

# NEWS FROM THE ESM PARTITION

## **New partition for submitting batch and interactive jobs dedicated solely to the ESM community**

- slurm queue “esm”
- use “-p esm” in your srun command or “#SBATCH --partition=esm” in your sbatch scripts
- 6 JUWELS nodes only, not suited for model simulations in production mode.
- Maximum wallclock time is 24 hours

## **Computing time usage of ESM projects**

- By end of 2019, the ESM project lost 24.7Mio core-h from it's computing time budget 05/2019 - 04/2020 (~8 – 50 % of the allocated core-h for the individual projects)
- The project „cesmtst“ used ~6Mio out of 26Mio core-h and lost ~8Mio core-h.
- cesmtst also is granted 500,000 core-h on JUWELS GPU nodes

# ESM USER ISSUES

## Survey Summer 2019:

- *A “pool directory” is needed for data exchange between different groups/ projects/ external partners (e.g. for synchronization with the DKRZ pool directory). The pool directory should be accessible from the compute nodes.*
  - Storage space for a dedicated “pool directory” can be applied for in terms of a data project.
  - User access/permissions to the data project resources can be managed via JuDoor.
  - Disk space on \$FASTDATA needs to be applied for to get access from the compute nodes.
- *More disk space on \$SCRATCH is needed.*
  - The quota on \$SCRATCH can be extended per request and on a case-by-case basis.
  - For large amounts of data (> 100 TByte) or for long term storage (archival, post-processing) ESM projects should apply for data projects providing access to additional disk space on \$FASTDATA, disk space connected to interactive nodes and cloud infrastructures (\$DATA), or the tape archive (\$ARCH).

# ESM USER ISSUES

- *Users have problems in compiling and running various ESM codes with both, Intel MPI and Parastation MPI.*
  - Problems with Intel MPI and Parastation MPI have been communicated to the vendors. Some issues are circumvented with the use of 2018 stages, others have solved with IntelMPI/2019.6.154
  - OpenMPI4.0 is ready for testing in Stages/Devel-2019a
- *Support for the installation of the “esm-tools” software package is requested.*
  - Slavko Brdar acts as point of contact (needs to be revised).
- *Job queue waiting times are perceived as unpredictable and generally too long. Priority of jobs is not transparent.*
  - It is intended that compute jobs have to be submitted to a common queue used by all users of the system. With this approach, users can potentially run jobs as large as the full machine.
  - Users can inquire about the prioritisation of their jobs using standard slurm commands or with the JSC tool “llview”

# ESM USER ISSUES

- *Nodes for post-processing with sufficient amount of memory and short queueing times are needed.*
  - JUWELS is equipped with 240 large memory compute nodes (“fat nodes”), providing up to 192 Gbyte of memory per node, that are particularly suited to run post-processing jobs.
- *Unlike other systems, JUWELS does not provide “shared” nodes for post-processing (to potentially reduce queueing time).*
  - Considering the size of JUWELS (Tier-0), node-sharing was not considered as a design option in its configuration.
  - Small post-processing jobs may run interactively on one of the 12 login nodes of JUWELS.
- *There are more down times on JUWELS compared to other systems.*
  - JUWELS was unfortunately affected by more downtimes during its first months. Partly, these downtimes were related to technical problems with the GPFS parallel filesystem.
  - Maintenance is foreseen about twice a month with notification to users well ahead. The duration depends on the demand of work to be done on the system.

# ESM USER ISSUES

- *Response times of JSC supercomputer support are too slow and the replies may be rather short or too general.*
  - JSC supercomputer support (sc@fz-juelich.de) for general technical problems.
  - Get in contact with project advisors assigned to the compute projects for more specific questions regarding ESM codes and workflows.
  - Point of contact for user support for the projects of the Helmholtz ESM community: Olaf Stein (o.stein@fz-juelich.de)

# UPCOMING EVENTS

**Helmholtz GPU Hackathon, Berlin, 9 – 13.3.2020**

<https://www.gpuhackathons.org/>

**HPC Tunathon, Jülich, ~25 –29.5.2020**

<http://fz-juelich.de/ias/jsc/2020/HPC-tuna>

To be discussed: Workshop/Hackathon on Gridtools (CLAW / other DSLs)