



Rapporti Tecnici INAF INAF Technical Reports

Number	372
Publication Year	2026-02-18
Acceptance in OA@INAF	2026-03-16T16:09:40Z
Title	PLEIADI@IRA CLUSTER: CALL4 USAGE REPORT
Authors	GANDOLFI, Matteo, BEDOSTI, Francesco
Publisher's version (DOI)	https://doi.org/10.20371/INAF/TechRep/372
Handle	http://hdl.handle.net/20.500.12386/46874

PLEIADI@IRA CLUSTER: CALL4 USAGE REPORT

Matteo Gandolfi, Francesco Bedosti

INAF-Istituto di Radioastronomia

ABSTRACT

This report documents the usage of the INAF Pleiadi@IRA cluster computing resources during Call 4 (August 2024 – February 2025). The PLEIADI partition delivered ~250k CORE hours supporting astrophysical simulations and data processing. The LOFAR partition consumed ~510k CORE hours for radio astronomy data processing. Key infrastructure improvements include the implementation of an automatic node power-down system, reducing idle power consumption by 61%.

1. INTRODUCTION

This report provides an overview of the Pleiadi@IRA cluster operations and resource consumption for Call 4.

1.1 PLEIADI@IRA

PLEIADI is a project by USC-C Computing of INAF-National Institute for Astrophysics, offering high-performance computing (HPC) and high-throughput computing (HTC) resources.

Individual researchers and teams belonging to research projects, European projects, PRIN, INAF mainstream projects, scientific missions, etc. that require computing can apply requesting the resources.

The Pleiadi infrastructures is distributed on the following sites: Bologna (IRA), Catania and Trieste (+ soon Palermo). [0] [1]

Pleiadi cluster at the INAF's Istituto di Radio Astronomia (Institute of Radio Astronomy) in Bologna, also called Pleiadi@IRA, is composed of 6 chassis with 12 nodes each, for a total of 72 nodes.

The cluster is managed with Slurm scheduler and is split in 2 between "normal" Pleiadi and Pleiadi-Lofar which is dedicated to Lofar parallel computing.

The Pleiadi@IRA cluster is deeply integrated in IRA's Computing Centre and could not function stand-alone, it uses the common networking, storage, software, power and cooling of the other IRA's computing facilities.

1.2 CALL TIMELINE

Call	Start Date	End Date	Duration
Call4	2024-08-01	2025-02-28	7 months

1.3 CLUSTER INFRASTRUCTURE

The Pleiadi@IRA cluster consists of 72 blade servers organized in 6 chassis (12 servers per chassis):

Partition	Chassis	Servers	Dedicated Use
pleiadi	1,2,6	36	USC-C allocations, astrophysical simulations
lofar	3,4,5	36	LOFAR radio astronomy data processing

Server specifications:

- CPU: 2× Intel Xeon E5-2697 v4 @ 2.30GHz (18 cores / 36 threads each)
- RAM: 256 GB DDR4 (some pleiadi nodes still at 128 GB during Call4)
- Total cores per server: 36 physical / 72 logical

2. RESOURCE ALLOCATIONS (PLEIADI PARTITION)

2.1 CALL4 ALLOCATIONS


Account	User	Project	Allocated
pleia02	Massimo Dall'Ora	The Milky Way Bulge view with Pleiadi	106,000
pleia04	Domenico Meduri	Global 3D direct numerical simulations of MHD instabilities in red giant cores	500,000
pleia05	Simón Ferrada Chamorro	The ALPACAS (ALL Purpose AstroChemical AnalySis) framework	250,000
pleia06	Styliani Boula	3D RMHD simulations: blazars	500,000
pleia14	Shourya Khanna	Modeling the Milky Way: kinematics, structure, & building a Selection Function for various stellar tracers.	100,000
pleia15	Angela Bongiorno	AGILE nelle PLEIADI	480,000

2.2 GUEST ALLOCATIONS (LOFAR PARTITION)

Account	User	Affiliation	Project
lofaruser0	Martin Hardcastle	Univ. Hertfordshire (UK)	LOFAR DR3 Processing

3. RESOURCE USAGE

3.1 GRAFANA REPORT



3.2 SUMMARY

Total cluster consumption: 774,121 CORE hours

Partition	CORE hours	% of Total
pleiadi	253,481	32.7%
lofar	510,122	65.9%
users (other)	10,517	1.4%

3.3 PLEIADI PARTITION USAGE

Account	User	Project	CORE hours	% Pleiadi
pleia15	Angela Bongiorno	AGILE nelle PLEIADI	188,657	74.4%
pleia02	Massimo Dall'Ora	Milky Way Bulge	39,445	15.6%
pleia14	Shourya Khanna	Milky Way modeling	20,017	7.9%
pleia08	Dario Borgogno	(sportello)	2,950	1.2%
pleia17	—	—	1,964	0.8%
pleia18	Alessandro Piombini	Thesis	344	0.1%
pleia04	Domenico Meduri	MHD simulations	104	<0.1%

3.4 LOFAR PARTITION USAGE

Account	User	CORE hours	% LOFAR
fdg	Francesco De Gasperin	346,225	67.9%
lofaruser0	Martin Hardcastle (guest)	163,897	32.1%

4. INFRASTRUCTURE IMPROVEMENTS

4.1 AUTOMATIC NODE POWER-DOWN SYSTEM (APRIL 2024)

Implemented automatic power management for idle compute nodes to reduce energy consumption.

Configuration:

- 8 blades always on for interactive sessions
- 56 blades can be powered down when idle
- Boot time for powered-down nodes: ~8 minutes

Tests executed on each chassis (12nodes):

State	Power per chassis
All blades on (idle)	~860 W
All blades off	~220 W
2 blades on	~360 W

Result: Total rack consumption reduced from ~4.6 kW to ~1.8 kW (61% reduction when idle)

4.2 SLURM ACCOUNTING MONITORING

Implemented Slurm accounting monitoring from common scheduler tools to Prometheus

Configuration:

The data is extrapolated using the common scheduler commands to see usage of the cluster per user, formatted and sent to Prometheus using the textfile-exporter plugin.

The data is then plotted with Grafana and this allows the operators to have a quick way to retrieve the usage of the cluster per user: select timeframe and check the delta per user.

5. TECHNICAL ISSUES AND SUPPORT

5.1 PLEIA04 - DOMENICO MEDURI (TICKET #180705)

Issue: Hybrid openMP+MPI crashes with magic-sph code

- Segmentation faults during runtime
- Memory allocation errors
- Tested with various --mem values (10000, 100000, 125000 MB)

Status: Open (as of February 2025)

Impact: User was unable to utilize the allocated 500,000 CORE hours (only 104 hours consumed).

5.2 PLEIA09 - SALVATORE COLOMBO (TICKET #525549)

Issue: PLUTO hydrodynamics code problems

- Cannot use more than 6 nodes
- Extremely slow file I/O (vtk files: 8000s vs expected 130s for 2GB files)

Resolution: User migrated to OACT (Catania) cluster which has established PLUTO support.

Notes: The debug on the IRA cluster continued after the allocation expired.

6. CONCLUSIONS

PLEIADI partition:

- Delivered 253,481 CORE hours (32.7% of cluster total)
- Primary consumer: AGILE nelle PLEIADI (Angela Bongiorno) - 74% of pleiadi usage
- Two users encountered technical issues limiting their utilization

LOFAR partition:

- Delivered 510,122 CORE hours (65.9% of cluster total)
- Supported LOFAR DR3 processing including international collaboration

Infrastructure:

- Automatic power-down system achieved 61% idle power reduction

7. APPENDIX

7.1 TICKET REFERENCES

Ticket	User	Issue	Status
#180705	Domenico Meduri	magic-sph crashes	Open
#525549	Salvatore Colombo	PLUTO I/O issues	Closed (OACT, opened in call 3)

7.2 LINK REFERENCES

- [0]: https://indico.ict.inaf.it/event/2870/contributions/20389/attachments/9211/18944/Pleiadi_USCVIII.pdf
- [1]: <https://www.ict.inaf.it/computing/pleiadi/>