

Table of Contents

Detected HW and node issues	2
CM nodes	2
DAM nodes	2
ESB nodes	2
SDV nodes	2
Software issues	2
Modular jobs failing	2
SLURM jobs	2
GPU direct usage with IB on ESB	2
GPU direct usage with Extoll on DAM	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: 2020-09-23

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.

Detected HW and node issues

CM nodes

- dp-cn[01,09,10]: nodes currently reserved for special use case during working hours
- dp-cn33: memory issues ([#2464](#))
- dp-cn49: configuration change required ([#2291](#))
- dp-cn50: node not reachable ([#2488](#))

DAM nodes

- dp-dam03: node currently reserved for special use case ([#2242](#))
- dp-dam04: showing low streams performance ([#2401](#))
- dp-dam05: node currently reserved for special use case
- dp-dam07: showing problems with its FPGA ([#2353](#))
- dp-dam[09,10]: nodes currently reserved for special use case during working hours

ESB nodes

- dp-esb05: node hangs in state "Idle+Completing"
- dp-esb11: wrong GPU Link Speed detected ([#2358](#))
- dp-esb24: CentOS8 Testbed ([#2396](#))
- dp-esb39: energy meter reading issues ([#2432](#))
- dp-esb52: energy meter reading issues ([#2433](#))
- dp-esb61: node not reachable ([#2469](#))
- dp-esb71: energy meter reading issues ([#2432](#))
- dp-esb73: energy meter reading issues ([#2433](#))

SDV nodes

- deeper-sdv01: node reachable via network, but marked as down in SLURM
- nfgw[01,02]: node reachable via network, but marked as down in SLURM
- knl01: NVMe issues ([#2011](#))
- ml-gpu02: memory issues reported with MCE ([#2489](#))

Software issues

Modular jobs failing

- users reported failing jobs that are doing MPI on more than one module using the gateways
- the problem is being investigated

SLURM jobs

- due to introduction of accounting with the start of the early access program there is some re-configuration of user accounts needed within SLURM to assign the correct QOS levels and priorities for the jobs
- this might lead to (temporary) failing job starts for certain users
- if you cannot start jobs via SLURM, please write an email to the support list: `sup(at)deep-est.eu`

GPU direct usage with IB on ESB

- currently only available via Developer stage, for testing load:

```
module --force purge
module use $OTHERSTAGES
module load Stages/Devel-2019a
module load GCC/8.3.0 or module load Intel
module load ParaStationMPI
```

- use PSP_CUDA=1 and PSP_UCP=1

GPU direct usage with Extoll on DAM

- new Extoll driver for GPU direct over Extoll currently being tested on the DAM nodes
- only available via Developer stage, for testing load:

```
module --force purge
module use $OTHERSTAGES
module load Stages/Devel-2019a
module load GCC/8.3.0 or module load Intel
module load ParaStationMPI
```

- expect performance and stability issues