


INAF



INAF
DS - ICT

MoU Framework

INAF – Cineca 2016-2019

INAF – Big Science

INAF has reached the scientific excellence in different fields: from computational astrophysics (theoretical modeling in Cosmology and Astrophysics) to ***large experiments and projects***.

The new frontiers in Astrophysics require and will require the massive use of infrastructures of HPC type.

The HPC systems represent today and in the near future computational fundamental facilities and these are today necessary to ***maintain the reached levels of excellence***

INAF – HPC in the past

INAF - CINECA old agreements

Computational Astrophysics, Visualization, Archiving and Post-Processing Data Analysis

-  1997 - 2001
- ➔ 2001 - 2005
- ➔ 2005 - 2007
- ➔ 2008 - 2010

In the last period important experiments have requested HPC computing resources (Tier-0 and Tier-1) for analysis of experimental data including comparisons with simulated data (Planck and Gaia missions). Specific agreements with the Cineca have been signed for these purposes

- ➔ 2010-2012 (Planck)
- ➔ 2012-2014 (Planck)
- ➔ 2013 -2021 (Gaia)

INAF – POR+PON

- ➔ POR 2005-2007 TRIGRID VL @INAF-OACT
- ➔ Cometa Consortium (PI2S2 project 2005-2009) (INAF-OACT+OAPA)
- ➔ Cosmolab Consortium (Cybersar project 2005-2009) (INAF-OACA)

MARCONI

TIER-0 @ CINECA

In 2015 the computing resources in Cineca were:

- Tier-0: **FERMI** (acquired in summer 2012)
- Tier-1: **GALILEO** (acquired in Jan 2015)
- Front-end, Viz, BigData: **PICO** (acquired in Nov 2014)

FERMI arrived at the expected end of its activity.

The Cineca governing bodies, aimed at supporting scientific research, approved a development plan, with an investment of Euro 50 million in two phases, from 2016 to 2020:

- 2 x 5 → 10 Pflops in 2016-2017
- 10 x 5 → 50 Pflops in 2019-2020




A tender was issued in 2015 and assigned Jan 2016 to **lenovo**

The system will be delivered in three phases:

- A1: April 2016 (BRD 2 PFs)
- A2: Sept 2016 (KNL 11 PFs)
- A3: July 2017 (SKL 5 PFs)

In total:

- 18 PFs peak performance
- 17PB raw storage



Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
A1 (half reserved to EUROfusion)						
2PFs	1512	2x Intel Broadwell 18cores @2.3GHz	128 GB	Intel OmniPath 2:1 100Gb/s	21	10 Front End Nodes (2xBDW 18c +128GB RAM)+ 2 MGMT nodes

A2

Core tot: 54.432
Core-h/anno=476.824.320

Core tot: 244.800
Core-h/anno=2.144.448.000

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
11 PFs	3600	Intel KnightsLanding 68cores @1.4 GHz	96 GB	Intel OmniPath 2:1 100Gb/s	50	4 Front End Nodes (2xBDW 18c +128GB RAM)+ 48 I/O nodes

A3 (great part reserved to EUROfusion)

Core tot: 60.480
Core-h/anno=529.804.800

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
5PFs	1512	2x Intel SkyLake 20cores @2.3GHz	192 GB	Intel OmniPath 2:1 100Gb/s	21	2 Front End Nodes + 2 MGMT nodes (2xSKL 20c +192GB RAM) (2xSKL 20cc +

Mou/Framework INAF - Cineca. Highlights

- Adaptation of existing programs and new development on the new Tier-0 systems generation, based on accelerators to avoid the *missing the excellence level reached*.**
- Framework agreement validity: 3 years. Economically valid for a year with explicitly renew and with equal economic conditions for a maximum of 3 years (additional 2 renewals)**
- Framework for Flagship Project. The Flagship Project, for the purposes of this framework agreement, are indicated by INAF (but must be accepted by CINECA)**
- Computational resources for a total of 50 million cpu hours/core KNL Marconi that however can be used on all the Cineca systems**
- Each supported Flagship Project will have reserved up to 5 % of computing resources. Cineca Consortium commits to double the resources assigned to them on the INAF budget. A maximum of 20% of the INAF budget can be assigned to the Flagship projects.**

Mou/Framework INAF - Cineca. Highlights

- ➔ **MoU Management Committee**: 3 Cineca and 4 INAF chaired by INAF. INAF: U. Becciani, G. Mulas, G. Murante, P. Rossi
- ➔ The Committee shall ensure the proper application of the agreement and the use of resources in an **efficient manner** and adjusted in accordance with the needs of researchers.
- Research Grant type PostDoc position** (hired by INAF) yearly renewable. This person will work at Cineca for the entire duration of the framework agreement. Call out now (<http://www.oats.inaf.it/index.php/en/job-opportunities.html>)
- ➔ The coordination of this unit is given by INAF
- ➔ The effective date of the agreement start is 1 May 2017

MoU INAF- Cineca

Calls for Applications: project classes

- ➔ First call was opened on 1st April 2017, deadline end of April (only for class A)
- ➔ 3 different classes of standard projects will be activated
 - ➔ A : 1 – 6 Million KNL core hours (regular calls twice per year)
 - ➔ B : 0.3 – 1 Million KNL core hours (rolling call always open)
 - ➔ Test : up to 0.3 Million KNL core hours (rolling call always open)
- ➔ Flagship projects: up to 5 Million KNL core hours (call always open)

MoU INAF- Cineca

Calls for Applications: project classes

- ➔ “Test” class: “light” proposal, call always open, examined and assigned quickly, intended for trial projects, i.e. for applications that were not tested yet on the available CINECA platforms. Not more than 1 million KNL hours can be assigned to each PI / project.
- ➔ Class A and B: production projects, intended for applications that have already been tested and are known to run well on CINECA platforms.
 - ➔ Class B is for “small” projects, up to 1 million KNL hours, and is almost as flexible as Test class: call always open, proposals examined quickly and, if accepted, resources made quickly available
 - ➔ Class A is for large projects, regular calls (and deadlines!) at least twice per year (1st call now open). Detailed justification of scaling properties and planned use of resources mandatory, to ensure that assigned resources can indeed be used within 1 year.
- ➔ Flagship projects: up to 5 million KNL hours, special class for HPC requests connected to internationally relevant INAF projects (declared by INAF, accepted by CINECA). Half of the hours “donated” by CINECA in exchange for visibility. Call always open.

MoU INAF- Cineca

Calls for Applications: how to apply

- ➔ All relevant information can be found at the URL
<https://www.ict.inaf.it/computing>
- ➔ All INAF personnel can log in to the above site using their IDEM credentials → if you don't have them yet, do request them at the URL
<http://www.ced.inaf.it/idem-info/idp-credenziali.html>
- ➔ On the ICT computing web site, select the “CINECA” tab (
<https://www.ict.inaf.it/computing/cineca/>) and then (visible only after logging in) click on the menu entry (on the left) “Richiesta risorse” (
<https://www.ict.inaf.it/computing/richiesta-risorse/>)
- ➔ Take the time to review the call for proposals (
<https://www.ict.inaf.it/computing/call-cineca/>), select the appropriate form, click on it
- ➔ Fill in the fields, submit, be patient. Be prepared to answer requests for clarifications on the proposal by the panel and/or referee(s)

MoU INAF- Cineca

How to access resources

After the proposal is accepted

- ➔ Computational resources are immediately made available to the user's Cineca account.... provided (s)he has one → if you don't, go to the URL <https://userdb.hpc.cineca.it/> and request one.
- ➔ Browse <http://hpc.cineca.it> select the hardware platform(s) you have obtained time on, read the online docs explaining how to access it/them
- ➔ If necessary, compile, tune, etc. the software you need to use. Take some time to review the most appropriate compilation flags, the most effective libraries to use. Request for support if you need it! In particularly complicated cases, the INAF user can visit CINECA to interact directly with CINECA personnel.
- ➔ Plan the use of the resources! Don't expect to be able to use all of your allocated time during the last month before expiration: beware the dreadful CINECA resource use linearisation!

MoU INAF- Cineca

Final remarks

All the INAF time MUST be used!

- ➔ If you have some idea that may make use of HPC resources, go for it!
- ➔ If you are in doubt, submit a “test” proposal first, use it to install, test, troubleshoot and optimise your code, in preparation for a production project. Keep in mind that on a massively parallel architecture scalability wins over efficiency
- ➔ As soon as you are in a condition to prepare the appropriate proposal, do so. The obvious next step after a test project is a C class project, and the call for these is always open
- ➔ Think BIG: one of the aims of this agreement is to train INAF users on top-level HPC machines so that they can then apply for large PRACE projects, which do require documented previous experience