



- RO-LCG 2018, Cluj-Napoca, 17-19 October 2018 -



# Advanced computing support for non-HEP research communities at IFIN-HH

**Ionuț Vasile, Dragoș Ciobanu-Zabet, Mihnea Dulea**

Department of Computational Physics and Information Technologies (DFCTI)  
Horia Hulubei National Institute for R&D in Physics and Nuclear Engineering  
(IFIN-HH)

Bucharest-Magurele, Romania



18/10/2018

CONDEGRID



# Introduction

- The Operations Centre of the Romanian National Grid Infrastructure (NGI-RO) currently provides a broad spectrum of computing resources and services to various non high-energy physics research communities.
- computing paradigms available: HPC (MPI/CPU and CUDA/GPGPU), grid computing, and cloud - which facilitates the provision, allocation and delivery of resources on demand for scientific researchers
- the sites managed by the Operations Centre support scientific areas of major interest in the national and European research strategies, such as the physics of intense laser – nuclear matter interaction, computational biology, nanomaterials technology and quantum optics devices.
- This talk overviews the recent developments regarding the support provided through the infrastructure of the GRIDIFIN and CLOUDIFIN sites



# A bit of history

- The development of the current advanced computing infrastructure has started more than a decade ago with the implementation of its first grid computing cluster for HEP research and its first HPC cluster in 2006.
- The infrastructure evolved periodically:



# European Open Science Cloud

- EOSC architecture includes: federated core of shared resources and multiple federated data infrastructures
- EOSC-Hub project has the goal to develop the initial federated core of shared resources
- IFIN-HH participates to the EOSC-Hub project through CLOUDIFIN, as provider of services for supporting existing or future user communities that will be featured in the EOSC-hub engagement plan as well as contributing with national services to the EOSC-Hub service catalog in accordance with the EOSC-Hub engagement rules





# Resources and services

## NGI-RO Operations Centre

### GRIDIFIN (NGI\_RO/HTC)

CREAM-CE (HTC)

Argus + MyProxy

SE - Storage (6 TB)

User Interface for any site, including CLOUDIFIN

Central LFC for non-HEP VOs

Accounting + Basic Monitoring

### External Services

VOMS for non-HEP communities

EGI Services: cASO, APEL-SSM, ARGO, ...

Top-BDII

RomanianGRID CA (ROSA - EuGridPMA)

### GRIDIFIN (HPC) for non-HEP communities

ARC-CE (HPC - MPI)

CREAM-CE (HPC - MPI)

Storage for HPC (65 TB)

CVMFS repository (testing)

Local LFC

Software repository (UMD, Linux)

### CLOUDIFIN IaaS cloud resources

Keystone-VOMS

APEL SSM

Block Storage 65TB

Cinder

Cloud controller

cASO

Glance

Compute Nodes

Nova

Neutron

Nova

Neutron

CloudKeeper

EGI AppDB VA Management

# CLOUD INFRASTRUCTURE

1<sup>st</sup> Romanian EGI certified cloud computing site CLOUDIFIN

Purpose:

- to provide Infrastructure-as-a-Service (IaaS) - virtual machines - to the national research and education community;
- to support within the EGI Federated Cloud the European ESFRI projects in which Romania is involved
- to participate to EOSC


CLOUDIFIN provides computing and storage resources for *fedcloud.egi.eu* (for EGI Federated Cloud), *eli-np.eu* and *biomed*.

And also *eli-laser.eu* (for the future ELI ERIC) and *ronbio.ro* (for computational biology).

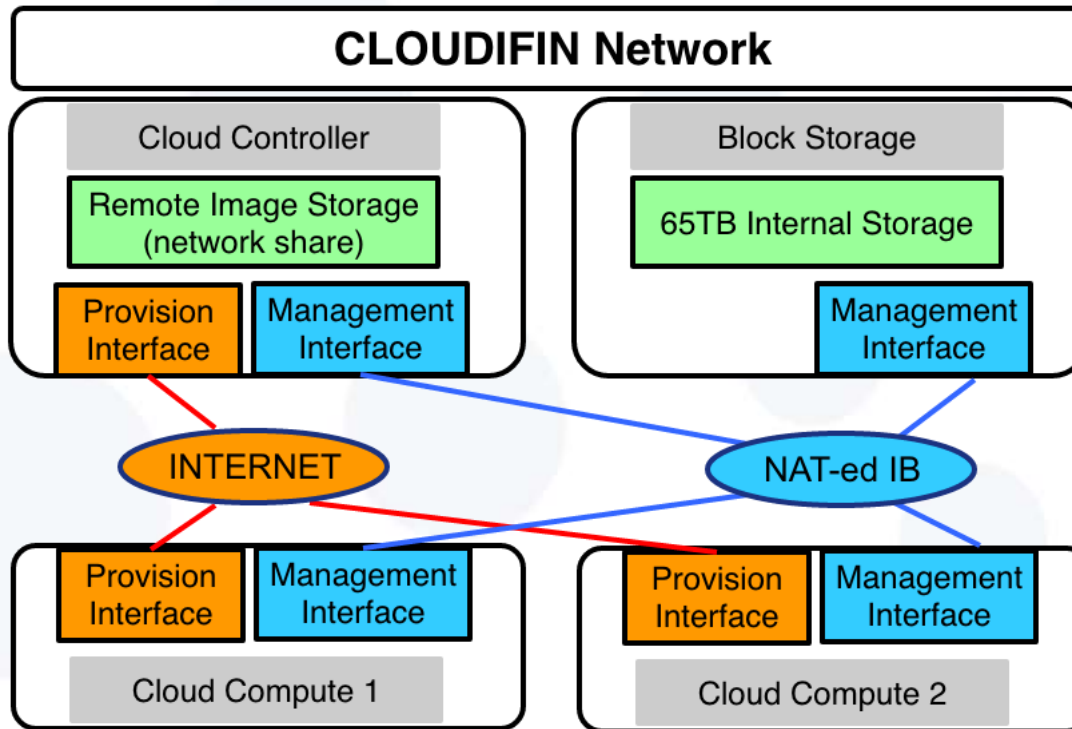
Upgraded to Centos 7.5 and OpenStack Pike.

GOCDDB registered services and hosts:

- Cloud controller (*cloud-ctrl.nipne.ro*):  
*eu.egi.cloud.vm-management.occi*, *eu.egi.cloud.accounting*
- Site BDII (*cloud-bdii.nipne.ro*): *eu.egi.cloud.information.bdii*



# CLOUD NETWORK INFRASTRUCTURE

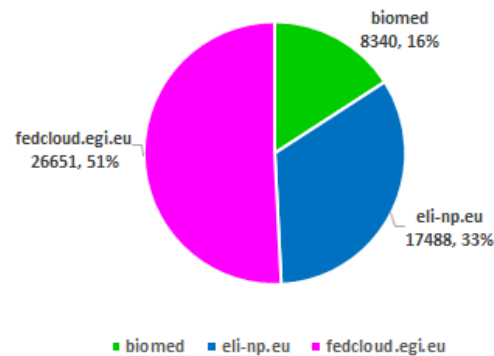


- Pool of public IPv4 IPs that are assigned on demand to the VMs.
- The IB network is used for installing software updates, management activities and data exchanges at service level (MQTT, Glance, Cinder, Neutron, Nova).
- Network shared Remote Image Storage which is used by Glance to store the disk images of the VAs from the EGI AppDB and the local disk images provided by the site to all the locally supported VOs. We only store the disk images endorsed by the supported VOs.

# CLOUDIFIN RESOURCES

## ❑ Cloud accounting data in CPU-hours

Wall time (hours) May 2017 - Sep. 2018



Pike infrastructure software:

- native OpenStack: Nova, Neutron, Glance, Cinder
- EGI integration: Keystone-VOMS, APEL-SSM, cASO and CloudKeeper.

Hardware resources (IB QDR interconnect):

- ❑ one Cloud Controller node (32GB RAM, 20 CPU-cores)
- ❑ one CloudKeeper node (8GB RAM, 8 CPU-cores)
- ❑ 7 Cloud Compute nodes (each with 64GB RAM and 16 CPU-cores)
- ❑ one Block Storage node (65.4TB in RAID6)



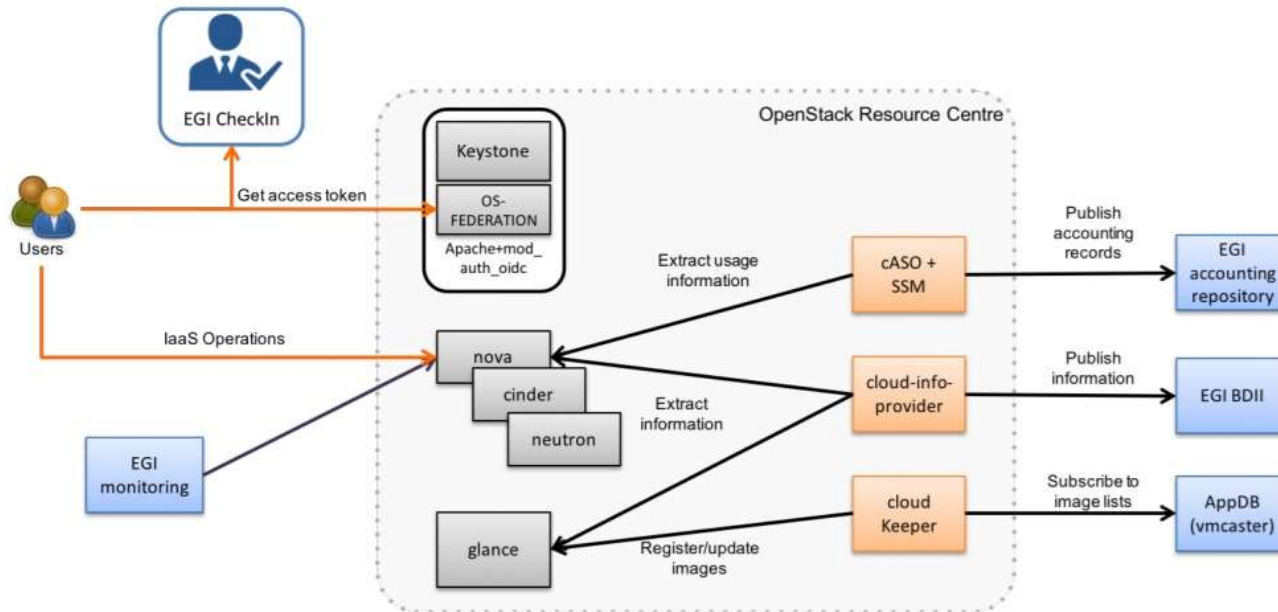


# EGI integration of CLOUDIFIN

## EGI FedCloud services

We installed EGI extensions on top of the OpenStack components.

- Keystone-VOMS Authorization plugin allow users with a valid VOMS proxy to access the OpenStack deployment
- cASO collects accounting data from OpenStack
- SSM sends the records extracted by cASO to the central EGI accounting DB
- BDII cloud provider registers the RC configuration and description through the EGI Information System to facilitate service discovery
- CloudKeeper provides access to the certified Virtual Appliances endorsed by the EGI VOs.



# HTC and HPC infrastructure

The GRIDIFIN site is based on CentOS 64bit Linux distribution and is composed of clusters with computing elements of different types: CREAM-CE and ARC-CE.

1. **CREAM-CE** is used on a virtualised HTC cluster for NGI\_RO services.
2. **CREAM-CE** is also used on an HPC cluster:
  - 28 worker nodes = 336 CPU-cores Intel Xeon X5650
  - IB QDR with OFED distribution
  - Torque + MAUI, both of them compiled with MPI and Infiniband support.
  - CentOS 6.x 64bit + UMD 3.x
3. **ARC-CE** is deployed in a virtualised environment:
  - 14 worker nodes = 336 CPU-cores Intel Xeon E5-2640
  - a couple of GPGPU nodes equipped with nVidia K80.
  - IB FDR with OFED distribution.
  - SLURM which offers native MPI/IB and CUDA support.
  - CentOS 7.x 64bit + UMD 4.x
4. **HPC** cluster:
  - 28 worker nodes = 336 CPU-cores Intel Xeon X5650
  - IB QDR with OFED distribution
  - Torque + MAUI, both of them compiled with MPI and Infiniband support.
  - CentOS 6.x 64bit
5. **GPGPU** cluster: 7x nVidia K80 + 2x nVidia P100 + older nVidia

# Software and Services

## User Interface:

- ❑ Implemented as a single point of entry to any GRID || HPC || CLOUD site
- ❑ configured to access the HTC and HPC part of GRIDIFIN
- ❑ configured to be a UI also for CLOUDIFIN.

## Scientific Software:

- GEANT, ROOT, EPOCH, FLUKA (eli-np.eu);
- VASP, (Tran)Siesta (gridifin.ro);
- NAMD, CHARMM, GROMACS, GAMESS-US, AMBER, OpenBabel, VINA... (ronbio.ro)

Libraries&tools: CUDA, OpenMPI, MVAPICH, Intel compilers, GNU compilers, etc.

**Linux repository:** mirrors for CentOS, Ubuntu, Debian, EGI CAS

**CernVM-FS private repository:** deployment of a local Stratum 0 repository which will eventually host the scientific software available to the *eli-np.eu* members. The CernVM-FS clients will be both in HTC and cloud environment.



# Romanian Node for Computational Biology (RoNBio)

RoNBio is an integrated system based on a **grid of HTC and HPC resources** that include those hosted at DFCTI, which is dedicated to the modeling and simulation of macromolecular structures. The system is accessible through a graphical frontend (applications portal) managed by DFCTI, and was developed in the framework of SIMBAGraN project.


It uses the **Taverna workflow management** system and automates procedures for the investigation of current research topics by means of programmable and reusable workflows.

The applications portal has a great potential for next generation sequencing analysis and other bioinformatics/biocomputing current research topics.

Grid jobs containing Taverna workflows are submitted and executed on the HTC/HPC clusters. The major advantage of Taverna is that it uses distributed services, optimizing the running of workflows, regardless of the local infrastructure.



# RoNBio workflow example





# User communities supported by GRIDIFIN

❑ GRIDIFIN provides computing and storage resources for three locally administered research communities (VOs):

- Experimental groups related of the Extreme Light Infrastructure – Nuclear Physics (ELI-NP) project

**eli-np.eu VO**

- National community of researchers in computational biology

**ronbio.ro VO**

- Researchers interested in condensed state physics and nanomaterial technology

**gridifin.ro VO**

❑ Hardware resources:

- 6 servers providing grid core services
- 672 CPU cores and 2500 GPU cores



# eli-np.eu VO

It is the only active VO for ELI & the most active VO on GRIDIFIN, with:

## Resource Centre GRIDIFIN — Total CPU time used (hours) by VO and Year (Custom VOs)

VO	2015	2016	2017	2018	Total
eli-np.eu	47,032	711,443	1,014,858	284,325	2,057,658

Allocated resources for scientific computing and data processing:

- 336 cores (Intel Xeon X5650 hex core), Infiniband QDR
- 336 cores (Intel Xeon E5-2640 12-core), Infiniband FDR

Data storage and management: 65 TB XFS in RAID6, shared through NFS over Infiniband QDR/FDR.

Community software run in the Grid, as MPI jobs:

- Epoch, Fluka, Geant, ROOT

Testing PIconGPU nVidia Tesla K80



# Summary, conclusions, future work

- ✓ Offer quality HPC, GRID and CLOUD services to the national research communities
- ✓ CLOUDIFIN participation to EOSC
- ✓ Support within the EGI Federated Cloud the European ESFRI projects in which Romania is involved
- ✓ The access of the resources (used by eli-np.eu and ronbio.ro) through CREAM-CE is becoming obsolete. Due to the complexity of the implementation and the fact that the middleware needs to be upgraded to the UMD 4 version, all the workernodes and the computing element will be migrated to ARC in the near future.
- ✓ Implement IPv6 in CLOUDIFIN (ready but not tested).
- ✓ Human resources involved 1.5 to 2.0. Need upgrades.



- RO-LCG 2018, Cluj-Napoca, 17-19 October 2018 -



**THANK YOU  
FOR YOUR ATTENTION!**

[itvasile@nipne.ro](mailto:itvasile@nipne.ro)



18/10/2018



17