

# Legion Quick Reference Sheet



## Access.

Legion can be accessed via SSH using either a command line SSH client or a windowed SSH client like PuTTY.

**Hostname:** legion.rc.ucl.ac.uk

For file transfers, users can either use SCP, or SFTP. Best transfer speeds will be attained by transferring via the login05 node:

**Hostname:** login05.external.legion.ucl.ac.uk

## Hardware.

ID	Class	Number	Cores	RAM	IB
login05	-	1	4	24 GB	x
login06 -- login09	-	4	12	24 GB	x
node-k01 -- node-p24	X	90	12	24 GB	✓
node-001 -- node-108	Y	50	12	24 GB	x
node-202 -- node-212	Z	4	12	48 GB	x
node-513 -- node-528	U	84	16	64 GB	x
node-601 -- node-603	T	3	32	1.5 TB	x

IB = Infiniband, the fast network for multi-node jobs.

## Scheduler.

Legion uses Sun Grid Engine to queue up job scripts to run later on the compute nodes. Users interact with this via commands on the login nodes.

Command	Description
qsub <i>script</i>	Submit a script to the scheduler
qstat	Show status of queued/running jobs
qstat -j <i>jobid</i>	Show detailed information about one job
qdel <i>jobid</i>	Remove a job from the scheduler
qrsh <i>options</i>	Request an interactive session on a test node

## Other Useful Commands.

Command	Description
man <i>command</i>	Show the manual page for a command
quotas	Show your home and scratch area quota status
groups	Show what groups you're in. May include a project group, and any program access groups

## Job Script Options.

The scheduler has many options to alter how and where your job can run, and what resources it will be allocated, shown below. As well as this, example job scripts for both generic use and for many specific applications are available from our website.

Flag	Effect
-N <i>jobname</i>	Set the name of the job in the queue to <i>jobname</i>
-P <i>projectid</i>	Set the project id to run the job under (only relevant for special access projects).
-pe mpi <i>cores</i>	Request cores to run distributed parallel jobs (MPI or other), with <i>cores</i> on more than one node if necessary. (Max: 256)
-pe smp <i>threads</i>	Request <i>threads</i> parallel threads on a single node. (Max: 32, sensible: 12 or 16)
-l h_rt= <i>hh:mm:ss</i>	Request an amount of time to be able to run the job for. (At the end of this time, the job will be killed if it is still running.)
-l mem= <i>bytes</i>	Request that processes be placed such that the job can use that much RAM per core requested.
-l tmpfs= <i>bytes</i>	Request the creation of a temporary area on the compute node that can store <i>bytes</i> bytes. Append "G" to request gigabytes. (Default: 10G)
-l gpu= <i>number</i>	Request 1 or 2 GPUs (per core if used with -pe) for your job. GPUs are only available on V-class nodes, so jobs requesting these will only run there (automatically).
-t <i>start-stop:step</i>	Specify that the job should be multiplied into an array of jobs, with \$SGE_TASK_ID in each set to a different value, starting at <i>start</i> , finishing with <i>stop</i> , and incrementing by <i>step</i> each time.
-cwd	Set the working directory the job starts in to the same directory it was submitted from.
-wd <i>directory</i>	Set it somewhere else.
-o <i>filename</i>	Direct the output stream into a given file. (Default: <i>jobname.ojobid</i> )
-e <i>filename</i>	Direct the error stream into a given file. (Default: <i>jobname.ejobid</i> )
-j y	Join together the output and error streams.
-hold_jid <i>jobid</i>	Hold the job until the job with id <i>jobid</i> has finished.
-ac allow= <i>node_class</i>	Only run on nodes of class <i>node_class</i> .

## Resource Limits.

- Jobs can use a single core for up to 3 days, or more cores/threads for up to 2 days.
- MPI jobs can use a maximum of 64 nodes.
- Each user may have up to 1000 jobs in the queue (an array of jobs counts as 1 for these purposes).

## Storage.

Jobs have access to three different areas of storage on Legion:

Location	Quota	Notes
/home/\$USER	50 GB	Read-only from compute nodes, writable from login nodes, backed up
/home/\$USER/Scratch	200 GB	Writable anywhere, not backed up
\$TMPDIR	(10 GB)	Defined per job, size requestable from job script (see -l tmpfs above).

## Module commands.

The `module` command sets environment variables to make it easier to run programs or use libraries. There is a default set of loaded modules on Legion which can be seen by running `module show default-modules`.

Command	Description
module	Give a brief list of the available module commands
module avail	Show all available modules.
module avail <i>text</i>	Show modules with names beginning with <i>text</i> .
module load <i>module</i>	Load <i>module</i> .
module unload <i>module</i>	Unload <i>module</i> .
module show <i>module</i>	Show the effect of loading a module.
module help <i>module</i>	Show a brief description of a module.

## Support.

Please contact [rc-support@ucl.ac.uk](mailto:rc-support@ucl.ac.uk) if you have any questions about the contents of this document, or any of our services or systems.

More documentation and support pages are available on our website, at: <http://www.ucl.ac.uk/research-computing/>.

## Last Updated.

This document was last updated on 2016-11-10.