

University of Iceland High Performance Computing

An introduction v2

Máni Maríus

June 2018



HÁSKÓLI ÍSLANDS
UPPLÝSINGATÆKNISVIÐ

- All of Gardar removed
- 46 new nodes added
- Support staff down to 1
- New support mail queue



- Garpur #1
 - Since 2015
 - 24/32 cores per node
 - 44 nodes + 3 GPU nodes
 - 128/256GB memory per node
 - 2x Tesla M2090 in each GPU node
- Garpur #2
 - Added late 2017
 - 32 cores per node
 - 192GB DDR4
 - interconnect: 50Gb/s Omnipath
- Jötunn
 - Since 2016
 - 24 cores per node
 - 4 nodes
 - 128GB memory per node




Name	Garpur	Garpur #2	Jötunn	Total
Year	2015	2017	2016	
Nodes	44	46	4	94
TFLOPS	37	98	4	139
kW	17.1	16.6	2	35.7

- 136 Garpur users
 - 47 unique users connected since june 1st
- 73 Jötunn users



Cluster usage by groups

Garpur Usage 2018 by groups



Garpur

- OpenHPC
- OS: Centos 7
- GCC & Intel compilers
- OpenMPI, IntelMPI, MPICH
- Python, R, Matlab
- VASP, GROMACS, PISM



Are you

- studying/working at an Icelandic University/Academic Institution?
- Doing a project supported by RANNIS?

→ then send an email to support-hpc@hi.is



What do I get with an account?

- SSH login
- Disk space
 - Home partition: 100GB +
 - Work partition: Unlimited¹
 - Jotunn disk space is more limited
- Unlimited CPU hours¹
- Support from us¹



¹Within resonable limits

- Account requests
 - support-hpc@hi.is
- Software requests
 - ihpc@hi.is
 - Request site is public
- News
 - ihpc@listar.hi.is
- Chat
 - ihpcgroup.slack.com



Cluster workflow

You should have received your login credentials by email.

- 1 Connect with ssh
 - ssh mani@jotunn.rhi.hi.is
- 2 Check cluster status
 - sinfo
 - squeue
- 3 Load modules *or* compile program on login node
 - module avail
 - module load ...
- 4 Create job file
- 5 Submit job to queue
 - sbatch myjob.sh
- 6 Check results
 - Use slurm directives to send email when job completes



Software on the cluster is provided in modules.

Use "module spider" to find all possible modules.

Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".

```
[mani@garpur-main ~]$ module avail
```

```
----- /opt/ohpc/pub/moduledepts/gnu-openmpi -----
boost/1.61.0      netcdf/4.4.1      netcdf-fortran/4.4.4  phdf5/1.8.17      scalasca/2.3.1      sionlib/1.7.0
fftw/3.3.4        netcdf-cxx/4.2.1  petsc/3.7.0           scalapack/2.0.2   scorep/3.0

----- /opt/ohpc/pub/moduledepts/gnu -----
R_base/3.3.1     gsl/2.2.1         hdf5/1.8.17          mpich/3.2         openblas/0.2.19     openmpi/1.10.4 (L)  pdtoolkit/3.22

----- /opt/ohpc/pub/modulefiles -----
gnu/5.4.0 (L)    ohpc              papi/5.4.3          prun/1.1

----- /opt/share/modulefiles -----
ADF/2017_r61726  gpaw/1.2.0        matlab/R2016a       proj4/4.9.2        tensorflow/1.1-cpu
ADF/2017_109 (D)  gromacs/5.1.4-w-plumed  ncl/6.4.0          pypy3/6.0          udunits/2.2.24
ase/3.13         gurobi/7.0.2     nco/4.6.1           python/intel-3.6   vasp/5.4.1-impi
cdo/1.8.2        intel/compiler/2017.2  openmpi/intel/2.1.0  python/2.7.13      vasp/5.4.1-tbdyn-impi
cp2k/5.1         intel/impi/2017.4    openmpi/intel/2.1.1 (D)  python/3.6.1 (D)  vasp/5.4.1
gcc/4.9.4        intel/mkl/2017.2    pism/0.7            siesta/gcc/4.1.b3  vasp/5.4.4-impi (D)
go/1.10.2        julia/0.6          pism/1.0 (D)        siesta/intel/4.0.1

Where:
L: Module is loaded
D: Default Module
```

Use "module spider" to find all possible modules.

Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".

Anything missing?

Anything outdated?

Send an email to ihcp@hi.is



Slurm Partitions

name	Time limit	nodecount	comment
short	4 hours	82	Default, limited to 1 node
normal	15 days	36	
long	30 days	36	
omnip	15 days	46	
himem	15 days	5	256GB memory nodes
himem-bigdisk	15 days	3	256GB memory and bigger disk
gpu	15 days	3	GPU Nodes

or use *sinfo* command



Interactive job

```
[mani@garpur ~]$ salloc -N 1 -p gpu --exclusive
salloc: Granted job allocation 203466
[mani@garpur ~]$ squeue -u mani
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
203466 gpu bash mani R 0:03 1 gpu-1
[mani@garpur ~]$ ssh $SLURM_NODELIST
Last login: Tue Jun 12 19:13:53 2018 from garpur.localdomain
[mani@gpu-1 ~]$ module load python/2.7.13
[mani@gpu-1 ~]$ module li
```

```
Currently Loaded Modules:
 1) prun/1.1 2) gnu/5.4.0 3) openmpi/1.10.4 4) ohpc 5) python/2.7.13
```

```
[mani@gpu-1 ~]$ nvidia-smi
Tue Jun 12 19:15:35 2018
```

```
-----+-----
| NVIDIA-SMI 375.26                Driver Version: 375.26                |
+-----+-----+-----+-----+-----+-----+
| GPU  Name      Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp     Perf          Pwr:Usage/Cap|      Mem-Usage |   GPU-Util  Compute M. |
+-----+-----+-----+-----+-----+-----+-----+
|   0   Tesla M2090      Off          | 0000:21:00.0  Off   |    0%      0
| N/A   N/A       P0        76W /  N/A | 0MiB / 5301MiB |      0%      Default |
+-----+-----+-----+-----+-----+-----+
|   1   Tesla M2090      Off          | 0000:22:00.0  Off   |    0%      0
| N/A   N/A       P0        77W /  N/A | 0MiB / 5301MiB |      0%      Default |
+-----+-----+-----+-----+-----+-----+
-----+-----
```

```
-----+-----
| Processes:                        GPU Memory
| GPU      PID  Type  Process name                               Usage
+-----+-----+-----+-----+-----+
| No running processes found
+-----+-----
-----+-----
```

```
[mani@gpu-1 ~]$ exit
logout
Connection to gpu-1 closed.
[mani@garpur ~]$ exit
exit
salloc: Relinquishing job allocation 203466
salloc: Job allocation 203466 has been revoked.
[mani@garpur ~]$ █
```



Typical slurm job workflow:

- 1 Decide how many nodes you need and on which partition (himem, default, gpu)
- 2 Create bash script with slurm directives
 - #SBATCH -J jobname
 - #SBATCH -N 2
 - #SBATCH -ntasks-per-node=2
 - #SBATCH -mail-user mani@hi.is
 - #SBATCH -mail-type=END
 - #SBATCH -array=0-15
- 3 Submit to queue
 - sbatch myjob.sh

```
[mani@garpur-main src]$ sbatch imb.sh  
Submitted batch job 7303
```

Further examples available on our website <http://ihpc.is>



Rules of thumb

- 1 Be respectful of others. Don't submit 10 jobs requiring 1 node each at once.
- 2 Allocate your job to 1 core, half a node or the whole node.
- 3 Tidy up unused disk space
- 4 Keep in mind resources other than CPU cores (e. g. memory)
- 5 If you know how long your job will run for, allocate only the needed walltime



System status

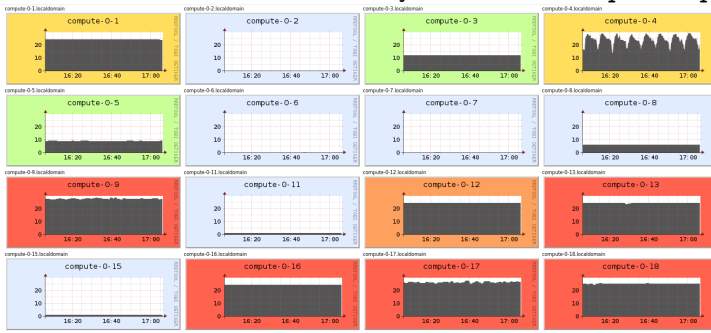
Check the status of the queue with

squeue

or

squeue -u mani

We also have a website with the system status <http://ihpc.is>



Reporting examples

```
[mani@garpur-main ~]$ sreport user top
```

```
-----  
Top 10 Users 2018-06-12T00:00:00 - 2018-06-12T23:59:59 (86400 secs)
```

```
Use reported in TRES Minutes
```

```
-----  
Cluster      Login      Proper Name      Account      Used      Energy  
-----  
garpur      jianyang   Jian Yang        jianyang     846483    0  
garpur      asod       Asmus O Dohn     asod         829440    0  
garpur      elena.pa+  Elena Papaleo    elena.papaleo 593512    0  
garpur      seb21     Sebastian Bohr   jesus-group   548301    0  
garpur      jdb5      Jan Burger       jesus-group   217493    0  
garpur      etayyebi Ebrahim Tayyebi etayyebi     152608    0  
garpur      aleksei   Aleksei Ivanov   aleksei      131781    0  
garpur      bab29     Barði Benedik+   bab29        119862    0  
garpur      christia+ Christian Bean    christianbean   108730    0  
garpur      nral     Nzar Rauf Abdu+  nral         46080     0
```

```
[mani@garpur-main ~]$ sreport job SizesByAccount user=mani START=2018-06-01
```

```
-----  
Job Sizes 2018-06-01T00:00:00 - 2018-06-12T23:59:59 (1036800 secs)
```

```
Time reported in Minutes
```

```
-----  
Cluster      Account      0-49 CPUs      50-249 CPUs      250-499 CPUs      500-999 CPUs      >= 1000 CPUs      % of cluster  
-----  
garpur      kerfis      39              0                  0                  0                  0                  100.00%
```



Any questions?
Send them to support-hpc@hi.is

