

# How to run a supercomputer

Joshua Kobschätzki

Extended abstract for CLT 2025 presentation proposal

Talk length: 15 to 20 minutes

The talk will mainly cover the Linux based infrastructure behind the [NHR@ZIB](#) cluster. The cluster is operated at [ZIB](#) in Berlin as part of the [NHR](#) and offers free<sup>1</sup> HPC compute resources to scientists.

The talk will provide a concise overview of the software used to run an HPC cluster, covering core system components like Slurm and Warewulf and their key responsibilities. It will also highlight network aspects to address the challenges and requirements of a heterogeneous HPC cluster.

The goal is to introduce the operation of a supercomputer from a system administration perspective, offering a high-level overview without delving too deeply into specific components. The target audience is individuals with basic knowledge of Linux and computer networks, such as hobbyists and enthusiasts.

The presentation may also include a brief section on how students and researchers can apply for a user account, granting access to 75K core hours per quarter to support research and studies.

This is my first time presenting this content, so a longer session would be beneficial. System architecture stuff of HPC systems are rarely shared at conferences, but I'm unsure how interesting HPC system architecture will be to CLT.

---

<sup>1</sup>The NHR project and NHR@ZIB are funded from public funds. The purpose of our project is to provide centralised HPC compute resources to support researchers with compute intensive tasks, like physics simulations. Projects from universities and other research institutions can apply for compute time and are granted compute time free of any charge.