



5th SIG-CISS meeting

# Building a Public Cloud from scratch: Get to 8K cores OpenStack

Cloud Engineering Team @ GRNET





# Cloud Engineering Team




Stamatis Katsaounis



Alexandros Soumplis



Giannis Polizois



Yanos Aggelopoulos



Aggelos Kolaitis



# Louros DC

A green data center





# Datacenter @ Louros



**09/2018**  
Non-production with  
numerous broken servers

Created by Neson Booth  
from Noun Project



**09/2019**  
Full in production with  
very few broken servers

Created by Neson Booth  
from Noun Project

## 10 Racks

### 198 VMC Lenovo RD550

- 40 cores @ 2.6G
- 192 GB RAM
- 2 x 280 GB SAS
- 2 x 10G NICs

## 5 Availability Zones

### 22 eVMC Lenovo RD550

- 40 cores @ 2.6G
- 384 GB RAM
- 2 x 280 GB SAS,  
2 x 380 GB SSD
- 2 x 10G NICs

## 5 Availability Zones

### 50 SC Lenovo RD550

- 40 cores @ 2.6G
- 128 GB RAM
- 2 x 280 GB SAS  
6 x 180 GB SSD  
12 x 4 TB NL-SAS
- 2 x 10G NICs

# Critical Design Decisions

# Principles



Stack decomposition and service isolation



Isolate Control Plane from Data Plane



Utilize metal with LXD containers and KVM pods



Follow the upstream path. Use techniques and do minor compromises to be fully compatible with upstream releases.



Resiliency is TOP priority. Redundancy, high availability, automatic failover



Automation for anything that has to be done twice



Scalability, scale out instead of scale up

**Be part of the  
community**

**build the  
community**

# What is OpenStack?





“OpenStack is a **cloud operating system** that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed and provisioned through APIs with common authentication mechanisms” [1]

[1] <https://www.openstack.org/software/>

# OpenStack Projects



**NOVA**

*an OpenStack Community Project*




**NEUTRON**

*an OpenStack Community Project*



**GLANCE**

*an OpenStack Community Project*



**KEYSTONE**

*an OpenStack Community Project*



**CINDER**

*an OpenStack Community Project*



**HORIZON**

*an OpenStack Community Project*



**WATCHER**

*an OpenStack Community Project*



**TELEMETRY**


*an OpenStack Community Project*



**DESIGNATE**

*an OpenStack Community Project*

# The Triplets



# The Outlier




**NOVA**  
an OpenStack Community Project



# Building Blocks





# Deployment



# Foundation Tools

## MaaS (Metal-as-a-Service)

The screenshot shows the MAAS web interface. At the top, there's a navigation bar with 'MAAS' and various menu items like 'Machines', 'Devices', 'Controllers', 'Pods', 'Images', 'DNS', 'AZs', 'Subnets', 'Settings', 'root', and 'Logout'. Below the navigation bar, it says 'Machines 299 machines available' and 'Add hardware'. There are tabs for '299 Machines' and '1 Resource pool'. A search bar is present with 'status:(=Deployed)'. Below that is a table of machines with columns for FQDN/MAC/IP, POWER, STATUS, OWNER/TAGS, ZONE/SPACES, CORES/ARCH, RAM, DISKS, and STORAGE. The table shows several machines in a 'Deployed' state, all running 18.04 LTS.

FQDN   MAC IP	POWER	STATUS	OWNER TAGS	ZONE SPACES	CORES ARCH	RAM	DISKS	STORAGE
lar1330.srv-louros.gr... 10.0.253.75 (+4)	On Ipmi	18.04 LTS	root foundation-node...	az3 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar0730.srv-louros.gr... 10.0.253.61 (+4)	On Ipmi	18.04 LTS	root prod, az2, found...	az2 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar2033.srv-louros.gr... 10.0.253.71 (+4)	On Ipmi	18.04 LTS	root os-str, prod, sc, a...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar2030.srv-louros.gr... 10.0.254.134 (+4)	On Ipmi	18.04 LTS	root az3, foundation-...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar2036.srv-louros.gr... 10.0.253.65 (+4)	On Ipmi	18.04 LTS	root sc, os-str, prod, f...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar1930.srv-louros.gr... 10.0.253.65 (+4)	On Ipmi	18.04 LTS	root foundation-node...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar2027.srv-louros.gr... 10.0.253.63 (+4)	On Ipmi	18.04 LTS	root prod, foundation...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar1936.srv-louros.gr... 10.0.253.62 (+4)	On Ipmi	18.04 LTS	root sc, az5, prod, os-...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB
lar1933.srv-louros.gr... 10.0.253.64 (+4)	On Ipmi	18.04 LTS	root sc, foundation-n...	az5 5 spaces	32 amd64	128 GiB	15	49.5 TB

## Juju

The screenshot shows the Juju web interface. At the top, it says 'Juju / admin / openstack'. Below that, it shows '62 applications' and '172 machines'. A sidebar on the left lists various applications with their counts: aodh (3), cellometer (3), cellometer-agent (160), ceph-mon (3), ceph-osd (33), ceph-radosgw (3), cinder (3), cinder-ceph (3), designate (3), designate-bind (2), elasticsearch (2), filebeat (224), glance (3), graylog (1), graylog-mongodb (1), hadcluster-aodh (3), and hadcluster-cinder (3). The main area displays a complex network diagram with nodes representing machines and applications, connected by lines representing network links.

# MaaS @ Louros



iPXE & DHCP



DNS & HTTP



HTTP Proxy



IPAM



KVM Pods

## ▼ Machines

Max

299

Machines

299

Physical

262

KVM hosts

7

Virtual Machines

37

Devices

112

## ▼ Networks

Spaces

16

Fabrics

14

VLANs

29

IPv4 subnets

17

Static & Dynamic IPs

590

Reserved IPs (total)

510



# Juju

## What is it anyway ?!

- Open source application modeling tool developed and maintained by Canonical Ltd
- Focus on applications and their relations
- Create a development environment on your laptop, then recreate that environment on the public cloud, onto bare metal servers or into a Kubernetes cluster
- Create **repeatable**, **systematic** and **secure** devops practices for all stages of your product's lifecycle
- Juju simplifies deployment, maintenance, scaling up and winding down

# Juju

ubuntu@cfp001:~\$ juju models  
Controller: louros-maas

Model	Cloud/Region	Type	Machines	Cores	Units
Controller	maas_cloud	maas	5	24	8
Kubernetes	maas_cloud	maas	100	400	23
<b>openstack*</b>	<b>maas_cloud</b>	<b>maas</b>	<b>231</b>	<b>6384</b>	<b>1859</b>
stgopenstack	maas_cloud	maas	83	1076	483

# Juju

Our production bundle is  
**~5500 lines !!!**

... and counting on

# Juju

Remains very simple to manage  
configuration for applications and  
relations

# Network Setup

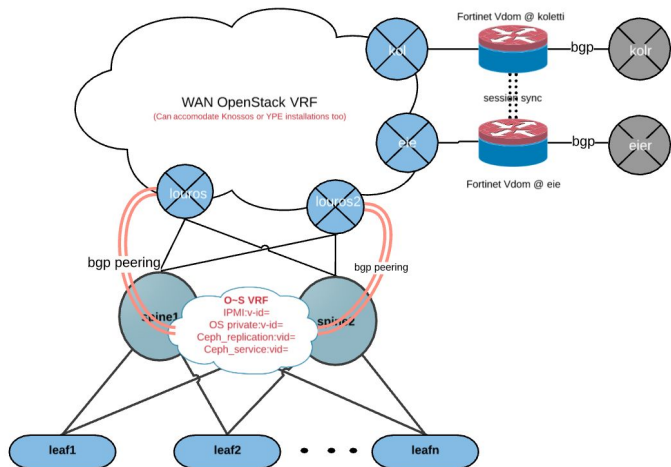


# Physical Network

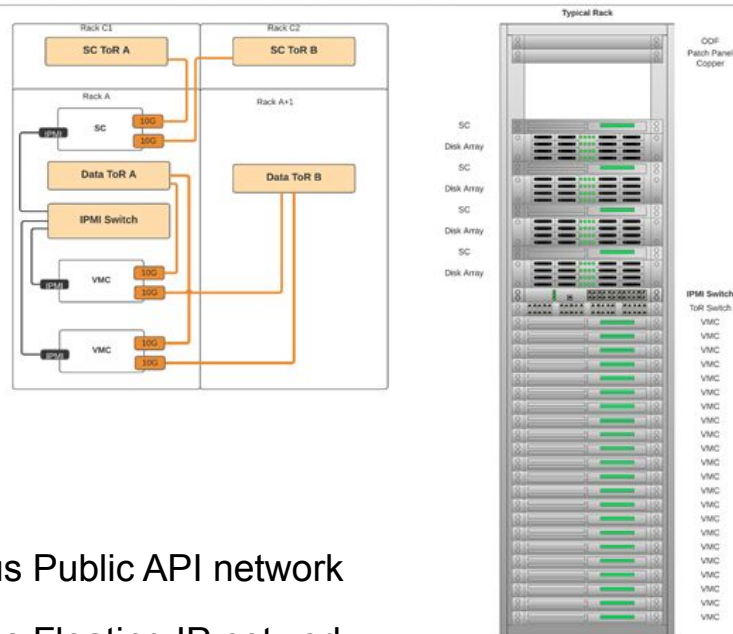
2 x 10 Gbps LACP bond per server

Jumbo Frames

Layer-2 access lists per VLAN



GRNET DATA CENTER - TYPICAL RACK & CONNECTIVITY



Autonomous Public API network

Autonomous Floating IP network

Private OAM / Internal networks

Internet access only through proxy

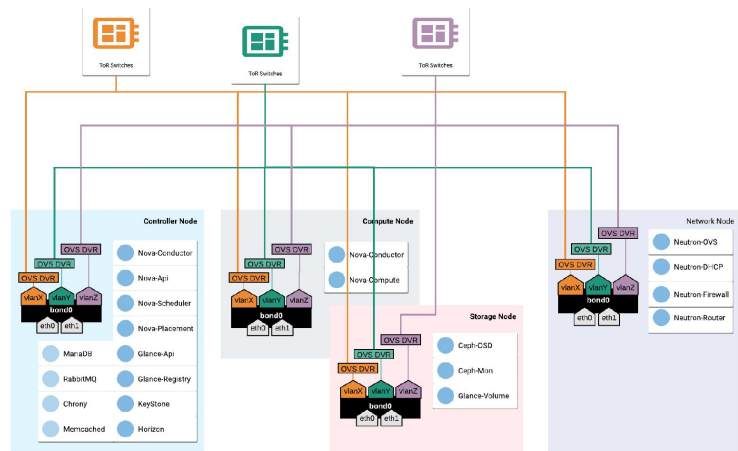
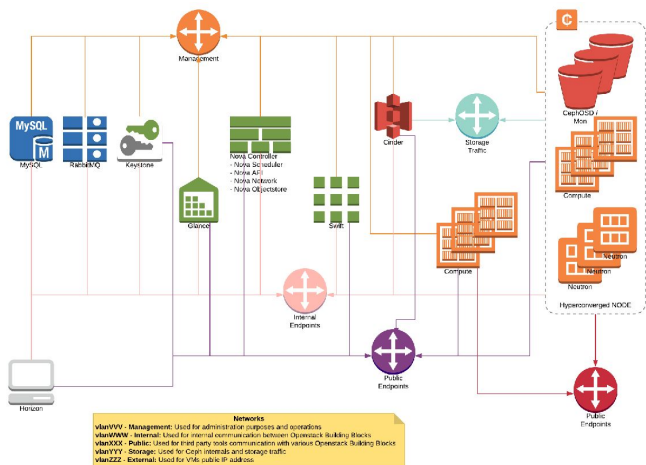


# Virtual Network

Dedicated Neutron gateways

Openvswitch & VxLANs

No DVR at the moment



Self Service tenant networks

Provider networks for Layer-2 VPNs

Be careful of MTUs !

# Observability

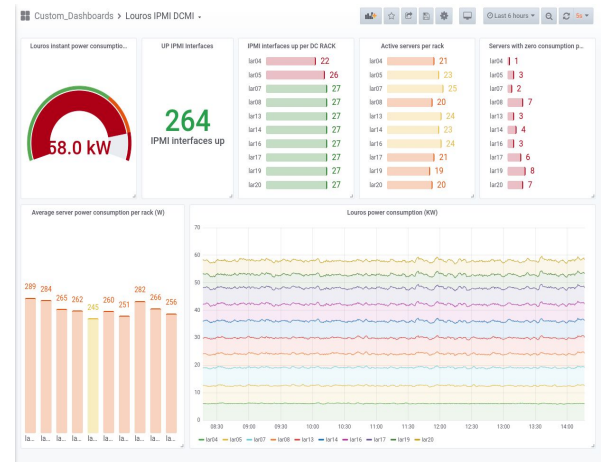
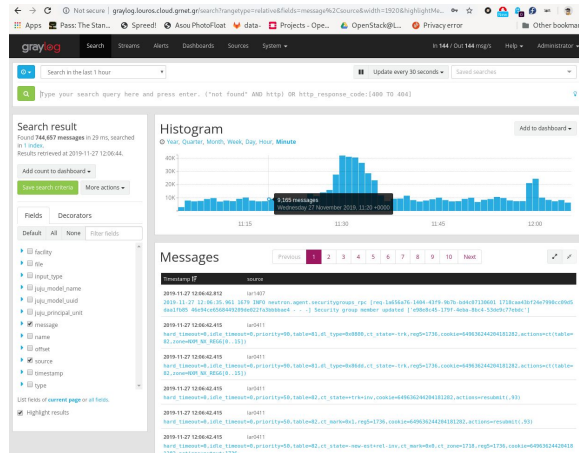
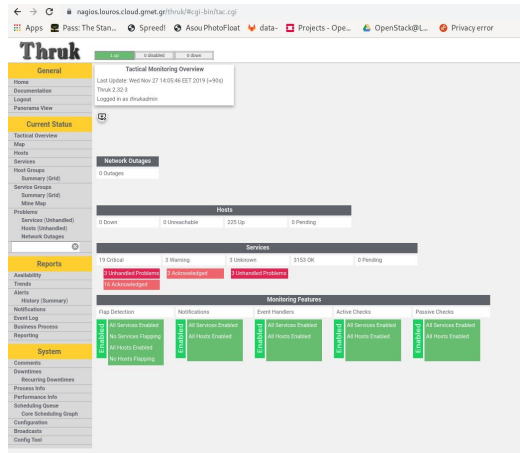


# Observability

Health Check  
Nagios

Logs  
Filebeat & Graylog

Metrics  
Prometheus & Grafana



# Observability

## Netdata

Real time metrics (1 sec granularity)

Very useful for debugging purposes

No need for long term metrics

Out of the box



# Ecosystem

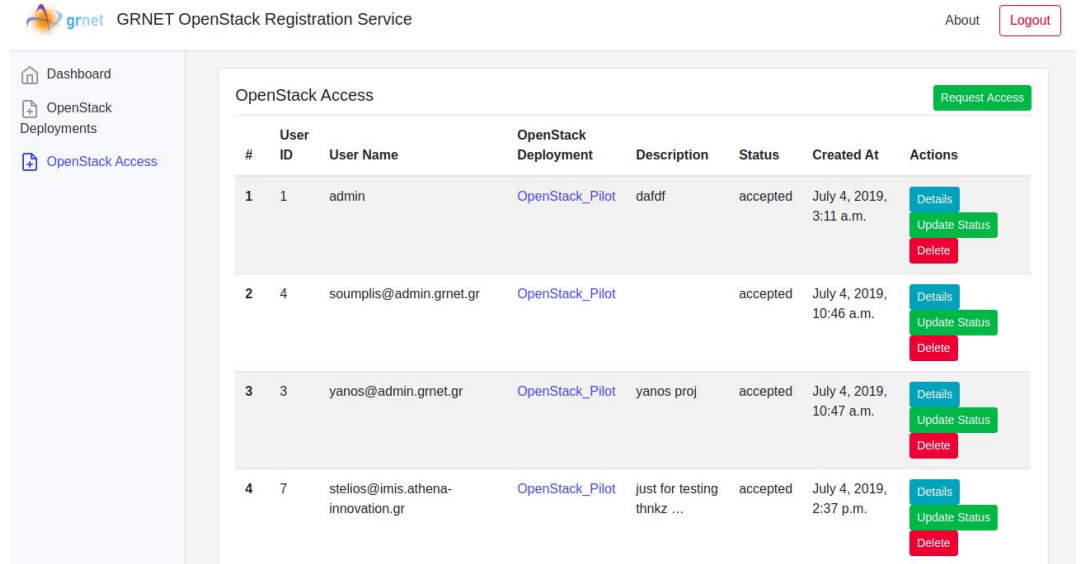


# Registration Platform

Django application

SAML integration

Ansible backend to create users in Keystone

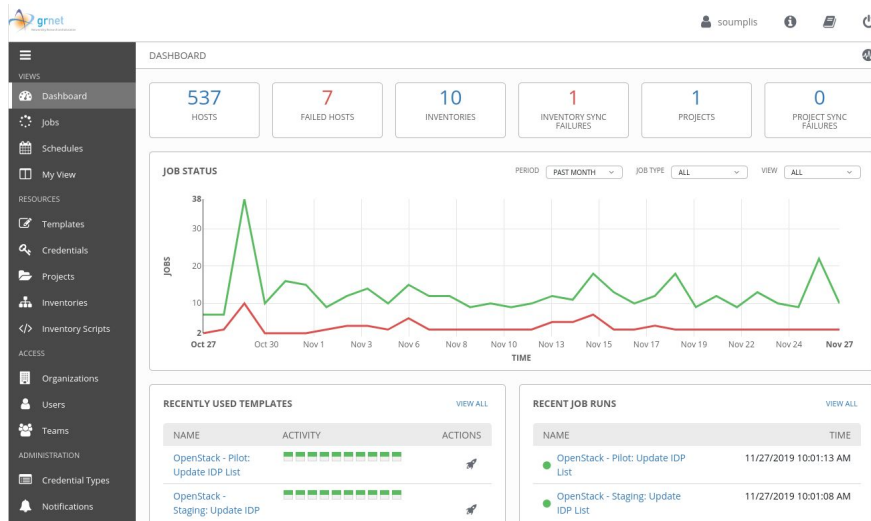


The screenshot displays the GRNET OpenStack Registration Service interface. The page title is "GRNET OpenStack Registration Service" and includes links for "About" and "Logout". A sidebar on the left contains navigation options: "Dashboard", "OpenStack Deployments", and "OpenStack Access". The main content area is titled "OpenStack Access" and features a "Request Access" button. Below this is a table listing access records.

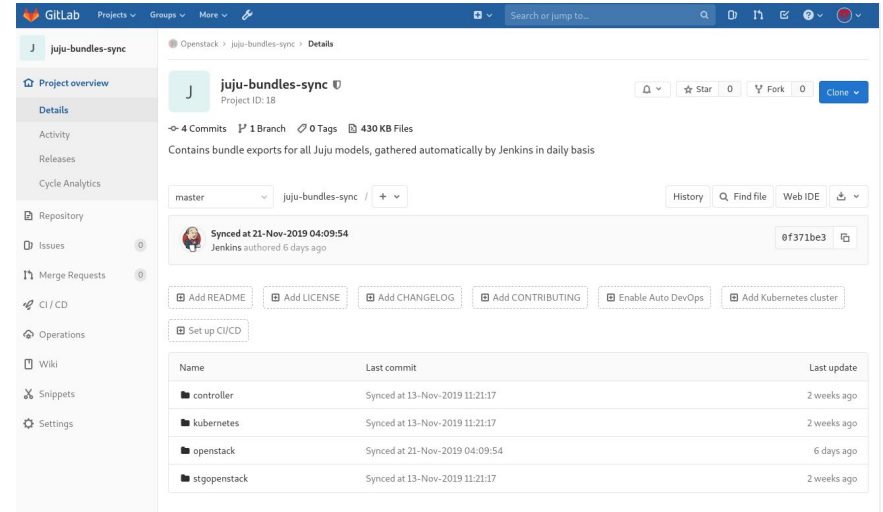
#	User ID	User Name	OpenStack Deployment	Description	Status	Created At	Actions
1	1	admin	OpenStack_Pilot	dafdf	accepted	July 4, 2019, 3:11 a.m.	<a href="#">Details</a> <a href="#">Update Status</a> <a href="#">Delete</a>
2	4	soumplis@admin.grnet.gr	OpenStack_Pilot		accepted	July 4, 2019, 10:46 a.m.	<a href="#">Details</a> <a href="#">Update Status</a> <a href="#">Delete</a>
3	3	yanos@admin.grnet.gr	OpenStack_Pilot	yanos proj	accepted	July 4, 2019, 10:47 a.m.	<a href="#">Details</a> <a href="#">Update Status</a> <a href="#">Delete</a>
4	7	stelios@imis.athena-innovation.gr	OpenStack_Pilot	just for testing thinkz ...	accepted	July 4, 2019, 2:37 p.m.	<a href="#">Details</a> <a href="#">Update Status</a> <a href="#">Delete</a>

# Configuration Management and Code Repositories

## Ansible & AWX



## GitLab



# Jenkins for Ops

Safe job execution  
(ex. Package update)

Complex jobs simplified  
(ex. NetBox validation)


Time consuming jobs  
(ex. Firmware upgrade)

netbox-validation-checks < 126

Branch: — 58m 1s No changes

Commit: — 7 hours ago Started by timer

Pipeline Changes Tests Artifacts Logout



Daily NetBox Validation / switch-connections-check - <1s

[Restart Daily NetBox Validation](#)

✓ > docker exec maas2netbox maas2netbox -c validate --log INFO -f switch\_connections — Shell Script

48m 48s

# Patchman

Central package control

Simple bash agent for clients

Read only for hosts

Easy filtering for OS, packages,  
repositories



Patchman Dashboard Hosts Repositories Packages Operating Systems Reports Django Admin Welcome, soumplis / Log out

Home

Patch Management Dashboard for tools.grnet.gr

323 Hosts require a reboot for security reasons (e.g. kernel update)

62 Hosts have bugfix-only updates pending

Hostname	Updates	Running Kernel	OS	Last Report	Reboot Status
backup.srv-louros.grnet.gr	8	4.15.0-60-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
cfp001	8	4.15.0-70-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-436add-26-lxd-14.srv-louros.grnet.gr	8	4.15.0-66-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-436add-26-lxd-15.srv-louros.grnet.gr	8	4.15.0-66-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-436add-26-lxd-16	8	4.15.0-66-generic	Ubuntu 18.04	Nov. 26, 2019, 2:19 p.m.	⊘
juju-436add-26-lxd-2.srv-louros.grnet.gr	12	4.15.0-66-generic	Ubuntu 18.04	Nov. 27, 2019, 1:34 p.m.	⊘
juju-436add-26-lxd-3.srv-louros.grnet.gr	7	4.15.0-66-generic	Ubuntu 18.04	Nov. 27, 2019, 1:33 p.m.	⊘
juju-436add-32-lxd-1.srv-louros.grnet.gr	8	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-436add-32-lxd-2.srv-louros.grnet.gr	12	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 1:35 p.m.	⊘
juju-436add-32-lxd-3.srv-louros.grnet.gr	7	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 1:33 p.m.	⊘
juju-436add-32-lxd-7.srv-louros.grnet.gr	8	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-436add-69-lxd-0.srv-louros.grnet.gr	12	4.15.0-55-generic	Ubuntu 18.04	Nov. 27, 2019, 1:35 p.m.	⊘
juju-436add-69-lxd-1.srv-louros.grnet.gr	7	4.15.0-55-generic	Ubuntu 18.04	Nov. 27, 2019, 1:35 p.m.	⊘
juju-436add-69-lxd-11.srv-louros.grnet.gr	8	4.15.0-55-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-436add-69-lxd-15.srv-louros.grnet.gr	8	4.15.0-55-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-980f92-92-lxd-16.srv-louros.grnet.gr	8	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-980f92-92-lxd-31.srv-louros.grnet.gr	8	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-980f92-92-lxd-39.srv-louros.grnet.gr	8	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-980f92-92-lxd-40.srv-louros.grnet.gr	8	4.15.0-58-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-980f92-94-lxd-10.srv-louros.grnet.gr	8	4.15.0-60-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
juju-980f92-95-lxd-12.srv-louros.grnet.gr	8	4.15.0-66-generic	Ubuntu 18.04	Nov. 27, 2019, 7 a.m.	⊘
lar0412.srv-louros.grnet.gr	8	4.15.0-55-generic	Ubuntu 18.04	Nov. 27, 2019, 6:57 a.m.	⊘
lar0415.srv-louros.grnet.gr	9	4.15.0-70-generic	Ubuntu 18.04	Nov. 27, 2019, 7:04 a.m.	⊙



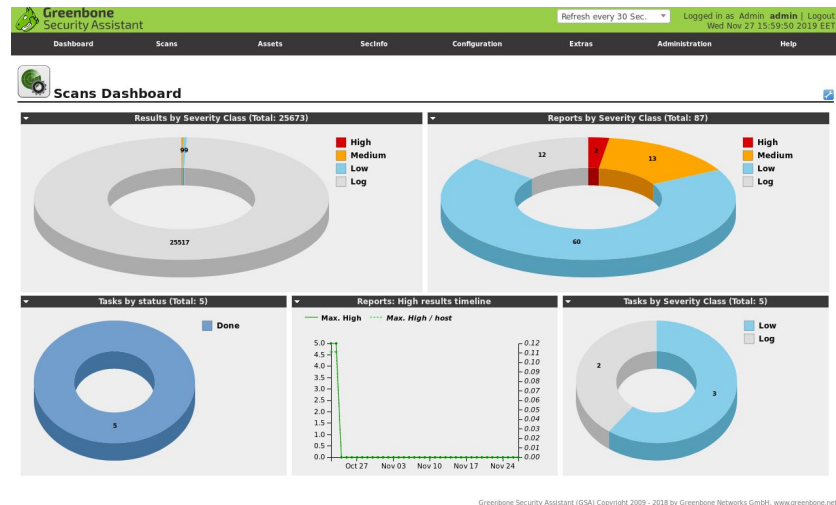
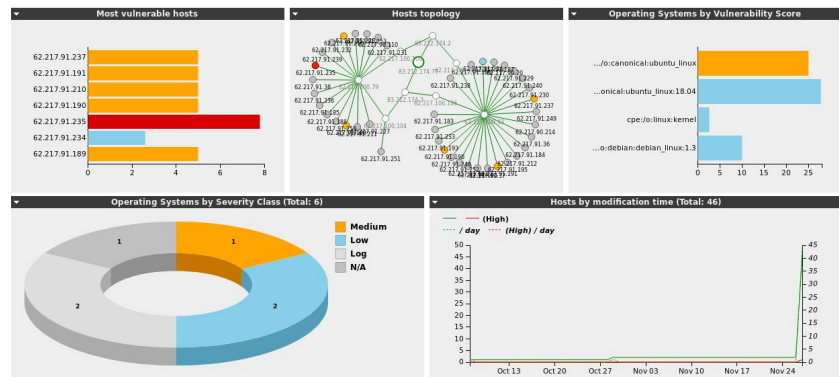
# Vulnerability Scanning

Public Perimeter - Discovery

Public Perimeter - Penetration Test



## Assets Dashboard



Greenbone Security Assistant (GSA) Copyright 2009 - 2018 by Greenbone Networks GmbH, www.greenbone.net

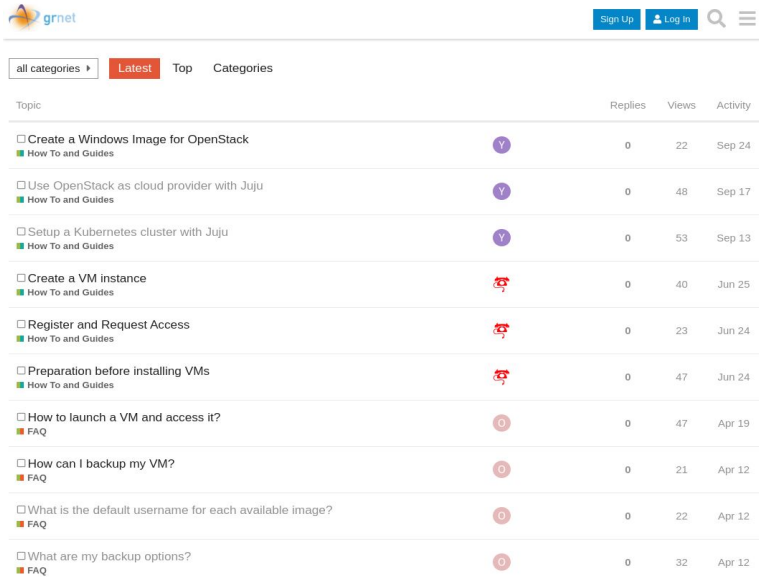
Public APIs - Discovery

Public APIs - Deep Inspection



# User support

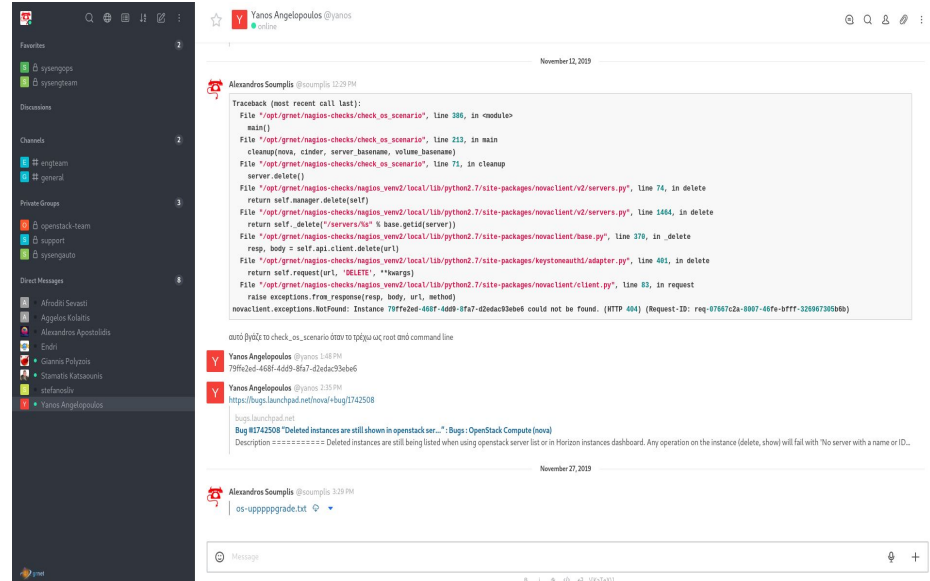
## Self Service - Tier 0



The screenshot shows the grnet website's self-service section. It features a navigation bar with 'Sign Up' and 'Log In' buttons, and a search icon. Below the navigation, there are tabs for 'all categories', 'Latest', 'Top', and 'Categories'. The main content area displays a list of articles with columns for 'Topic', 'Replies', 'Views', and 'Activity'. Each article includes a title, a 'How To and Guides' tag, a status icon (Y or D), and a date.

Topic	Replies	Views	Activity
<input type="checkbox"/> Create a Windows Image for OpenStack How To and Guides	0	22	Sep 24
<input type="checkbox"/> Use OpenStack as cloud provider with Juju How To and Guides	0	48	Sep 17
<input type="checkbox"/> Setup a Kubernetes cluster with Juju How To and Guides	0	53	Sep 13
<input type="checkbox"/> Create a VM instance How To and Guides	0	40	Jun 25
<input type="checkbox"/> Register and Request Access How To and Guides	0	23	Jun 24
<input type="checkbox"/> Preparation before installing VMs How To and Guides	0	47	Jun 24
<input type="checkbox"/> How to launch a VM and access it? FAQ	0	47	Apr 19
<input type="checkbox"/> How can I backup my VM? FAQ	0	21	Apr 12
<input type="checkbox"/> What is the default username for each available image? FAQ	0	22	Apr 12
<input type="checkbox"/> What are my backup options? FAQ	0	32	Apr 12

## Instant Help - Tier 1



The screenshot shows a Slack chat conversation. On the left, a sidebar lists channels like #sysmops and #sysmgmt, and a list of direct messages including 'Alexandros Apostolidis'. The main chat area shows a message from Alexandros Souplis at 12:29 PM. The message contains a traceback from a Python script, a command to check a scenario, and a link to a bug report. A response from Yanos Angelopoulos at 1:48 PM provides a link to the bug report and a note that deleted instances are still shown in the OpenStack dashboard. A second message from Alexandros Souplis at 3:29 PM shows a file named 'os-upppgrade.txt'.

```
Traceback (most recent call last):
  File "/opt/grnet/nagios-checks/check_os_scenario", line 386, in <module>
    main()
  File "/opt/grnet/nagios-checks/check_os_scenario", line 213, in main
    cleanup(memo, cinder, server_baseName, volume_baseName)
  File "/opt/grnet/nagios-checks/check_os_scenario", line 71, in cleanup
    server.delete()
  File "/opt/grnet/nagios-checks/nagios_vemv2/local/lib/python2.7/site-packages/novaclient/v2/servers.py", line 74, in delete
    return self.manager.delete(self)
  File "/opt/grnet/nagios-checks/nagios_vemv2/local/lib/python2.7/site-packages/novaclient/v2/servers.py", line 1464, in delete
    return self._delete("servers/%s" % base.getId(server))
  File "/opt/grnet/nagios-checks/nagios_vemv2/local/lib/python2.7/site-packages/novaclient/base.py", line 376, in _delete
    resp, body = self.api.client.delete(url)
  File "/opt/grnet/nagios-checks/nagios_vemv2/local/lib/python2.7/site-packages/keystonemiddleware/adaptor.py", line 401, in delete
    return self.request(url, "DELETE", **kwargs)
  File "/opt/grnet/nagios-checks/nagios_vemv2/local/lib/python2.7/site-packages/novaclient/client.py", line 83, in request
    raise exceptions.from_response(resp, body, url, method)
novaclient.exceptions.NotFound: Instance 79ff2e6d-468f-46d9-8fa7-d2edac35be6 could not be found. (HTTP 404) (Request-ID: req-87667c2a-8907-46fa-8f7b-32869739596b)
```

curl \$(cat /dev/urandom | tr -dc 'a-z0-9' | fold -n 40 | xargs | sh) to check\_os\_scenario dir to /opt/ux-root and command line

**Yanos Angelopoulos** @yanos 1:48 PM  
79ff2e6d-468f-46d9-8fa7-d2edac35be6

**Yanos Angelopoulos** @yanos 3:25 PM  
https://bugs.launchpad.net/nova/+bug/1742508

bugs.launchpad.net  
**Bug #1742508 "Deleted instances are still shown in openstack ser...": Bugs: OpenStack Compute (nova)**  
Description ===== Deleted instances are still being listed when using openstack server list or in Horizon instances dashboard. Any operation on the instance (delete, show) will fail with 'No server with a name or ID...


November 27, 2019

**Alexandros Souplis** @souplis 3:29 PM  
os-upppgrade.txt

# Giving back to Community




# Community - Our contributions




Created by Josh Sorosky  
from Noun Project

## 17 Upstream Patches

- 2 @ OpenStack projects
- 8 @ OpenStack charms
- 3 @ Nagios charms
- 2 @ Juju / MaaS core
- 2 @ Prometheus IPMI exporter



Created by Josh Sorosky  
from Noun Project



Created by Josh Sorosky  
from Noun Project

## 19 Upstream Bug Reports

- 6 @ OpenStack projects
- 5 @ OpenStack charms
- 3 @ Nagios charms
- 3 @ Juju / MaaS core
- 2 @ Other

## 3 Projects @ GRNET GitHub

- Ansible Inventory Server
- MaaS to NetBox
- MaaS / Juju Toolkit

Thank You !

