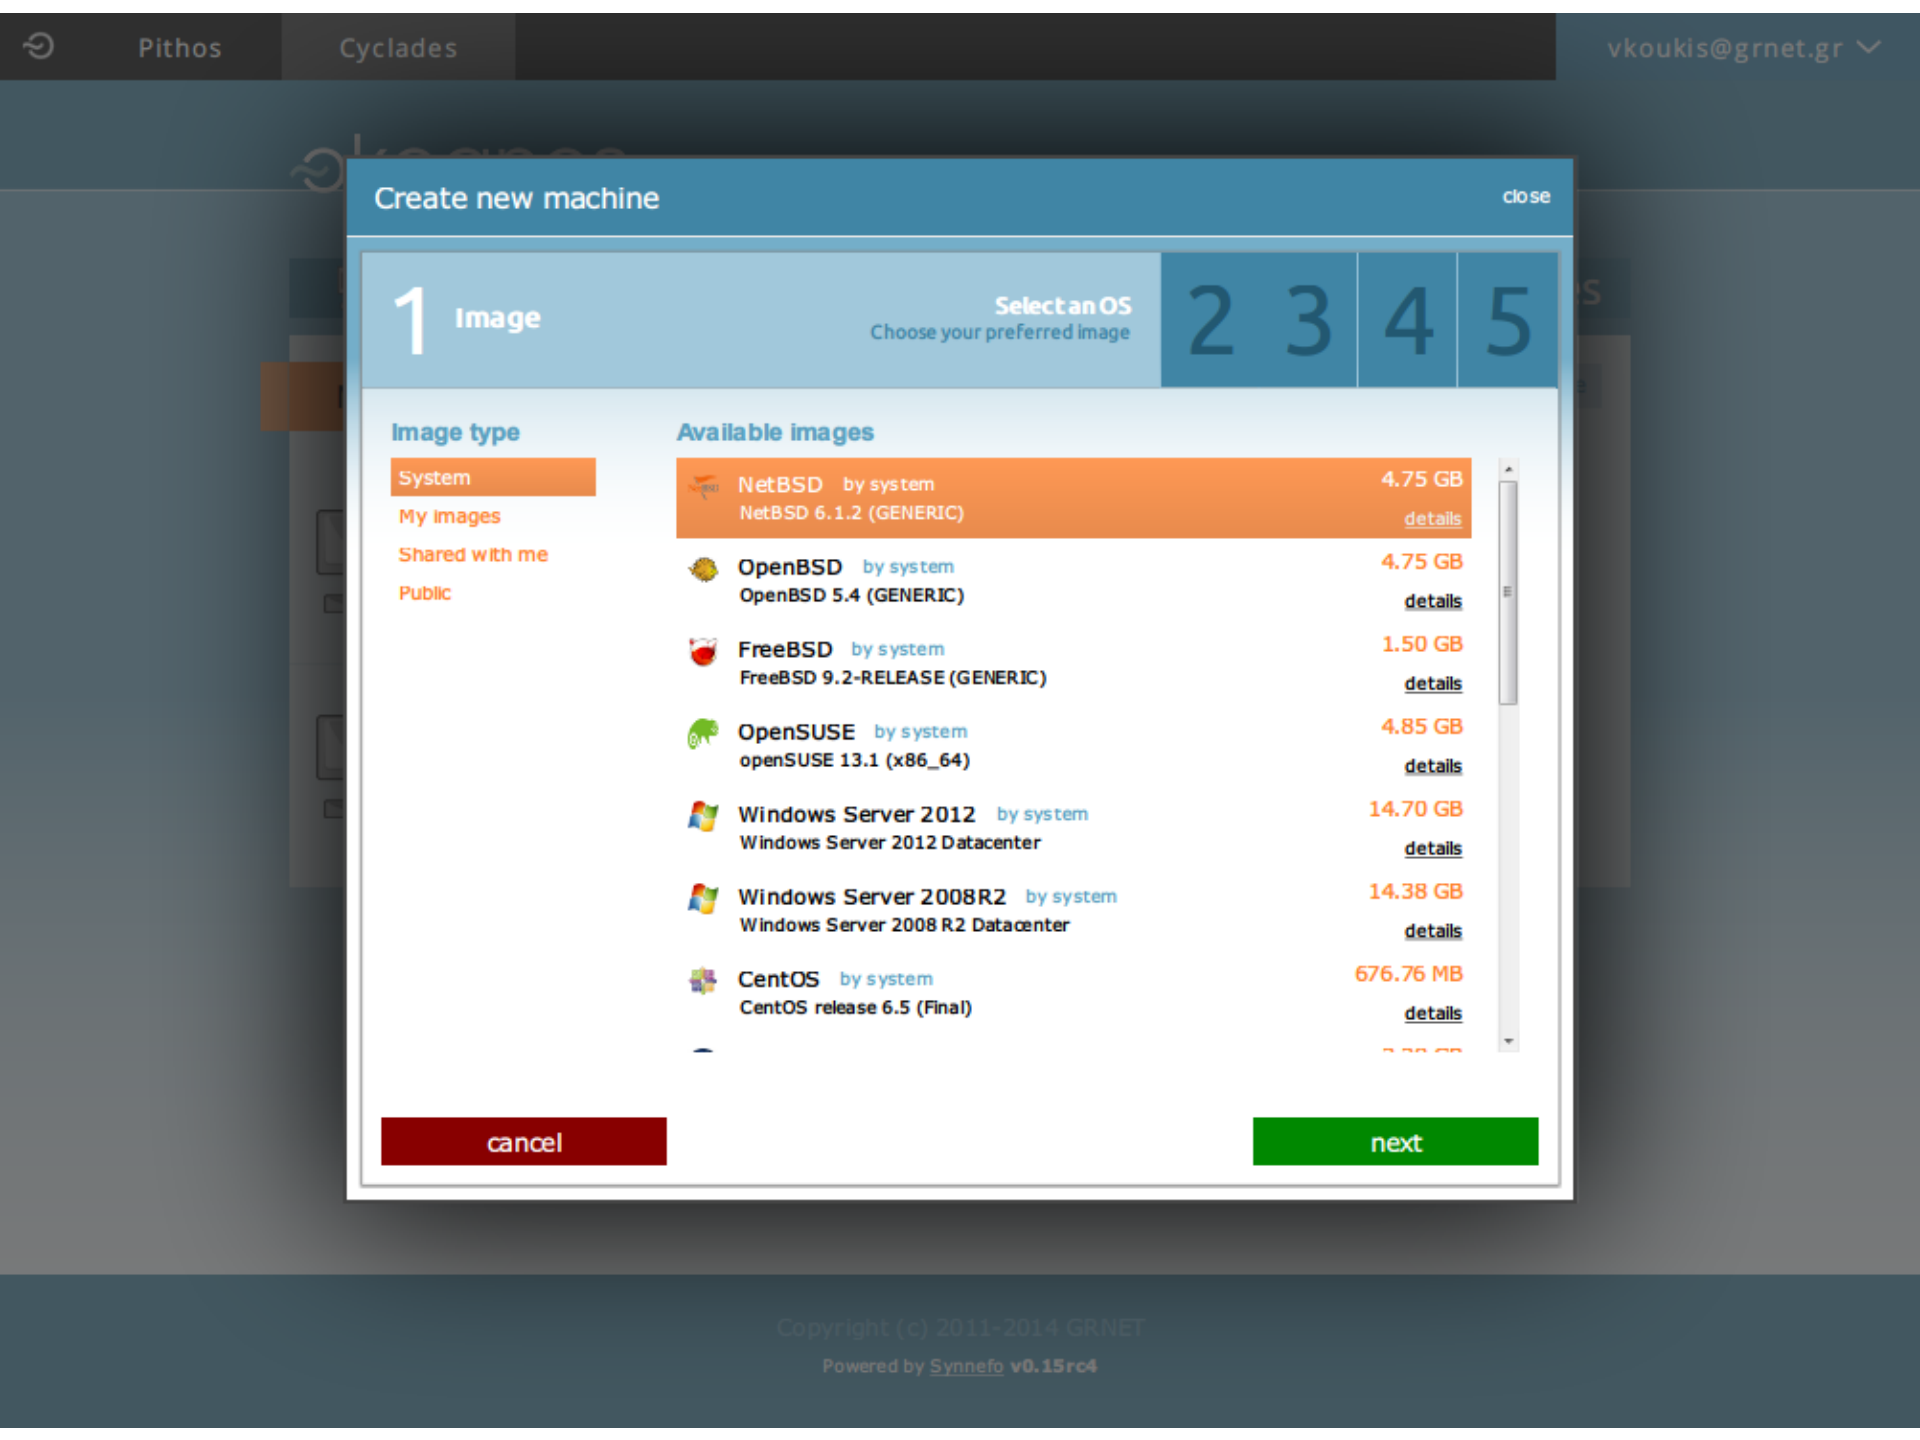
Pithos
Cyclades
vkoukis@grnet.gr
Create new machine
close
1 Image
Select an OS
Choose your preferred image
2 3 4 5
Image type
System
My images
Shared with me
Public
Available images
NetBSD by system
NetBSD 6.1.2 (GENERIC)
4.75 GB
details
OpenBSD by system
OpenBSD 5.4 (GENERIC)
4.75 GB
details
FreeBSD by system
FreeBSD 9.2-RELEASE (GENERIC)
1.50 GB
details
OpenSUSE by system
openSUSE 13.1 (x86_64)
4.85 GB
details
Windows Server 2012 by system
Windows Server 2012 Datacenter
14.70 GB
details
Windows Server 2008R2 by system
Windows Server 2008 R2 Datacenter
14.38 GB
details
CentOS by system
CentOS release 6.5 (Final)
676.76 MB
details
cancel
next
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

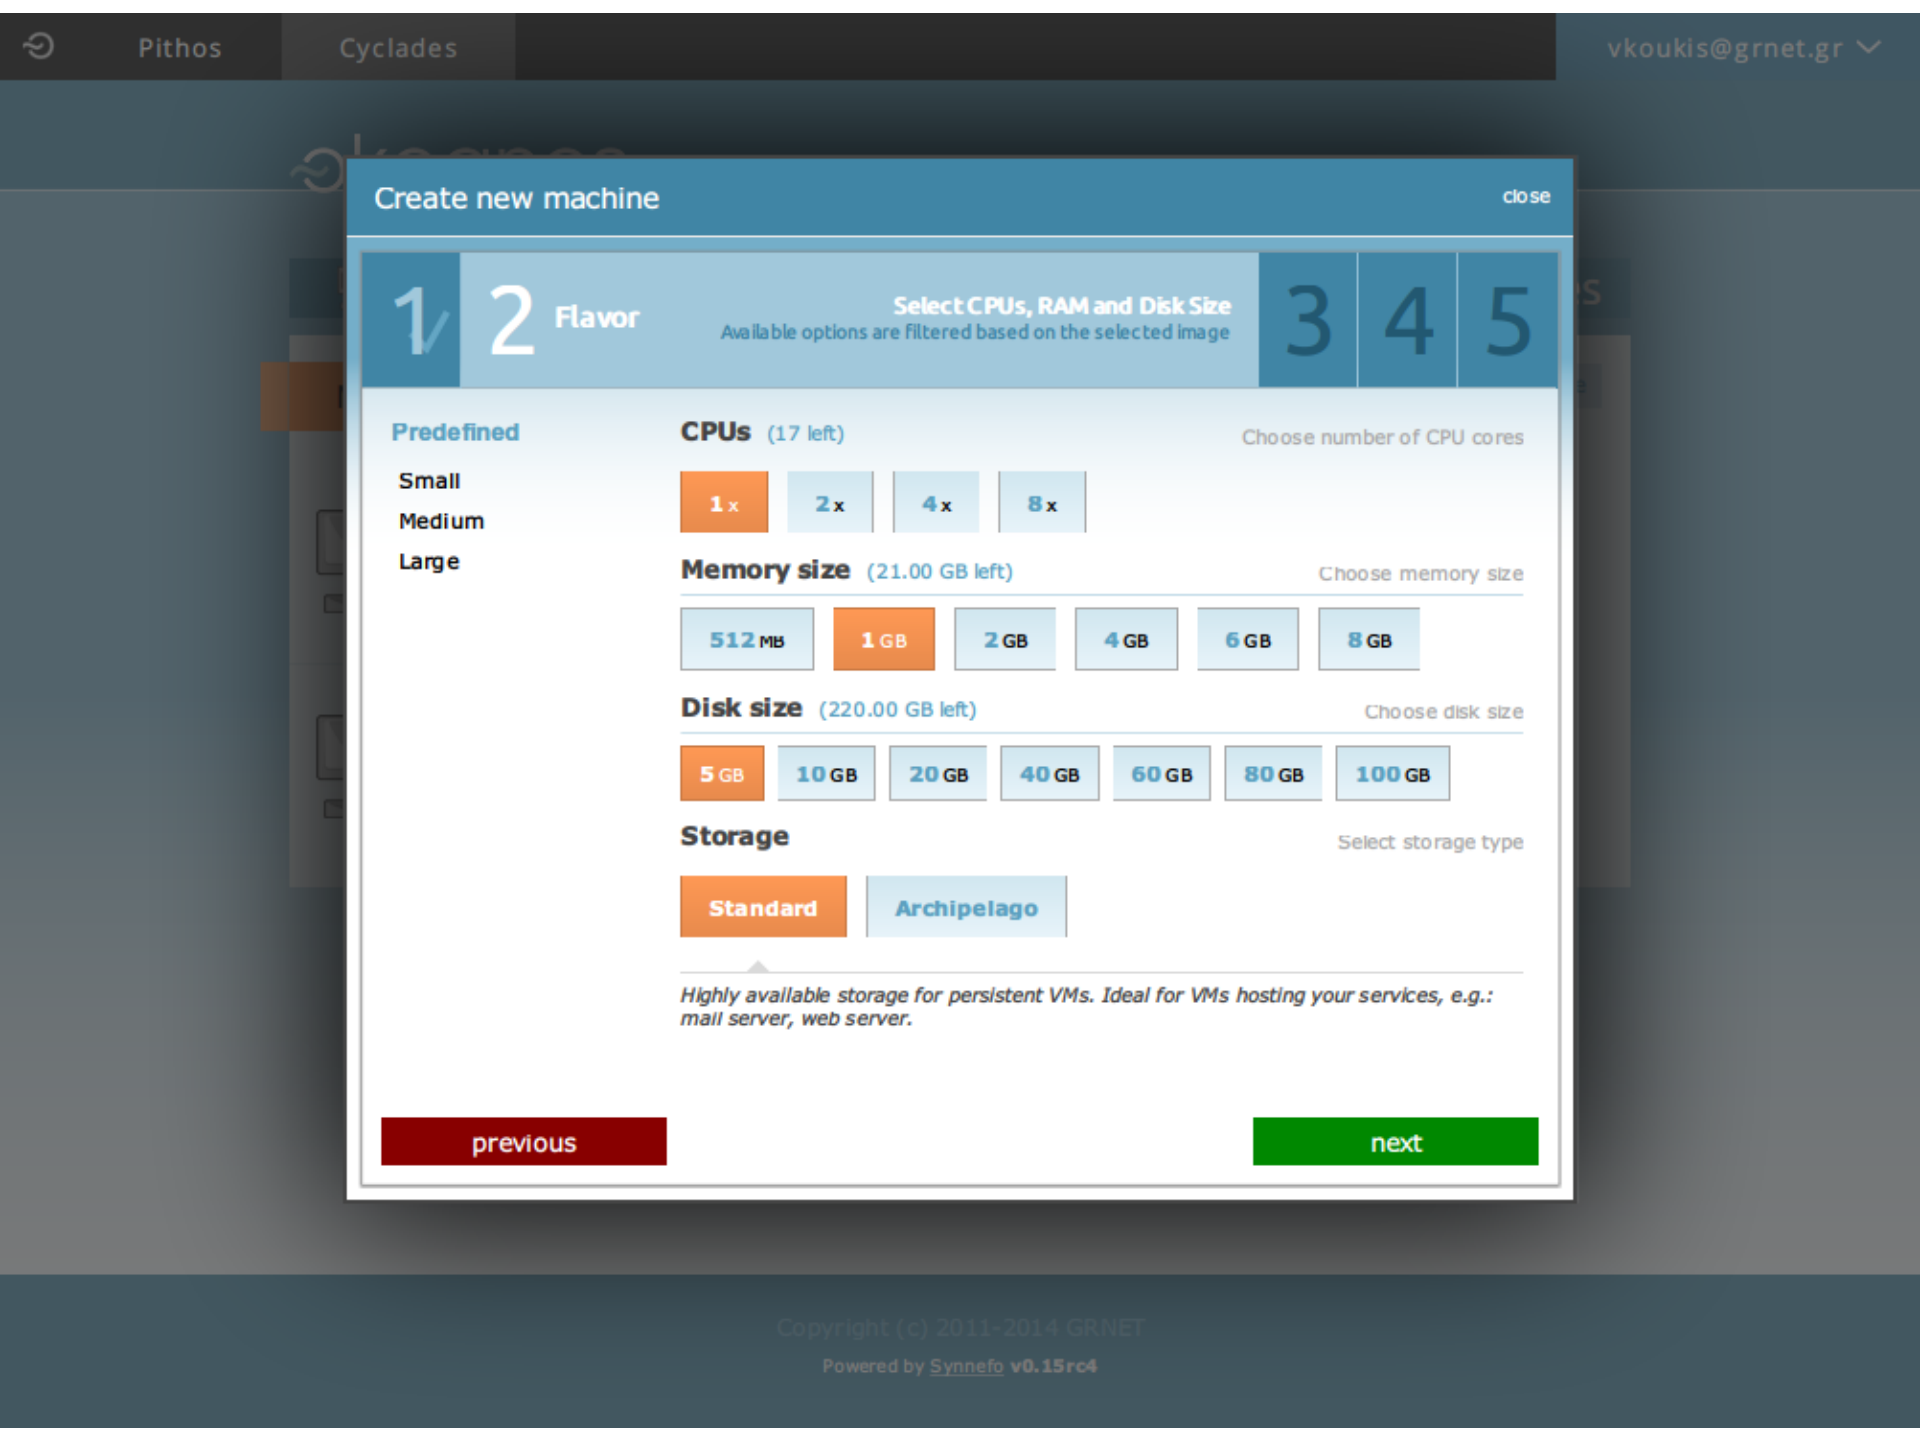
Pithos
Cyclades
vkoukis@grnet.gr
Create new machine
close
1
2 Flavor
Select CPUs, RAM and Disk Size
Available options are filtered based on the selected image
3
4
5
Predefined
Small
Medium
Large
CPUs (17 left)
Choose number of CPU cores
1 x
2 x
4 x
8 x
Memory size (21.00 GB left)
Choose memory size
512 MB
1 GB
2 GB
4 GB
6 GB
8 GB
Disk size (220.00 GB left)
Choose disk size
5 GB
10 GB
20 GB
40 GB
60 GB
80 GB
100 GB
Storage
Select storage type
Standard
Archipelago
Highly available storage for persistent VMs. Ideal for VMs hosting your services, e.g.: mail server, web server.
previous
next
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

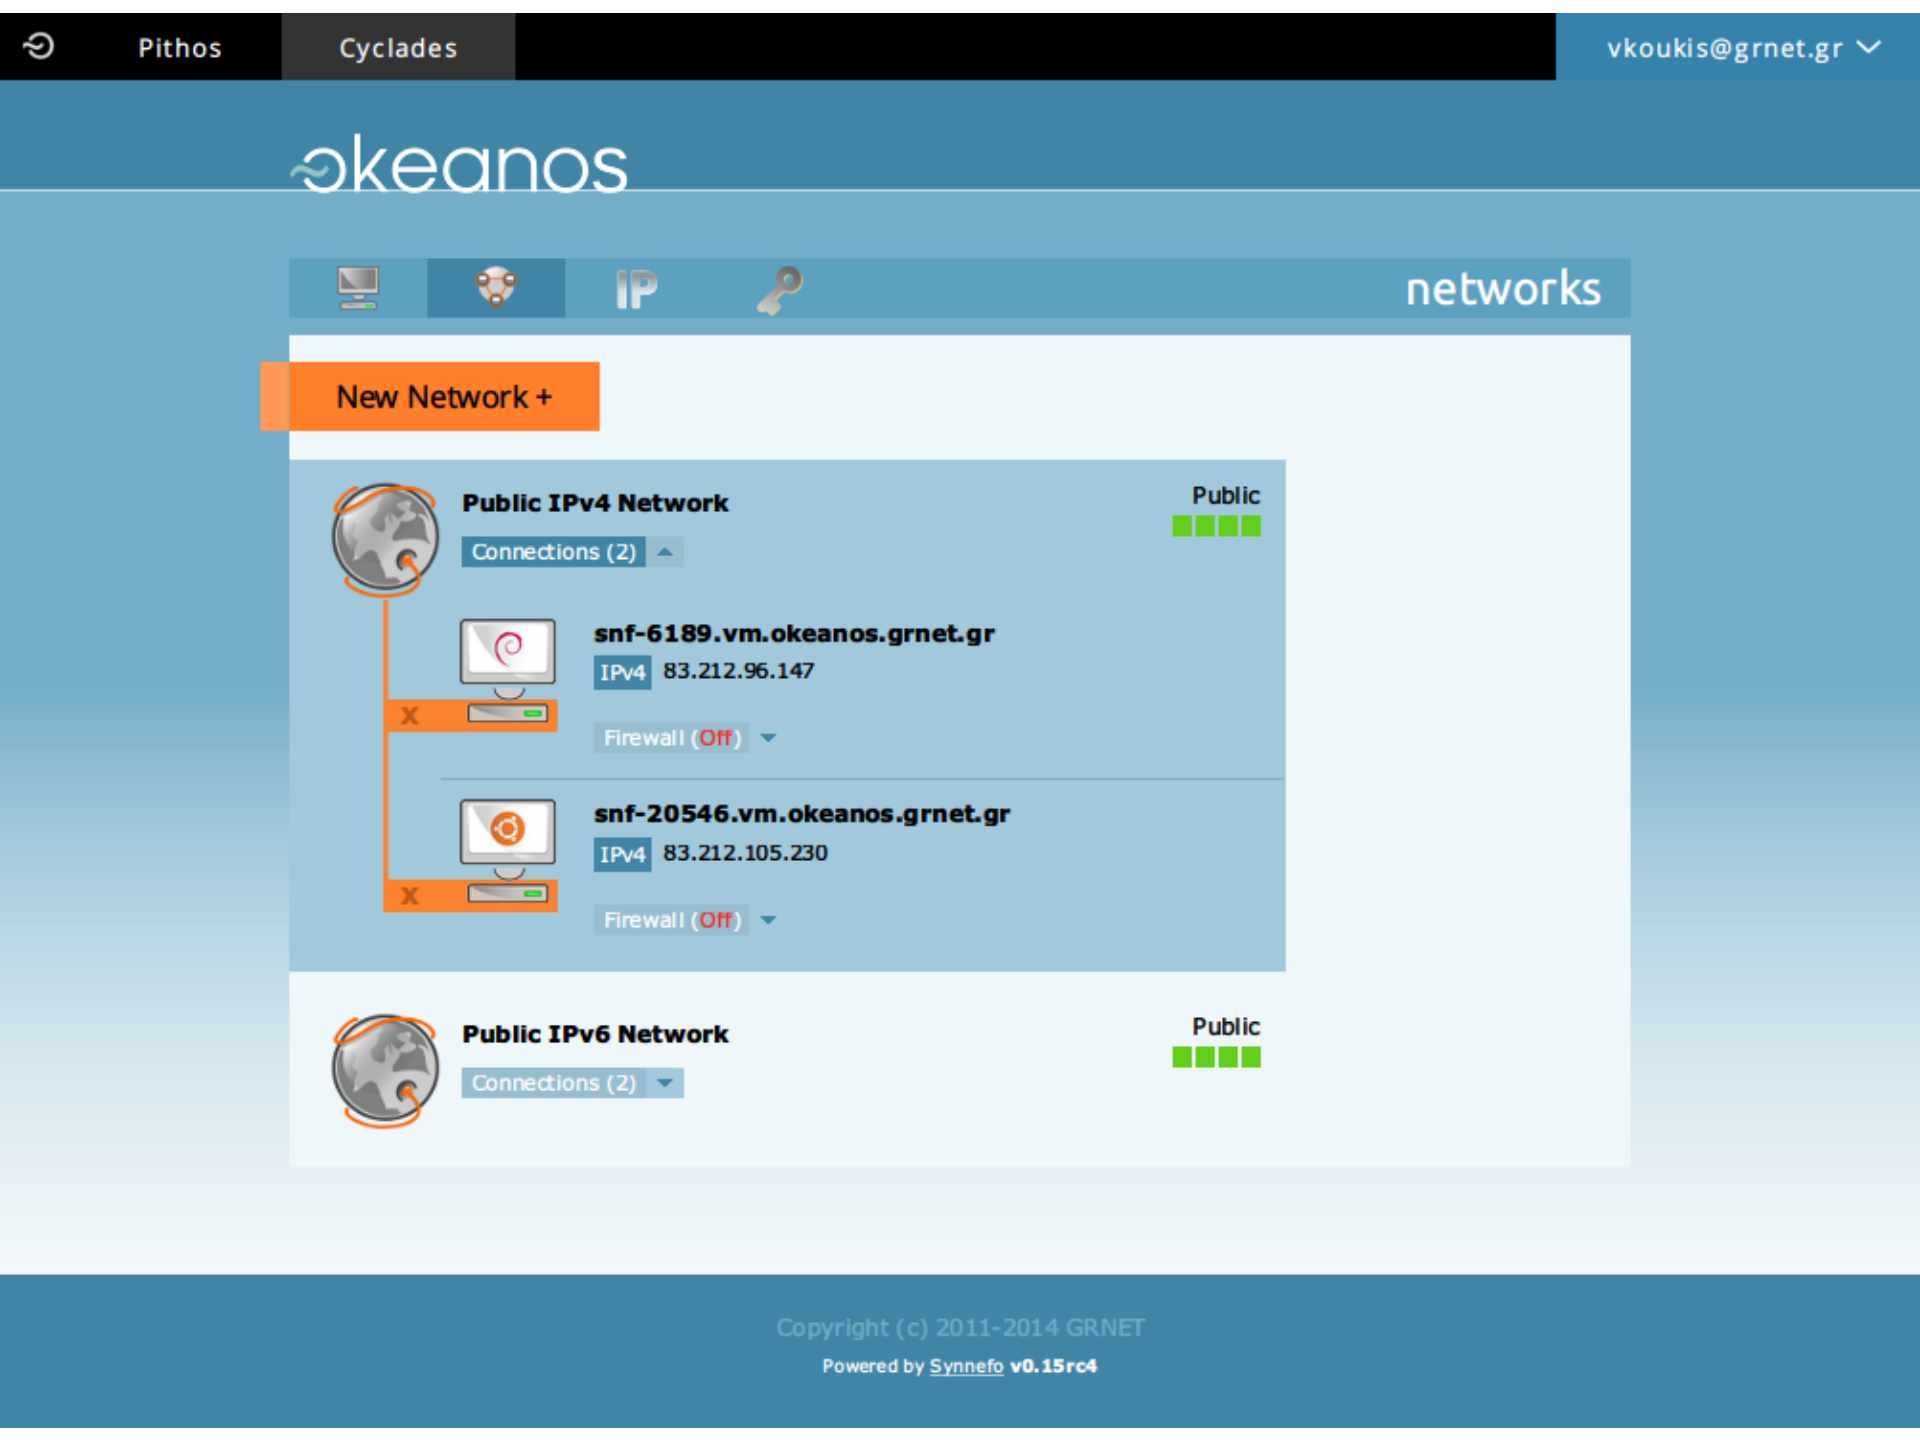
Pithos
Cyclades
vkoukis@grnet.gr
okeanos
networks
New Network +
Public IPv4 Network
Public
Connections (2)
snf-6189.vm.okeanos.grnet.gr
IPv4 83.212.96.147
Firewall (Off)
snf-20546.vm.okeanos.grnet.gr
IPv4 83.212.105.230
Firewall (Off)
Public IPv6 Network
Public
Connections (2)
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

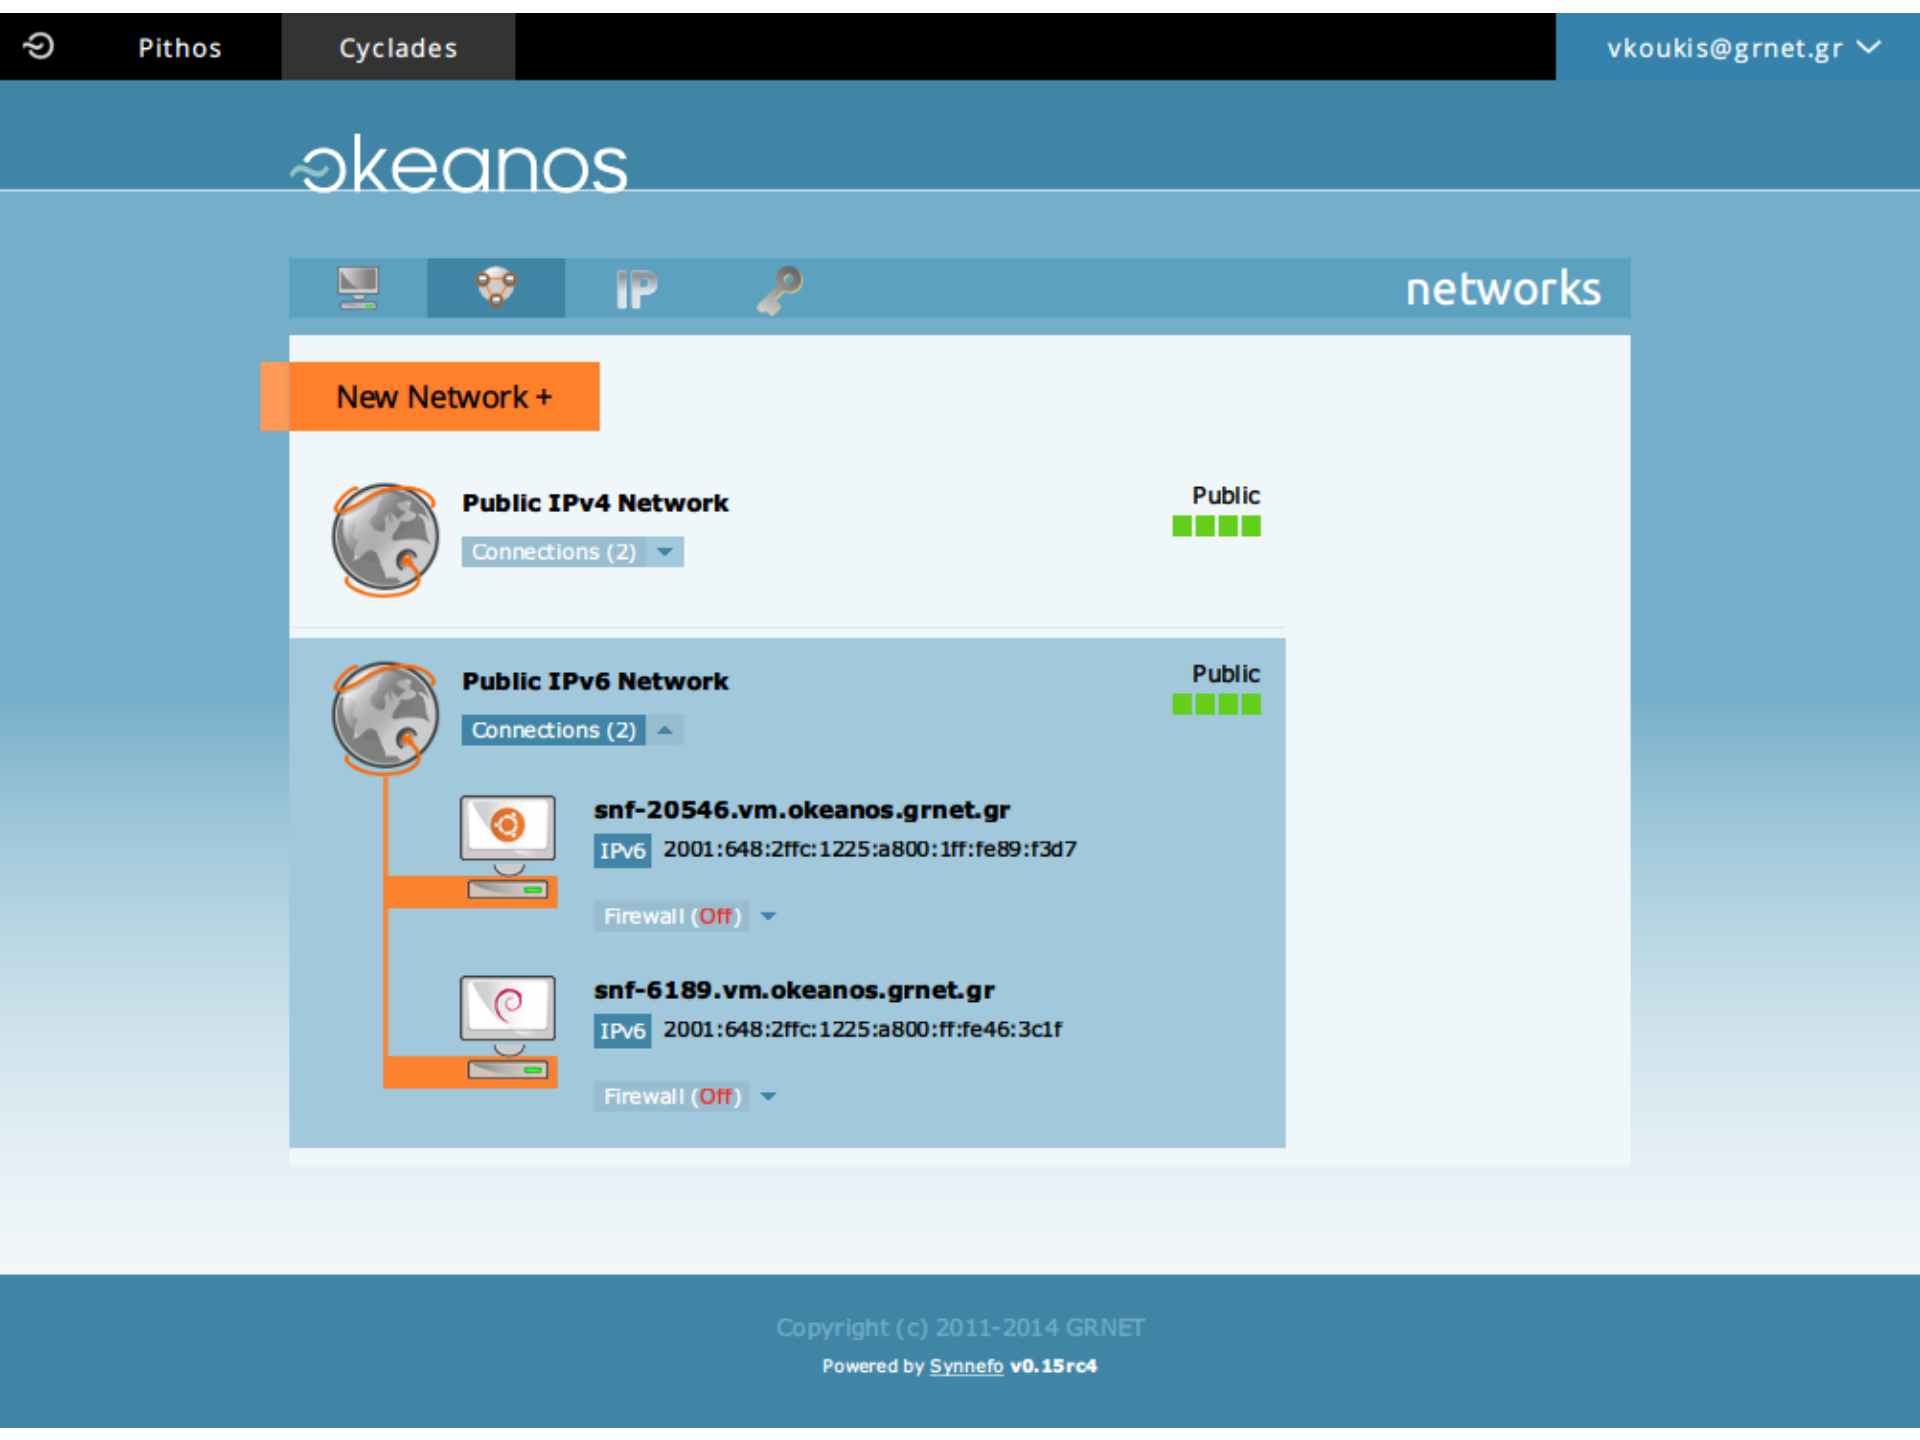
Pithos
Cyclades
vkoukis@grnet.gr
okeanos
IP
networks
New Network +
Public IPv4 Network
Connections (2)
Public
Public IPv6 Network
Connections (2)
Public
snf-20546.vm.okeanos.grnet.gr
IPv6 2001:648:2ffc:1225:a800:1ff:fe89:f3d7
Firewall (Off)
snf-6189.vm.okeanos.grnet.gr
IPv6 2001:648:2ffc:1225:a800:ff:fe46:3c1f
Firewall (Off)
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

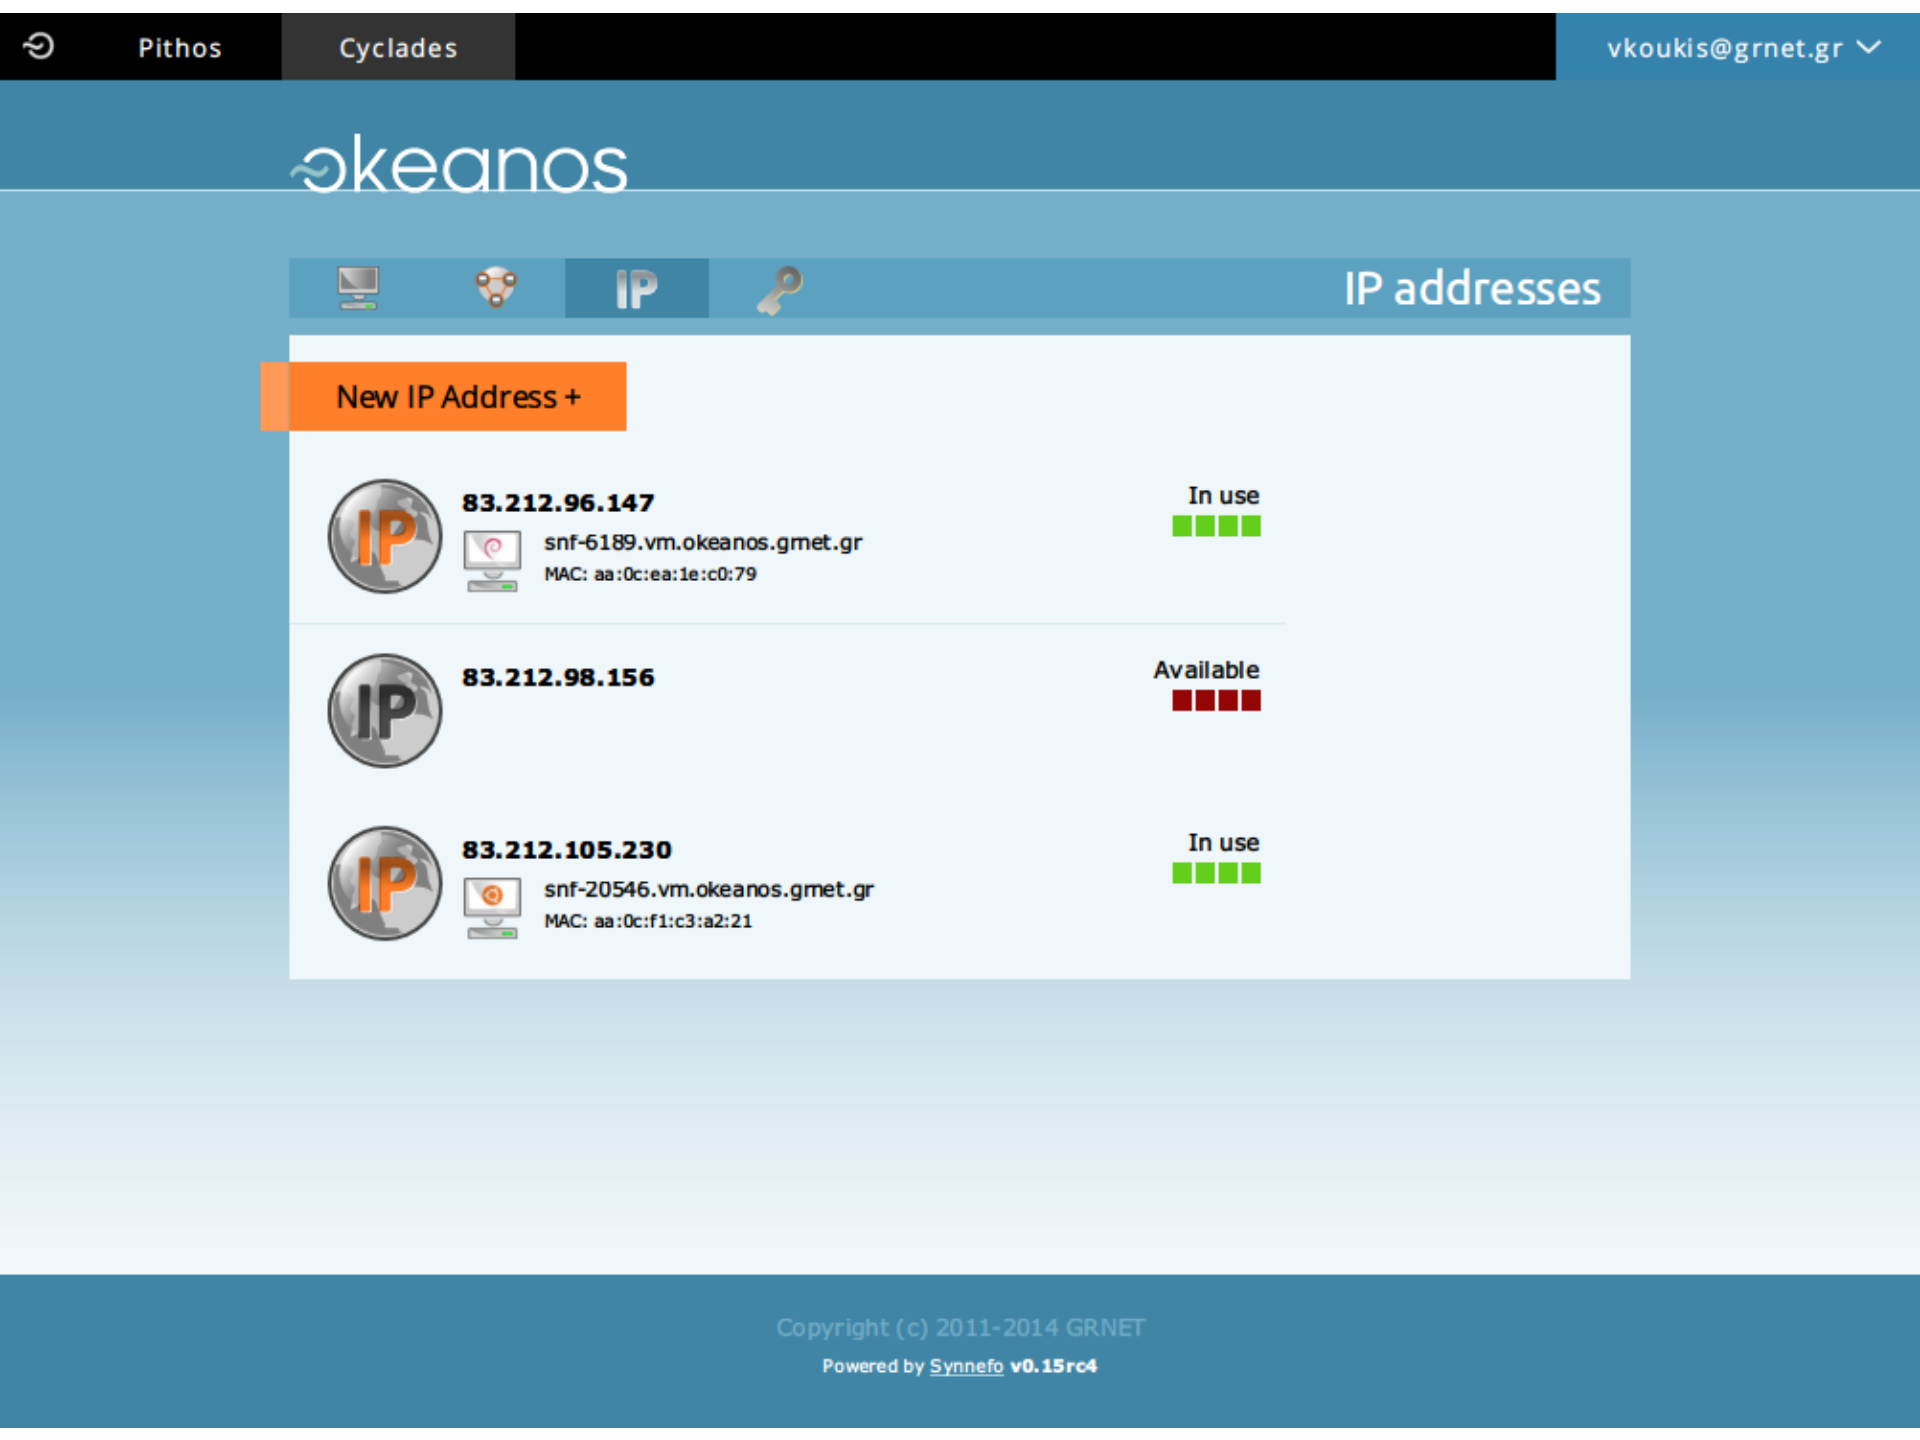
Pithos
Cyclades
vkoukis@grnet.gr
okeanos
IP
IP addresses
New IP Address +
83.212.96.147
snf-6189.vm.okeanos.grnet.gr
MAC: aa:0c:ea:1e:c0:79
In use
83.212.98.156
Available
83.212.105.230
snf-20546.vm.okeanos.grnet.gr
MAC: aa:0c:f1:c3:a2:21
In use
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

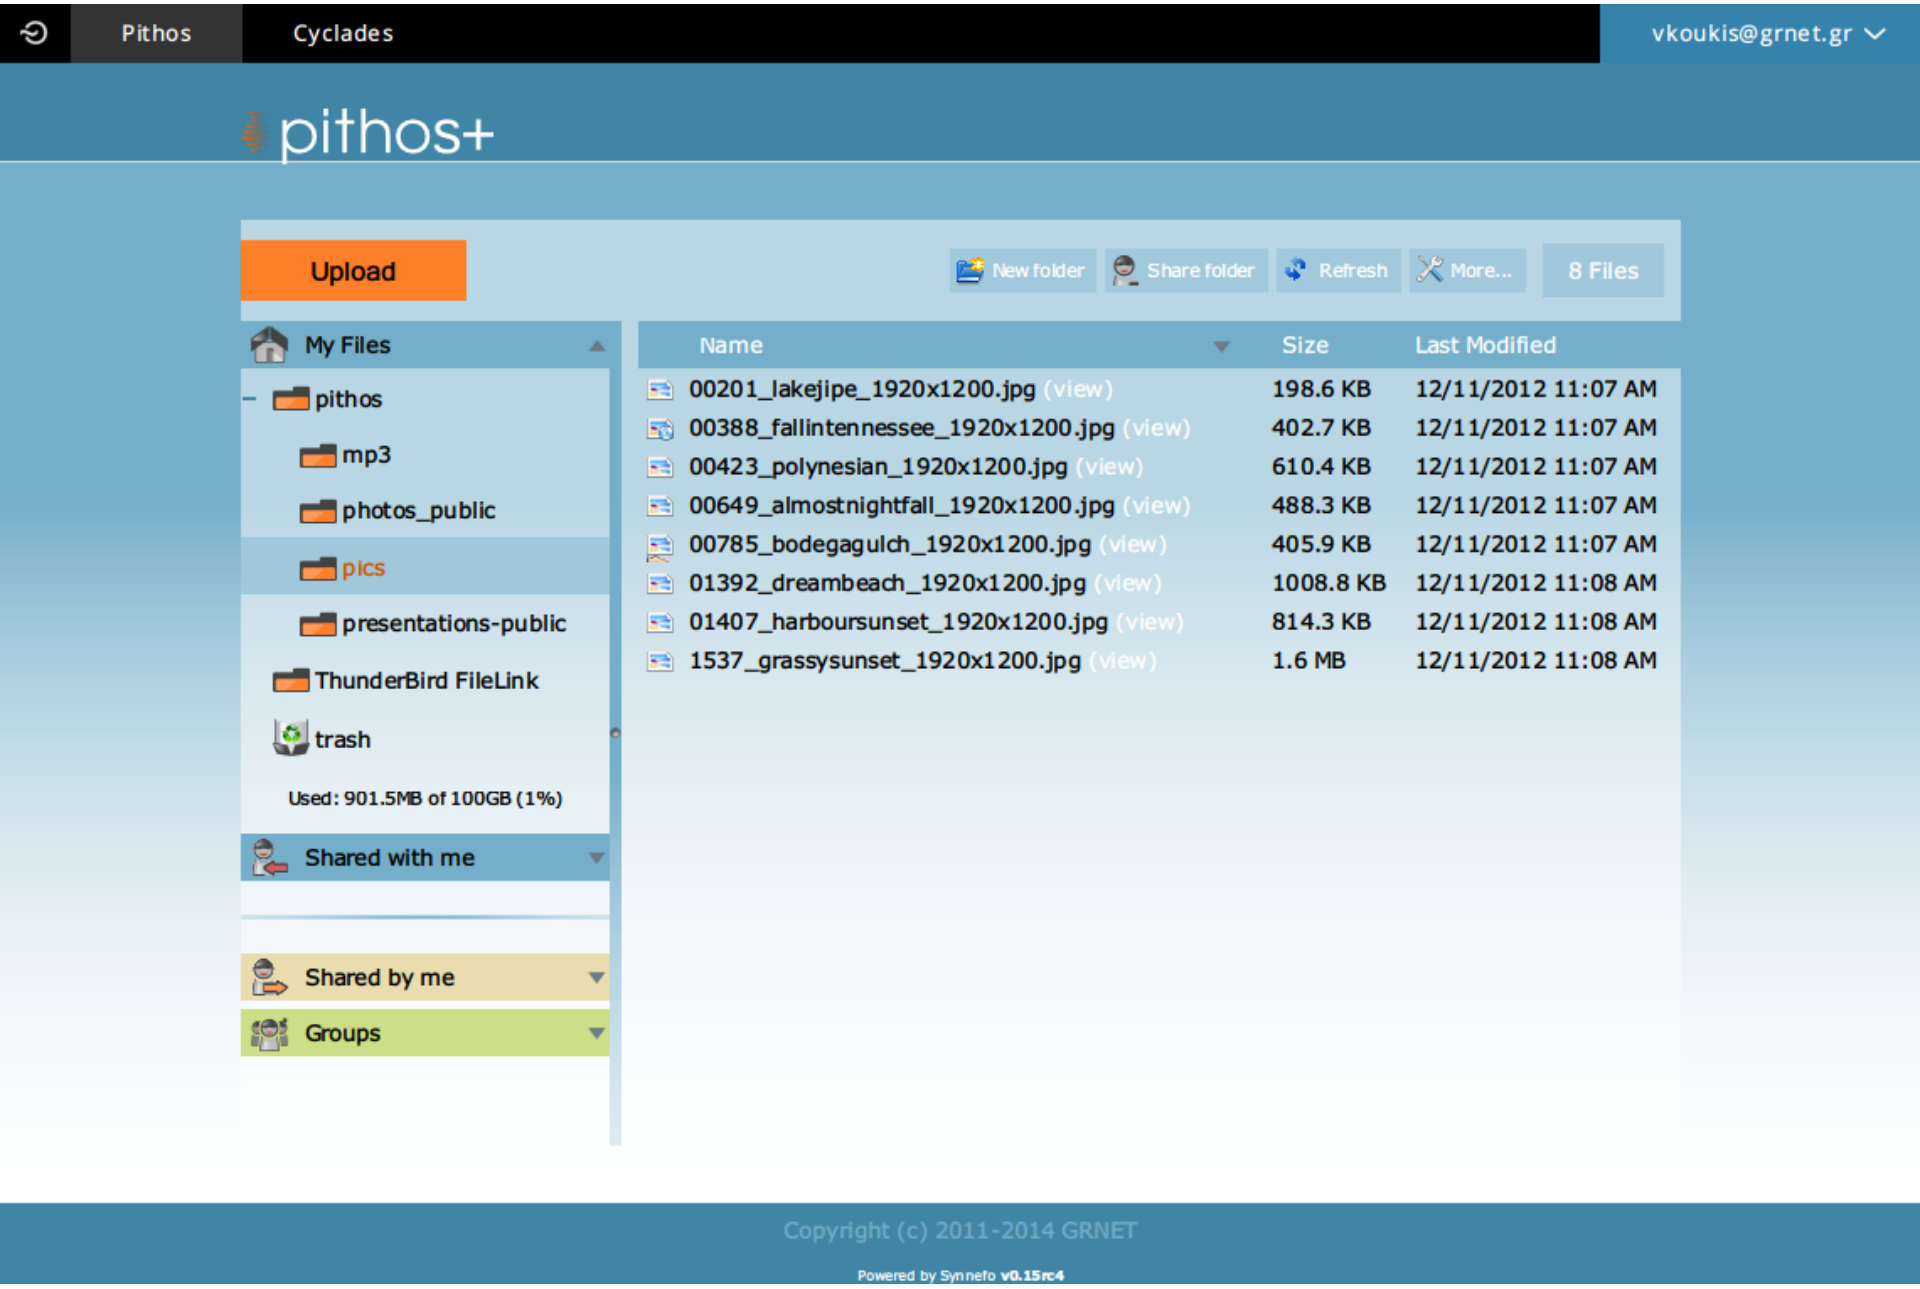
Pithos
Cyclades
vkoukis@grnet.gr
pithos+
Upload
New folder
Share folder
Refresh
More...
8 Files
My Files
pithos
mp3
photos_public
pics
presentations-public
ThunderBird FileLink
trash
Used: 901.5MB of 100GB (1%)
Shared with me
Shared by me
Groups
Name
Size
Last Modified
00201_lakejipe_1920x1200.jpg (view) 198.6 KB 12/11/2012 11:07 AM
00388_fallintennessee_1920x1200.jpg (view) 402.7 KB 12/11/2012 11:07 AM
00423_polynesian_1920x1200.jpg (view) 610.4 KB 12/11/2012 11:07 AM
00649_almostnightfall_1920x1200.jpg (view) 488.3 KB 12/11/2012 11:07 AM
00785_bodegagulch_1920x1200.jpg (view) 405.9 KB 12/11/2012 11:07 AM
01392_dreambeach_1920x1200.jpg (view) 1008.8 KB 12/11/2012 11:08 AM
01407_harboursunset_1920x1200.jpg (view) 814.3 KB 12/11/2012 11:08 AM
1537_grassysunset_1920x1200.jpg (view) 1.6 MB 12/11/2012 11:08 AM
Copyright (c) 2011-2014 GRNET
Powered by Synnefo v0.15rc4

okeanos dashboard

LOGIN

Sign up

If you are a student, professor or researcher you can login using your academic account.

ACADEMIC LOGIN

Classic login (username/password)

vkoukis@grnet.gr

••••••••••••••

SUBMIT

Forgot your password?

RESOURCE USAGE

Storage Space

901.48 MB out of 100.00 GB Storage Space

1%

System Disk

80.00 GB out of 300.00 GB System Disk

27%

CPUs

5 out of 22 CPUs

23%

RAM

5.00 GB out of 26.00 GB RAM

19%

Virtual Machines

2 out of 12 Virtual Machines

17%

Private Networks

0 out of 15 Private Networks

0%

Public IPs

3 out of 9 Public IPs

33%

synnefo

grnet

Integration with Synnefo

GanetiCon 2014

cven@grnet.gr