

SINFO



Ques in SLURM are called partitions. Each available partition has certain settings regarding available resources, maximal job time, etc. To see some information on available partitions run `sinfo` on Neumann.

A sample output is printed here:

```
$ sinfo
PARTITION    AVAIL  TIMELIMIT  NODES  STATE NODELIST
container    up    infinite    4    drain c[028,110,114,152]
container    up    infinite   148    alloc
c[001-002,005-022,027,029-109,111-113,116-128,133-150,155-166]
container    up    infinite    19    idle
c[003-004,023-026,129-132,151,153-154,167-172]
gpu          up    infinite    2    idle c[502-503]
big*         up 4-20:00:00    2    drain c[028,114]
big*         up 4-20:00:00   135    alloc
c[005-022,027,029-044,047-064,067-084,089-106,111-113,116-128,133-150,155-166]
big*         up 4-20:00:00    6    idle c[167-172]
sw01_short  up    2:00:00    1    drain c110
sw01_short  up    2:00:00   13    alloc
c[001-002,045-046,065-066,085-088,107-109]
sw01_short  up    2:00:00    6    idle c[003-004,023-026]
[...]
```

What you can see here is list of partitions in different states.

- *idle* nodes are currently unused. Sometimes they are kept free by SLURM to start the next job.
- *alloc* (-ated) are nodes when they are being used by a job.
- *drain* or *draining* are nodes which are removed or are to be removed from the available nodes in a partition

A short explanation of the partitions:

Partition	Explanation
big	This is the default que. Most calculation are run here. At max a job can run 4,8 days with a maximum of approx. 140 nodes. However, only need as many as you really need and the fewest possible. See parallel efficiency .
sw01_short	This que is for short jobs. During the day the maximum time for a job is 1h. In the evening hours the maximum time is increased up to approx. 12h. That way test jobs can be run in working hours and computations can use free nodes efficiently at night. The next day, the partition will most often be free again (at least only 1h jobs will run). If you want to run 1h jobs take the short partition only

Partition	Explanation
short	This partition is nearly the same as sw01_short. It uses the same nodes + 2 extra nodes. This partition keeps the 1h limit at all times (afaik). Maybe this partition will replace sw01_short at some point.
sw04_longrun	This partition runs jobs for about 14 days at max. Use this if you cannot save your simulation in between. However, avoid using this que. The waiting times are long because of that long max computation times. Most often you can do the same task with big, too.
longrun	Same as sw04_longrun. Probably it will replace it at some point.
gpu	In this partition there are 2 nodes which have only 8 CPUs per node. However, each node has a dedicated high-performance graphic card
container	This partition just collects all resources. It is only of interest for the administrator.

The nodes taken out of the partitions can be listed with the option -R

```
$ sinfo -R
REASON          USER      TIMESTAMP           NODELIST
possible_bad_mem root      2017-06-21T14:07:05 c004
pcie-errors     root      2017-06-16T11:30:04 c110
hardware.ecc errors root      2016-04-25T12:10:06 c114
bad_mem         root      2017-08-23T14:55:13 c142
Not responding  root      2017-08-27T18:41:21 c031
possible_bad_mem root      2017-06-21T14:15:33 c009
```

The normal view of sinfo might be difficult to understand. Add the following alias to your .bashrc file (in your home directory).

```
alias si='sinfo -o "%12P %.10A %.5D %.4c %.6mMB %.11l %3p %5h" | grep -v -e container -e extra'
```

This will print sinfo in the following format:



```
$ si
PARTITION      NODES(A/I)  NODES  CPUS  MEMORYMB  TIMELIMIT  PRI
SHARE
gpu            0/2         2      16   254000MB  infinite  1
NO
big*          135/6       144    16   254000MB  4-20:00:00 1
NO
sw01_short    11/8        20     16   254000MB  2:00:00   1
NO
sw04_longrun  11/0        12     16   254000MB  14-20:00:00 2
NO
sw09_urgent   0/7         8      16   254000MB  1:00:00   2
NO
urgent        46/6        54     16   254000MB  9:00:00   3
```



NO

Now you can see easily how many nodes are available in each partition and how many are currently unused (*l*).

From:

<https://wikis.ovgu.de/hpc/> - HPC

Permanent link:

<https://wikis.ovgu.de/hpc/doku.php?id=guide:sinfo>

Last update: **2017/09/27 20:48**

