




# Access, disk spaces, accounting

SCAI User Support

*MARCONI USER GUIDE*

*<https://wiki.u-gov.it/confluence/display/SCAIUS/UG3.1%3A+MARCONI+UserGuide>*





# How to get projects for MARCONI

**ISCRA Projects:** computing time allocation on the basis of a research project.

<http://www.hpc.cineca.it/services/iscra>

**LISA Projects:** computing time allocation on the basis of a research project, with current research bodies in Lombardia.

<http://www.hpc.cineca.it/services/lisa>

**PRACE Projects:** computing time allocation at European level on the basis of a research project.

<http://www.prace-project.eu/HPC-access?lang=en>

**Agreements:** several Italian Research Institutions have special agreements for computing provision

**Try Project:** small computing time allocation to performe testing activity



# How to get access credentials

## FOR A NEW USER:

🔧 Registration to our UserDB site:

- 📌 [Go to https://userdb.hpc.cineca.it](https://userdb.hpc.cineca.it)
- 📌 Create new user

🔧 Associaton to a valid project, as a:

- 📌 "Collaborator"
- 📌 "Principal Investigator" (PI)



# How to get access credentials

- 🕒 Log in to userdb
- 🕒 Follow “**HPC access**” menu link to become an HPC user:
  - 🕒 Documents
  - 🕒 Institution Info
  - 🕒 Personal Info
- 🕒 Receive two emails with the hpc credentials (username and password)
- 🕒 Change the password as soon as possible by typing the command “passwd”

The new password can take up to 2 hours to work



# How to login

📌 **SSH – to access the cluster from your pc**

`ssh <username>@login.marconi.cineca.it`

- 📌 by giving the password
- 📌 by a public key (without using the password)

📌 create the key on your pc:

```
$ ssh-keygen
```

```
$ ls -l .ssh
```

```
-rw----- 1 ... 668 Apr 26 14:56 id_rsa
```

```
-rw-r--r-- 1 ... 601 Apr 26 14:56 id_rsa.pub
```

📌 copy the key to the destination cluster

```
scp id_rsa.pub <username>@login.marconi.cineca.it:~
```

```
$ cat id_rsa.pub >> ~/.ssh/authorized_keys
```



# How to login

## 📌 RCM – to access in a graphical way

📌 Download RCM client and launch it

<https://hpc-forge.cineca.it/svn/RemoteGraph/branch/multivnc/build/dist/Releases/?p=817>

📌 Log in to the cluster through “new login” button

📌 Create a graphic session through “new display” button → **\$DISPLAY**

## USE CASES

📌 Access to Remote Desktop

📌 Run a GUI application (e.g., totalview):

```
qsub -v DISPLAY=`hostname`$DISPLAY job.sh
```



# How to login to transfer data

📌 **SCP** – to copy a small amount of data

```
scp <file> <username>@login.marconi.cineca.it:/.../
```

📌 **RSYNC** – to copy a large amount of data

```
rsync --timeout=600 -r -avzHS --bwlimit=80000 --block-size=1048576 --  
progress <data_path_from>  
username@login.marconi.cineca.it:<data_path_to>
```

📌 **GridFTP** – to copy a large amount of data

X509 certificate or Cineca certificate (mailto: superc@cineca.it)



# Local Disk Spaces

## USER SPACES

📍 **\$HOME** → /marconi/home/userinternal/<username>

- 🔗 Permanent (dependent of the life of the user) and backed-up
- 🔗 Quota = 50 GB by default
- 🔗 For storing source code, executables, configuration files or important input files
- 🔗 For compiling your program
- 🔗 Mounted on the login and compute nodes

📍 **\$CINECA\_SCRATCH** → /gpfs/scratch/userinternal/<username>

- 🔗 Temporary (files older than 30 days automatically deleted)
- 🔗 Not backed-up
- 🔗 No quota
- 🔗 For production testing, for temporary output files
- 🔗 Mounted on the login and compute nodes

## PROJECT SPACE

📍 **\$WORK** → /gpfs/work/<project\_name>

- 🔗 Permanent (deleted six months after the end of the corresponding project)
- 🔗 Backed-up and parallel filesystem (GPFS)
- 🔗 1 Tb quota by default
- 🔗 For production activity
- 🔗 Mounted on the login and compute nodes





# Local Disk Spaces

## \$WORK

- “chprj” command
  - l list your PROJECTS
  - d <project\_name> set your default project for \$WORK

- Unix file permissions:

```
$WORK →/gpfs/work/<default project name>  
Owner: PI  
UNIX group: project_name  
drwxrwx--- 29 PI project_name 4096 May 17 15:11 .
```

All collaborators of the project and the PI can write into \$WORK

```
$WORK/subdir or file  
Owner: subdir creator (PI or collaborator)  
UNIX group: interactive  
drwxr-xr-x 29 Collaborator Interactive 4096 May 17 15:11 .
```

In order to sharing personal data between all collaborators of the project:

```
chgrp -R subdir project_name # change unix group  
chmod -R 770 subdir # add rwx permissions to group
```



# Shared Disk Spaces

You need to ask for this kind of resource explicitly, it does not come as part of a project  
(mailto: [superc@cineca.it](mailto:superc@cineca.it))

## USER SPACE

‣ \$TAPE → /gss/gss\_work/tape/userexternal/<username>

- conceived for saving “personal” data on magnetic media
- shared among platforms
- Quota=500 GB
- mounted on the login nodes
- mounted on the compute nodes only of PICO cluster

## PROJECT SPACE

‣ \$DRES

- conceived for saving “project” data on magnetic media
- shared among platforms and projects
- mounted on the login nodes
- mounted on the compute nodes only of PICO cluster
- FS: normal filesystem access on high throughput disks
- ARCH: magnetic tape archiving with a disk-like interface via LTFS
- REPO: smart repository based on iRODS



# cindata

\$ cindata

-----Asynchronous data report-----

USER	AREAID	AREADESCR	USER USAGE				TOTAL USAGE		
			FRESH	SPACE	QTA	QTA%	SPACE	MAX	MAX%
sbuenomi	galileo_hpc-prod	/galileo/prod	-3hou	0	--	--%	790G	1T	77.2%
sbuenomi	work-old_tape	<a href="#">/gss/gss_work/tape</a>	-3hou	0	--	--%	29T	400T	7.4%
sbuenomi	work-cin_staff	/gss/gss_work/cin_staff	-3hou	2M	--	--%	10T	30T	34.3%
sbuenomi	galileo_work-cin_staff	<a href="#">/gpfs/work/cin_staff</a>	-3hou	377M	--	--%	<b>4.5T</b>	<b>10T</b>	<b>45.2%</b>
sbuenomi	galileo_work	/gpfs/work/	-3hou	390M	--	--%	139T	299T	46.5%
sbuenomi	galileo_hpc	/galileo/	-3hou	445M	--	--%	6.7T	8.2T	81.5%
sbuenomi	galileo_hpc-home	<a href="#">/galileo/home</a>	-3hou	<b>445M</b>	<b>50G</b>	<b>0.9%</b>	5.5T	--	--%
sbuenomi	work-DRES_cin_fs	<a href="#">/gss/gss_work/DRES_cin_fs</a>	-3hou	4.3G	--	--%	<b>3.9T</b>	<b>4.9T</b>	<b>79.8%</b>
sbuenomi	work	/gss/gss_work/	-3hou	4.3G	--	--%	1.1P	1.4P	84.0%
sbuenomi	galileo_scr	<a href="#">/gpfs/scratch/</a>	-3hou	<b>54G</b>	--	--%	249T	299T	83.4%
sbuenomi	galileo_scr	/gpfs/scratch_old/	-8wee	730G	32T	2.2%	227T	284T	79.8%



# Accounting

## saldo

#####

### \$ saldo -b

account	start	end	total (local h)	localCluster Consumed(local h)	totConsumed (local h)	totConsumed %	monthTotal (local h)	monthConsumed (local h)
cin_staff	20110323	20200323	400000004	496643	24525028	6.1	3649635	9860
cin_external	20150319	20201231	2000	297	297	14.9	28	0
cin_tmpaccG	20150325	20201231	0	0	0	0.0	0	0

In the CINECA system it is possible to have more than 1 budget (“account”) from which you can use time.

The accounts available to your UNIX username can be found from the saldo command.



# Accounting

## saldo

#####

\$ saldo -r

```

.....
.....
20160427  amarani0  cin_priorit      3:12:48      4
20160503  amarani0  cin_staff        0:04:34      1
20160504  amarani0  cin_staff        2:03:36      1
20160509  amarani0  cin_priorit     27:45:04      2
20160513  amarani0  cin_priorit      6:12:16      1
20160517  amarani0  cin_priorit      0:00:48      2
  
```

```

-----Total from 201401 to 201612-----
      username  account      localCluster  num.jobs
                          Consumed/h

      amarani0  cin_staff      734:18:35      80
      amarani0  cin_priorit    1657:15:47     162
-----
                          Total      2391:34:22     242
  
```



# Accounting

## saldo

#####

\$ saldo -r -a cin\_priorit

```

.....
.....
20160606 amarani0 cin_priorit      0:00:04      1
20160606 mippolit cin_priorit      0:05:20      1
20160607 ccavazzo cin_priorit      3:12:00      1
20160607 mippolit cin_priorit      8:19:44      5
-----Total from 201401 to 201612-----
username account      localCluster  num.jobs
Consumed/h

nbesker1  cin_priorit      115:18:45      3
mguarra1  cin_priorit      11:41:08      12
nspalla1  cin_priorit      34825:15:01    59
faffinit  cin_priorit      28:07:36      22
amarani0  cin_priorit      1657:15:47    162
mippolit  cin_priorit      342:40:16     48
mcestari  cin_priorit      648:15:28     19
cpadrin0  cin_priorit      15211:21:52   88
ccavazzo  cin_priorit      36278:39:15  228
aemerson  cin_priorit      1029:16:50    53
-----
Total      90147:51:58    694
  
```



## FERMI accounts to MARCONI

📌 The budgets that will be moved from FERMI to MARCONI will be recalculated by applying the following factor:

1 h Marconi = 5-6 h Fermi

Marconi budget = Fermi budget / 5 or 6



# Accounting **Billing policy**

- 🔧 No billing for **serial** work (command line or batch → queue serial)
- 🔧 For **batch** jobs the billing is based on "elapsed time" and "effective number of cores" (reserved, not used!) by the batch job. The memory request is also taken into account!

**accounted hours = WallClockTime x ReservedCores**





# Accounting

## Budget linearization

☛ Monthly quota defined for each account:

$$\text{monthTotal} = (\text{total\_budget} / \text{total\_no\_of\_months})$$

☛ As long as the budget is consumed (**monthConsumed**), the jobs submitted from the account will gradually lose priority, until the monthly budget is fully consumed.

☛ On MARCONI, there is a finer graduation of the linearization effect on the priority, as the linearization number is a .1 decimal ranging from 1.0 to 0.0, and depends on the percentage of the monthly quota consumed

☛ The daily budget consumption is updated one time for day (2:00 pm)