



OSDC.de
OPEN SOURCE DATA
CENTER CONFERENCE

Open Source Data Center Conference 2014

08. – 10. April, Berlin

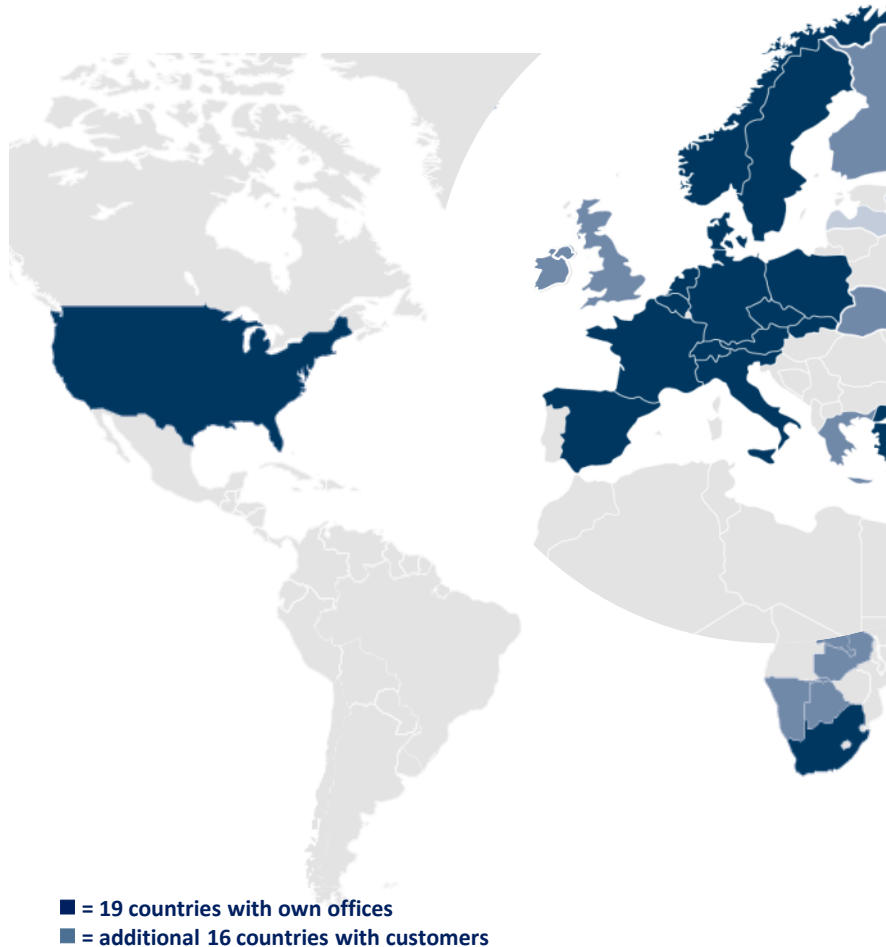
Introduction to Synnefo

open source cloud stack

Thomas Schend, Systems Architect, CompuGroup Medical AG

- **Introduction**
- What do we need from IaaS?
- Introducing to Synnefo
- Components
- Ganeti
- Demo
- Q & A

successfully worldwide in 35 countries for approx. 385,000 customers



Indicators

	2012	Guidance 2013
Customers	385.000	
Group sales	€ 451 Mio.	€ 458 – 463 Mio.
EBITDA	€ 105 Mio.	€ 97 – 100 Mio.
employees	3.700	

Objectives

- Best provider of IT solutions for higher quality and efficiency in health care
- Worldwide leader for MIT (Medical Information Technology)
- High profitability and an average of 30% growth p.a.

- Systems Architect, responsible for Infrastructure Design
- OSS Enthusiast
- LPIC-2, MCITP, VCP, CCA, DCIE, ...

Contact: thomas.schend@gmail.com

- Introduction
- **What do we need from IaaS?**
- Introducing to Synnefo
- Components
- Ganeti
- Demo
- Q & A

- Reliable
- Scalable (Backend and VMs)
- Easy to manage
- Long lived workloads, not all VMs are EC2 instances ;-)

“VMs are not cattle, they are pets”

Vangelis Koukis, Technical Lead, Synnefo

- Introduction
- What do we need from IaaS?
- **Introducing to Synnefo**
- Components
- Ganeti
- Demo
- Q & A

- Open Source
- Production-quality IaaS
- made for persistent VMs
- commodity hardware – No SAN
- Large-scale Installations
- simple UI for the end-user
- uses Ganeti as a Cluster Manager

- Design started late 2010
- Production since July 2011
- ~oceanos* public cloud running currently ~8700 VMs

**<http://oceanos.grnet.gr>*

Synnefo Layers

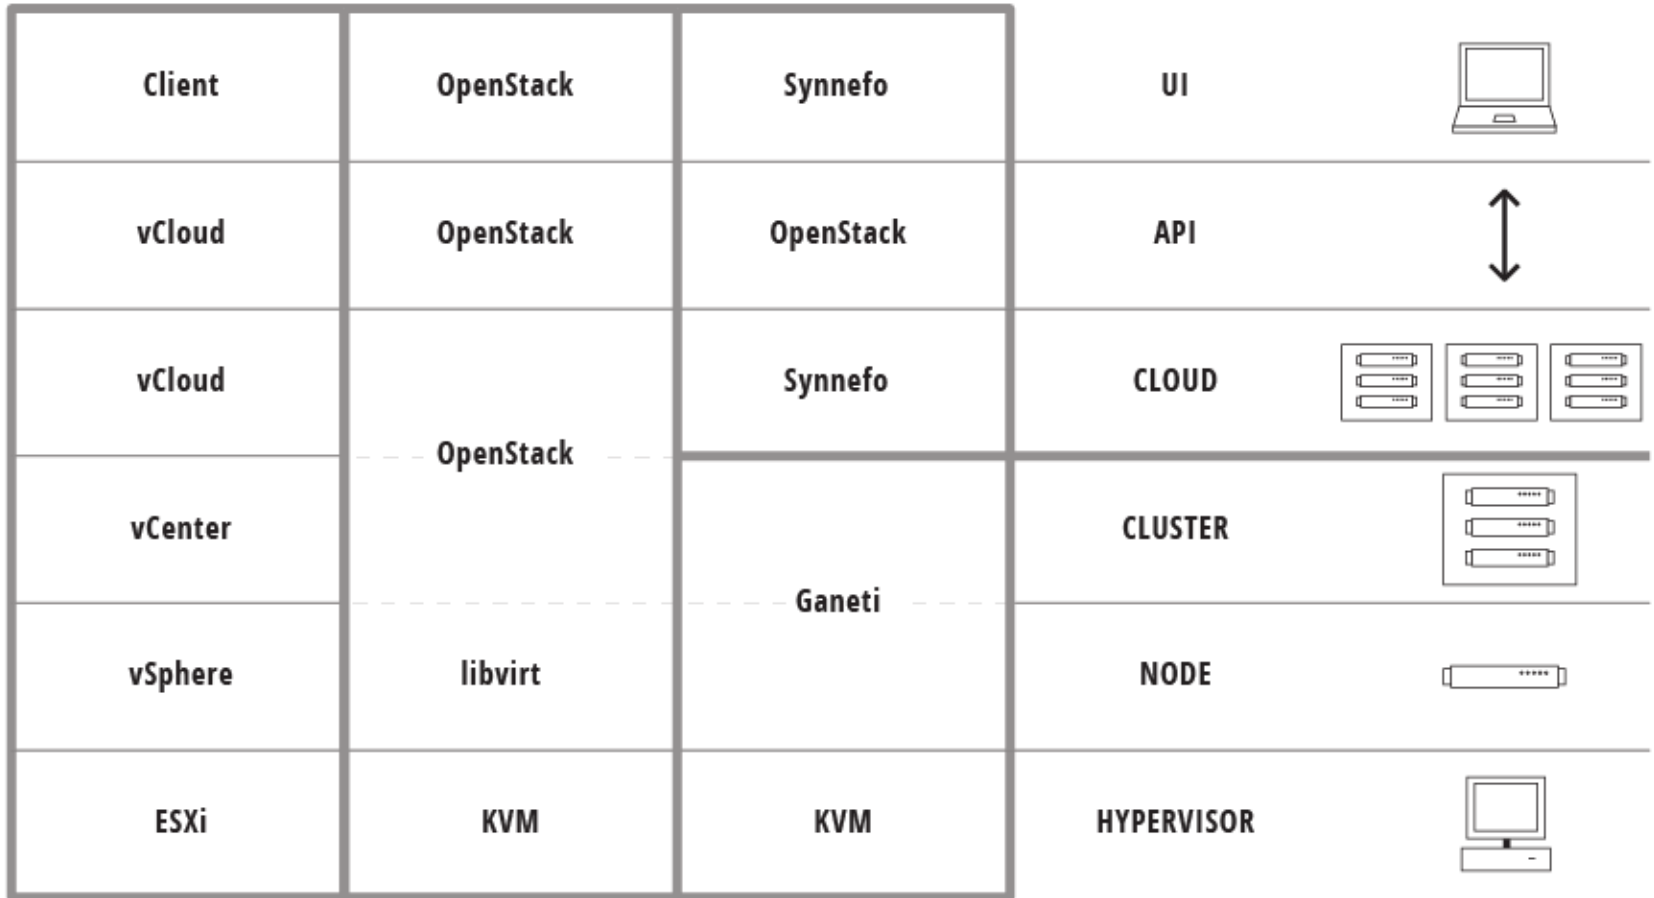


Image provided by GRNET

Synnefo birdseye view

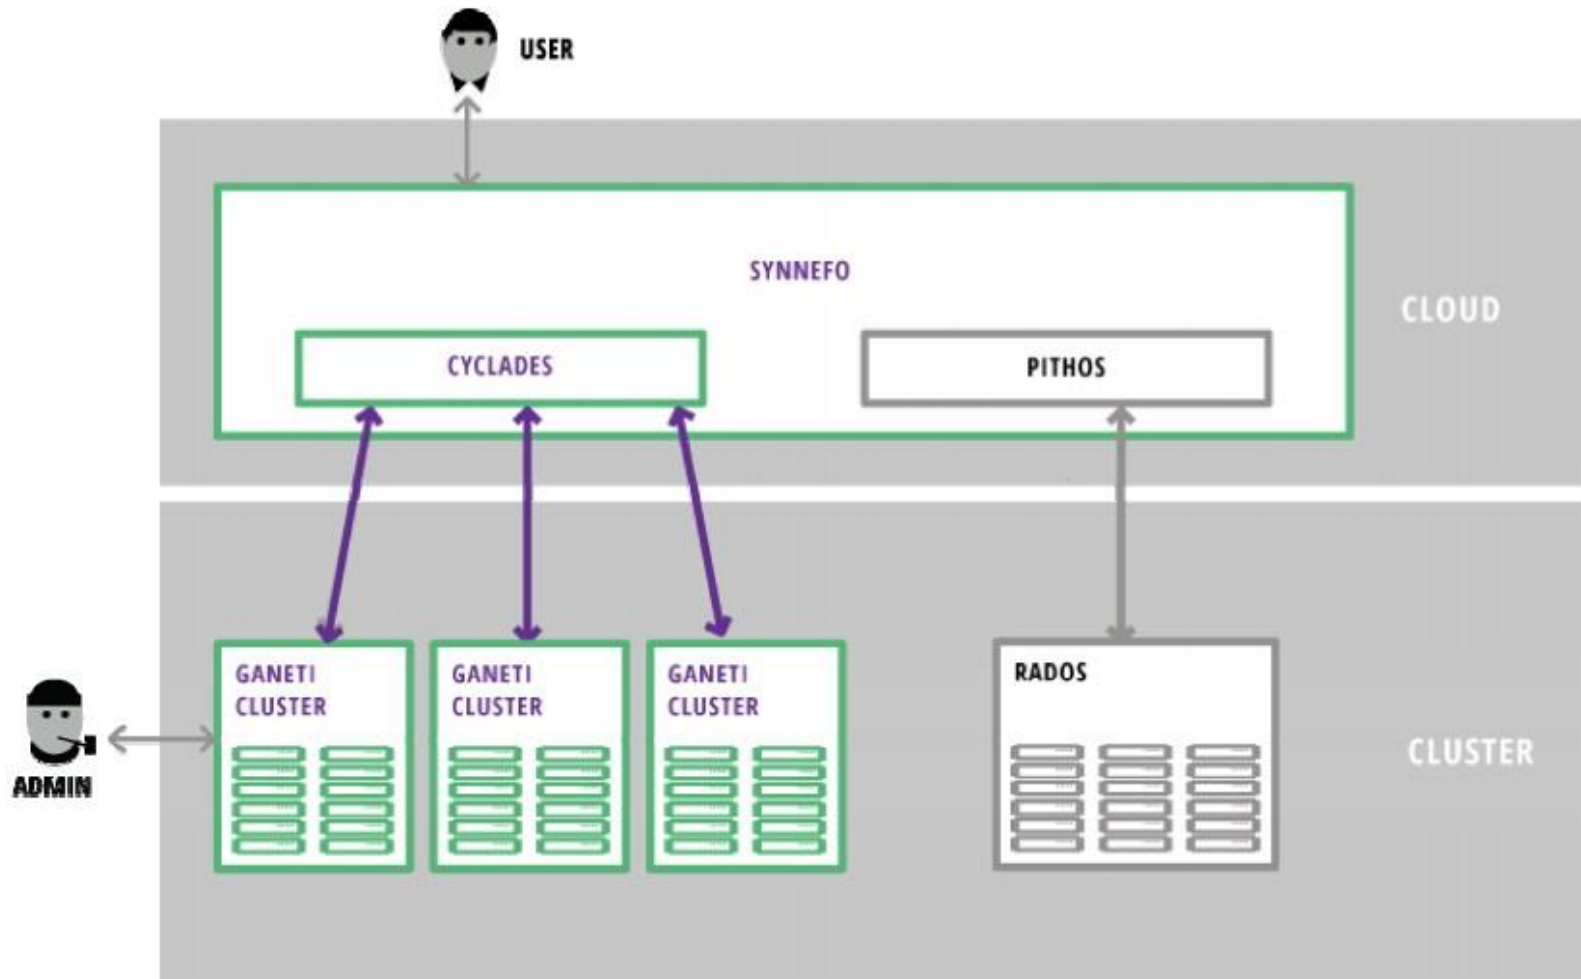


Image provided by GRNET

- Build with Ganeti as a backend
- Astakos: Identity/Account services
- Cyclades: Compute/Network/Image/Volume services
- Pithos: File/Object Storage service
- snf-manage commandline tool
- Basic helpdesk tool
- snf-burnin / snf-burnin-run
- kamaki commandline client
- Code on github

- Introduction
- What do we need from IaaS?
- Introducing to Synnefo
- **Components**
- Ganeti
- Demo
- Q & A

Synnefo Architecture I

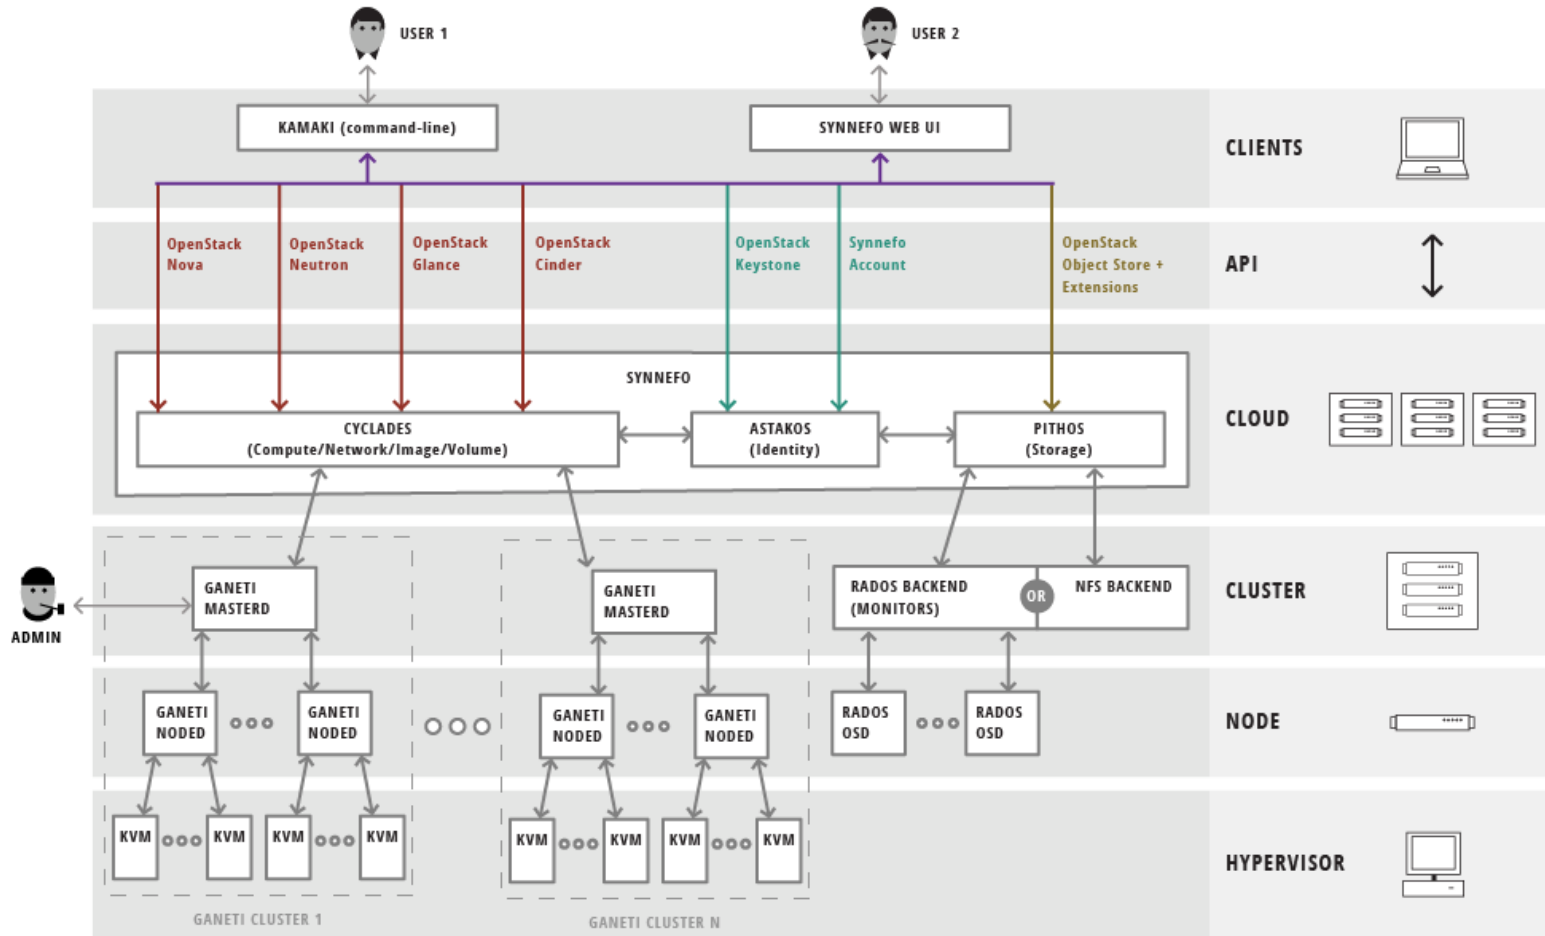


Image provided by GRNET

- Identity service
- Provides Dashboard
- Track Resource usage
- Project Management
- Multiple Auth methods
 - Twitter / Google / LinkedIn / LDAP / Active-directory

- Network management
- L3 Routed/L2 VLANs/L2 isolation with mac filter
- nfdhcpd

- Archipelago storage abstraction
- ceph/NFS/Gluster backend
- very fast, thin cloning
- Snapshots

- Object storage service
- Efficient syncing
- Hashing before upload
- Deduplication in backend
- Web UI written in GWT
- Clients for Windows, MacOS and IOS
- Sharing public and private

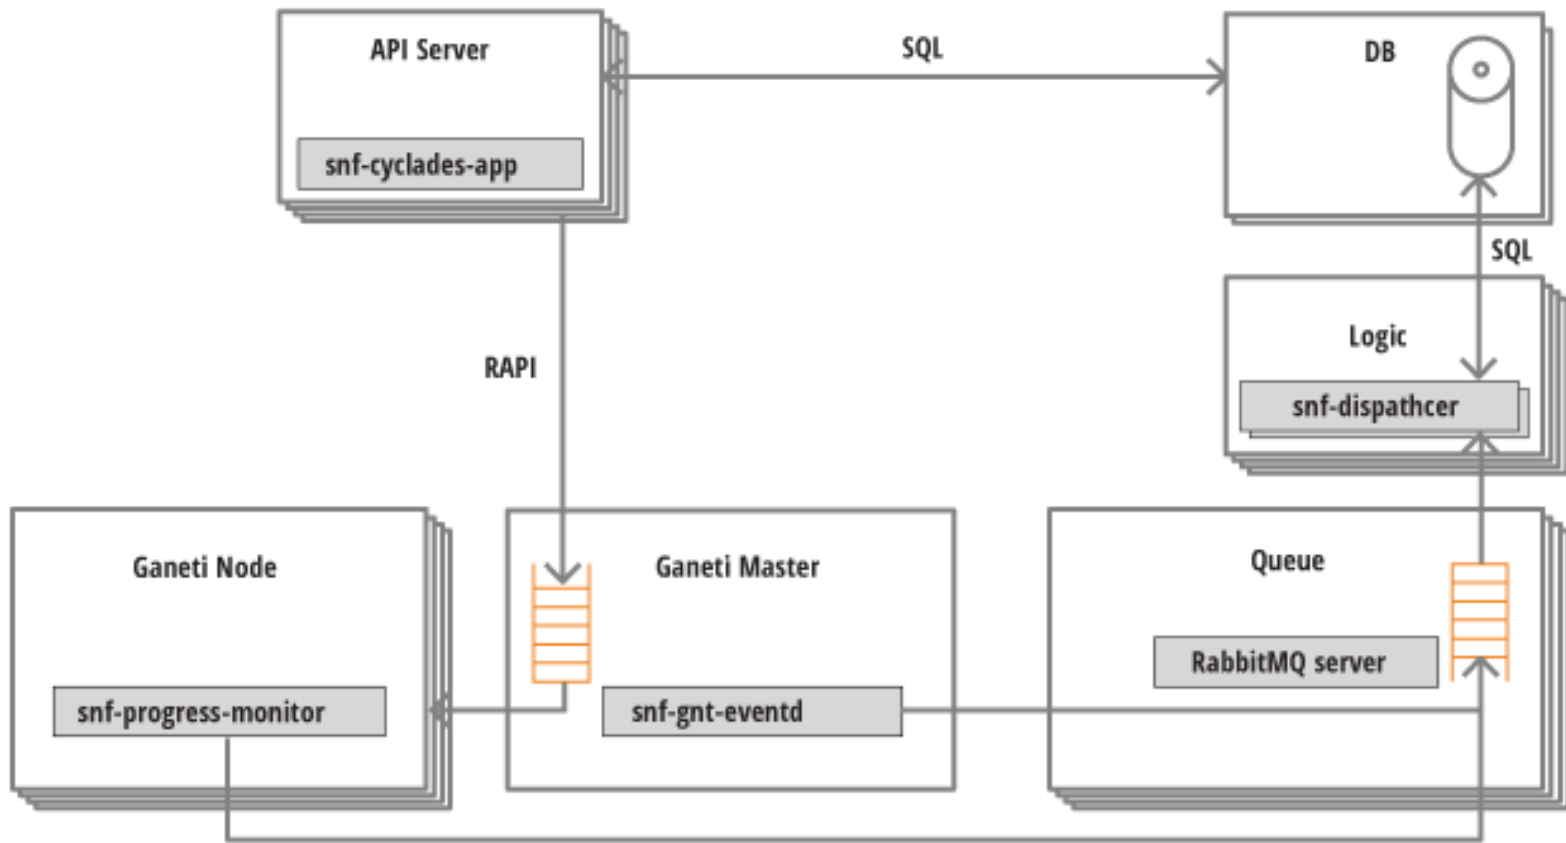


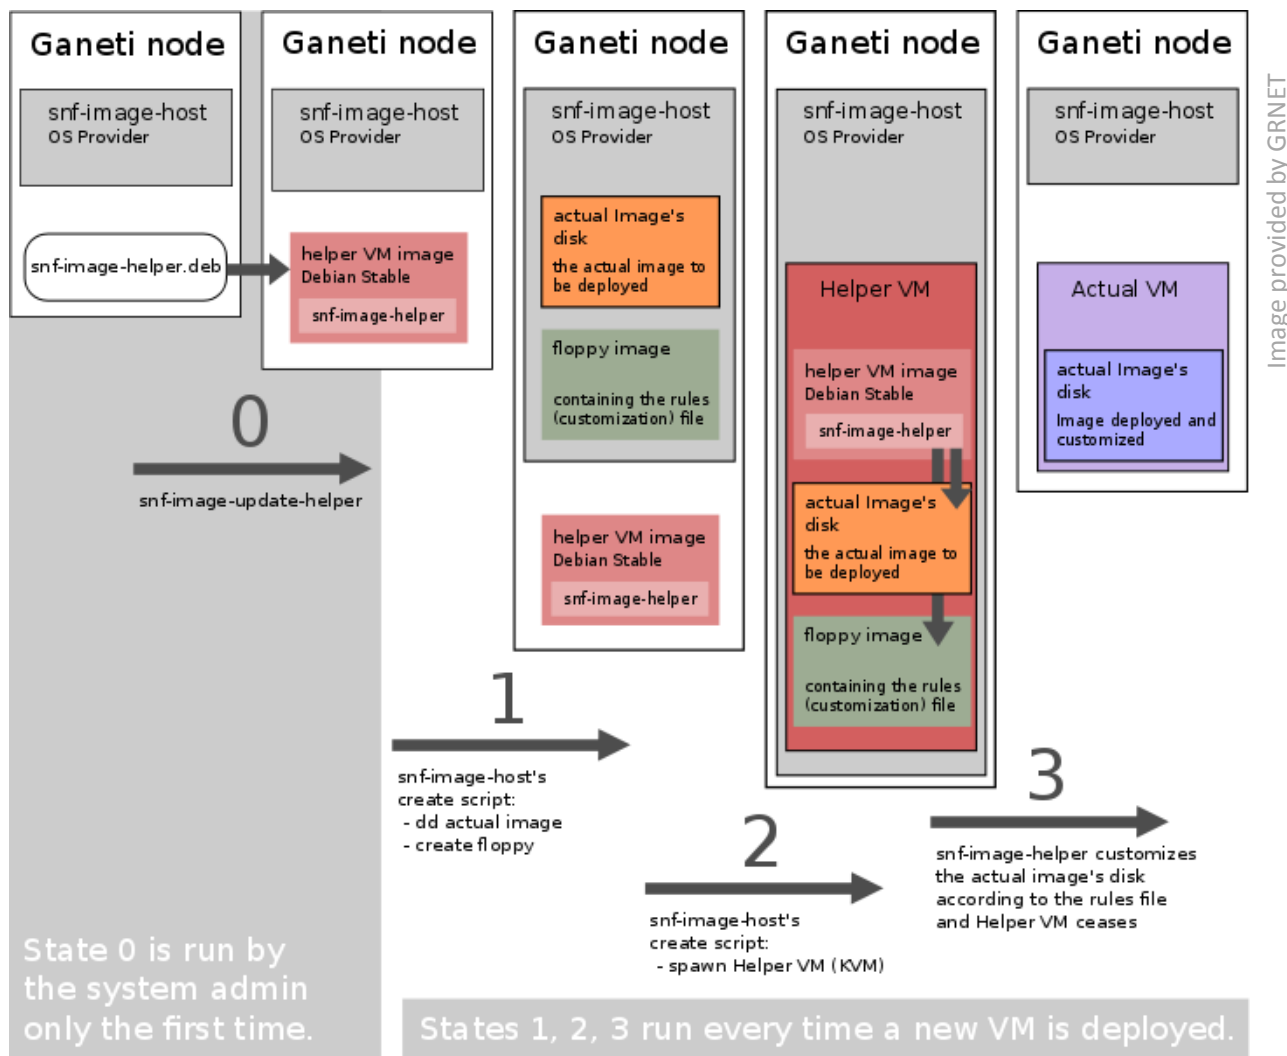
Image provided by GRNET

- written in python / django
- apache/gunicorn
- postgres database
- snf-gnt-eventd / snf-progress-mon
- rabbitmq
- snf-dispatcher
- snf-vncauthproxy
- build for wheezy, can run fully virtualized
- speaks openstack API but does not share any code

- Secure image deployment of untrusted images
- Isolated from the ganeti node
- Supports Linux, Windows and BSD

- customization tasks
 - Assign hostname
 - Delete SSH keys
 - Change password
 - Resize Filesystem
 - Run unattended install (windows)
 - Inject files

Synnefo snf-image deployment process



- Introduction
- What we need from IaaS?
- Introducing to Synnefo
- Components
- **Ganeti**
- Demo
- Q & A

- Very mature, production ready
- opensourced since August 30, 2007 by google
- uses Mature open source tools LVM,DRBD,python,haskell
- Different hypervisors (Xen,KVM,LXC)
- Low overhead for cluster management



- Usage at google
 - small Clusters in Offices (OIB)
 - large Clusters in Datacenters for virtual desktops
- mostly using XEN
- Not used for user facing services

- Scale from 1-200 nodes and beyond
- Live migration without shared storage
- Noderoles: master, master candidate, regular nodes
- Nodegroups
- OVF import tool
- Backup with included tool
- Job queue
- Commands start with gnt-

- OS Interface to provision new instances
- snf-image secure image deployment
- ganeti-instance-debootstrap, image, ganeti-os-defs
- Htools to balance clusters, plan capacity and do maintenance

- supported file,plain,sharedfile,ceph,NFS,Gluster
- Ext storage provider with plugins for Archipelago, HP EVA,,IBM Storwize

- Hooks for every action (pre/post)
- Tags for every object
- Exclusion tags to keep machines apart from each other
- Deploy flash in every node for high performance (flashcache,bcache,dm-cache)
- OOB management interface
- Ganeti watcher

- No need for special hardware (SAN)
- But SAN/Shared Storage is supported
- Manage cluster as an entity not single nodes (libvirt)
- Few dependencies
- No database
- Easy recovery in case of hardware outage
- No SPOF which can affect the reliability of VMs
- Very good documentation

- Deployment automation puppet, chef, cfengine, ..
- Monitoring icinga/nagios/check_mk/...
- Openvswitch LACP + 10GBe
- Autorepair
- Packages for all major OS debian, ubuntu, rhel/centos

- cli with bash completion
 - LUXI
 - RAPI
- Different web UIs using the RAPI
 - Ganeti web manager
 - Ganetimgr
 - Synnefo

`gnt-cluster init / info / verify / command / copyfile`

`gnt-node add / list / failover / evacuate`

`gnt-instance add / info / replace-disks`

`gnt-job list ...`

`gnt-network ...`

`gnt-backup ...`

`...`

- written in haskell

`hail`

`hbal`

`hspace`

`harep`


`hinfo`

`hroller`


...

Typical Ganeti Cluster


```
g n t - c l u s t e r   i n i t   [ - s   i p ]  
g n t e n a b d e d a d d y p e r s i s p t s = k v m  
g n t e n a b d e d a d d   [ - s   i p ]   n o d e 3  
g n t - c l u s t e r   v e r i f y  
g n t - i n s t a n c e   a d d   - t   d r b d   -
```



Node01



Node02

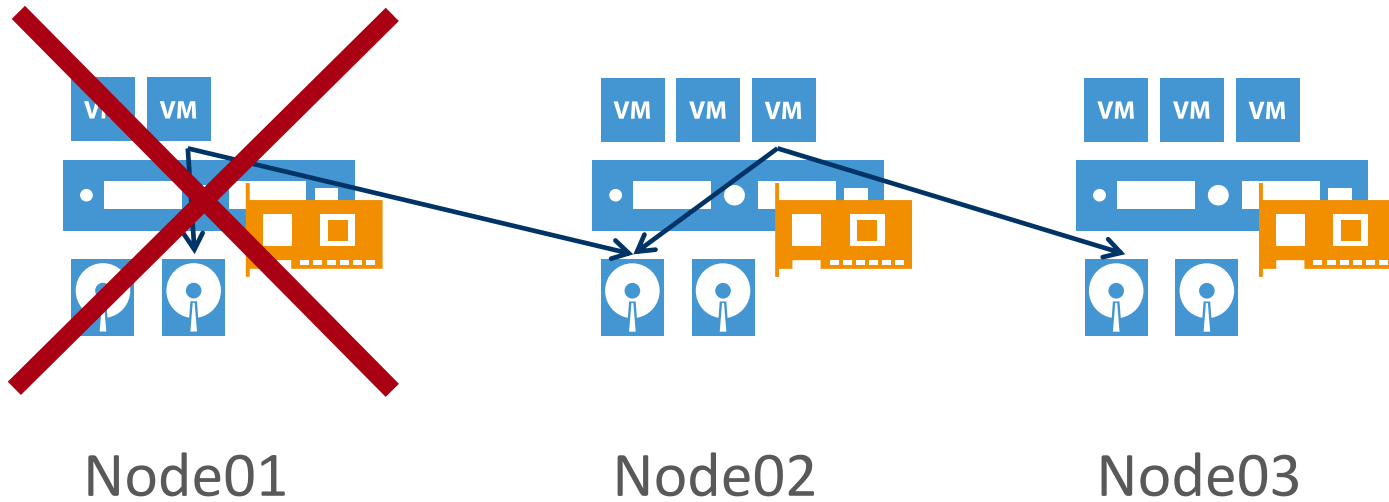


Node03

Veeam virtualization stencils

Ganeti Failover and Repair

```
g n t - n o d e   m o d i f y   - 0  
g e $ - n o d e 0 1 f a i l o v e r   - -  
@ g n e n e d e o e s à s t a t e y - l   h a i l  
n o d e 0 1
```



Veeam virtualization stencils

“Ganeti didn't have any generally known meaning before this project. And now it means kick-ass cluster management”

Martin Beauchamp, on the Ganeti Users List

- Introduction
- What do we need from IaaS?
- Introducing to Synnefo
- Components
- Ganeti
- **Demo**
- Q & A

- Login to Astakos
- syncing and Pithos UI and App
- Create VMs with Cyclades UI
- Snapshot Feature
- OOB Console
- Hotplug of NICs and Networks
- Floating IPs
- Clean-up

- Introduction
- What do we need from IaaS?
- Introducing to Synnefo
- Components
- Ganeti
- Demo
- Q & A

Try it!

<http://www.synnefo.org>

<https://github.com/grnet/synnefo>

<https://code.google.com/p/ganeti/>

Q & A