

Wikiprint Book

Title: Public/User_Guide/PaS

Subject: DEEP - Public/User_Guide/PaS

Version: 70

Date: 18.05.2024 06:48:38

Table of Contents

Login node	3
Detected HW and node issues	3
Cooling issues	3
CM nodes	3
DAM nodes	3
ESB nodes	3
SDV nodes	3
Software issues	3
Cuda and Rocky 8.6	3
nvidia driver mismatch	3
nvidia profiling tools	4

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: 2023-10-11

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

Please refer to the [Project 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-cn03: [#2374](#) - Kernel update
- dp-cn25: [#2769](#) - SEL Problems
- dp-cn50: [#2769](#) - Testbed for beegfs

DAM nodes

- dp-dam02: [#2874](#) - Reserved for FPGA tests
- dp-dam16: [#3231](#) - ECC Memory

ESB nodes

- dp-esb[17]: [#3196](#) - Energy meter not working

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 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:

```
srunc --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