

Table of Contents

Login node	2
Detected HW and node issues	2
Cooling issues	2
CM nodes	2
DAM nodes	2
ESB nodes	2
SDV nodes	2
Software issues	2
MODULEPATH	2
Cuda and Rocky 8.6	2
nvidia driver mismatch	3
nvidia profiling tools	3

This page is intended to give a short overview on known issues and to provide potential solutions and workarounds to the issues seen.

Last update: 2023-03-16

Please, use the support mailing list `sup(at)deep-sea-project.eu` to report any issues

To stay informed, please refer to the [News page](#). Also, please pay attention to the information contained in the "Message of the day" displayed when logging onto the system. The system status is reported on [?JSC status](#) as well.

Login node

- Time limit for user processes enforced on deepv login: **Processes will be killed after 24 hours** In case of problems, please contact niessen@?

Detected HW and node issues

Cooling issues

- pump in JSC cooling loop is running in manual mode: frequently running HPL jobs (with low priority) to create some load (waste heat)
 - HPL jobs can be killed on demand: in case of problems (your jobs being blocked by HPL runs), please contact j.kreutz@? or niessen@?

CM nodes

- dp-cn25: Thermal issues within chassis slot ([#2769](#))

DAM nodes

- dp-dam02: reserved for FPGA tests
- dp-dam13: failing healthcheck: memory_not_reclaimable
- dp-dam16: testbed

ESB nodes

- dp-esb[07]: wrong BIOS settings ([#2881](#))
- dp-esb[17]: IB HCA issues ([#3140](#))
- dp-esb[75]: Easybuild testbed ([#3094](#))

SDV nodes

- deeper-sdv cluster nodes (Haswell) have been taken offline: deeper-sdv[01-16]
 - not included in SLURM anymore
 - deeper-sdv[09-10] used for testing (please contact j.kreutz(at)fz-juelich.de if you would like to get access)
- knl01: serves as golden client for imaging only
- dp-sdv-esb[01,02]: will only be powered on demand

Software issues

MODULEPATH

- MODULEPATH variable might get overwritten when switching stages
- leads to various modules not being detected / found correctly
- re-setting the MODULEPATH manually might solve the issue, e.g. for the 2022 stage, please try:

```
export MODULEPATH=/usr/local/software/skylake/Stages/2022/modules/all/Compiler/sidecompiler/GCCcore/11.2.0:/usr/local/s
```

Cuda and Rocky 8.6

New CUDA drivers on the compute nodes. In case of problems, please manually prepend your `LD_LIBRARY_PATH` (first for libcuda, second for libcublas, fft, etc.):

```
ln -s /usr/lib64/libcuda.so.1 .
ln -s /usr/lib64/libnvidia-ml.so.1 .
LD_LIBRARY_PATH=./usr/local/cuda/lib64:$LD_LIBRARY_PATH srun <srun_args> <exe> <exe_args>
```

nvidia driver mismatch

- loading CUDA module and trying to run `nvidia-smi` (or any application trying to use the GPU) leads to

```
Failed to initialize NVML: Driver/library version mismatch
```

- workaround is to unload the driver module: `ml -nvidia-driver/.default`
- for further information, please also see [?here](#)

nvidia profiling tools

- to launch the tools on a compute node using X-Forwarding another SSH session is needed:

```
srun --forward-x -p dp-esb -N 1 -n 1 --pty /bin/bash -i
ssh -X -J <your account>@deep.zam.kfa-juelich.de <your account>@<the node you received>
```

- you will still see a warning "OpenGL Version check failed. Falling back to Mesa software rendering.", but the profiling tool (e.g. `nsight-sys`) should start up