Hâpy

OpenNebula GNU/Linux distribution for two french ministries

Daniel Dehennin

Pôle de Compétences Logiciels Libres

OpenNebula TechDay Paris 2015

Pôle de Compétences Logiciels Libres

FOSS and agility in french Minister of National Education

- Main mission is to maintain the EOLE GNU/Linux meta-distribution
- CeCILL / GPL software licensing
- Agile consulting for other development teams

The EOLE GNU/Linux meta-distribution

EOLE: GNU/Linux meta-distribution

Ensemble Ouvert Libre et Évolutif

- Local project in educational constituency of Dijon in 2000
- National project in 2001 to protect students and administrative datas
- Based on Mandrake Linux
- Switched to Ubuntu GNU/Linux in 2007

Turnkey solutions for national education

One ISO to rule them all

From elementary to high school

- Zéphir: Centralised server management
- Amon: Firewall, proxy and IPSec VPN
- Sphynx: VPN concentrator
- Horus: Administrative staff Samba server
- Scribe: Student communication and file server
- AmonEcole: Merge Amon and Scribe functionalities with containers
- Eclair: LTSP server
- Seshat: Centralised MTA and web SSO
- Thot: Centralised LDAP

Highly adaptable

Patch and extend

- Variables declared in XML files
- Python Cheetah templates of configuration files
- Pre/post scripts (ex: populate database)

Adapt or create your own derivatives to fit your needs

MEDDE derivatives

Variation and new servers

Ministère de l'écologie, du développement durable et de l'environnement (MEDDE) created their own derivatives

- eSSL: variant of firewall Amon (2009)
- eCDL: NT domain controller (2011)
- eSBL: file server, plugged on eCDL (2011)

Largely deployed

Even on a boat, not on the map ;-)

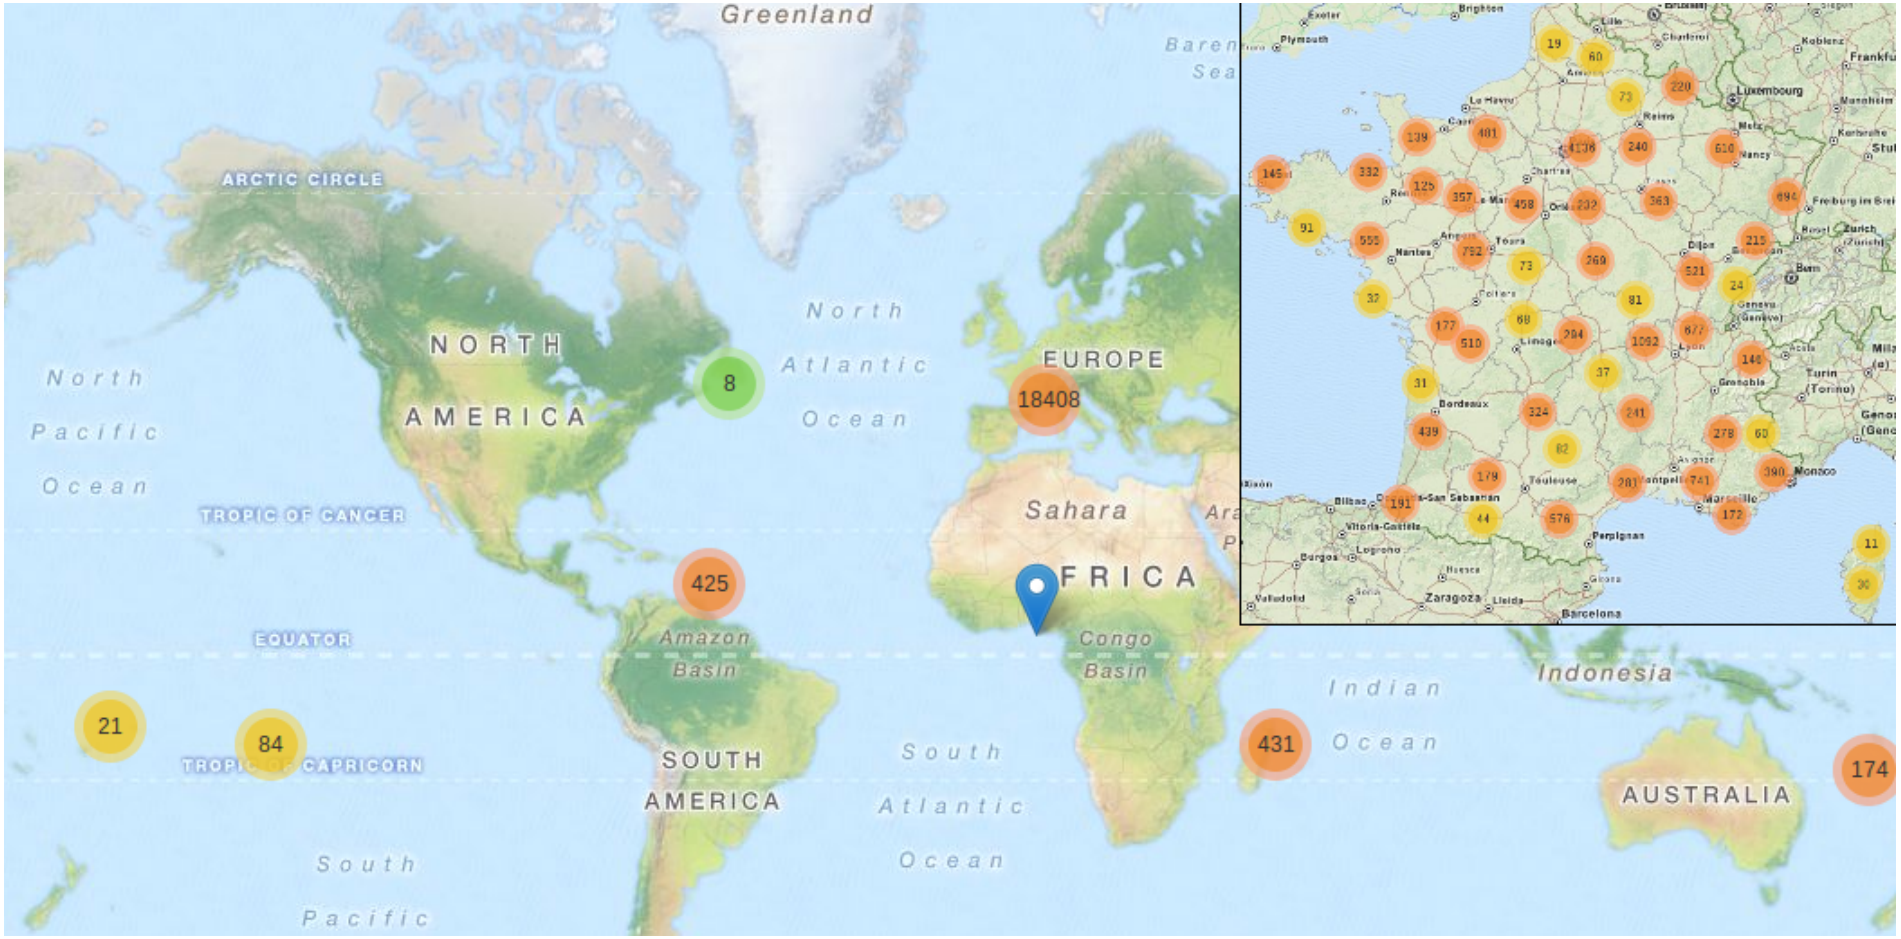

Bare metal to configured server in 30 mn

3 little steps for humans

1. Installation from ISO
2. Configuration
3. Deployment

Each step can be done by different people

Automatic installation from hybrid ISO

OEM like install

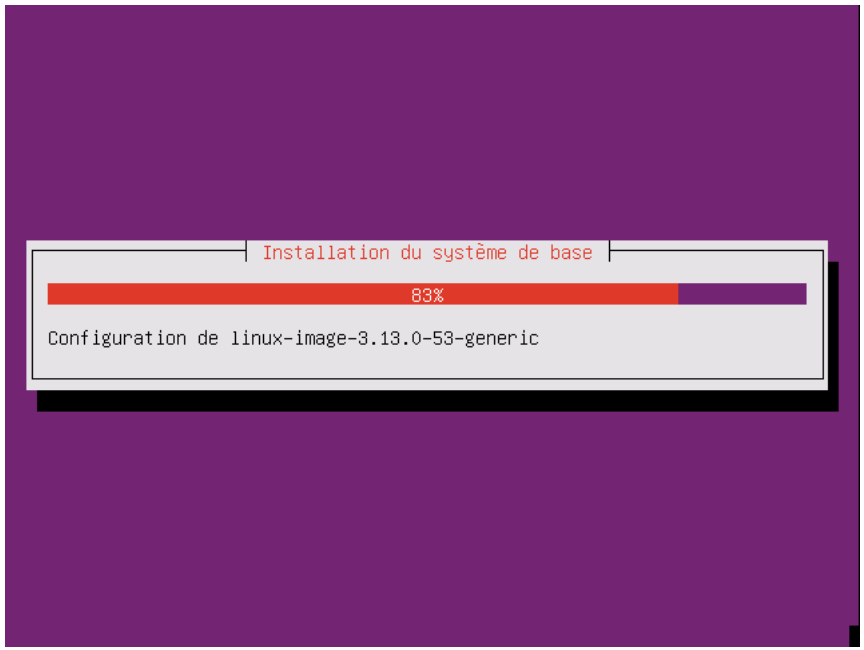

A web interface for configuration

Local on the server or central on Zéphir

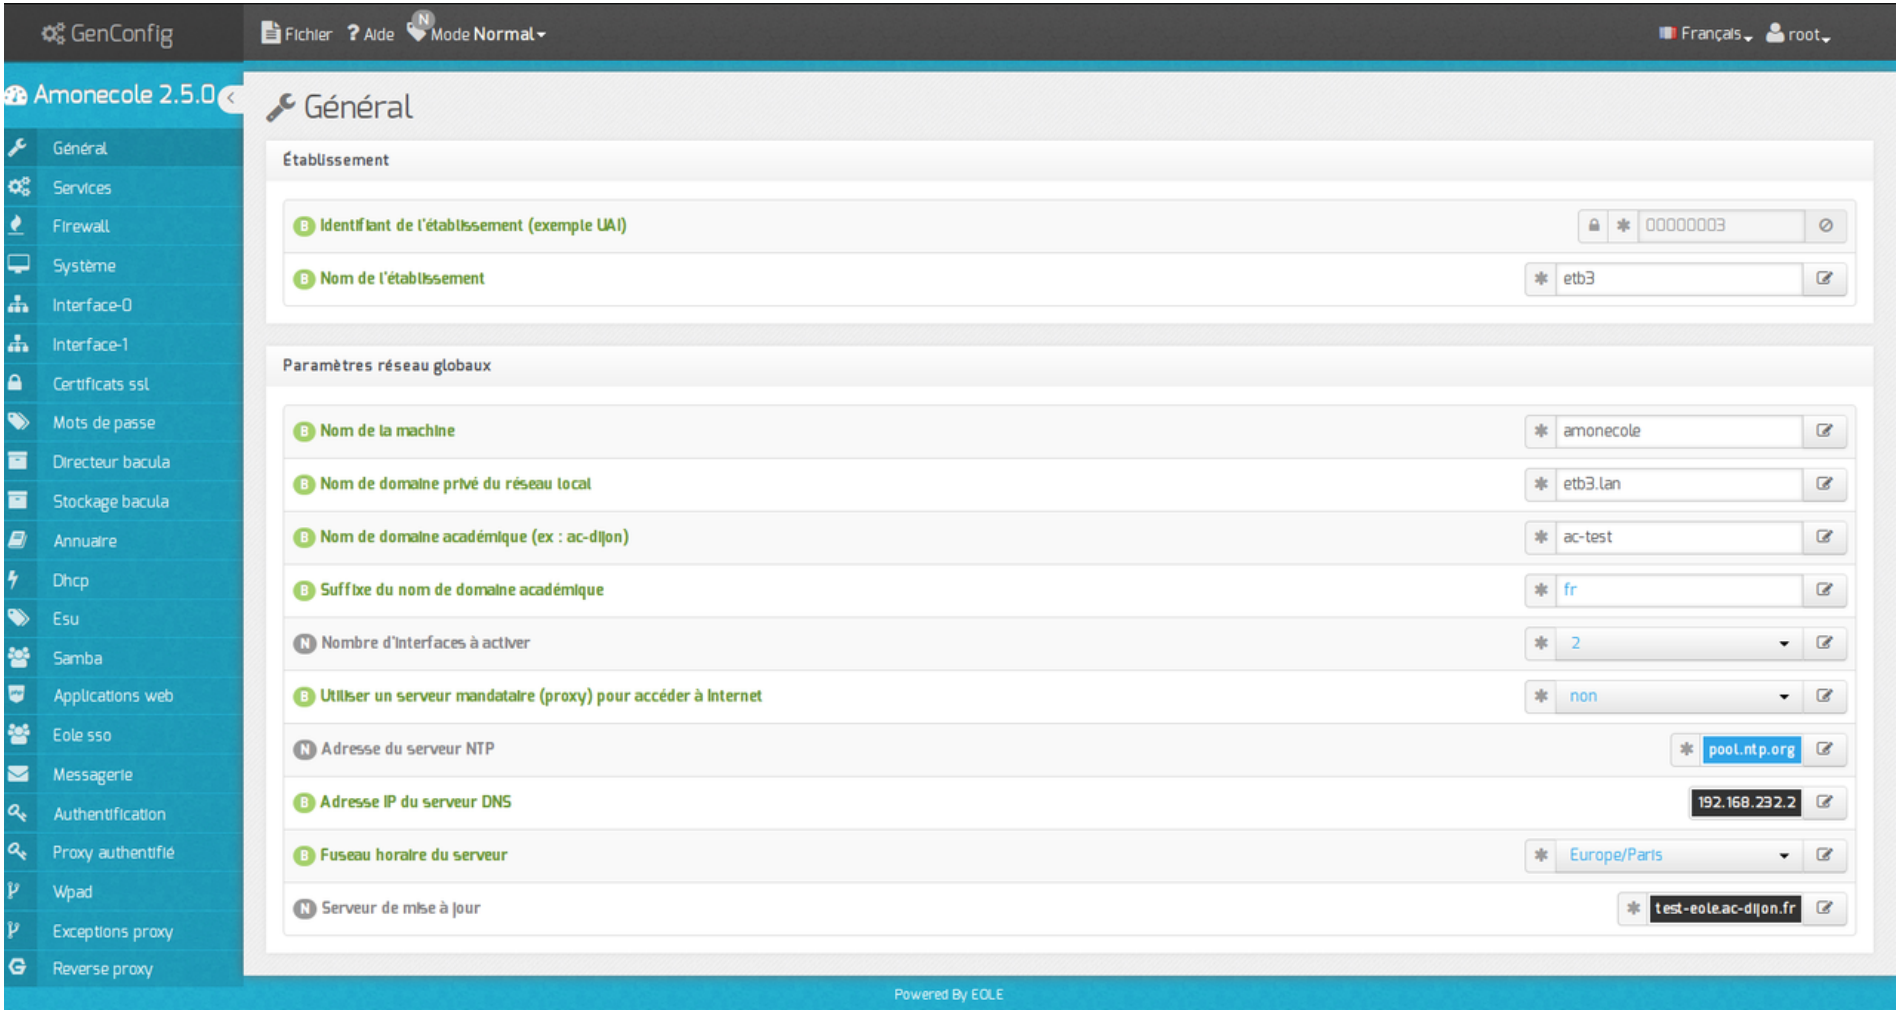

Instantiate the server

Generate config files and start services

1. Register the server on Zéphir
2. Retrieve the configuration
3. Run *instance*

Manage and monitor the server

command line or web interface

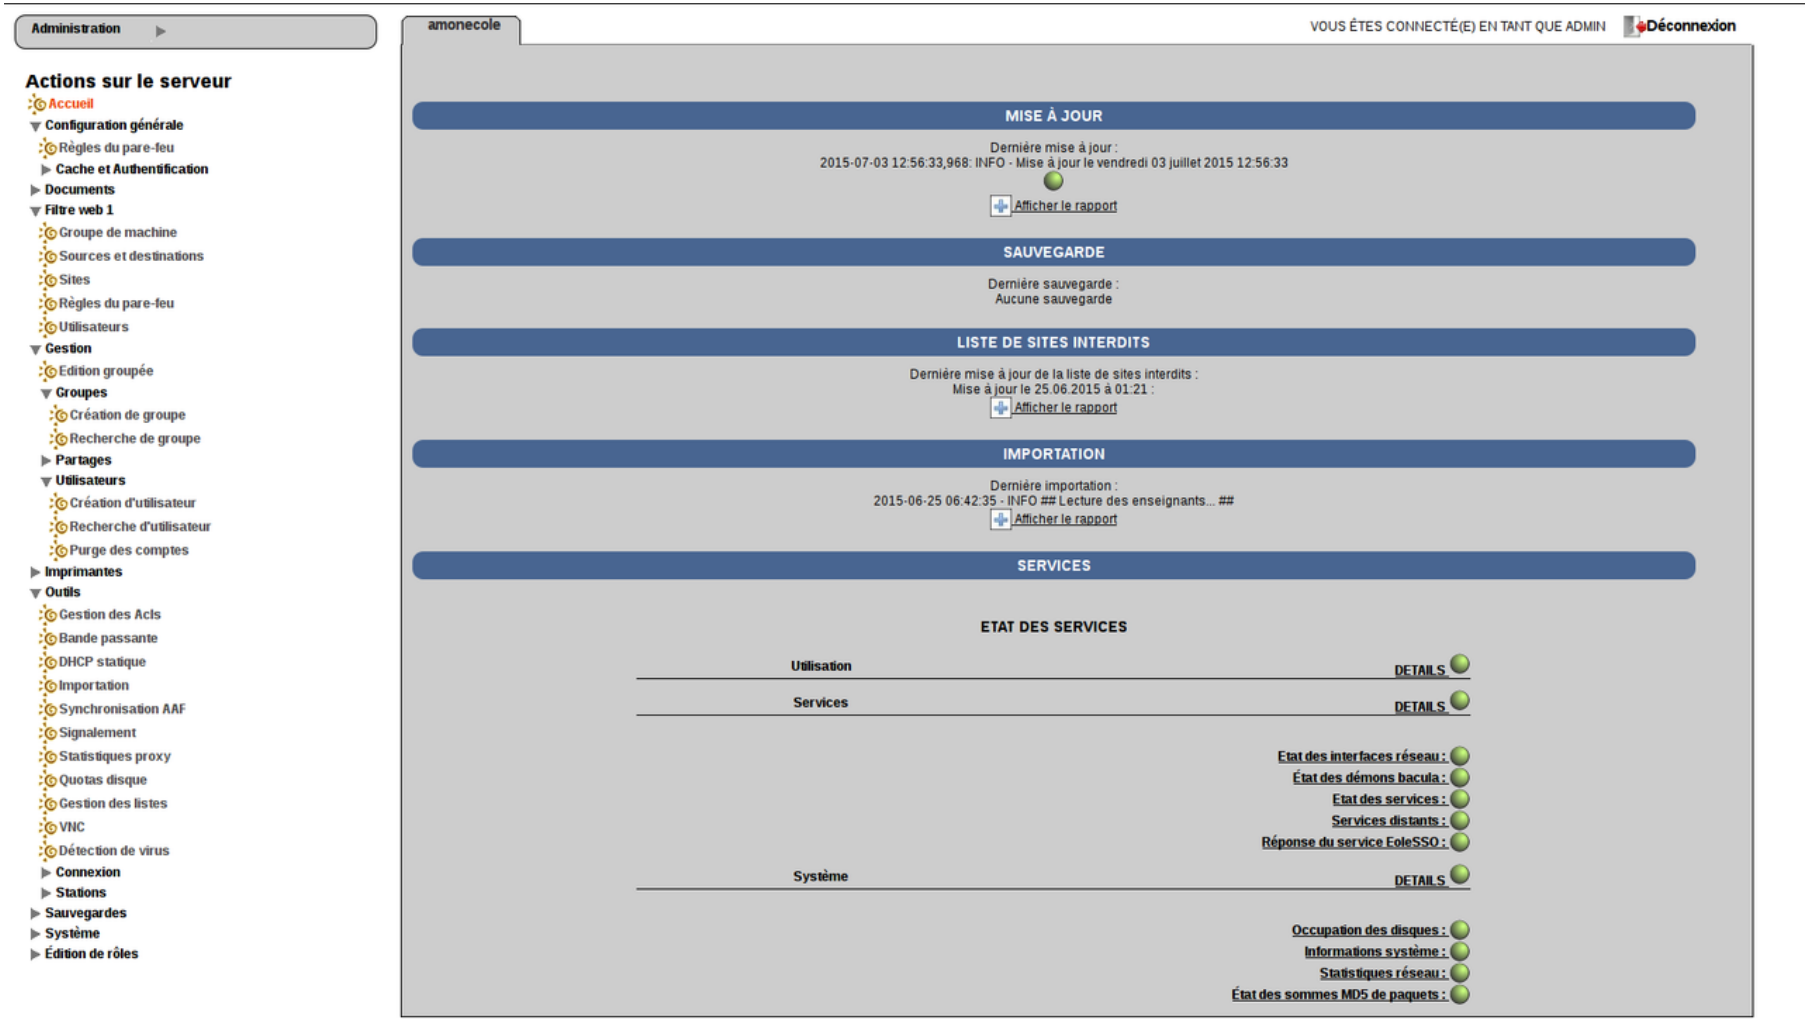

Road to a new galaxy

Why did we get to OpenNebula?

Bare metal elastic limit is too low

- Testing our OS was done on physical desktop computers
- Some “lucky” developers could have at most 2 VMs on their workstation

EOLE development needed elasticity

Looking for virtualisation infrastructure

Many choices: too big, not enough flexible or immature

2012: two new quite powerful workstations $\Rightarrow$ testing party

- Proxmox needed a reboot to add a new network
- Archipel barely emerged
- Ganeti was promising
- OpenStack was already too much

Start with OpenNebula 3.8

Complete virtual infrastructure per user

- Standard network names
- VLAN isolated networks
- One user == one gateway
- Generated with Jenkins

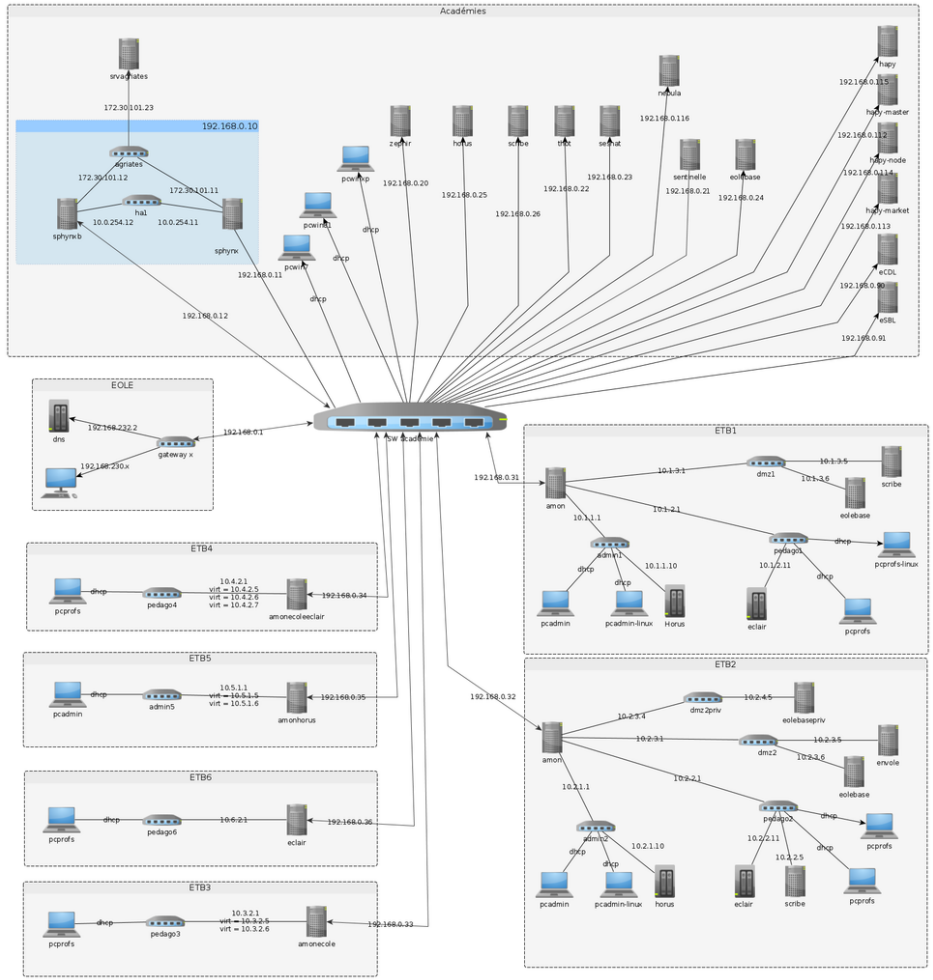

Sharing VM templates requires avoiding *UNAME* on networks

Pilot OpenNebula from Jenkins

Continuous integration of OS

- Check installation from ISO
- Check default configurations
- Check daily upgrade to find broken packages
- Check user database import

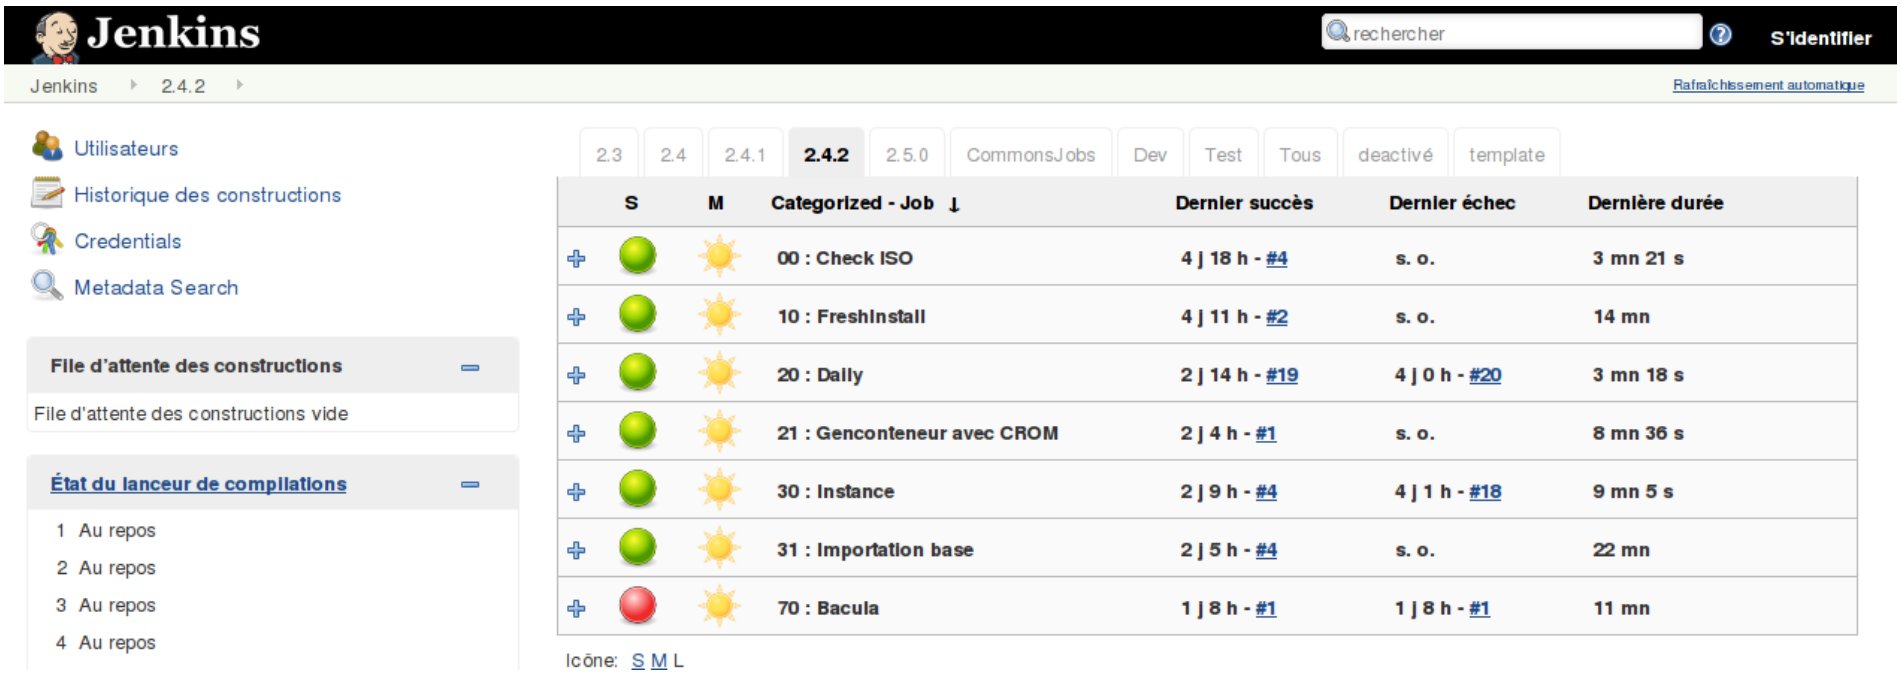

Jenkins jobs produce ready to use VMs

Reduce environment setup time

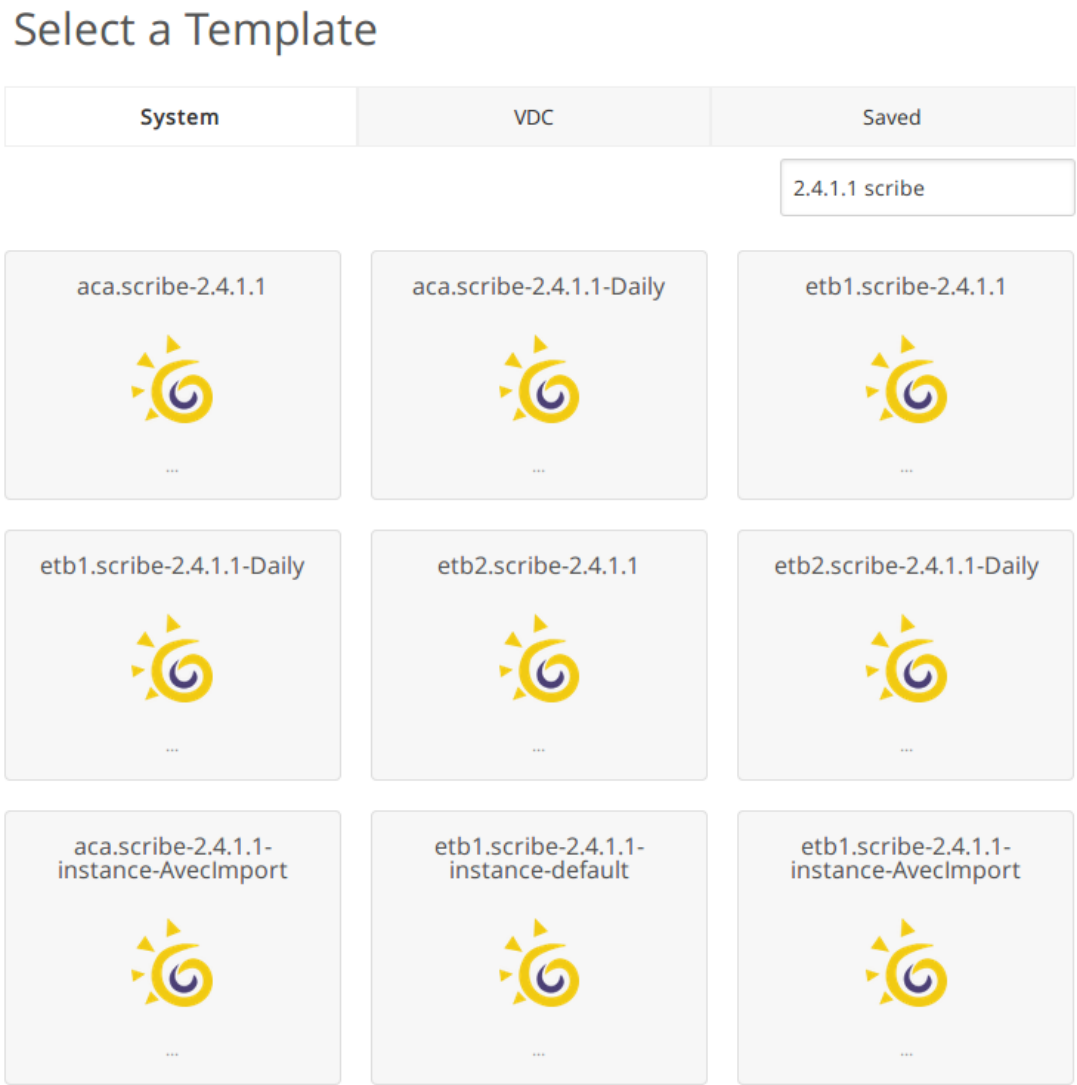

Remove physical limitations

Test beds was burning

- The two dedicated workstations was fine for testing
- Bumped memory to 2x32GB

NFS access on workgroup NAS was too slow

Two 24 CPU 96GB RAM Blades with 3TB GFS2 on SAN

Road to Hâpy-ness

2013: talk at our annual seminary

Teasing inside

General presentation of OpenNebula

High interest from our community to add it on the ISO

First steps to OpenNebula distribution

The community wanted it, the community did it

Common elaboration between

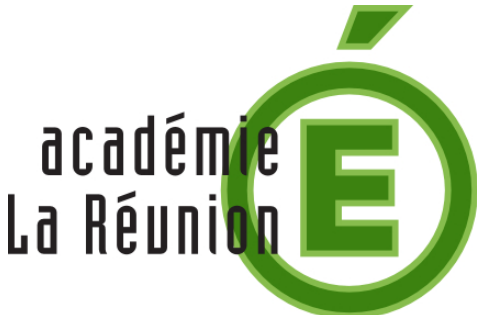

Hypervisor

Virtual network

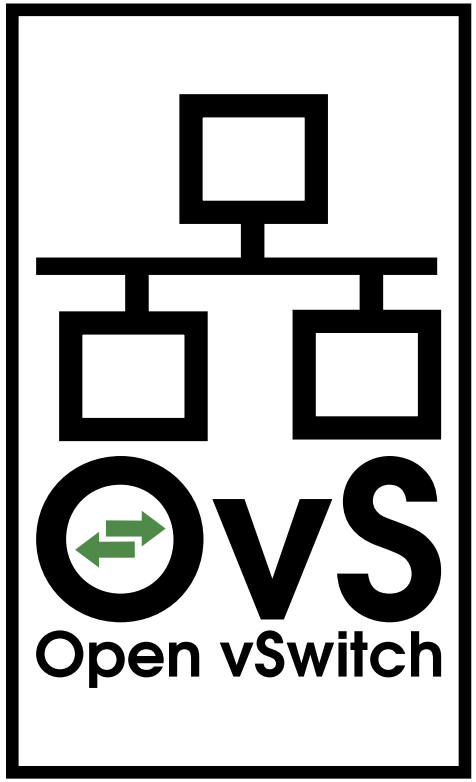

2014: presentation of the work in progress

Show must go on

- Only the first part was funded
- Single node OpenNebula
- Preparation of virtual networks from Zéphir
- Preparation of datastores from Zéphir
- Automated deployment of EOLE servers based on Zéphir

Engagement of MEDDE to finalise it

2015: Hâpy new distributions

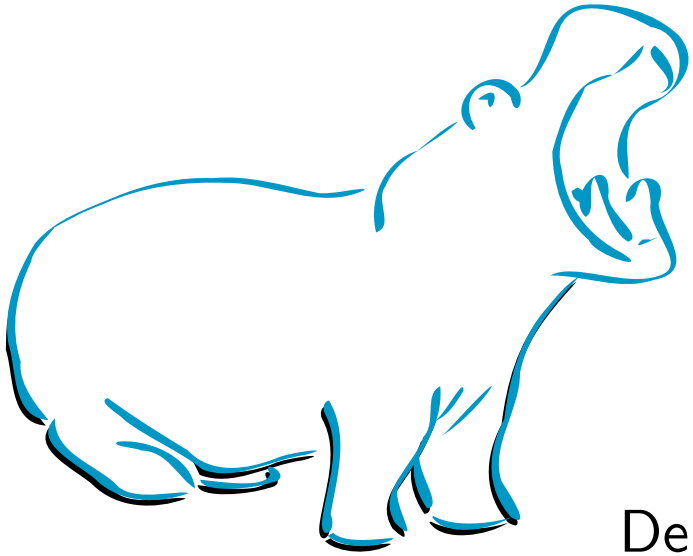

Deification of annual flooding of the Nile

- Integrated on ISO
- Missing community feedback ⇒ tagged experimental

Several distributions depending on the working mode

Single node server

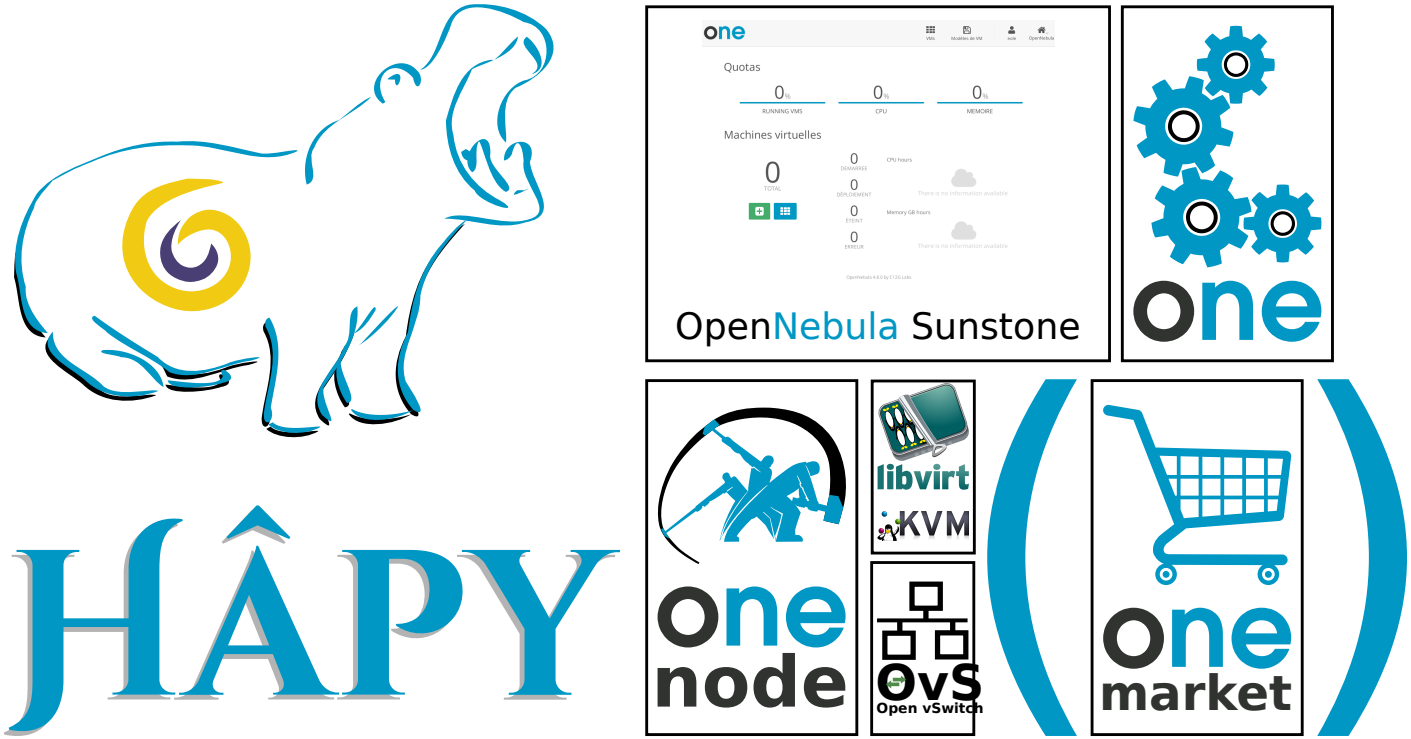

Cluster: the frontend

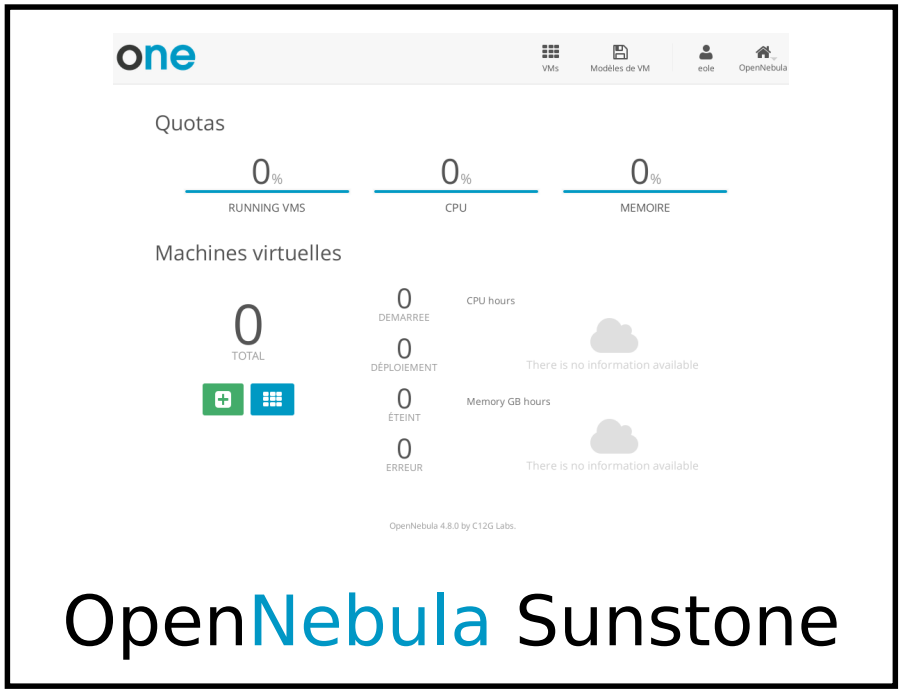

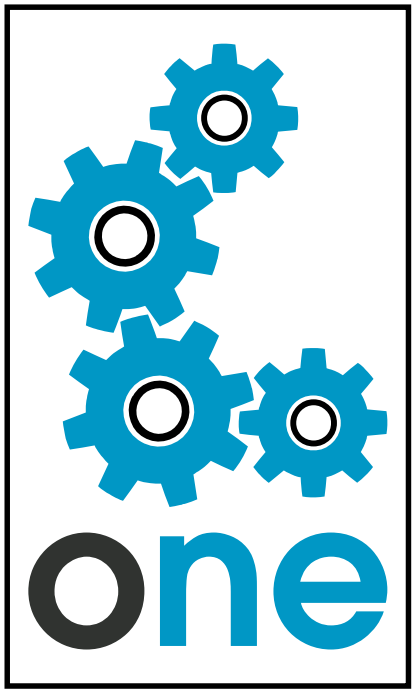

Cluster: the nodes

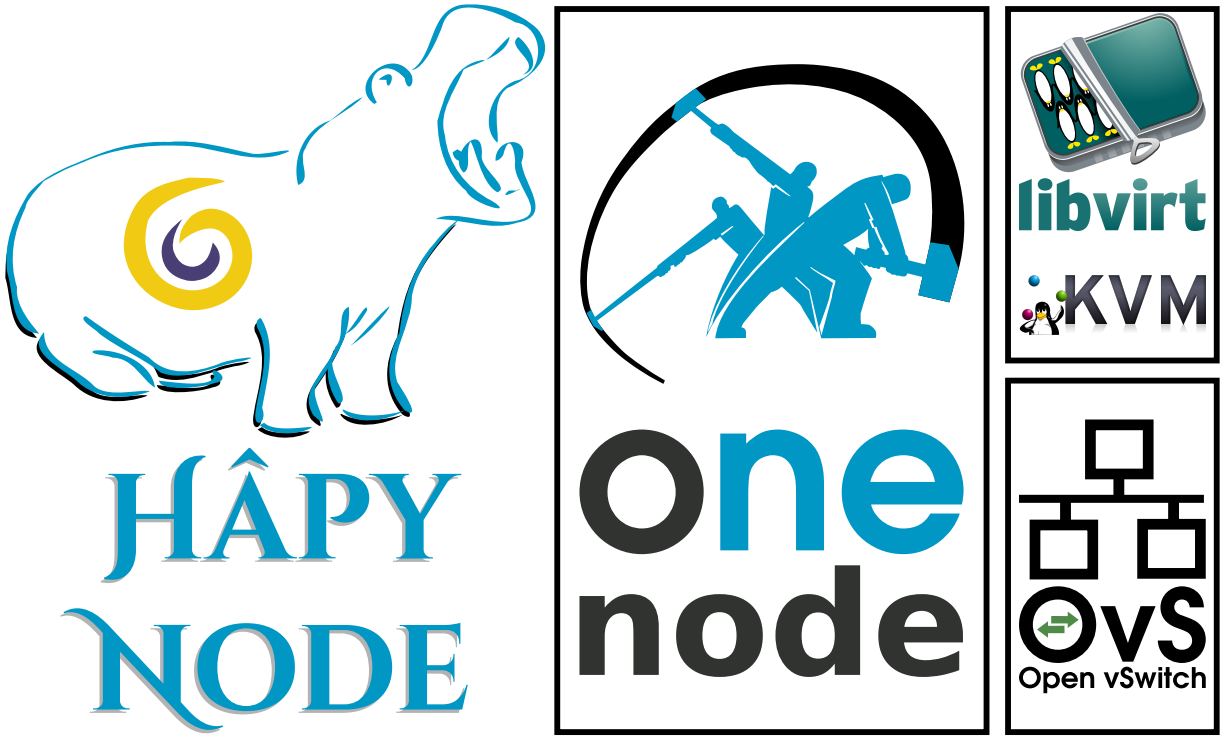

Hâpy deployment

La Réunion is the hâpy leader

- 10 schools deployed
- 120 schools in 3 years
- 4 virtual machines per Hâpy (Amon, horus, 2 proprietary OS)
- 12 cores CPU, 64GB RAM, 1.2TB SAS internal disks
- 2TB for backup (VMs + ONE database)

Summary

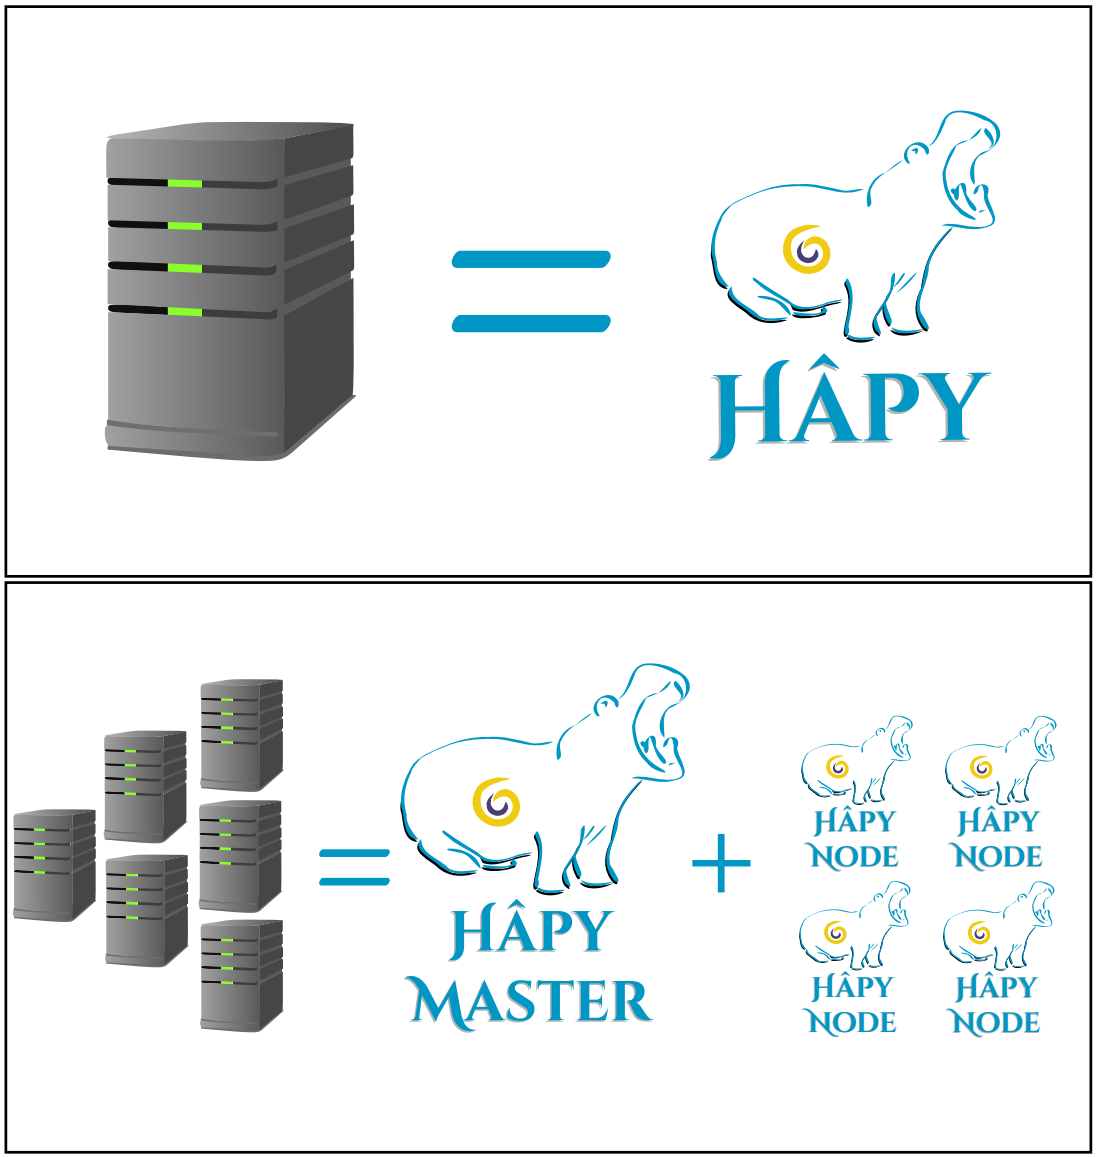

Evolutivity

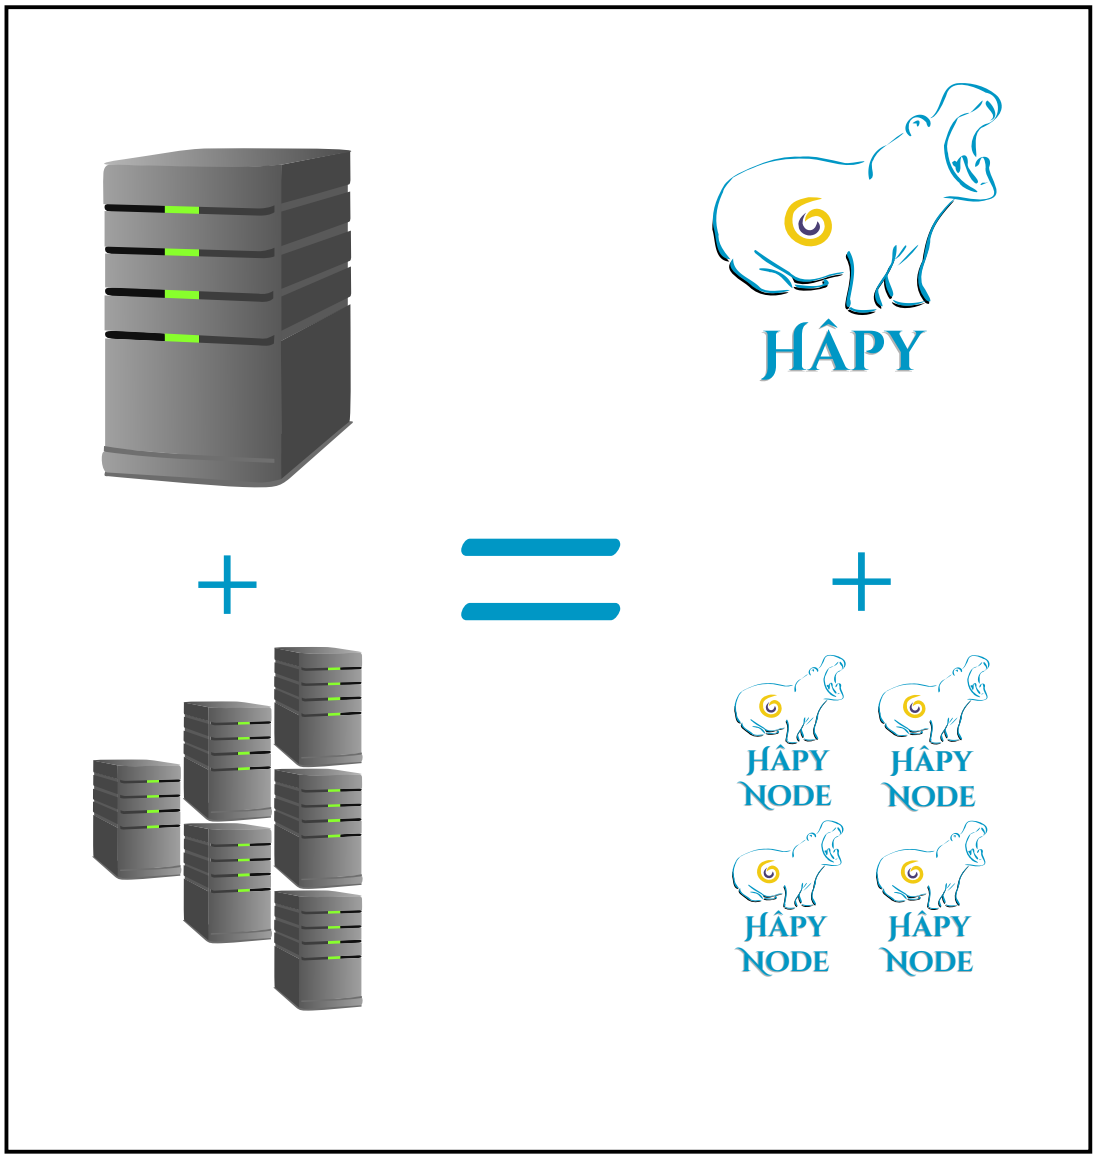

Outlook

- Integrate latest OpenNebula version
- Automatic migration of database on upgrade
- Use MySQL by default
- Support a distributed file system
- Better market support

Manage hundred of remote OpenNebula from a central console

Questions?

Thanks

Many thanks to the FOSS community for all the great software. So few things would exists without them.

This talk was realised with the help of the following libre software:

- Composition system LaTeX TeX Live
- The most powerful text editor available today GNU/Emacs
- The Awesome window manager
- The Universal Operating System Debian GNU/Linux




Licence

The slides are licensed under Creative Commons BY-NC-SA 2.0-FR

- Attribution
- Non Commercial
- Share alike
- France

You can obtain a copy of the license

by Internet

http://creativecommons.org/licenses/by-nc-sa/2.0/fr

by snail mail

Creative Commons
444 Castro Street, Suite 900 Mountain View,
California, 94041, USA.