

HPC working seminar for physicists

Scientific Computing Department at HIM

Dr. Dalibor Djukanovic

Dr. Peter-Bernd Otte

bi-weekly meeting – 10.5.2022



Today's Topics


1. Himster 2: Presentations and People
2. best of tips and tricks (Himster 2 usage, singularity, CI, Jupyter)
3. your questions / discussion / requests to the maintainers

- compact in time (20mins + user questions/discussion).
- bring people together tackling the same problems
- minutes: <https://www.hi-mainz.de/research/computing/hpc-working-seminar/>

Presentations 2021 & 2022

- AG Denig BES III: Detector Simulation & Data Analysis
 - Christoph Florian Redmer, Riccardo Aliberti, Yasemin Schelhaas, Thomas Lenz, Max Lellmann, Yuping Guo, Tong Liu, Yaqian Wang, + MSc + BSc
- MESA: Beam Dynamics
 - Sebastian Taubert
- EMP Panda: Detector Simulation
 - Luigi Capozza, Dong Liu, Sahra Wolff, Alexander Greiner, Julian Moik
- AG Bacca Nuclear Theory Group: Precision calc. of few-body and many-body systems
 - Bijaya Acharya, Joanna Sobczyk, Limone Li Muli, Francesca Bonaiti
- SPEC-F-Hyp Panda: Detector Simulation & Data Analysis
 - Sebastian Bleser, Michael Bölting, Martin Christiansen, Falk Schupp, Marcell Steinen
- THFL: Nucleon form factors from lattice QCD
 - Miguel Salg

Mogon 2 ↔ Himster 2



HPC-Gate
(PrivacyIdea)

head nodes 1..3

home directory

LUSTRE

Mogon 2:
80% of compute nodes & interconnect

GPU
nodes

Himster 2:
320 compute nodes & Fat tree interconnect

Further shared:
- maintenance
- user software
- slow control

shared

Non shared hardware

Software (1/2): Modules on Himster

14.9.2021

- `echo $MODULEPATH`
- `export MODULEPATH=$MODULEPATH:/cluster/him/modulefiles/`
- `module avail`
 - Remarks: D = standard module, aliases exist, L = loaded
- `module avail 2>&1 | wc -l`
- `module load tools/Singularity/`
- `module list`

Jupyter on headnode with plain python

28.9.2021

usage:

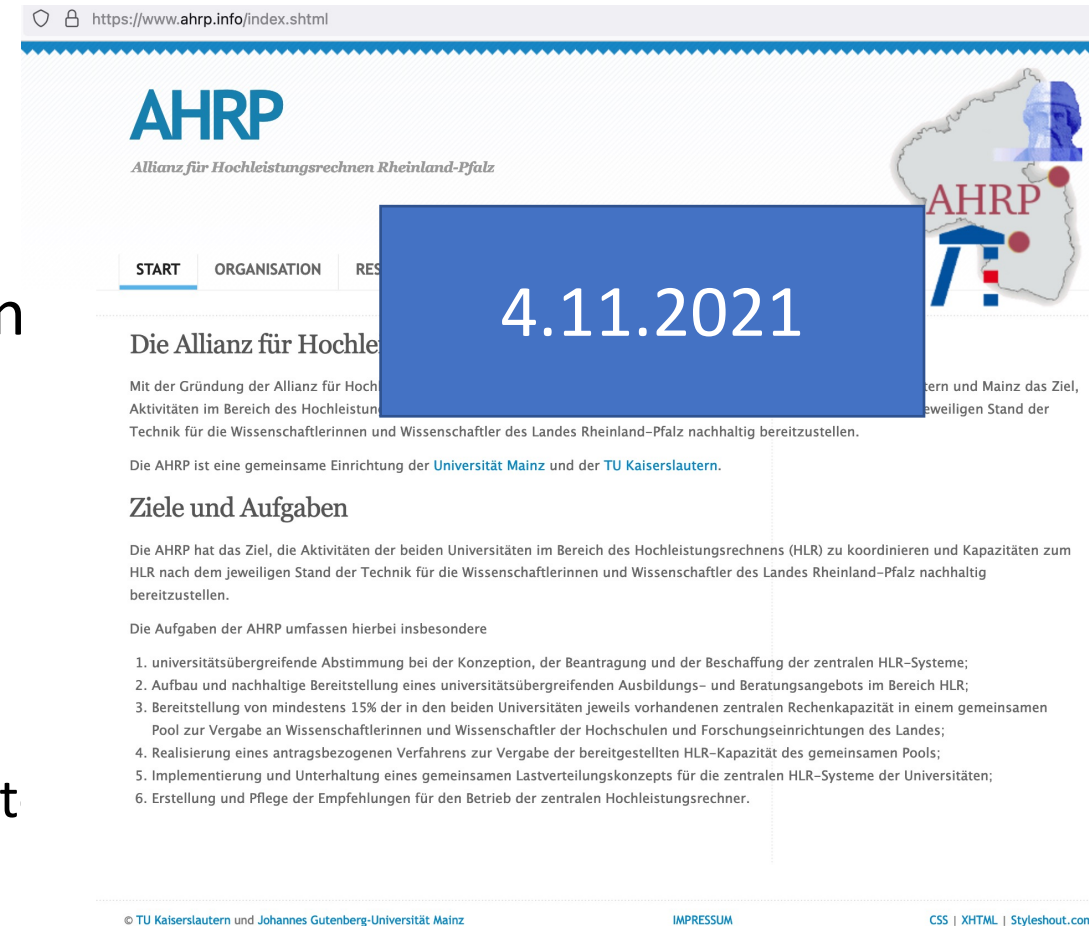
1. ssh himster2
 2. [pbotte@login23 ~]\$ source testjupyter/bin/activate
 3. (testjupyter) [pbotte@login23 ~]\$ jupyter notebook
 4. Open locally: <http://localhost:12345>
 - Enter the code presented in terminal
- Caution:
 - others might already use port 8888.
 - If port already in use, change config file and the port forward in SSH.

JupyterHub @ ZDV

- <https://jupyterhub.zdv.uni-mainz.de>
- 5*4 CPUs with 5*12GB RAM
- Intended for teaching
- For research: Solution with Mogon 2 will follow

Apply for GPU nodes

- GPU nodes only available on Mogon2, belong to University
- <https://www.ahrp.info/>
 - → Ressourcen
 - Antrag auf Nutzung eines rheinland-pfälzischen Hochleistungsrechners
 - (ENGLISH VERSION: Application for Using a Rhineland-Palatinate High Performance Comput
- Fill out application, together with 1-2 sides of project description
- Technical contact person: fill in my name (but contact me first)



The screenshot shows the homepage of the AHRP (Allianz für Hochleistungsrechnen Rheinland-Pfalz) website. The URL in the browser is <https://www.ahrp.info/index.shtml>. The page features a navigation menu with 'START', 'ORGANISATION', and 'RES'. A large blue overlay with the date '4.11.2021' is positioned over the main content area. The main heading is 'Die Allianz für Hochleistungsrechnen Rheinland-Pfalz'. Below this, there is a paragraph describing the alliance's mission and a list of goals and tasks. The footer contains copyright information for TU Kaiserslautern and Johannes Gutenberg-Universität Mainz, along with links for 'IMPRESSUM', 'CSS | XHTML | Styleshout.com', and 'TU Kaiserslautern und Johannes Gutenberg-Universität Mainz'.

AHRP
Allianz für Hochleistungsrechnen Rheinland-Pfalz

START ORGANISATION RES

4.11.2021

Die Allianz für Hochleistungsrechnen Rheinland-Pfalz

Mit der Gründung der Allianz für Hochleistungsrechnen Rheinland-Pfalz sollen die Aktivitäten im Bereich des Hochleistungsrechnens (HLR) an der Johannes Gutenberg-Universität Mainz und der TU Kaiserslautern koordiniert werden und Mainz das Ziel, den HLR nach dem jeweiligen Stand der Technik für die Wissenschaftlerinnen und Wissenschaftler des Landes Rheinland-Pfalz nachhaltig bereitzustellen.

Die AHRP ist eine gemeinsame Einrichtung der [Universität Mainz](#) und der [TU Kaiserslautern](#).

Ziele und Aufgaben

Die AHRP hat das Ziel, die Aktivitäten der beiden Universitäten im Bereich des Hochleistungsrechnens (HLR) zu koordinieren und Kapazitäten zum HLR nach dem jeweiligen Stand der Technik für die Wissenschaftlerinnen und Wissenschaftler des Landes Rheinland-Pfalz nachhaltig bereitzustellen.

Die Aufgaben der AHRP umfassen hierbei insbesondere

1. universitätsübergreifende Abstimmung bei der Konzeption, der Beantragung und der Beschaffung der zentralen HLR-Systeme;
2. Aufbau und nachhaltige Bereitstellung eines universitätsübergreifenden Ausbildungs- und Beratungsangebots im Bereich HLR;
3. Bereitstellung von mindestens 15% der in den beiden Universitäten jeweils vorhandenen zentralen Rechenkapazität in einem gemeinsamen Pool zur Vergabe an Wissenschaftlerinnen und Wissenschaftler der Hochschulen und Forschungseinrichtungen des Landes;
4. Realisierung eines antragsbezogenen Verfahrens zur Vergabe der bereitgestellten HLR-Kapazität des gemeinsamen Pools;
5. Implementierung und Unterhaltung eines gemeinsamen Lastverteilungskonzepts für die zentralen HLR-Systeme der Universitäten;
6. Erstellung und Pflege der Empfehlungen für den Betrieb der zentralen Hochleistungsrechner.

© TU Kaiserslautern und Johannes Gutenberg-Universität Mainz IMPRESSUM CSS | XHTML | Styleshout.com

Archiving Service on Himster 2/Mogon 2 (1/2)

DMP (data management plans) necessary for founding program

- [Experts: https://researchdata.uni-mainz.de](https://researchdata.uni-mainz.de)

9.11.2021


IRODS

- Archiving service connected to Himster 2
- No space limit (if >1TB inform support first)
- <250TB/user -> free
- combine files until they are larger than 10GB
- Idea:
 - write once, read from time to time
 - share your raw data together with paper (measurement or analysis VMs)

Tbit Link HIM <-> GSI

GSI Test Lustre File System mounted in Mainz, same data locally. IOR: 1.0GBytes/s write, 1.1 Gbytes/s read

23.11.2021



Containerised Analysis

- Easiest way to freeze your analysis and run it again at any later time
- Same analysis on all machines
(HPCs, around the world, workstation)
- easy start for new students

- Singularity (no Docker @ HPC)
- Try it first on your office computer
 - https://sylabs.io/guides/3.0/user-guide/quick_start.html
- Recent example:
 - BDSIM for MAMI: <https://gitlab.rlp.net/-/snippets/3010>

Worked out example: BDSIM with Singularity

- physical setup is described in files on your home directory.
- Analysis framework (Root, Geant, ..., BDSIM) sits in singularity-image: /lustre/miifs05/scratch/him-acid/singularity/bdsim.sif

- Call:

```
#!/bin/bash
#SBATCH -A m2_him_exp           # Specify allocation to charge against
#SBATCH --partition=himster2_exp. # Queue name 'smp' or 'parallel' on Mogon II
#SBATCH --time=24:00:00.       # Run time (hh:mm:ss)

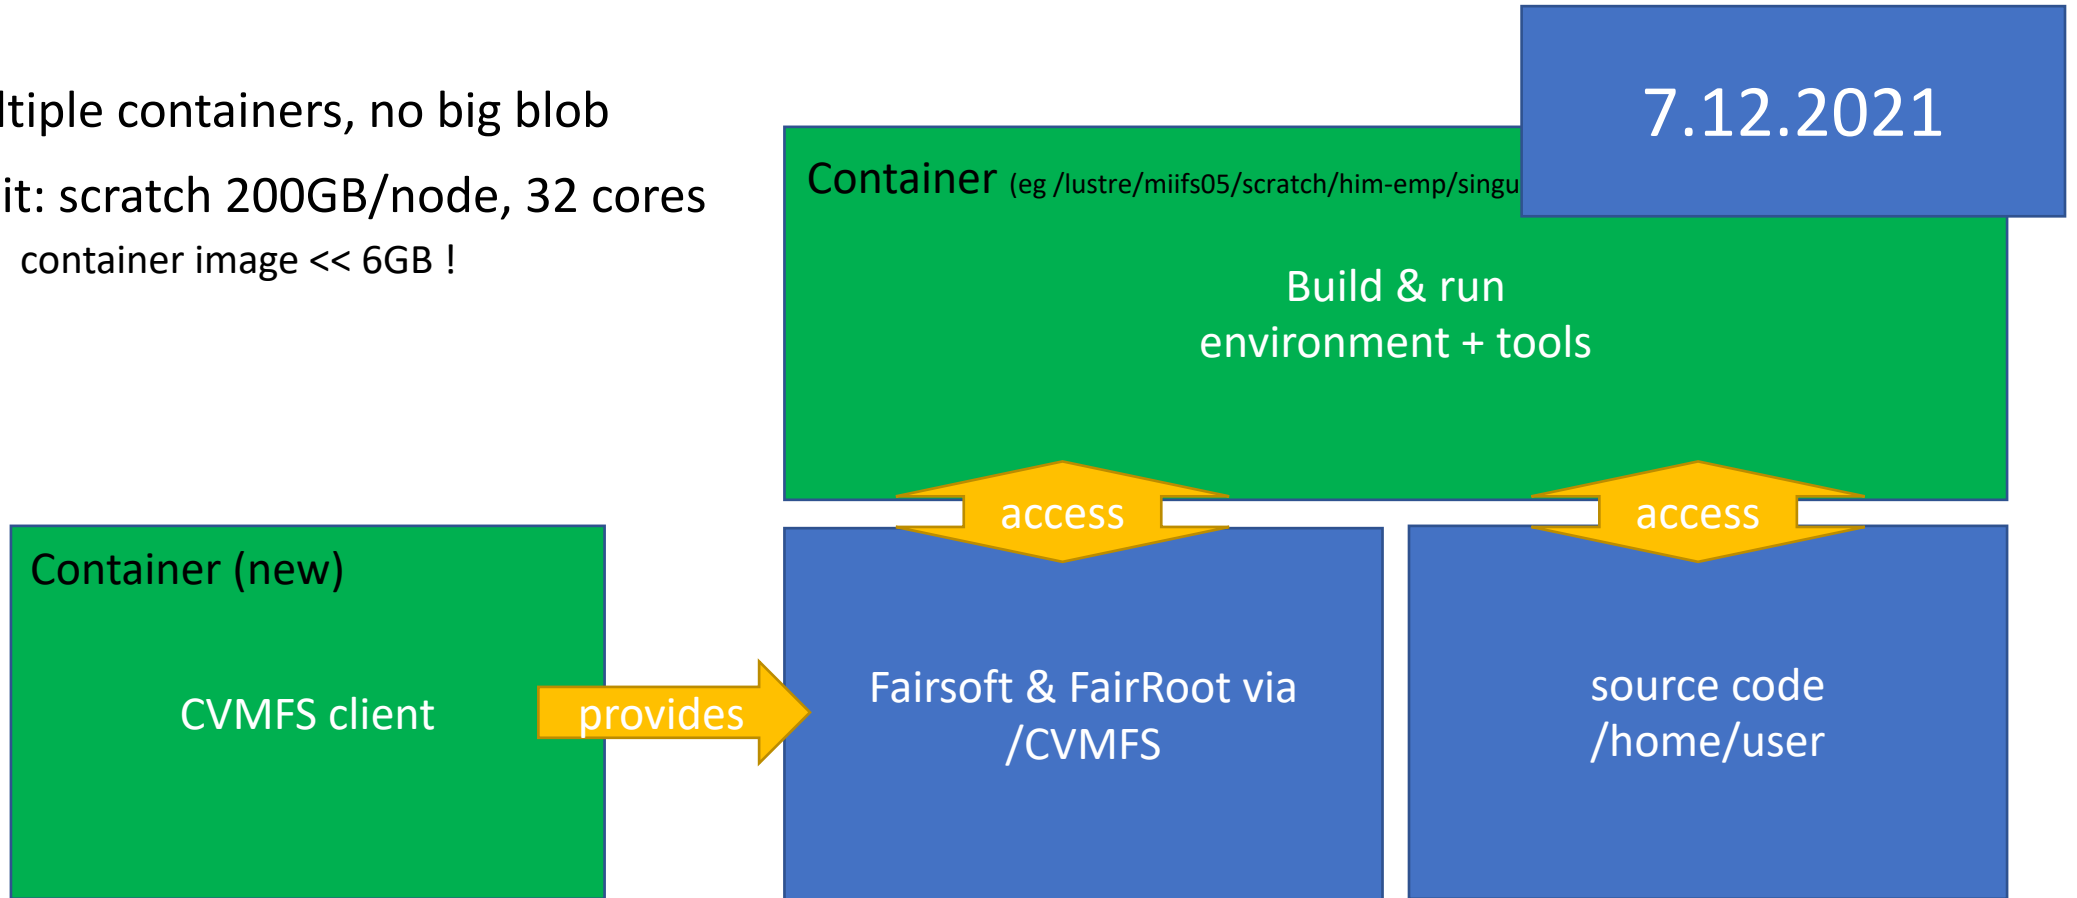
#Load the Singularity module
module load tools/Singularity

#if image is >250MB, change the TMP dir to prevent a overfull /tmp directory on node
SINGULARITY_TMPDIR=/localscratch/${SLURM_JOB_ID}/singularity_tmp/
export SINGULARITY_TMPDIR
mkdir -p $SINGULARITY_TMPDIR

singularity exec /lustre/miifs05/scratch/him-acid/singularity/bdsim.sif /bin/bash -c \
'source /local-tmp/bdsim-build/bin/bdsim.sh && source /usr/local/bin/geant4.sh && \
cd ~/bdsim-test/ && \
bdsim --file=positronBeamline.gmad --batch --ngenerate=1000 --outfile=output'
```

Containerised Analysis

- multiple containers, no big blob
- Limit: scratch 200GB/node, 32 cores
 - container image << 6GB !



Presentation by Roman Klasen on 7th December for PandaRoot analysis

CI with gitlab.rlp.net

- GitLab:
 - extensive web service for source code management
 - git, wiki, issue tracking, continuous integration & deployment
- details: <https://www.zdv.uni-mainz.de/gitlab/>
- Limits:
 - maximal 100 projects/user
 - max. 10 GB /project
 - gitlab runner:
 - 3 VMs, each 2vCPU and 4GB RAM
 - 1h timeout
 - Shared with all on campus
- Status: 27.1.2022

15.3.2022

Simple test

- Tests automatically run, when file “.gitlab-ci.yml” exists.

Fails:

```
image: alpine:latest
```

```
test_simple:
```

```
script:
```

```
- exit 1
```

works:

```
image: alpine:latest
```

```
test_simple:
```

```
script:
```

```
- exit 0
```


Python test (1/2)

- Use gitlab example
- Your repository gets automatically included

```
image: python:latest

# Change pip's cache directory to be inside the project directory since we can
# only cache local items.
variables:
  PIP_CACHE_DIR: "$CI_PROJECT_DIR/.cache/pip"

# If you want to also cache the installed packages, you have to install
# them in a virtualenv and cache it as well.
cache:
  paths:
    - .cache/pip
    - venv/

before_script:
  - python --version # For debugging
  - pip install virtualenv
  - virtualenv venv
  - source venv/bin/activate

test:
  script:
    - python hello.py
```

Python test (2/2)

- Dependent on return value of hello.py

CI fails:

```
import sys  
print("Hello")  
sys.exit(1)
```

CI works:

```
import sys  
print("Hello")  
sys.exit(0)
```

Advanced example: CI for BES3

Ingredients:

1. Docker image (CERN CentOS 7 + build environment + BES3-CVMFS)
<https://gitlab.rlp.net/bes3-mainz/CI/bossdocker/>
2. CI definition file builds the docker image (standard)
3. CI definition file in analysis repository (build generator +)
https://gitlab.rlp.net/bes3-mainz/generator/phokhara/-/blob/CI_integration/.gitlab-ci.yml

Work by Riccardo Aliberti & Thomas Lenz

Alternative Workflow manager

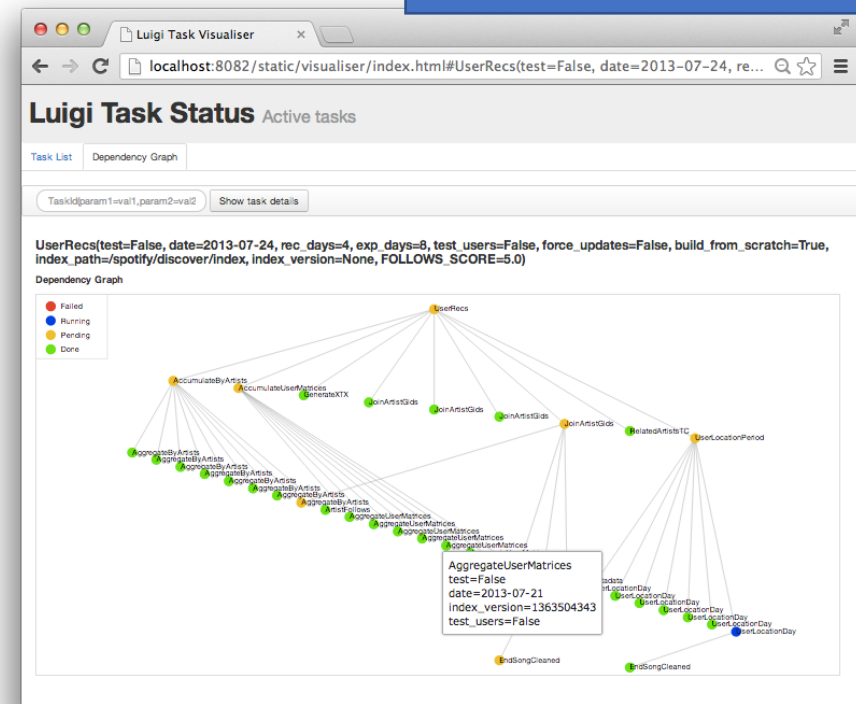


<https://github.com/spotify/luigi>

pip install luigi

A screenshot of the Luigi Task Status web interface. The top navigation bar is green and contains 'Luigi Task Status', 'Task List', 'Dependency Graph', 'Workers', and 'Resources'. Below this, there are several summary cards: 'PENDING TASKS' (3), 'RUNNING TASKS' (1), 'BATCH RUNNING TASKS' (0), 'DONE TASKS' (134), 'FAILED TASKS' (3), and 'UPSTREAM FAILURE' (1). A table below shows a list of tasks, with one task 'LoadTable' highlighted. The table has columns for Name, Details, Priority, Time, and Actions.

29.3.2022




Dependency graph manger

26.4.2022

SLURM scheduler: Multifactor Priority


https://slurm.schedmd.com/priority_multifactor.html



SLURM scheduler: Backfilling

Performed only when jobs with higher prio are not affected

https://slurm.schedmd.com/sched_config.html



Hints (2/2)

- Use devel queue from Mogon2
- Reserve only as much resources as necessary
- Reserve resources with salloc, use them later several times with srun

- Ask for a (recurring) reservation
- Alter SLURM parameters?
- Developer node on Himster 2 exp?

your questions / discussion /
requests to the maintainers?

Next meeting, 24th May 2022

- hand in your topics!

- Check out www.hi-mainz.de/tfp22

TOOLS FOR PHYSICISTS 2022 WERKZEUGE FÜR PHYSIKER

Be prepared for the real lab work - know how to tackle the problems.
10 independent hands-on topics. Get in touch with the pros in their field.
Focusing on thesis starters (Bachelor, Master, PhD), Postdocs welcome.



Technisches Zeichnen
Mo, 25.4. 14:15, Konferenzraum HIM



Mathematica
We, 15.6. 14:15, HIM Conference Hall



Code Optimization
We, 4.5. 14:15, HIM Conference Hall




3D Printing and Designing Basics
We, 22.6. 13:00, HIM Conference Hall



Boost your Analysis with High
Performance Computing
We, 11.5. 14:15, HIM Conference Hall



PCB Design with KiCAD
We, 29.6. 14:15, HIM Conference Hall




Statistics
We, 18.5. 14:15 + 25.5. + 1.6., HIM Conference Hall



EM Noise & Interference in
Measurement Setups
We, 6.7. 14:15, HIM Conference Hall



Estimates, Analogies, Storytelling and
Handwaving for Physicists
We, 8.6. 14:15, HIM Conference Hall



Geant 4
We, 13.7. 14:15 + 20.7., HIM Conference Hall

Visit the course webpage and register today! www.hi-mainz.de/tfp22
Registration mandatory, limited seats.
Organised by Dr. Peter-Bernd Otte (HIM)



Scan me