

Energy Measurement

For the CM and ESB modules there is a fine grained energy measurement in place using Megwares energy meters attached to the compute nodes of these modules. There are different ways to get information about energy consumption for your jobs (and the nodes). Preferred methods are:

SLURM `sacctl` command

This is probably the easiest way to get energy consumption for your interactive and batch jobs. Once your job has finished, you can use the `sacct` command enquire the Slurm database about its energy consumption. For further information and an example on how to use the command for energy measurements, please see [accounting](#).

DCDB (coming soon)

The "Datacenter Database" very frequently (every 10 seconds) stores measured values from node and infrastructure sensors including power and energy consumption of the compute nodes. This allows for a very fine grained analysis of consumed energy for your jobs, e.g. by specifying precise time stamps / ranges and by providing access to measured data from different components like CPU, GPU, memory etc. instead of making available an accumulated value only. On the other hand it offers a convenient way for analysis of SLURM jobs.

The DCDB can be queried from the login node using the DCDB client tools. A user guide that gives some more details on the database and explains how to use it for energy measurements will be attached to this page in the next days.

Use `/sys` files

The energy meters provide their measured values through the `/sys` filesystem on the nodes using different files. To query the overall energy (in Joules) a node has consumed so far, you can use the `energy_j` file. You should integrate readings in your SLURM job script before and after you `srun` your commands to measure consumed energy by your commands (applications): Unit=[Joules]

CM Module

```
srun sh -c 'if [ $SLURM_LOCALID == 0 ]; then echo ${SLURM_NODEID}: $(cat /sys/devices/platform/sem/energy_j); fi'
```

ESB Module

There are two energy meters present for the ESB nodes, one for the CPU blade and one for the GPU part:

```
srun sh -c 'if [ $SLURM_LOCALID == 0 ]; then echo ${SLURM_NODEID}: $(cat /sys/devices/platform/sem.1/energy_j); fi'
srun sh -c 'if [ $SLURM_LOCALID == 0 ]; then echo ${SLURM_NODEID}: $(cat /sys/devices/platform/sem.2/energy_j); fi'
```

To get the consumed energy by a multi-node job you have to accumulate all the values which makes it quite cumbersome, but for running on single nodes (maybe even in an interactive session) reading out current values directly from the files might be quite useful.