Table of Contents

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: 2022-11-18

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.

Detected HW and node issues

CM nodes

- dp-cn25: SEL ProblemsFW issues ([#2769]())
- dp-cn30: Image update needed ([#2991]())
- dp-cn[47-50]: BeeOnd testbed

DAM nodes

- dp-dam02: reserved for FPGA tests
- dp-dam16: testbed

ESB nodes

- dp-esb[07]: testbed
- dp-esb[11]: memory issues ([#2857]())
- dp-esb[31]: GPU issues ([#2949]())
- dp-esb[47]: SEL Problems ([#2998]())
- dp-esb[75]: Image update needed

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

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 unload the driver module: `ml -nvidia-driver/.default`
- for furhter 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 profling tool (e.g. `nsight-sys`) should start up