

Customized resources

Customized resources and queues for specific research groups and organizations

- [LNEC users](#)
 - [Login Machine](#)
 - [Directories and filesystem](#)
 - [MATLAB](#)
 - [Job Queues](#)
 - [Special resources](#)
 - [Dedicated resources](#)
 - [Old wiki documentation](#)
- [LIP users](#)
- [All Users](#)

LNEC users

The information contained on this chapter is complementary information specific for LNEC users that is not covered on the general guidelines

LNEC users

Login Machine

INCD-Lisbon

ui6-lnec.ncg.ingrid.pt

Directories and filesystem

INCD-Lisbon

The medusa cluster at INCD-Lisbon has the following Lustre filesystems. Lustre is a shared filesystem available both in the computing nodes and in the login nodes. Lustre is best suited for the storage of large/huge files.

- For more specific information regarding the Lustre deployment at INCD see the [Lustre section](#)

Summary

Name	Purpose	Availability	Quota	TimeLimit	Backup
/home/lnec/user	User default home	always	2.7TB total for all users	permanent	no
/data/lnec	Group data large data files [4].	always	5.0 TB for all users	temporary (can be deleted after 6 months)	no
/exper-sw/lnec	Install software for groups	on request	20GB/group	permanent	no

MATLAB

- The MATLAB runtime compiler is available to the LNEC Medusa users. This software allows to use the Matlab compiled code in other machines without having Matlab installed.
- To use MATLAB in the Medusa cluster: **load <name_of_the_module>**

```
matlab/R2012a
```

```
matlab/R2012a.i686
```

```
matlab/R2012b
```

- References[\[link\]](#)
 - Agradecimentos / Créditos: Carlos M. Lima Azevedo

LNEC users

Job Queues

INCD-Lisbon

Job queues available for LNEC users.

List of Queues

Name	Jobs max elapsed time	access	Memory	Max #cores[1]	Comments
medusa	18 days	everyone	6 GB/Core	64	default queue for LNEC users
fast_medusa	24 hours	restricted[2]	3 GB/Core	32	specific queue for short priority jobs
hpc	72h	everyone	6 GB/Core	64	default queue for all users no special privileges

[1] Maximum number of cores a user can request

[2] access based on evaluation and upon request

Special resources

Some LNEC users(*) have dedicated servers (gorgon201 and gorgon301) both accessible from medusa and fast_medusa queues. To use these nodes add one of the following lines to the submitted script.

- Access to **gorgon201.ncg.ingrid.pt**

```
#$ -l Inecit=y
```

- Access to **gorgon301.ncg.ingrid.pt**

```
#$ -l Inecintel=y
```

(*) on request

Dedicated resources

- Cluster MEDUSA - Gorgon201: Servidores TYAN

1. Servidores TYAN FT48B8812 em rede Ethernet (1 Gbit/s)
2. 1 servidor, 32 cores, 8 GB de RAM por core (256 GB de RAM por servidor)
3. 1 disco de 280 GB
4. Servidores instalados com Scientific Linux 6, versão x86_64

```
processor      : 0
vendor_id     : AuthenticAMD
cpu family    : 21
model         : 1
model name    : AMD Opteron(TM) Processor 6220
stepping      : 2
cpu MHz       : 1400.000
cache size    : 2048 KB
physical id   : 0
siblings      : 8
core id       : 0
cpu cores     : 4
apicid        : 32
initial apicid : 0
fpu           : yes
fpu_exception : yes
cpuid level   : 13
wp            : yes
flags         : fpu vme de pse tsc msr pae mce cx8 apic mtrr pge mca cmov pat pse36 clflush mmx
                fxsr sse sse2 ht syscall nx mmxext fxsr_opt pdpe1gb rdtscp lm constant_tsc rep_good nonstop_tsc
extd_apicid
                amd_dcm aperfmperf pni pclmulqdq monitor ssse3 cx16 sse4_1 sse4_2 popcnt aes xsave avx lahf_lm
cmp_legacy
                svm extapic cr8_legacy abm sse4a misalignsse 3dnowprefetch osvw ibs xop skinit wdt lwp fma4
nodeid_msr    topoext
                perfctr_core cpb npt lbrv svm_lock nrip_save tsc_scale vmcb_clean flushbyasid decodeassists
pausefilter pftthreshold
bogomips      : 6000.41
```


TLB size : 1536 4K pages
clflush size : 64
cache_alignment : 64
address sizes : 48 bits physical, 48 bits virtual
power management: ts ttp tm 100mhzsteps hwpstate [9]

- Cluster Medusa - Servidores GATEWAY

1. Servidores Gateway GR380 F2 em rede Ethernet (1 Gbit/s)
2. 1 servidor, 12 cores, 4 GB de RAM por core (32 GB de RAM por servidor)
3. 1 disco de 500 GB
4. Servidores instalados com Scientific Linux 6, versão x86_64

```
processor[]: 0
vendor_id[]: GenuineIntel
cpu family[]: 6
model[]: 45
model name[]: Intel(R) Xeon(R) CPU E5-2620 0 @ 2.00GHz
stepping[]: 7
cpu MHz[]: 2001.000
cache size[]: 15360 KB
physical id[]: 0
siblings[]: 12
core id[]: 0
cpu cores[]: 6
apicid[]: 0
initial apicid[]: 0
fpu[]: yes
fpu_exception[]: yes
cpuid level[]: 13
wp[]: yes
flags[]: fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush dts acpi mmx fxsr sse
sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc arch_perfmon pebs bts rep_good xtopology
nonstop_tsc aperfmperf pni pclmulqdq dtes64 monitor ds_cpl vmx smx est tm2 ssse3 cx16 xtpr pdcm pcid dca
sse4_1 sse4_2 x2apic popcnt tsc_deadline_timer xsave avx lahf_lm ida arat epb xsaveopt pln pts dts
tpr_shadow vnmi flexpriority ept vpid
bogomips[]: 4000.18
clflush size[]: 64
cache_alignment[]: 64
address sizes[]: 46 bits physical, 48 bits virtual
power management:
```


LNEC users

Old wiki documentation

The old Medusa wiki page can be found [here](#)

LIP users

The information contained on this chapter is complementary information specific for LIP users that is not covered on the general guidelines

All Users

ui6-lnec.ncg.ingrid.pt

The medusa cluster at INCD-Lisbon has the following Lustre filesystems. Lustre is a shared filesystem available both in the computing nodes and in the login nodes. Lustre is best suited for the storage of large/huge files.

{{@https://wiki.incd.pt/link/153#bkmrk-servidores-gateway-g}}

1. Servidores Gateway GR380 F2 em rede Ethernet (1 Gbit/s)
2. 1 servidor, 12 cores, 4 GB de RAM por core (32 GB de RAM por servidor)
3. 1 disco de 500 GB
4. Servidores instalados com Scientific Linux 6, versão x86_64