

FERMI @ CINECA: access and usage

e.rossi@cineca.it



www.cineca.it




- Configuration
- Access
- Data storage
 - File systems
 - Archiving
- Documentation, announcements and support




Access to FERMI and other related

CINECA
Consorzio Interuniversitario



CINECA in Top 500





FERMI configuration

- 10 racks BlueGene/Q
 - 10K Compute nodes (16 cores & 16 GB memory)
 - 160K cores
 - 2.097 PFlops peak performance
 - 1.725 PFlops sustained (82.2%)
 - 96 I/O nodes
 - ✓ 8 racks with 8 i/o nodes
 - ✓ 2 racks (11 and 31) with 16 I/O nodes
- 8 front-end nodes
 - 4 login nodes – round robin served
 - 4 nodes for pre/post-processing and data moving.
Available through “serial” LL class
 - Power6 nodes (p520), 8 cores, 4.695 GHz, 16GB memory
- 2 Service nodes
 - Reserved to system management and procedures
- Filesystems
 - \$HOME 100 TB
 - \$CINECA_SCRATCH 900 TB → 3.000 TB
 - \$CINECA_DATA 200 TB → 1.400 TB




Access to FERMI

- Same username on all HPC systems in CINECA (FERMI, PLX, ...)
- **Usernames** are personal and grant for interactive access;
Accounts refer to projects active on one or more systems and grant for batch access.
- Connect to login-nodes via ssh, using your username/passwd
`ssh <username>@login.fermi.cineca.it`
- Automatically assigned to one of 4 login-nodes

Filesystems

- The users can store data in three different file systems and one archive system (cartridges)



- \$HOME: (100 TB)
- \$CINECA_SCRATCH (0.9 PB → 3.0 PB)
- \$CINECA_DATA/PROJECT (0.2 PB → 1.4 PB)
- cart*_*



Filesystem: \$HOME

- Daily backup
- 100TB, 50GB quota,
- For regular used data files and applications
- Quota information with `cindata`




- No backup, daily cleanup for files not used for 30 days (not yet in place)
- Files of interest should be promptly copied to other supports
- 0.9 PB → 3.0 PB, no quota
- Fastest and largest filesystem: for data used or generated during the computation (Block size= 4MB → 8MB)
- Usage information with `cindata`

- Shared among all HPC systems
- No backup, no cleaning procedure
- Available only on front-end nodes (not mounted on compute nodes)
- 250 TB, initial quota of 200GB, may be increased on request
- Usage and quota information with `cindata`
- In the near future we will move towards a project-based repository, much larger and possibly based on IRODS (march 2013)

- You can archive large amount of data by using the **cart** * commands. The first time you need a special authorization, so please contact our HPC support (superc@cineca.it).
- You can acquire as many "volumes" as you want and then store data on them. When you do not need your archived data anymore, you can delete single files or whole volumes.
- Since it is more efficient to store few large files than many small ones, we suggest to use the **tar** command to aggregate many small files if needed.

<code>cart_new</code>	create a new volume
<code>cart_dir</code>	show all the defined volumes/files in a given volume
<code>cart_put</code>	save a file in the volume
<code>cart_get</code>	retrieve a file from the volume
<code>cart_del</code>	delete files or full volumes
<code>cart_acl</code>	grant access to a given volume for a specific user
<code>cart_dir_access</code>	show the volumes/files of a different owner
<code>cart_get_access</code>	retrieve a file from a volume of a different owner



Listing usage and quota info

cindata

```
[erossi00@fen07 ~]$ cindata
```

-----ASynchronous Data report-----									
-----			-----USER USAGE-----			-----TOTAL USAGE-----			
USER	AREADESCR	FRESH	SPACE	QTA	QTA%	SPACE	MAX	MAX%	


erossi00	/shared/data/	-2hou	64K	100G	0.0%	177T	189T	93.7%	
erossi00	/fermi/home	-2hou	376M	--	--%	1.2T	--	--%	
erossi00	/fermi/	-2hou	376M	50G	0.7%	1.4T	85T	1.6%	
erossi00	/gpfs/scratch/	-2hou	20G	--	--%	325T	854T	38.1%	







documentation

- www.hpc.cineca.it/

**SCAI**
SuperComputing Applications and Innovation

SuperComputing Applications and Innovation

Home | [Contacts](#) | [CINECA](#)



ABOUT US

RESOURCES

SERVICES


FOR USERS

TRAINING

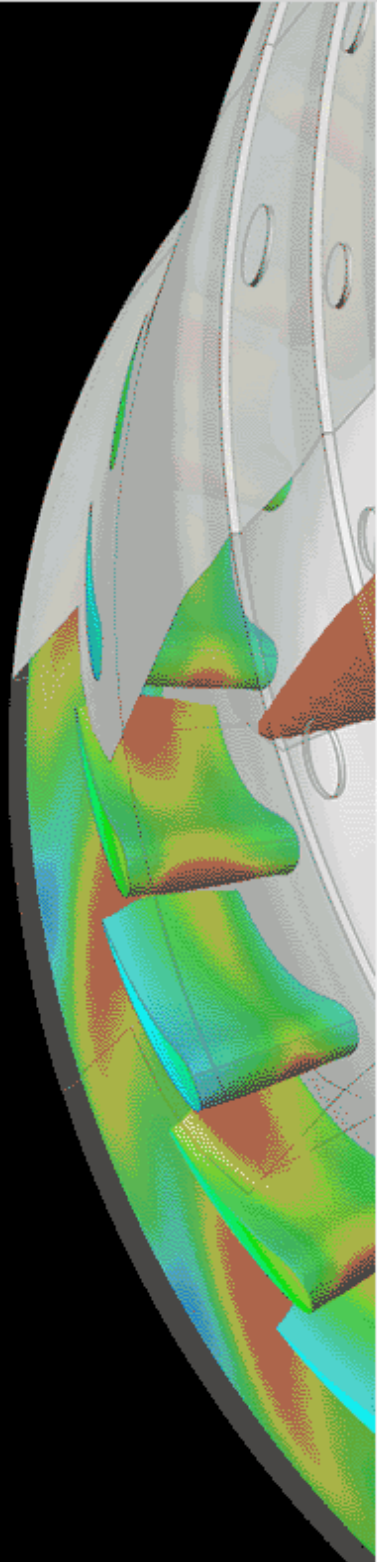
PROJECTS

For users

- » [My portal](#)
- » [Getting started](#)
- » [Get in touch](#)
- » [Help desk](#)
- » [Documentation](#)



[Help desk](#)



FERMI status

PLX status

For users

- » My portal
- » Getting started
- » Get in touch
- » Help desk

Documentation

HPC User Guide
2012
FAQ
Other Documents
Services
One-page manuals
Quick help

Home > For users

Documentation

New documents:

- **Software Benchmarks**
Available soon.
- **Fermi reference guide**
User guide to use our new BlueGene/Q system FERMI.

In the menu:

- **HPC User Guide (new)**
A comprehensive guide to the scientific HPC systems in CINECA, with how-to information for the users.
- **FAQ**
Frequently Asked Questions related to popular topics among the users.
- **Other Documents**
Documents of particular relevance, organised by recently issued and by HPC systems.
- **One-page Manuals**
Short manuals about topics as "vim" or "dbx", collected for you by the User Support team.
- **Services**
Documentation about middleware services, including data management and remote visualization
- **Quick Help**
A collection of quick links to the most requested help topics.

Help desk




Service Closed

Center news

23/11/2012

Power outage





Come and learn the use of FERMI, the new CINECA BG/Q system!

1 2 3

News

Special Features

20/11/2012

FERMI reconfirmed in the Top500 top-ten!

The Italian supercomputing system, installed at CINECA, reconfirmed in the top-ten of November 2012 Top500 ranking list




[Read more](#)

19/11/2012

CINECA at SC12 in Salt Lake City!

For the first time Cineca was one of the exhibitors in the major High Performance Computing conference worldwide.



[Read more](#)

19/11/2012

ETP4HPC at SC12

The European Technology Platform for High Performance Computing was presented at SC12



[Read more](#)

[more News](#)

Call for proposals



ISCRA
Next A&B CALL
OPEN: 13 February 2013
CLOSE: 26 March 2013
ClassC: 28 February 2013



PRACE 6th Call for Projects
23 October 2012/ call closed!



The project ends on 31/12/2012 and no further calls are foreseen.

Training

- 12-13 November HPC Surgery: debugging, profiling and optimization of a scientific code (PATC event)

List of all training events [here](#)

Systems status

FERMI status Up

PLX status Up

For users

- My portal
- Getting started
- Get in touch**
- Help desk
- Documentation

Help desk



Service Closed

Center news

23/11/2012

Power outage

23/11/2012

Reboot of FERMI login node fen01

20/11/2012

FERMI shut down for management operation

14/11/2012

Call for proposals for DECI-10 call for Tier-1 resources

» more Center News

Information and announcements

Home > For users

Get in touch

How to get Center's announcements (HPC-news)

We manage a mailing list (HPC-news) for posting announcements, scheduled downs, software updates, any problems and so on, about our HPC computing resources. It is advisable for HPC users to be included in that list!

You can subscribe (or unsubscribe) to HPC-news by sending an email from the address you want to subscribe. You can consult the archive browsing the archive web site.

To subscribe to HPC-news:

send a mail to listsenv@list.cineca.it
in the body --> "subscribe hpc-news"
in the subject --> any string...

To unsubscribe to HPC-news:

send a mail to listsenv@list.cineca.it
in the body --> "unsubscribe hpc-news"
in the subject --> any string...

To consult the archive:

for the recent archive go to → [Center News](#)
for the full archive go to → <http://list.cineca.it/archives/hpc-news.html>

FERMI status



PLX status



Consultant on-duty

- Monday to Friday 9-13 14-18
- Ticket Trac system

Advanced HPC Support

- su appuntamento e programmato
- working with users (enabling, scaling, etc.)

Portal: **hpc.cineca.it**

E-mail: **superc@cineca.it**

HPC User's Mailing list: **HPC-News**

- Task-local files should be avoided
 - Large number of MPI process, I/O becomes critical
 - (lucido 71 di "application stack and ...")
- Use the correct filesystem:
 - \$CINECA_SCRATCH has the best performance
- Pay attention to block alignment
 - Best performance if I/O is done in chunks equal in dimension to the block size of the filesystem (8MB for CINECA_SCRATCH)
- Use Parallel I/O when possible
 - MPI I/O, HDF5, netCDF, ...