

Wikiprint Book

Title: System usage

Subject: DEEP - Public/User_Guide/SDV_Cluster

Version: 6

Date: 09.03.2025 16:10:57

Table of Contents

System usage	3
Filesystems and local storage	3
Using /nvme	3
Multi-node Jobs	3

System usage

The system can be used through the SLURM based batch system that is also used for the DEEP-EST Cluster Module and (most of) the remaining compute nodes. You can request cluster nodes on the SDV with an interactive session like this:

```
srun --partition=sdv -N 4 -n 8 --pty /bin/bash -i
srun ./hello_cluster

Hello world from process 6 of 8 on deeper-sdv07
Hello world from process 7 of 8 on deeper-sdv07
Hello world from process 3 of 8 on deeper-sdv05
Hello world from process 4 of 8 on deeper-sdv06
Hello world from process 0 of 8 on deeper-sdv04
Hello world from process 2 of 8 on deeper-sdv05
Hello world from process 5 of 8 on deeper-sdv06
Hello world from process 1 of 8 on deeper-sdv04
```

When using a batch script, you have to adapt the `—partition` option within your script: `—partition=sdv`

Filesystems and local storage

The home filesystems on the SDV are provided via GPFS/NFS and hence the same as on the DEEP System. The local storage system of the SDV running BeeGFS is available at

```
/sdv-work
```

This is NOT the same storage being used on the DEEP-EST system. Both, the DEEP-EST System and the DEEP-ER SDV have their own local storage. On the DEEP-EST nodes it is mounted to `/work` on the deeper-sdv nodes it can be found in `/sdv-work`. The DEEP-EST storage in `/work` can be accessed, but with 1 GbE only. Hence, it should not be used for performance relevant applications since it is much slower than the local storages (see NVMe) and the DEEP-ER SDV storage (`/sdv-work`).

Using /nvme

During job startup, all files of non-privileged users within `/nvme` are removed. If you want to keep your files across consecutive jobs on a particular SDV node, add a list of filenames to `$HOME/.nvme_keep`:

```
/nvme/tmp/myfile.A
/nvme/tmp/myfile.B
```

This will keep the files `/nvme/tmp/myfile.{A,B}` across two or more job runs in a row.

Multi-node Jobs

The SDV Cluster nodes are connected via EXTOLL tourmalet (100 Gb/s). Please use

```
module load extoll
```

to run jobs on multiple nodes.