

Building a GNU/Linux distribution with DevOps in mind

Daniel DEHENNIN

Pôle de Compétences Logiciels Libres

OpenNebulaConf 2016

CC BY-SA 4.0



Pôle de Compétences Logiciels Libres

FOSS and agility in french Minister of National Education

- Original mission \Rightarrow EOLE GNU/Linux meta-distribution
- CeCILL / GPL software licensing
- Agile consulting for other development teams

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

Started with OpenNebula 3.8

First uses

local workflow on central servers

- Each user was responsible for its own infrastructure
- Team documentation with conventions

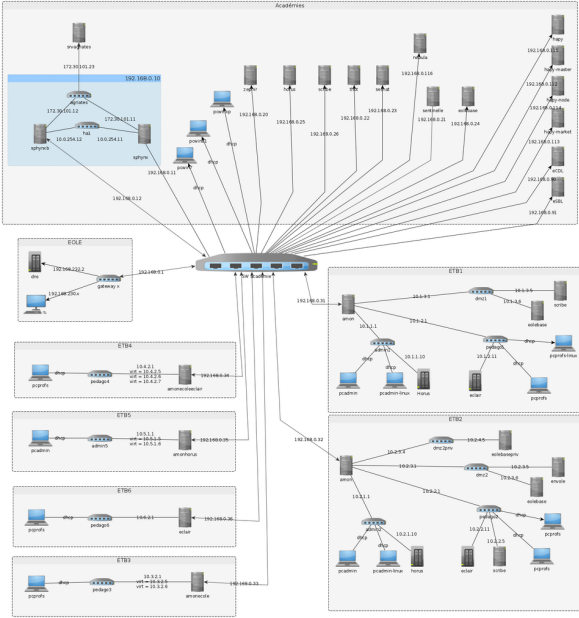
Everybody was admin

Load on higher skilled users

Emerged use cases and needs

- Short lifetime VMs
- Factorise VM templates and images
- Reduce usage cost for users
- Reproducible environments for QA
- Use production like environments for dev

Complete virtual infrastructure per user



Complete virtual infrastructure per user

workaround #2125

- VLAN isolated networks
- Standard network names
- One set of networks per user
- One user = one gateway

Sharing VM templates requires avoiding *UNAME* on networks

Automatic user environment generation

manual creation of 21 times 25 networks is not an option

- custom contextualisation of the gateway (per user IP)
- generate/update virtual networks

Pilot OpenNebula from Jenkins

continuous integration of OS

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

Produce VM images and templates at each step

Pilot OpenNebula from Jenkins

continuous integration of OS

The screenshot shows the Jenkins web interface. At the top, there's a navigation bar with the Jenkins logo, a search bar, and a 'log in' link. Below the navigation bar, there's a sidebar on the left with menu items: 'People', 'Build History', 'Forge Eole', 'Disk Usage', and 'Credentials'. The main content area is divided into two sections. The top section is the 'Build Queue', which is currently empty, showing 'No builds in the queue.' The bottom section is the 'Build Executor Status', which shows a list of 10 executors, all in an 'Idle' state. To the right of these sections is a large table displaying the build history. The table has columns for version (2.3 to 2.6.0), status (S, W), categorized job names, last success/failure times, and last duration. The jobs listed include 'Check ISO', 'FreshInstall', 'FreshInstallFromUbuntu', 'Daily', 'Instance', 'Importation base', 'Check Instance', 'Creole Lint', 'Module tests', 'Certificats', 'Pylint', 'Upgrade Auto', and 'Sauvegarde'.

2.3	2.4	2.4.1	2.4.2	2.5.0	2.5.1	2.5.2	2.6.0	Dev	En erreur	Test	Tous	deactive	infra	template
S	W	Categorized - Job	Last Success	Last Failure	Last Duration									
+	🟢	☀️ 00 : Check ISO	6 days 16 hr - #28	N/A	36 sec									
+	🟢	☀️ 10 : FreshInstall	6 days 10 hr - #22	6 days 19 hr - #16	13 sec									
+	🟡	☀️ 11 : FreshInstallFromUbuntu	5 mo 23 days - #14	6 mo 6 days - #17	15 min									
+	🟢	☀️ 20 : Daily	14 hr - #106	5 days 14 hr - #103	2 min 48 sec									
+	🔴	☀️ 30 : Instance	13 hr - #175	13 hr - #45	11 min									
+	🟢	☀️ 31 : Importation base	13 hr - #123	21 days - #108	17 min									
+	🟢	☀️ 50 : Check Instance	12 hr - #112	16 hr - #113	10 min									
+	🟢	☀️ 51 : Creole Lint	11 hr - #113	18 days - #17	1 min 35 sec									
+	🔴	☀️ 52 : Module tests	11 hr - #93	11 hr - #92	3 min 59 sec									
+	🟡	☀️ 53 : Certificats	N/A	N/A	N/A									
+	🟡	☀️ 54 : Pylint	4 mo 16 days - #1	N/A	4 min 14 sec									
+	🟡	☀️ 60 : Upgrade Auto	N/A	N/A	N/A									
+	🟢	☀️ 70 : Sauvegarde	10 hr - #102	7 days 10 hr - #64	4 min 22 sec									

Pilot OpenNebula from Jenkins

some numbers













- \approx 120 VMs per night
- \approx 90k VMs since 2014
- \approx 3TB of qcow2 images (25TB virtual)

Jenkins jobs produce ready to use VMs

reduce environment setup time

Template

etb1 2.6.0 ALL Labels

<p>etb1.eolebase-2.6.0rc2</p>  <p>...</p> <p>system</p>	<p>etb1.scribe-2.6.0rc2</p>  <p>...</p> <p>system</p>	<p>etb1.scribe-2.6.0rc2-Daily</p>  <p>...</p> <p>system</p>	<p>etb1.scribe-2.6.0rc2-AveclmportSconet-Auto</p>  <p>...</p> <p>system</p>
<p>etb1.eclairdmz-2.6.0rc2</p>  <p>...</p> <p>system</p>	<p>etb1.eclairdmz-2.6.0rc2-Daily</p>  <p>...</p> <p>system</p>	<p>etb1.eclairdmz-2.6.0rc2-Instance-Auto</p>  <p>...</p> <p>system</p>	<p>etb1.eolebase-2.6.0rc2-Daily</p>  <p>...</p> <p>system</p>
<p>etb1.horus-2.6.0rc2</p>  <p>...</p> <p>system</p>	<p>etb1.horus-2.6.0rc2-Daily</p>  <p>...</p> <p>system</p>	<p>etb1.dcpedago-2.6.0rc2</p>  <p>...</p> <p>system</p>	<p>etb1.dcpedago-2.6.0rc2-Daily</p>  <p>...</p> <p>system</p>

Development is so simple

until an IA will make our work

Preparing the coding session

- 1 Start an infrastructure
- 2 Create a branch of your repository
- 3 Clone the repository on the VM

Development is so simple

until an IA will make our work

Hack until it works

- ① Code locally on your workstation
- ② Pull the changes on the VM
- ③ `make install`
- ④ Test and cycle to ① until it works

Development is so simple

until an IA will make our work

Test like a user

- 1 Cleanup your local branch
- 2 Merge, push and make a package
- 3 Cleanup the VM
 - DELETE-RECREATE (< 5.0)
 - Revert to initial disk snapshot (≥ 5.0)
- 4 Upgrade OS \Rightarrow new package
- 5 Test

Jenkins jobs will install the new packages during next night

QA campaign

Automation is not the panacea

- Squash-TM
- Targeted environments
- Critical features

Physical limitations

test bed was burning

≈ 40k VMs

- The two dedicated workstations was fine for testing
- Workload was memory bounded ⇒ bumped to 2x32GB

NFS access on workgroup NAS was too slow

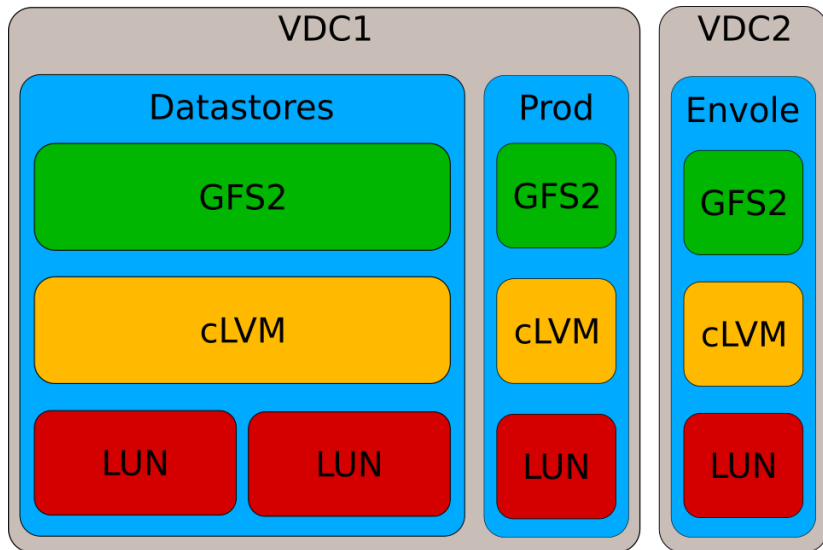
New hypervisor nodes

blade runners

- First VDC
 - Two servers
 - 48 cores
 - 384GB RAM
- Second VDC
 - One server
 - 8 cores
 - 42GB RAM

Storage on SAN

the corosync/pacemaker/cLVM/GFS2 sandwich



Storage on SAN

Everybody has a gun and too many want to use it

corosync/pacemaker can be challenging

Hot/cold storage \Rightarrow I/Os not striped on all LUNs

The future

- Replace home made orchestration code with SaltStack
- Host other teams
- Ceph storage
- Docker uses

The killer features of OpenNebula

you can't make us use something else without them

- Low load on the team
- We can adapt to our use cases and **contribute**
 - Features 3/17
 - Bugs 4/46

Muchas gracias **OpenNebula Systems**

Thanks

Many thanks to the FOSS community for all the great software. So few things would exist 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-SA 4.0**

- Attribution
- Share alike

You can obtain a copy of the license

by Internet

<http://creativecommons.org/licenses/by-nc-sa/4.0>

by snail mail

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