

System usage

The DEEP-EST Data Analytics Module (DAM) can be used through the SLURM based batch system that is also used for (most of) the Software Development Vehicles (SDV). You can request DAM nodes (`dp-dam[01-16]`) with an interactive session like this:

```
srun -N 4 --tasks-per-node 2 -p dp-dam --time=1:0:0 --pty /bin/bash -i
[kreutz1@dp-dam01 ~]$ srun -n 8 hostname
dp-dam01
dp-dam01
dp-dam02
dp-dam02
dp-dam03
dp-dam03
dp-dam04
dp-dam04
```

When using a batch script, you have to adapt the partition option within your script: `--partition=dp-dam`

Using Cuda

To compile and run Cuda applications on the Nvidia V100 cards included in the DAM nodes, it is necessary to load the CUDA module:

```
[deamicisl@deepv ~]$ ml CUDA
[deamicisl@deepv ~]$ ml

Currently Loaded Modules:
  1) GCCcore/.8.3.0 (H)   2) binutils/.2.32 (H)   3) nvidia/.418.40.04 (H,g)  4) CUDA/10.1.105 (g)

Where:
  g:  built for GPU
  H:  Hidden Module
```

Filesystems and local storage

The home filesystem on the DEEP-EST Cluster Module is provided via GPFS/NFS and hence the same as on (most of) the remaining compute nodes. The local storage system of the DAM running BeeGFS is available at

```
/work
```

The file servers are reachable through the 40 GbE interface of the DAM nodes.

This is NOT the same storage being used on the DEEP-ER SDV system. Both, the DEEP-EST prototype system and the DEEP-ER SDV have their own local storage.

It's possible to access the local storage of the DEEP-ER SDV (`/sdv-work`), but you have to keep in mind that the file servers of that storage can just be accessed through 1 GbE ! Hence, it should not be used for performance relevant applications since it is much slower than the DEEP-EST local storages mounted to `/work`.

There is node local storage available for the DEEP-EST DAM node (2 x 1.5 TB NVMe SSD), but configuration is to be done for those devices.

Multi-node Jobs

Attention: Since the Extoll network is not in place yet multi-node MPI Jobs are currently disabled.