bw|HPC – S5

KIT
Karlsruhe Institute of Technology

UNIVERSITÄT MANNHEIM

UNIVERSITÄT HEIDELBERG
ZUKUNFT
SEIT 1386

universität ulm

Universität Konstanz

universität freiburg

EBERHARD KARLS
UNIVERSITÄT TÜBINGEN

UNIVERSITY OF HOHENHEIM

University of Stuttgart
Germany

ESSLINGEN UNIVERSITY

Cluster Usage – Basics

Samuel Braun
KIT, SCC

Baden-Württemberg
Ministerium für Wissenschaft,
Forschung und Kunst

KIT | NHR
National High-Performance Computing Center

Outline

- Overview of cluster usage
- Workflow(s)
- Do’s and Dont’s, Policies
- How to write tickets
- IDE usage

What you will learn in this course and why do you need it?

- Summary on cluster usage
- Typical steps on the cluster
- What you are allowed to do and what you should avoid

- **Guidance on understanding the more in-depth topics**

- **Material: https://indico.kit.edu/event/5462/**

Overview of cluster usage

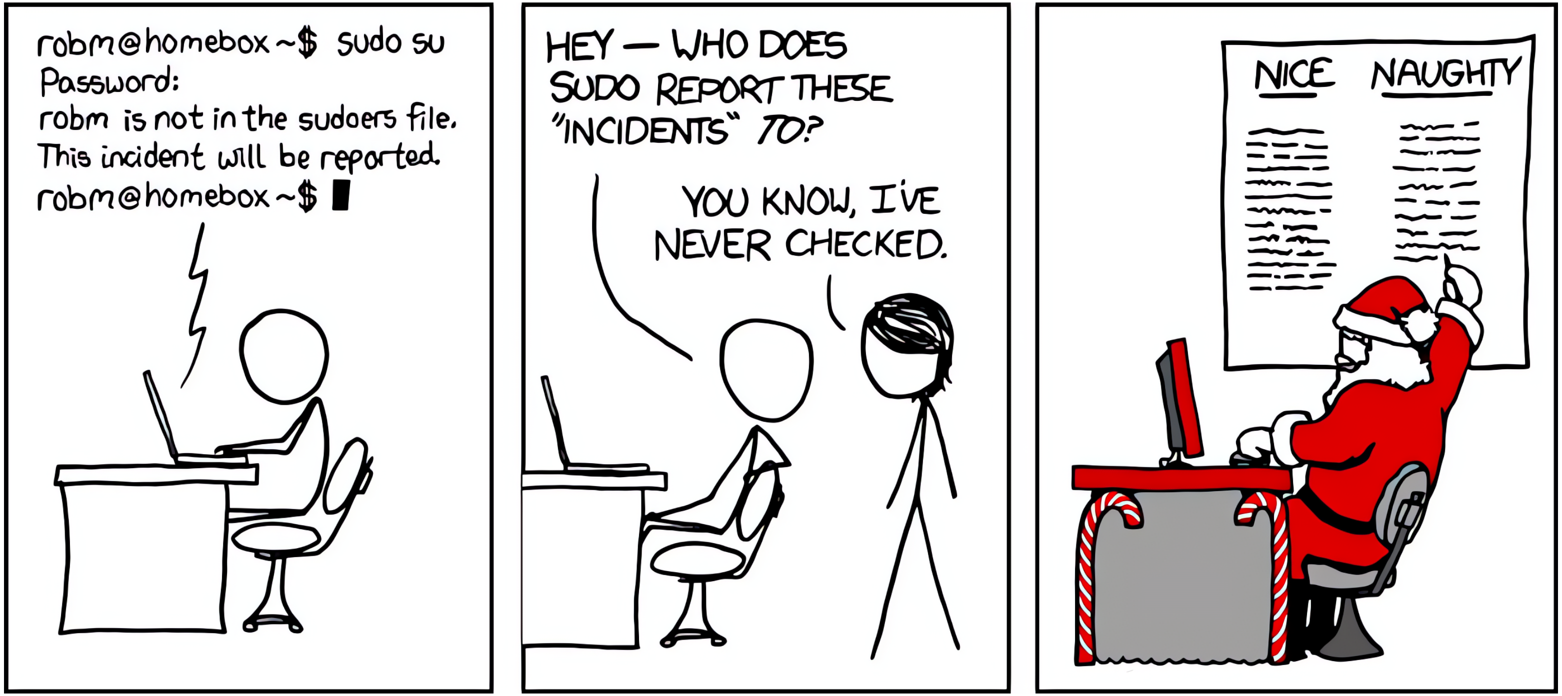

```
robm@homebox~$ sudo su
Password:
robm is not in the sudoers file.
This incident will be reported.
robm@homebox~$ ▮
```

Source: https://xkcd.com/838/

Overview of cluster usage

- Using the cluster involves more than just knowing details on
  - Login procedure
  - File systems
  - Software
  - Batch system
  - Submission of compute time applications
- It also involves understanding
  - Concepts/environment
  - Workflows
- **Goal:** Guidance on understanding/classifying the more in-depth topics of this event

Workflow(s)

- **What you do once/rarely**
  - Get familiar with cluster
  - Know/prepare login method(s)
  - Get required software
  - Acknowledgement in publications
- **What you do regularly**
  - Login/logout
  - Copy data to/from cluster
  - Pre-/post-processing
  - Manage compute jobs
  - Manage data storage

Workflow(s): Get familiar with cluster

- Awareness of which cluster you're using
  - bwUniCluster 3.0
  - JUSTUS 2
  - NEMO 2
  - …
  - HoreKa
- Documentation
  - https://wiki.bwhpc.de
  - https://www.nhr.kit.edu/userdocs/

- Browse the cluster-specific documentation
- Read at least the „Getting Started“
  - Follow the links
- Be aware of Do’s and Dont’s, and Policies

Workflow(s): Know/prepare login method(s)

- SSH – default, **recommended** login method
  - TOTP + Service password
  - SSH-keys: interactive login w/o password
- Jupyterhub
- IDEs like VS Code
  - … is actually SSH

- Recommendations
  - Use password manager
  - Don’t reuse passwords
  - Set up more than one source for OTPs

Workflow(s): Get required software

- Default source: Software-Modules (Lmod)
  - `module spider` etc.
- Compile/install own software
  - `HOME` directory: allowed and intended
  - **You** are responsible
  - No/limited support by us
- Container
  - Apptainer or Enroot

- Python
  - Use Python native venv’s + pip
  - Avoid conda (depending on context)
  - **Do NOT use Anaconda services/channels**
  - You are responsible (!)
- Software from gitlab / github
  - `git clone ...`
  - Allowed and intended
  - You are responsible

Workflow(s): Acknowledgement in publications

- Usage of bwHPC and NHR@KIT ressources is free of charge
- You must acknowledge the ressources in publications, status reports
- Report your papers
  - publications@bwhpc.de
  - publications@nhr.kit.edu

- https://www.nhr.kit.edu/userdocs/horeka/acknowledgement/
- https://wiki.bwhpc.de/e/BwUniCluster3.0/Acknowledgement

Workflow(s): Login/logout

- Login
  - You are in the HOME directory
  - You are on one of the login nodes
  - No computation here!
- Tmux sessions only for allowed activities!

- Logout
  - Do not forget to properly logout
  - Regularly check, if some processes keep running, kill them
  - Check for remaining processes e.g. from VS Code. Kill them.
  - Properly end Jupyter sessions
    - `scancel` also works

Workflow(s): Copy data to/from cluster

- `scp` and `rsync`
- Avoid drag & drop
  - MobaXterm
  - Jupyterhub
- Software/Configurations → HOME
- Data → Workspaces
- Tmux session or compute job for long running copy tasks

Workflow(s): Pre-/post-processing

- **Light** pre-/post-processing
  - **Login node**
  - Low compute intensity
  - Short runtime
  - Not I/O heavy
- Compilation of code
  - Login node only for **fast** compile jobs
  - `make -j`: No!
  - Be aware: `pip install` sometimes also does compilation
- More demanding activities
  - **Compute nodes**
- Visual pre-/post-processing
  - VNC-session
  - Server-client tools, e.g. ParaView
  - X11-Forwarding (Linux, MobaXterm)

Workflow(s): Manage compute jobs

- Know the queues
- Interactive jobs: `salloc`
- Batch jobs: `sbatch`
  - Maintain your batch scripts
  - Don’t blindly recycle old batch scripts

- Keep track of ongoing and completed jobs
- Check running jobs
  - `squeue`
  - By email
    Make sure the email address you provided is valid!
- Waiting times are normal
  - Please NO ticket after 2 hours of waiting time
  - Possibly adjust your resource allocations

Workflow(s): Manage data storage

- **HOME**
  - Software
  - Config-files
- **Workspaces**
  - All of your (big) data
  - Keep an eye on the lifetime
  - Renew on time
  - Delete data and workspaces if processing is done
- Additional file systems
  - As per requirements

- Be aware of which file system features backups / snapshots
  - Always take **data loss due to technical errors** into account

Do’s and Dont’s, Policies

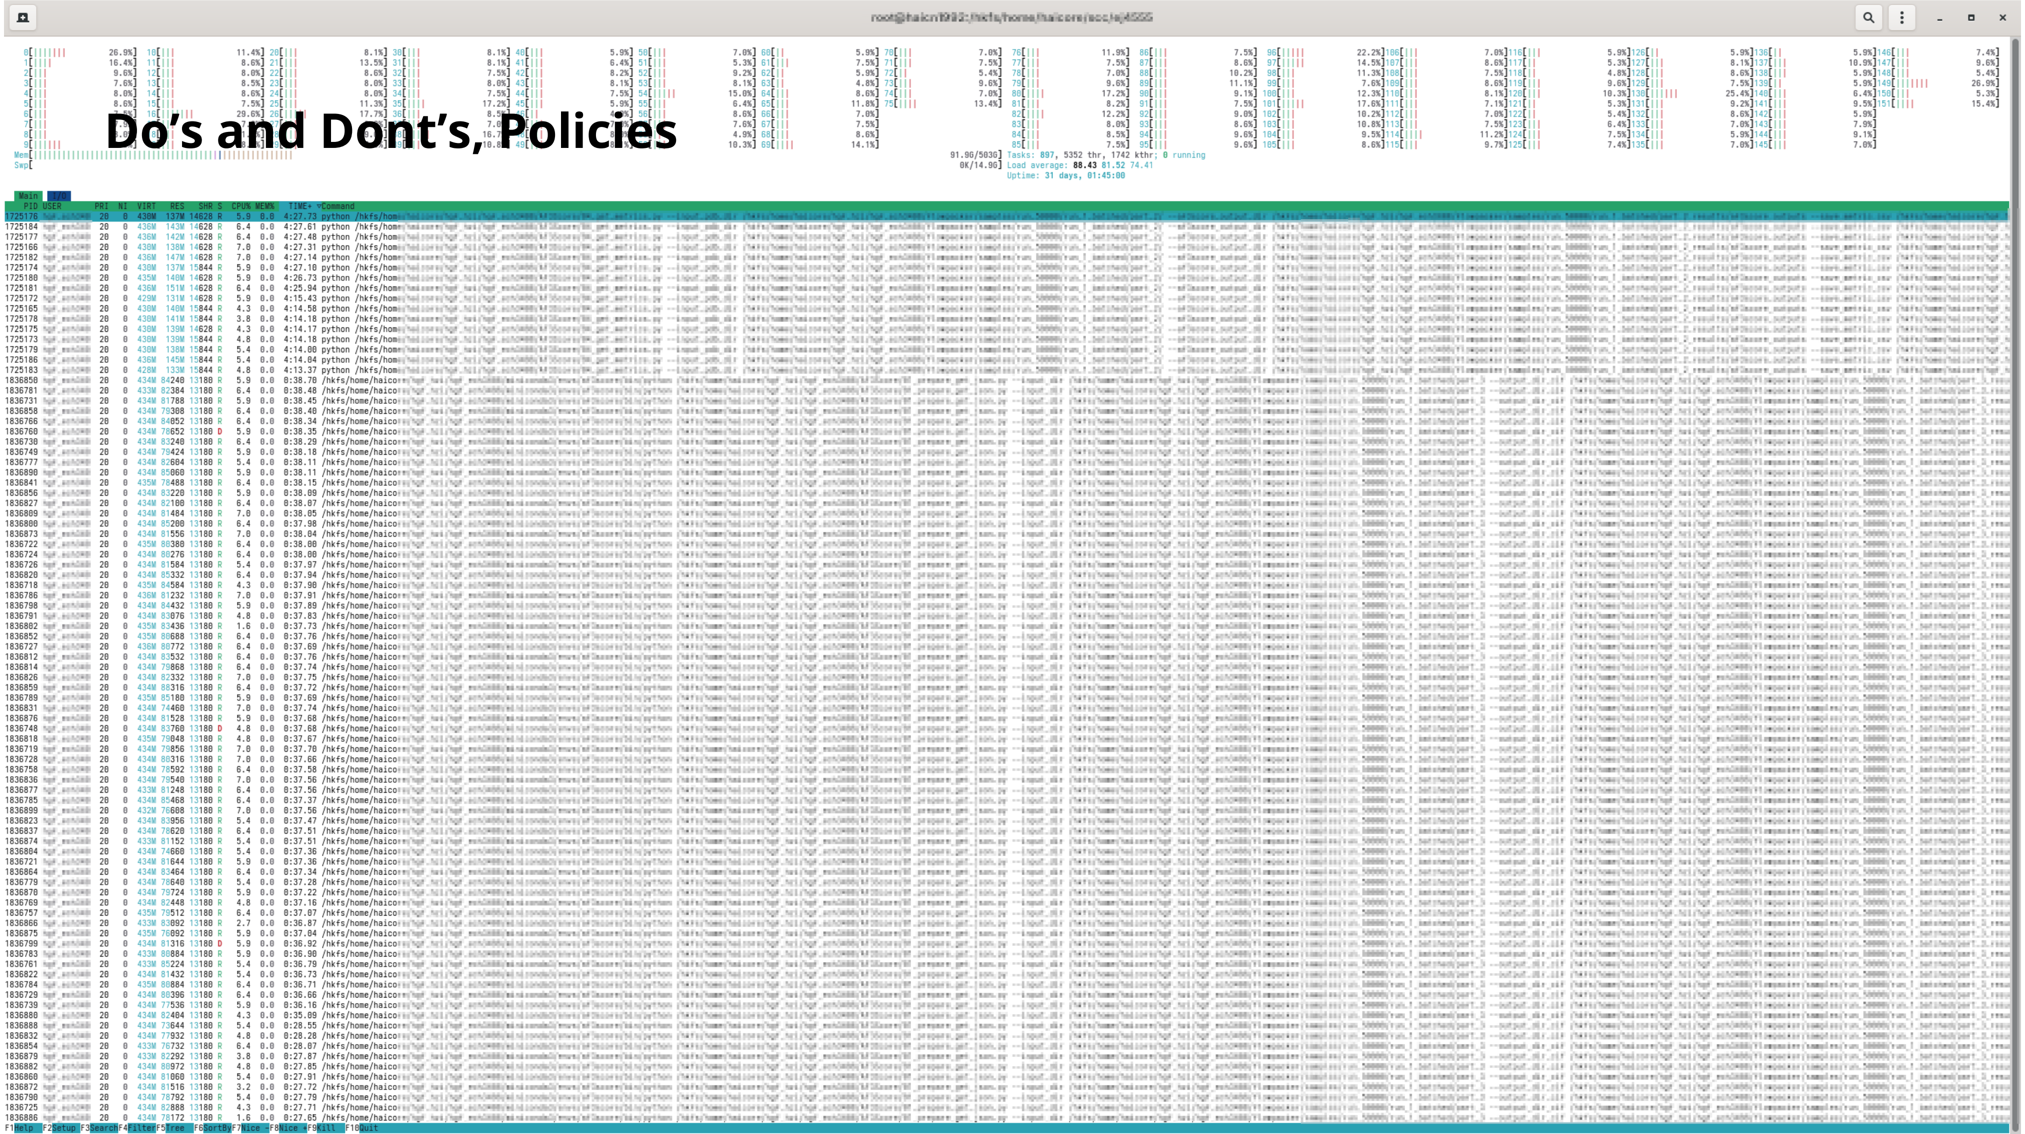

```
Mem[                91.9G/503G] Tasks: 897, 5352 thr, 1742 kthr; 0 running
Swp[                  0K/14.9G] Load average: 88.43 81.52 74.41
                                Uptime: 31 days, 01:45:00

  Main    I/O
    PID USER       PRI  NI  VIRT   RES   SHR S  CPU% MEM%   TIME+  Command
1725176 [redacted]  20   0  430M  137M 14628 R   5.9  0.0  4:27.73 python /hkfs/hom[redacted]
1725184 [redacted]  20   0  436M  143M 14628 R   6.4  0.0  4:27.61 python /hkfs/hom[redacted]
1725177 [redacted]  20   0  436M  142M 14628 R   6.4  0.0  4:27.48 python /hkfs/hom[redacted]
1725166 [redacted]  20   0  430M  138M 14628 R   7.0  0.0  4:27.31 python /hkfs/hom[redacted]
1725182 [redacted]  20   0  436M  147M 14628 R   7.0  0.0  4:27.14 python /hkfs/hom[redacted]
1725174 [redacted]  20   0  430M  137M 15844 R   5.9  0.0  4:27.10 python /hkfs/hom[redacted]
1725180 [redacted]  20   0  435M  140M 14628 R   5.9  0.0  4:26.73 python /hkfs/hom[redacted]
1725181 [redacted]  20   0  436M  151M 14628 R   6.4  0.0  4:25.94 python /hkfs/hom[redacted]
1725172 [redacted]  20   0  429M  131M 14628 R   5.9  0.0  4:15.43 python /hkfs/hom[redacted]
1725165 [redacted]  20   0  430M  140M 15844 R   4.3  0.0  4:14.58 python /hkfs/hom[redacted]
1725178 [redacted]  20   0  430M  141M 15844 R   3.8  0.0  4:14.18 python /hkfs/hom[redacted]
1725175 [redacted]  20   0  430M  139M 14628 R   4.3  0.0  4:14.17 python /hkfs/hom[redacted]
1725173 [redacted]  20   0  430M  139M 15844 R   4.8  0.0  4:14.18 python /hkfs/hom[redacted]
1725179 [redacted]  20   0  430M  138M 15844 R   5.4  0.0  4:14.00 python /hkfs/hom[redacted]
1725186 [redacted]  20   0  436M  145M 15844 R   5.4  0.0  4:14.04 python /hkfs/hom[redacted]
1725183 [redacted]  20   0  428M  133M 15844 R   4.8  0.0  4:13.37 python /hkfs/hom[redacted]
1836850 [redacted]  20   0  434M 84240 13180 R   5.9  0.0  0:38.70 /hkfs/home/haico[redacted]
1836781 [redacted]  20   0  433M 82384 13180 R   6.4  0.0  0:38.48 /hkfs/home/haico[redacted]
1836731 [redacted]  20   0  434M 81788 13180 R   5.9  0.0  0:38.45 /hkfs/home/haico[redacted]
1836858 [redacted]  20   0  434M 79308 13180 R   6.4  0.0  0:38.40 /hkfs/home/haico[redacted]
1836766 [redacted]  20   0  434M 84052 13180 R   6.4  0.0  0:38.34 /hkfs/home/haico[redacted]
1836760 [redacted]  20   0  434M 78652 13180 D   5.9  0.0  0:38.35 /hkfs/home/haico[redacted]
1836730 [redacted]  20   0  434M 83240 13180 R   6.4  0.0  0:38.29 /hkfs/home/haico[redacted]
1836749 [redacted]  20   0  434M 79424 13180 R   5.9  0.0  0:38.18 /hkfs/home/haico[redacted]
1836777 [redacted]  20   0  434M 82604 13180 R   5.4  0.0  0:38.11 /hkfs/home/haico[redacted]
1836890 [redacted]  20   0  434M 85060 13180 R   5.9  0.0  0:38.11 /hkfs/home/haico[redacted]
1836841 [redacted]  20   0  435M 78488 13180 R   6.4  0.0  0:38.15 /hkfs/home/haico[redacted]
1836856 [redacted]  20   0  434M 83220 13180 R   5.9  0.0  0:38.09 /hkfs/home/haico[redacted]
1836827 [redacted]  20   0  434M 82100 13180 R   6.4  0.0  0:38.07 /hkfs/home/haico[redacted]
1836809 [redacted]  20   0  434M 81484 13180 R   7.0  0.0  0:38.05 /hkfs/home/haico[redacted]
1836800 [redacted]  20   0  434M 85200 13180 R   6.4  0.0  0:37.98 /hkfs/home/haico[redacted]
1836873 [redacted]  20   0  434M 81556 13180 R   7.0  0.0  0:38.04 /hkfs/home/haico[redacted]
1836722 [redacted]  20   0  435M 80380 13180 R   6.4  0.0  0:38.00 /hkfs/home/haico[redacted]
1836724 [redacted]  20   0  434M 80276 13180 R   6.4  0.0  0:38.00 /hkfs/home/haico[redacted]
1836726 [redacted]  20   0  434M 81584 13180 R   5.4  0.0  0:37.97 /hkfs/home/haico[redacted]
1836820 [redacted]  20   0  434M 85332 13180 R   6.4  0.0  0:37.94 /hkfs/home/haico[redacted]
1836718 [redacted]  20   0  435M 84584 13180 R   4.3  0.0  0:37.90 /hkfs/home/haico[redacted]
1836786 [redacted]  20   0  436M 81232 13180 R   7.0  0.0  0:37.91 /hkfs/home/haico[redacted]
1836798 [redacted]  20   0  434M 84432 13180 R   5.9  0.0  0:37.89 /hkfs/home/haico[redacted]
1836791 [redacted]  20   0  434M 83076 13180 R   4.8  0.0  0:37.83 /hkfs/home/haico[redacted]
1836802 [redacted]  20   0  435M 83436 13180 R   1.6  0.0  0:37.73 /hkfs/home/haico[redacted]
1836852 [redacted]  20   0  435M 80688 13180 R   6.4  0.0  0:37.76 /hkfs/home/haico[redacted]
1836727 [redacted]  20   0  436M 80772 13180 R   6.4  0.0  0:37.69 /hkfs/home/haico[redacted]
1836812 [redacted]  20   0  434M 83532 13180 R   6.4  0.0  0:37.76 /hkfs/home/haico[redacted]
1836814 [redacted]  20   0  434M 79868 13180 R   6.4  0.0  0:37.74 /hkfs/home/haico[redacted]
1836826 [redacted]  20   0  434M 82332 13180 R   7.0  0.0  0:37.75 /hkfs/home/haico[redacted]
1836859 [redacted]  20   0  434M 80316 13180 R   6.4  0.0  0:37.72 /hkfs/home/haico[redacted]
1836789 [redacted]  20   0  435M 85180 13180 R   5.9  0.0  0:37.69 /hkfs/home/haico[redacted]
1836831 [redacted]  20   0  434M 74460 13180 R   7.0  0.0  0:37.74 /hkfs/home/haico[redacted]
1836876 [redacted]  20   0  434M 81528 13180 R   5.9  0.0  0:37.68 /hkfs/home/haico[redacted]
1836748 [redacted]  20   0  434M 83760 13180 D   4.8  0.0  0:37.68 /hkfs/home/haico[redacted]
1836818 [redacted]  20   0  435M 79048 13180 R   4.8  0.0  0:37.67 /hkfs/home/haico[redacted]
1836719 [redacted]  20   0  434M 79856 13180 R   7.0  0.0  0:37.70 /hkfs/home/haico[redacted]
1836728 [redacted]  20   0  434M 80316 13180 R   7.0  0.0  0:37.66 /hkfs/home/haico[redacted]
1836758 [redacted]  20   0  434M 78592 13180 R   6.4  0.0  0:37.58 /hkfs/home/haico[redacted]
1836836 [redacted]  20   0  434M 79540 13180 R   7.0  0.0  0:37.56 /hkfs/home/haico[redacted]
1836877 [redacted]  20   0  433M 81248 13180 R   6.4  0.0  0:37.56 /hkfs/home/haico[redacted]
1836785 [redacted]  20   0  434M 85468 13180 R   6.4  0.0  0:37.37 /hkfs/home/haico[redacted]
1836899 [redacted]  20   0  432M 76608 13180 R   7.0  0.0  0:37.56 /hkfs/home/haico[redacted]
1836823 [redacted]  20   0  434M 83956 13180 R   5.4  0.0  0:37.47 /hkfs/home/haico[redacted]
1836837 [redacted]  20   0  434M 78620 13180 R   6.4  0.0  0:37.51 /hkfs/home/haico[redacted]
1836874 [redacted]  20   0  433M 81152 13180 R   5.4  0.0  0:37.51 /hkfs/home/haico[redacted]
1836804 [redacted]  20   0  434M 74660 13180 R   5.4  0.0  0:37.36 /hkfs/home/haico[redacted]
1836721 [redacted]  20   0  434M 81644 13180 R   5.9  0.0  0:37.36 /hkfs/home/haico[redacted]
1836864 [redacted]  20   0  434M 83464 13180 R   6.4  0.0  0:37.34 /hkfs/home/haico[redacted]
1836779 [redacted]  20   0  434M 78640 13180 R   5.4  0.0  0:37.28 /hkfs/home/haico[redacted]
1836870 [redacted]  20   0  434M 79724 13180 R   5.9  0.0  0:37.22 /hkfs/home/haico[redacted]
1836769 [redacted]  20   0  434M 82448 13180 R   4.8  0.0  0:37.16 /hkfs/home/haico[redacted]
1836757 [redacted]  20   0  435M 79512 13180 R   6.4  0.0  0:37.07 /hkfs/home/haico[redacted]
1836866 [redacted]  20   0  433M 83092 13180 R   2.7  0.0  0:36.87 /hkfs/home/haico[redacted]
1836875 [redacted]  20   0  435M 76092 13180 R   5.9  0.0  0:37.04 /hkfs/home/haico[redacted]
1836799 [redacted]  20   0  434M 81316 13180 D   5.9  0.0  0:36.92 /hkfs/home/haico[redacted]
1836783 [redacted]  20   0  433M 80884 13180 R   5.9  0.0  0:36.90 /hkfs/home/haico[redacted]
1836761 [redacted]  20   0  433M 85224 13180 R   5.4  0.0  0:36.79 /hkfs/home/haico[redacted]
1836822 [redacted]  20   0  434M 81432 13180 R   5.4  0.0  0:36.73 /hkfs/home/haico[redacted]
1836784 [redacted]  20   0  435M 80884 13180 R   6.4  0.0  0:36.71 /hkfs/home/haico[redacted]
1836729 [redacted]  20   0  434M 80396 13180 R   6.4  0.0  0:36.66 /hkfs/home/haico[redacted]
1836739 [redacted]  20   0  434M 77536 13180 R   5.9  0.0  0:36.16 /hkfs/home/haico[redacted]
1836880 [redacted]  20   0  434M 82404 13180 R   4.3  0.0  0:35.09 /hkfs/home/haico[redacted]
1836888 [redacted]  20   0  434M 73644 13180 R   5.4  0.0  0:28.55 /hkfs/home/haico[redacted]
1836832 [redacted]  20   0  434M 77932 13180 R   4.8  0.0  0:28.28 /hkfs/home/haico[redacted]
1836854 [redacted]  20   0  433M 76732 13180 R   6.4  0.0  0:28.07 /hkfs/home/haico[redacted]
1836879 [redacted]  20   0  433M 82292 13180 R   3.8  0.0  0:27.87 /hkfs/home/haico[redacted]
1836882 [redacted]  20   0  434M 80972 13180 R   4.8  0.0  0:27.85 /hkfs/home/haico[redacted]
1836860 [redacted]  20   0  434M 81060 13180 R   5.4  0.0  0:27.91 /hkfs/home/haico[redacted]
1836872 [redacted]  20   0  434M 81516 13180 R   3.2  0.0  0:27.72 /hkfs/home/haico[redacted]
1836790 [redacted]  20   0  434M 78792 13180 R   5.4  0.0  0:27.79 /hkfs/home/haico[redacted]
1836725 [redacted]  20   0  434M 82888 13180 R   4.3  0.0  0:27.71 /hkfs/home/haico[redacted]
1836886 [redacted]  20   0  434M 78172 13180 R   1.6  0.0  0:27.65 /hkfs/home/haico[redacted]
F1Help  F2Setup F3SearchF4FilterF5Tree  F6SortByF7Nice -F8Nice +F9Kill  F10Quit
```

Do’s and Dont’s, Policies

- A cluster is **used** and **shared** by **many people at the same time**
- You're not alone on the cluster
- Ressources are limited
    - CPUs, GPUs, Memory
    - Storage
    - Bandwidth(s)
    - Compute time w/in project/group

Do’s and Dont’s, Policies

- **Policies** (cf. https://wiki.bwhpc.de/e/BwUniCluster3.0/Policies)
  - File system quotas
  - …
  - **Allowed Activities on Login Nodes**
    - To guarantee usability for all the users of clusters you must not run your compute jobs on the login nodes.
    - Compute intensive jobs must be submitted to the queuing system.
    - **Any compute job running on the login nodes will be terminated without any notice.**
    - Any long-running compilation or any long-running pre- or post-processing of batch jobs must also be submitted to the queuing system.
- **In the event of repeated or serious violations, we will suspend the account. We will contact the user via a ticket assigned to them.**

Do’s and Dont’s, Policies

- **Dont’s**
  - Calculations on login nodes
  - `watch squeue / sinfo_t_idle`
    This puts heavy load on the workload manager
  - Tmux sessions apart from allowed activities
  - Exessive VS Code activities w/ many open work sheets, processes, AI-helpers
  - Dangling VS Code processes

- **Do’s**
  - Keep an eye on your resource usage
  - Compute „intensity“: `htop`
  - Disk usage: `du`
  - Manually login and check for dangling VS Code residuals

How to write tickets

- **Which cluster** you are referring to?
- Avoid sending **screenshots**
- You may write in **German**
- Choose a **meaningful subject**
- **Structure** your message
- Ask a **question** or make a **request**
- **Be kind** (and sometimes patient)

❤️

- https://bw-support.scc.kit.edu
- https://nhr-helpdesk.scc.kit.edu

IDE usage

- Visual Studio Code
  - Remote SSH
  - Mind security implications
    - code tunnel (**please don’t**)
    - AI-helpers
  - Mind ressource usage
    - AI-Helpers
    - Many open work sheets, processes

- Keep an eye on your resource usage
  - `htop` on login nodes
- Remote development/debugging, heavy prossesing
  - Use compute nodes

Weitere Informationen

Die Initiative „Baden-Württemberg High Performance Computing (bwHPC)“ bietet Wissenschaftler*innen in Baden-Württemberg Zugang zu verschiedenen Hochleistungsrechnern – ideal für datenintensive Forschung, komplexe Berechnungen oder KI-Entwicklung. Auch Studierende, Promovierende und Forschende aus Baden-Württemberg können diese landesweite Infrastruktur und die zugehörigen Support- und Schulungsangebote nutzen.
bwHPC betreibt mehrere spezialisierte Clusterrechner, die in Fachbereiche aufgeteilt sind:

- **bwUniCluster 3.0** (Karlsruher Institut fürTechnologie): Grundversorgung für alle Disziplinen
- **JUSTUS 2** (Universität Ulm): Theoretische Chemie, Quanten- und Festkörperphysik
- **NEMO 2** (Universität Freiburg): Neurowissenschaften, Elementarteilchenphysik, Mikrosystemtechnik, Materialwissenschaften
- **Helix** (Universität Heidelberg): Strukturelle Biologie und Systembiologie, Medizinwissenschaft, weiche Materie, Computational Humanities, Mathematik und Informatik
- **BinAC 2** (Universität Tübingen): Bio- und Medizininformatik, Pharmazie, Astrophysik und Geowissenschaften

Die Registrierung erfolgt clusterspezifisch und setzt ein sogenanntes „Entitlement“ sowie ein genehmigtes Rechenvorhaben voraus. Eine Übersicht zum Anmeldeprozess finden Sie im bwHPC-Wiki.

Webseite

https://www.bwhpc.de

Wiki

https://wiki.bwhpc.de

Mail

office@bwhpc.de