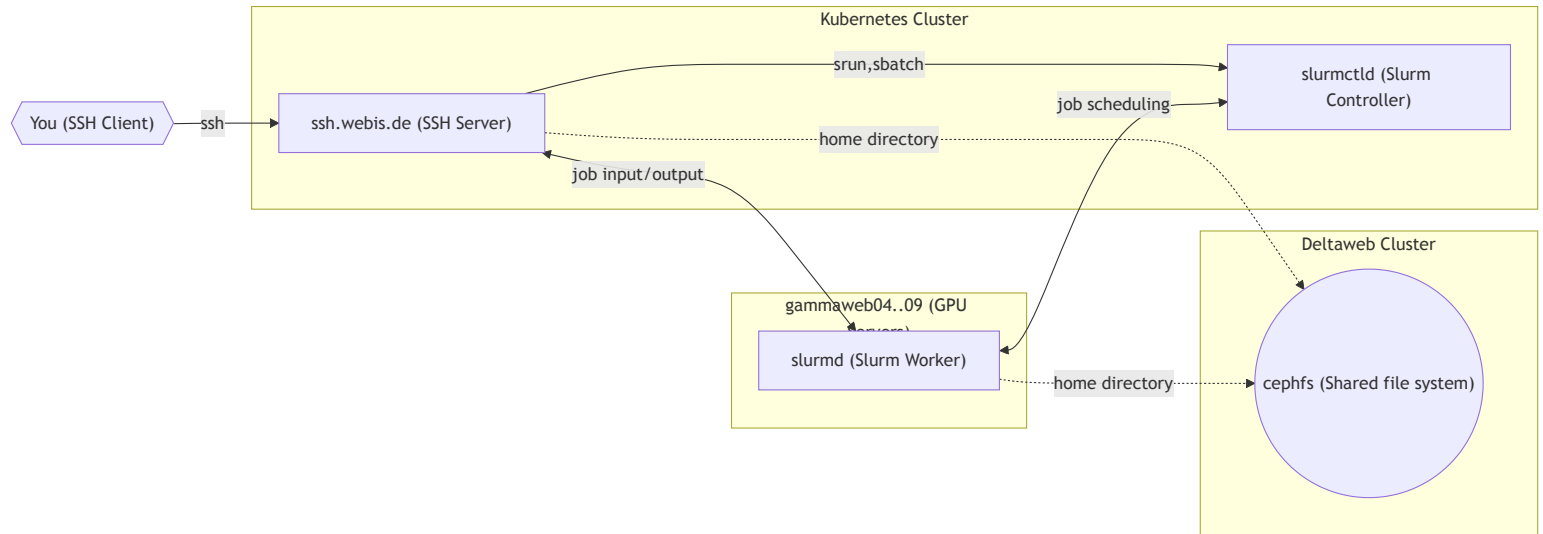


Hands-On Introduction to Slurm on the Webis GPU Cluster

Content

- ❑ Background Information
- ❑ Finding Help
- ❑ Getting Started with the Cluster and PyCharm

Background Information



Cluster components:

- ❑ **SSH gateway (`ssh.webis.de`):**
Runs an SSH server that you can connect to; do not run any jobs directly on this node.
- ❑ **Slurm worker (`gammawebXX`):**
Accepts jobs scheduled from the SSH gateway.

- ❑ **Worker-local file system** (`/var/tmp/$USER`):
Plenty of fast temporary storage accessible from your Slurm worker. Will be available as long as your job runs.
- ❑ **Shared file system CephFS** (including the home directory `/mnt/ceph/storage/data-tmp/current/$USER`):
Slow persistent but temporary storage. Can be used to transfer data between jobs and into the cluster (via the SSH gateway).

Finding Help

1. Webis FAQ [faq.webis.de]
2. Webis Knowledge Base [kb.webis.de]
3. (As applicable) Google
4. Discord (including the error log)

Getting Started

- ❑ Set up and connect to the VPN [\[kb.webis.de\]](https://kb.webis.de)
- ❑ Generate an SSH key pair [\[faq.webis.de\]](https://faq.webis.de)
- ❑ Enable and configure SSH authentication in GitLab [\[faq.webis.de\]](https://faq.webis.de)
- ❑ (Optional) Configure `/.ssh/config`
- ❑ Read and understand the Slurm user guide [\[kb.webis.de\]](https://kb.webis.de)
- ❑ Run a Slurm job on gammaweb

PyCharm

- Synchronizing files between a local repository and the CephFS home directory [[jetbrains.com](https://www.jetbrains.com)]