



ESM ISSUES AND DISCUSSION, NEWS FROM THE ESM PARTITION

OLAF STEIN, JSC, SDL CLIMATE SCIENCE

Mitglied der Helmholtz-Gemeinschaft



THE JUWELS ESM PARTITION

- Computing time granted to projects with participation of HGF E&E institutions
- Strong participation of Helmholtz ESM project partners
- Since 11/2019 projects related to PL-ExaESM are also eligible
- Grant committee: Daniela Jacob (GERICS), Stefan Kollet (FZJ/IBG-3) and Thomas Jung (AWI)



Note: The effective computing power per core-hour of JUWELS differs for each module. Thus, resources need to be converted to floating point operations in units of 10^{18} (EFLOP).

JUWELS Cluster

- ESM Partition holds 20% share of total computing time since 05/2018
- ~170 Mio Core-h or 52500 EFLOP granted each year (plus ~3 Mio Core-h or 9000 EFLOP on NVIDIA V100 GPUs)

JUWELS Booster

- ESM Partition holds 10% share of total computing time since 11/2020
- ~33 Mio Core-h or 200000 EFLOP granted each year on NVIDIA A100 GPUs

ESM PARTITION: PARTNERS AND PROJECTS



Partners

HGF E&E

AWI, GEOMAR, GFZ, UFZ,
KIT, FZJ: IEK-7, IEK-8, IBG-3

HGF AST

DLR-PA, DLR-SC

HGF INF

FZJ: JSC, KIT

Projects

Currently: 22 projects (5 with resources on Booster)

Project cesmtst for collaborative efforts and testing (→ PI Daniel Caviedes Voullieme)

ESM PARTITION: COMPUTING TIME ALLOCATION AND USAGE

	Allocation 11/2020-10/2021	Allocation 5/2021-4/2022	Allocation 11/2021-10/2022
Cluster CPUs	69 Mio Core-h or 21500 EFLOP	105 Mio Core-h or 32500 EFLOP	61.5 Mio Core-h or 19000 EFLOP
Cluster GPUs	550,000 Core-h or 1650 EFLOP	760,000 Core-h or 2300 EFLOP	940,000 Core-h or 2850 EFLOP
Booster	9.7 Mio Core-h or 58000 EFLOP	12 Mio Core-h or 72000 EFLOP	7.3 Mio Core-h or 43800 EFLOP

Since 2018: 573 Mio Core-h approved or 355000 EFLOP (~90% of requested resources)

Actual data usage on GPFS file systems:

- \$DATA: 3.5 PByte (max. 5 Pbyte)
- \$SCRATCH: 1.2 PByte (max. 2.5 Pbyte)
- \$PROJECT: 158 TByte (max. 350 Tbyte)
- \$ARCHIVE: 1.5 Pbyte

Publications (peer-reviewed)

29 acknowledgements on Google Scholar

NEWS FROM JSC & ESM PARTITION

Data on \$DATA now accessible from compute nodes

- A few JUWELS Cluster/Booster compute nodes can access data residing on \$DATA directly
- for your slurm jobs use: `salloc --constrained=largedata`

New Status page for JSC resources

- <https://status.jsc.fz-juelich.de/>

llview job reporting

- https://www.fz-juelich.de/ias/jsc/EN/Expertise/Support/Software/LLview/jobreporting_node.html

JSC Service Status		
- Cluster Systems		
 JURECA Booster	Next Maintenance at 01. Oct. 2022, 09:00:00	 JURECA DG
 JUBUF HPC	Next Maintenance at 01. Mar 2022, 09:00:00	 JUWELS Booster
 JUWELS Cluster		 HDF-ML
 JUZEAT	Next Maintenance at 22. Mar 2022, 10:00:00	 JUDAC
- File Systems		
 \$FASTDATA		 SHOME
 \$SCHATCH I		 SHOME EXT
 \$ARCHIVE	Next Maintenance at 05. Mar 2022, 14:00:00	 \$DATA
 \$SCHATCH II		
- Services		
 JUBUF CLOUD		 JUSTCOM
 Backup		 Cloud Object Storage
 HDF Cloud		 JuDoor
 Jupyter JSC		 UNICORE
- Support		
 JSC Support		
- Icons guide		
 Service is operating normally		
 Service is operating normally. There is some additional information available		
 Service is lightly degraded, which should not affect most users		
 Service is degraded, which can impact some users		
 Service is heavily degraded, which will impact most users		
 Service is currently unavailable		
 JÜLICH Forschungszentrum		
Service Status Information Legend Feedback		

NEWS & ISSUES

Object Store in XCST

- store, manage, and access large amounts of data via HTTP
- Supported object storage access protocols: OpenStack Swift and S3
- not possible to store data via POSIX and access them via the object storage interface or vice versa

Data quota

- Quota listed on \$DATA listed with `du` or `jutil project dataquota` currently counts double the size of data actually stored due to internal data replication
- Quota values on \$SCRATCH are currently untrustworthy, sometimes even negative quota occur

Backup

- \$HOME, \$PROJECT, \$ARCHIVE have full external TSM backup
- \$DATA provides snapshots instead of backup: State is saved at a certain point in time: daily/weekly/monthly (last three retention)
- \$FASTDATA provides snapshots instead of backup: daily/weekly/monthly (last one/one/two retention)
- Storage team is working to create a full backup from snapshots

UPCOMING EVENTS

Interactive High-Performance Computing with Jupyter, 5.-7.4. 2022

<https://www.fz-juelich.de/SharedDocs/Termine/IAS/JSC/DE/Kurse/2022/ptc-interactive-hpc-2022.html>

High-performance computing with Python, 20.-24.6. 2022

<https://www.fz-juelich.de/SharedDocs/Termine/IAS/JSC/DE/Kurse/2022/ptc-hpc-python-2022.html>

Directive-based GPU programming with OpenACC, 26.-28.10. 2022

<https://www.fz-juelich.de/SharedDocs/Termine/IAS/JSC/DE/Kurse/2022/gpu-openacc-2022.html>