

Contact & Help

Need help running your job on the cluster? Need a question answered about a particular tool or cluster resources?

- Please contact the cluster team at cluster-help@luis.uni-hannover.de **using the email address associated with your cluster account.**
- You can also call us on the phone +49 511 762 791000 (be aware, though, that some problems are complex and we sometimes need to look through logs and think about what we see, too).
- **Sie können uns natürlich auch auf Deutsch kontaktieren** (we get a lot of inquiries in English, even from obviously German-speaking users; so we just want to mention that both German and English are equally suitable to contact our support. We maintain the documentation in English since many users come from an international background and almost anything on a cluster system is in English, anyway).

However, please read this documentation first and check whether your question is already answered in the [FAQ](#) or possibly even a page specifically written for the issue at hand. If you find the documentation lacking, let us know what's missing so we can update and improve it. We will also be happy to integrate useful workflows users develop e.g. using software installed on the cluster, since we are no experts in these products.

Whenever you have a question, please try to help us help you by providing the following information:

- **Do NOT use “reply” to an old mail to create a new inquiry.** Create a new one, choose a suitable subject line, and do not address it to an individual person, but to the common cluster-help-Adress mentioned above. Otherwise, the topic will be both obsolete and completely different to what you really want, and on top, it will be attached to an old case. You will **not** get a timely response if we think you are just responding to an introductory course invitation, or if you personally address someone who is currently not at their desk. And your case will get sorted to the bottom of the queue in at least some of the to-do lists we have due to its old age — and thus easily get overlooked.
- **State your username** (nh...), **Job ID(s)** and **locations on the system** (directories, hostnames etc.)
- **NEVER send us your password.** We do not need it. **Sending your password will lead to your account being locked.**
- **One problem per case:** Nobody at the LUIS knows everything. If you pack several different topics into one case, fewer persons are able to answer them all, even if three out of your five questions would be easy to reply to. So the most difficult question in an inquiry will usually influence response time. One question at a time will usually get you answers much quicker. *If* things could depend on each other, then by all means describe them in one case. Some things are complex.
- **Details known about the problem, e.g.**
 - JobID(s) of the job(s) that have problems
 - what command(s) in what sequence lead to the problem? What did you do? In which directory?
 - **Have you checked your quota (use the command `checkquota`)?** (really!)
 - the time something happened; be precise if you can: log files are big and sometimes there are 1000 entries every minute

- the batch script (as an attachment or - even better - its location on the cluster)
- output from the program, if available, like myjob.o12345 and myjob.e12345 (as an attachment or - even better - its location on the cluster)
- what did you already try to solve the problem yourself?
- for batch jobs:
 - did you request sufficient memory for your job?
 - in case of problems running a multi-node job: did you already try a single-node job? What was the outcome? Is it even reasonable to run your job over several nodes when it would fit on a single node (intra-node communication between processes may be a lot faster than inter-node communication)?
 - is the combination of cores, nodes and memory reasonable or even possible on the system? Did you check the [list of nodes](#) available to find your setup, or did you just copy a very old setup for a partition that does not even exist any more from a former colleague of a former colleague?
 - make sure you do not change/delete the files we need to trace your problem while we are at the case. The best method is to provide a minimal setup in a separate directory that we can use. The simpler and clearer your bug report, the easier it is to isolate the problem. A very good case shows “this runs ok” in comparison to “this does not” (e.g. a smaller job setup, less single/multi nodes, memory request, different parameter, different environment variable...)
- for questions related to file transfers or graphical logins: **how do you access the cluster?** Weblogin? PuTTY? X2Go? SSH? What is your workstation's OS?
- what is unusual or unexpected in the output files you already have?
- bus errors and segmentation violations usually are not a system problem. They occur when a program tries to access memory locations or hardware that does not exist, which may happen e.g. when the software has problems with proper use of pointers or when it does not deal properly with unexpected input data.
- Screenshots help, but consider doing copy&paste in plain text for textual information
- **After you got your response, do not just reuse the case to ask for something completely different.** The person who replied to your first case may not know the answer to your new problem, but the case remains assigned to the same agent. One case per problem.
- In case you solve your problem yourself in the time it takes us to answer your case, please tell us. Ideally, you also tell us *how* you solved it: another user may have the same problem one day, and we do not know everything.
- **We will NOT do your Homework for you.** That means we expect you to at least read your own log files and check with the documentation and the FAQ on this site. We also expect you to do your own thinking.

Examples for typical cases that make it almost impossible for us to help without asking further questions

- **“The cluster is running very slow”** *What are you trying to do? On what node(s)? How do you log in (describe the method, ssh, X2Go, OpenOnDemand)? How do you quantify “slow”? What part of the cluster is “slow”, a file system?*
- **“I cannot log in”** *What is your account name? Did you succeed to log in previously, or is this the first time? How do you try to log in (ex. ssh, X2Go)? Did you check that you are within the LUH network, either in an institute, or are you using the VPN? What is the operating system of the machine you come from? When exactly (hh:mm) did you try it? Do you even get a password prompt or similar? Which messages, if any, did you get? Did you try an alternative method, like e.g. ssh instead of the web site? Did you check your disk quota?*

- **“My job does not start”** *What do you mean by “does not start” - 2 minutes or 3 weeks? Where is the job script or the job ID, respectively? Which account? Are you using a resource that is scarce in the cluster, like large-memory nodes or GPU nodes?*
- **“My jobs crashed”** *Which job IDs, what account? Did you check that you are not over quota (checkquota)? Did the same jobs run previously without errors, and if so, what changed?*
- **“Some part of a complex software suite like Ansys, Abaqus, Comsol, Matlab etc. does not work as I expect”** *What exactly did you do to create the problem? Please also include a description of how you started the software, which version/modules you loaded, what commands you entered, where you clicked to start what. We are in no way experts in every software that's available on the cluster and can only provide basic support for individual applications on a best-effort basis.*

From:

<https://docs.cluster.uni-hannover.de/> - **Cluster Docs**

Permanent link:

https://docs.cluster.uni-hannover.de/doku.php/guide/how_to_get_support

Last update: **2026/03/30 17:49**

