

FAQ: Frequently asked questions

Formal issues

I forgot my password!

→ Go to your project management (usually that's the people of your own institute) who will help you with that. We can not and will not reset your password, that is entirely the project management's responsibility.

I get mail from the cluster, but I am not working at the LUH any more - please remove me from the list!

→ Talk to the person at your former institute who provided you with your account credentials.

Accounts on the cluster are not created by us ("Scientific Computing"), but by the person responsible for the management of the institute's compute project they have with the LUIS. They use an interface called BIAS for this, which is also maintained by a different group than us, and we only receive the accounts from there. So your account in the project your institute has with LUIS needs to be deactivated when you leave the institute, and that automatically removes you from our mailing lists.

In addition, the email addresses that these removal requests come from are usually NOT on any of our lists. They only get email because messages addressed to their former institute's mail accounts are being forwarded and the account has not been properly deactivated/removed.

Connecting to the cluster, disk quota problems

I can not connect to the cluster!

→ First check whether you get any messages at all or if the system appears totally silent. If that is the case, check whether you are possibly "outside" of the LUH network. You'd then first need to connect via VPN.

→ If you can connect to the cluster from a shell (command line) via ssh, but are denied login graphically via X2Go, chances are high that you are over your quota (maximum disk space allocation) in the HOME directory and/or your grace period has expired. Use the command "checkquota" on the command line to verify this and or see the corresponding FAQ about "Disk Quota Exceeded".

→ If you just deleted files in your HOME directory without thinking or respecting any of the hints on the file systems page, you may first need to log into the cluster via ssh and password, then have a look around at what you destroyed and recreate your environment. If you still need help, make sure

to follow the guidelines in the how to get help section of this documentation. Do your own homework and tell us what you already did to minimize the problems you created and where you need help. Have a look at the next FAQ as well. For those who really do not read the documentation: You should NEVER delete your .profile and .bashrc files nor the .ssh or the .x2go directory, and you should ONLY fiddle with their contents if you at least somehow know what you are doing (or are willing to learn and have a backup-shell already open to test your new setup with a third shell...).

Disk Quota Exceeded errors OR Upload of file fails OR Can not create file / directory

→ Use the command “checkquota” on the command line to find out which filesystem is affected. Read the [documentation about the file systems provided by the cluster and how to use them](#) to understand which one should be used for what purpose. Then delete/move files to get below your quota limits and try again.



DO NOT simply delete files or directories in your HOME before understanding what they do, in particular DO NOT simply delete directories like .ssh, .x2go or .bashrc.

Have a look in other directories like .cache/sessions and directories that evidently look like temporary/scratch locations. You will need to analyze how you put these files here. The initial setup of a new account only needs a few kilobytes, and it is very easy to stay below a few megabytes in your HOME and run perfectly fine; so large amounts of data will have been put there by YOU, usually using conda or pip). To learn how to change the location Conda is installing to, read the corresponding section in [Conda tips](#). Also have a look at the next FAQ, “where actually are these files?”.

The command checkquota tells me that I am over my quota, but where actually are these files?

→ use the command `du -mad1 | sort -n` in the file system that checkquota tells you it sees the problem. You can quickly change to the file system using the environment variables we automatically set for you, `cd` moves to your \$HOME, `cd $BIGWORK` and `cd $PROJECT` to the other two major file systems.

Explanation: `du` = “check disk usage”, `-m` = “show megabyte units” (to enable numerical sorting afterwards), `-a` = “show all files, including those starting with a dot” (those are otherwise “hidden”, and some applications like e.g. Ansys use dot-directories e.g. for temporary files), `-d1` = “only show files up to a depth of one (this level)” (otherwise the command would descend down the directory tree). `|` = pass the output of `du` to `sort`, which numerically (`-n`) sorts what it gets.

I need more space on BIGWORK, please increase my quota!

→ Please state *how much* storage you need, *why* you need it, for *how much time* and, of course, the account name that needs the space. We will be happy to consider reasonable requests. If you later on notice that the first guess was not sufficient, simply contact us again.

The default quota is a compromise between the needs of individual users and the available disk space. The limits are set so about 90 % of our users should hopefully never encounter any difficulties. It also exists

1. to prevent job scripts or users from accidentally filling up the BIGWORK directory due to a mistake,
2. to ensure that people clean up their scratch files from time to time, since disk storage is a limited resource shared by all,
3. as a reminder that one needs to always make backup copies of important results and data to own storage, since it is neither possible to provide a backup for BIGWORK (due to the fast-changing nature of that file system and due to performance considerations), nor is there a guarantee that this file system will not become corrupted and loose data one day,
4. to reserve some space enabling us to consider requests for more space for those users who really need it for some period of time. This is also the reason why we want to know for what time you need the extended quota.

I need millions of files in BIGWORK for my project, please increase my inode quota!

→ No.

Unfortunately, we can't do that.

Lustre in the current configuration is good scaling large bandwidths, but very weak with many very small files, since that is something the metadata server (MDS) of the file system has to carry alone. This file system is just not suitable for many small files (much "classic" file system hardware using mechanical hard disk drives is severely bottlenecked when it comes to IOPs; we hope the situation will get better with the next filesystem in a few years, which may be comprised of SSDs that can handle orders of magnitude more IOPs). The result would be both slow performance and also decreased overall system stability for all, multiplied by the number of people who nowadays request really huge numbers of files "for LLM training". We are happy to see that the cluster *also* is quite able to carry these workloads. But strictly speaking anything that constantly needs many small files on storage media is not classic HPC any more.

A possible workaround is to pack all the small files into one larger dataset, either with a suitable format like HDF5, or using good old tar to create one large and possibly "striped" archive file on BIGWORK (see [File systems in the cluster](#), particularly the sections about \$TMPDIR and "Lustre file system and stripe count", how to do/optimize this). Then extract that file at the start of your job, ideally to a fast local scratch SSD drive directory like the one \$TMPDIR points to. Newer nodes in the cluster are usually equipped with such a drive, cf. our [computing hardware table](#) to find out which ones are. You'll probably need to include a `#SLURM -partition= ...` line within your job script to ensure you'll only land on suitable nodes. Nodes marked "NVME" in the table have the newest and usually fastest drives. "SSD" is also good; stay away from "HDD".

My job starts, but then suddenly aborts

→ Use the `checkquota` command to verify that you are not over your soft limit ("Quota"). If you are, check that the grace period has not yet expired. Issues are usually marked in red instead of a green line. In case the grace period (the time you may overstep your soft limit) is over, you'll see a simple dash "-", otherwise, the remaining time is displayed. The dash is only relevant if you *are* over quota. We recommend to also check the other quota-related FAQ.

I want to access some server somewhere on the internet from within my jobs on the compute nodes

→ By policy, compute nodes cannot access the internet outside the computing cluster. Exceptions need to belong to the LUH network. If you need such an exception, contact cluster support stating IP address, port number(s), protocol(s) and account name(s) that should be allowed to use the exception as well as a contact person, the reason and duration of the exception. See https://docs.cluster.uni-hannover.de/doku.php/resources/computing_hardware

I get kicked out of my connection when trying to load a larger model or to compute something

→ you probably try to run something directly on the login nodes, which are unsuitable for computations and anything larger.

To prevent people from doing this (and thus impede the work of others), there is a limit of 1800 cpu seconds on the login nodes, after which processes automatically get terminated by the system. 1800 cpu seconds should be enough for months of shell access, but will get used up quickly when you try to calculate something on multiple cpu cores. Read the docs how the cluster is intended to be used. As a first start, check out our OOD-Portal on <https://login.cluster.uni-hannover.de> and submit an interactive job. It is never a good idea to try anything more than small tests, editing or a (small) compilation on a login node, since they are shared between all users (simply try the command `w` on the shell prompt to see how many others are working on the same machine).

I get kicked out of my connection when trying to transfer large amounts of data

→ see last question directly above. Using `scp` or `rsync` to transfer files is the correct way, but the encryption used to secure transport may need an amount of cpu that can exceed the 1800 seconds. For that, we provide a node expressly dedicated for data transfer to/from the cluster, transfer.cluster.uni-hannover.de

When trying to connect to the cluster, I get: "Missing font family helvetica, OpenGL might have problems on this machine", or other strange graphics problems

→ You are probably using X2Go, and you did not check the "Full Installation" checkmark. Please reinstall X2Go and make sure all fonts are installed. Graphics problems can be quite nasty, since there are so many possible configurations. It often also depends on how you access the cluster (ssh/MobaXTerm, Cygwin/X, X2Go, Open OnDemand,...), from what kind of operating system (Linux, Mac, Windows...?) you start, and of course the software you want to use. Be sure to include that information, and be sure to tell us what you do step by step to give us a chance to possibly reproduce things. So please NOT something like "when I connect to the cluster, everything is broken", BUT "I connect as account nhxxxxx to the cluster using X2Go, of which I have a full installation as suggested in the documentation, and after I load this module and start software xyz, I get the following output (screenshot, log file...)". You'll need to help us help you.

Optimizing your work, "my job does not start, something's wrong with the cluster!"

My jobs do not run, even though some nodes appear to be free - there must be something wrong with the cluster!

Usually there isn't.

Scheduling is more complex than most people imagine. It's a mixture of rules, priorities, exceptions and available resources that is hit by a mélange of very different jobs. For some hints on that, consult our Slurm pages, chapter [How the scheduler works](#).

The short answer usually is: just because a particular node matching the requirements of your job appears to be free at a certain time, this does not automatically mean that the scheduler will pick or even consider *your* job as the next to run.

If you did not even check whether something is free and just want to complain about you waiting, please first check your position in the queue, and also consider whether you possibly requested a rare resource like a GPU or a combination of resources that can only be satisfied by few nodes at all. We will not prioritize people just because they are creating cases. We will also not start a deeper search for possible errors in the batch system when you do not even tell us the slightest details, such as job IDs, your username and the time when you looked at the system. We see many jobs that can not immediately run and even sometimes jobs that can not run at all, but only in few cases it's a problem with the system.

You should also consider that free nodes may first need to boot up even if your job can be considered immediately. We automatically power down nodes that are not used for more than 23 minutes (yes, that's intentionally a prime number...). That means if your job appears to be stuck waiting, it may just be waiting for a node to become ready (usually about 10 minutes).

If you and several of your colleagues experience long waiting times that can not be explained, open a

case, stating example jobIDs and *when* (the time and date) you looked at the system. State the commands used to diagnose the problem. Sometimes, there's really something that can be optimized, which may only recently have turned up because some other user submitted a new configuration of jobs. Or a node is stuck and that somehow affects other components in the cluster. Slurm is good, but not perfect, and sometimes the interplay of many factors in a complex cluster with many different partitions, users, reservations and priorities creates unexpected results. Do not be afraid to ask, we are in this together.

But we *really* only have a *very* limited amount of *very expensive* GPU cards. Just mentioning...

What can I do to get more jobs to run or to decrease wait times in the job queue?

- check whether you request an adequate wall time in your jobs (`-time=...`). The longer the time limit, the longer you will wait until the scheduler finds a window for your job (in general).
- if you are able to submit shorter jobs, you may fit into gaps between longer-running jobs.
- if you can submit jobs that need less than 12 hours, you may benefit from the so-called FCH clusters that we run for various institutes. They usually are reserved during office times (08.00-20.00 on work days), so they become available during the night adding a lot of extra resources to the cluster. The same goes for jobs requesting less than 60 hours during weekends.
- if you are unhappy with the performance of a job, try running it on a specific partition and adapt it to the nodes in that partition. Check how many cores there are on a socket (`lscpu` or `cat /proc/cpuinfo` on that node) or how many NUMA domains there are, respectively, and try to first “fill sockets”, then fill nodes, only thereafter request multiple nodes.
- if you are an expert or want to become one, consider checking the output of `lstopo` (becoming available e.g. after using the command `module load foss`). That shows you the hardware topology of a node, including cores, packages, cache levels, networking hardware etc. Try the command first on a login node to see what the output looks like. You'll need to be aware that we run all jobs submitted to SLURM in containers (“cgroups”), so `lstopo` will only show the parts you requested from the batch system, and if you only requested part of a node, that's what you will see.
- usually a job that runs on one single node will run faster than a job using the same amount of cpus distributed over several nodes due to less complexity in message passing. We see jobs that request several nodes, but only a few cpu cores on each node. That is usually both inefficient and also somewhat impolite to other users, who have optimized their jobs to requested one full node. Please refrain from just requesting 100 cpus “on any nodes, in any distribution” if possible. Consider to request full nodes instead — in order to avoid leaving only fragmented nodes for your colleagues.
- requesting fewer cpus per node while requesting more nodes may possibly get your job started quicker, but it may then need more time to complete in total than if you were to run a job on a single node, since data transfers between nodes (inter-node) are usually slower than those within one node (intra-node). So that's usually inefficient for all (also see last item).
- use `clusterinfo -v[v]` to find out what cpu architectures the nodes have. Depending on which cpu features your computation can use, you may or may not benefit from using a newer cpu (some keywords: AVX, cache size, the flags section of `lscpu`). Some jobs use about the same amount of time running on an older cpu, since they are mainly memory-bound or need to frequently access files — so why not use the shorter waiting times of an older partition that is less used? Some simple tests will easily tell you whether you benefit at all from a newer cpu.
- the command `scontrol show partition [<partitionname>]` will show you which

partitions are available in general and which nodes belong to a specific partition. `scontrol show node [<nodename>]` shows resources available on all nodes (warning: long) or on a specific node. Of course, you can also [have a look at our table](#).

- think about what bottlenecks your job has and how many cores will still result in a speed up. Painting a small room using three workers may speed up the process, but 1000 workers in a small room will just step on each others' toes. Check Amdahl's law.
- check whether the bottleneck could be I/O (or memory) instead of cpu. "A cluster is a tool to turn a cpu-bound problem into an I/O-bound problem". Check where your jobs take the most time, and know your workload.
- in case you need a GPU to run, check whether you accidentally request a specific type of GPU (like "v100"), when you just need "any" GPU. The more relaxed your constraints are, the easier it becomes for the scheduler to satisfy them, and since there are many types of GPUs in the cluster, specifying exactly one type will make you wait much longer.

I have access to a partition of the FCH service (nodes my institute has bought and which are reserved exclusively for agreed times), so my jobs should run. But I still have to wait. There really IS something wrong with the cluster!

→ first check whether you are working in the middle of the night. FCH reservations usually cover working days, a typical example would be from 08.00-20.00 Mo-Fr. That means that in the middle of the night or on weekends, you have no direct (absolute) priority. You still have a very high indirect priority because other users can only use your institute's partition when it's not exclusively reserved and thus during the week can not run jobs there that would take longer than 12 hours to complete. Since during daytime, your priority is absolute, and since you do not have that limit, you could, of course, have a job on your partition that runs well into the night, thus shortening the time available to others, against which they can do absolutely nothing. That means that your maximum waiting time will still usually just be a few hours (that is, if none of your colleagues with the same priority submit jobs to the same partition). BUT you may occasionally encounter the phenomenon that both your job and a (higher-priority, usually due to long waiting times) job not eligible to access your institute's reservation may have to wait until the current daily reservation has expired, IF you submit a job that requests a longer wall time than the reservation includes (if you submit an 8-day-job, you would not expect it to immediately start just because you have a reservation for, say, 12 hours, would you?). So your priority is only "absolute" as long as your jobs can complete within your reservation window, but it is NOT when your jobs do not fit that window any more - you just have a higher chance that the scheduler still decides your job should be preferred nevertheless immediately.

→ since the last SLURM upgrade (03/2026), you'll need to explicitly request - - `reservation=<reservationname>` in your jobs **for jobs that request a longer walltime than the duration of that reservation** in order to get them scheduled on your institute's partition. For the time being, that also means that you'll need to possibly fix jobs that you submitted in the web portal (Open On Demand, via your browser pointing to <https://login.cluster.uni-hannover.de>) on the command line (ssh to `login.cluster.uni-hannover.de`) with the command

```
scontrol update jobid=<yourjobid> --reservation=<your_fch_reservation_name>
```

To see your own jobs in the queue, you could use `squeue -me`. To find out which reservation(s) you can use, you could either guess the name (using your own institute's name, prefix used for nodes in your institute's partition), or use something like

```
scontrol show reservation | grep -A 10 "^ReservationName=" | egrep -B 5 "Users=.*$(id -u)|Accounts=.*$(id -g)"
```

(look for the name after "ReservationName=").

Since we can not determine which jobs you intend to have using your institute's reservation(s) and which ones you want to submit into the normal queue, and since the OOD interface does not yet offer a convenient mechanism to directly check which reservations you have access to, you'll need to adjust this yourself for the time being.

→ check whether there is a reservation for a scheduled lecture that prevents you from accessing the nodes. The command `scontrol show reservations` will help.

→ if you submit jobs that request longer walltimes than the exclusive reservations set for your institute, you are outside "absolute" priority (and you will need to explicitly state `reservation=<reservationname>` in your job script to get you job to run on your institute's partition). FCH reservations may be seen as blocking all "external" jobs requesting longer walltimes than the non-reserved time windows available in that partition, thus massively reducing your jobs' competition for priority. But you'll still be in the queue with "other" jobs that request less than 60 hours (weekend reservation window) or even less than 12 hours (night window). Usually the scheduler will massively prefer your jobs even if they request longer run times, but in some cases you will submit a 200-hour-job Friday afternoon that only starts sometime on Monday. Because the scheduler already has some higher-priority jobs of other users in the queue that it scheduled for the week end, your job was just submitted and thus has a low priority, and it also would continue to run outside your reservation window. In these cases the external higher-priority job(s) will wait together with your job in the queue, then start almost exactly after your daily reservation ends, run throughout the weekend, and thereafter your job will run. The same may sometimes happen during the week with short, less-than-12-hour jobs of other users. Your priority is absolute during your reservation windows, for jobs that fit into the windows. Test it by submitting a job that requests less than the remaining time until the closing of your daily reservation.

tl;dr: you must not expect just ANY job to start immediately just because you have access to an FCH partition. You can expect jobs to start immediately that fit into your reservation. But if you submit long-running jobs, they won't always start to run immediately just because a (shorter) reservation is still currently active, and in some cases resources may sit unused since none of the jobs can currently run until your reservation ends, due to priorities.

I get a message about the (external) email address I use to get status mails about my jobs

→ Sending lots of similar emails to external freemail providers leads to the uni-hannover.de Domain being blocked by those sites, which in turn leads to almost nobody external receiving mail from LUH accounts or machines any more. Therefore, we ask you to use your institute's email account (the one ending in @xxxx.uni-hannover.de), and we also ask you to NOT setup an automatic forward to a Google, GMX or similar external freemailer account.

My job stays in Cancelled or Completing state for a long time (days)

→ please let us know stating the usual details. Sometimes a node really crashes or gets stuck, and failure to cancel or complete a job over several hours or even days usually is a sign for a problem.

Problems with jobs or specific software requirements

What is the meaning of "Script is written in DOS/Windows text format"?

→ [transferring_data#converting_jobscripts_written_under_windows](#)

Something in this very complex and quite expensive software-suite (that is installed on the cluster and typically used by experts) does not work as I want it, I fail to mention any details and I expect YOU to solve MY problem NOW

→ we will invariably point you to our [How to get Support](#) page and request you do your own homework.

My jobs have crashed, and the job status says something about "out of memory"

→ Use the command

```
sacct -j <myjobid> -o JobID,NodeList,ReqMem,AveRSS,MaxRSS,MaxVMSize,State
```

to find out how much memory you requested in your job and how much memory your job tried to use. If the number under MaxRSS is near or even above ReqMem, you should request more memory in your job, plus a little reserve.

VM = virtual memory. That usually is the memory that is allocated (but not necessarily used).

RSS = resident set size, that is memory that is allocated AND at least somehow used — usually in pieces of 4 kB, which is the standard memory page on Linux/x64.

When requesting memory, try to always leave at least 4GB of memory to the operating system on *each* node your job will run on. You will also benefit if the system has memory for buffers etc. The [table of available resources](#) shows the physical memory, of which you'll need to subtract these 4GB.

Attention: the sacct command above only tells you about the system memory usage of your job. If you run into OOM on a graphics card, you'll still have to adapt your usage, but that will not show using sacct.

If some of your jobs survive despite demanding too much memory, it is due to the limits not being enforced every few microseconds, since that would cost performance as well. So be kind to your colleagues on the same system. Due to the fact that things run in real-time, they may very well be

the first victims at times if you occupy more than your share. So please take OOM seriously, such an event may even lead to a system crashing quite unexpectedly. See <https://lwn.net/Articles/104185/> (do not read if you are afraid of flying, and rest assured that things do *not* work like that in airplanes).

My jobs sometimes crash, depending on which partition they run on, and I see "illegal instruction messages" in my job logs

→ This may be due to you trying to run software that was compiled to use newer cpu features (like AVX512) than are available on the partition the job ran on — like e.g. the software was compiled for skylake, but the job ran on haswell. Use the command `clusterinfo -v` on a login node of the cluster to check the cpu generations available.

When compiling software yourself, keep in mind that we have several login nodes from various cpu generations. See [#build_software_from_source_code](#) [Modules & Application Software, Build Software from Source Code](#). The cpu architecture of the login node you compile on may automatically determine which platform your installer/make/... compiles for — some packages assume you want to install the software on the same architecture you want to run on, and this may of course fail if you use a newer login node to compile/install than the partition your jobs later run on. `hostname` shows you the name of the node you are on, `lscpu` shows you details about the cpus in place.

Use `#SBATCH --partition=` to only allow certain partitions, and use the table under [Hardware specifications of cluster compute nodes](#) to find out which cpu generation the major partitions offered run.

My jobs sometimes crash when they run on the enos-partition

The enos partition is the only partition that is equipped with OmniPath instead of InfiniBand as the network interface. An MPI application may have problems running there, but you can usually get things to work by adding the following line to your job script:

```
[[ $HOSTNAME =~ ^enos-.* ]] && export MPI_IB_PKEY=0x8001
```

My jobs do not start immediately, even when I know the resources should be available (e.g. using an FCH partition)

As a contribution to saving power, we use an energy saving option in Slurm that automatically shuts down nodes that would be idle for more than a few minutes. Slurm boots these nodes once they are needed. During boot, these nodes show the status "POWERING_UP" and jobs waiting for the nodes receive the state "CF/CONFIGURING".

When compiling, I get an error "unsupported option --add-needed"

That usually is a problem with using the "Gold" instead of the "BFD" linker (there's usually both flavours on a system). With gcc, the option to use is `-fuse-lld=bfd`.

Ansys/2025.1 crashes with "Unexpected error: [...] addins could not be loaded"

In case Ansys/2025.1 crashes with a message "Unexpected error: The following required addins could not be loaded: Ans.SceneGraphChart.SceneGraphAddin. The software will exit.", try `module load foss/2021b Mesa/.21.1.7` first before starting `runwb2`. The problem seems to be specific to Ansys/2025.1, which seems to have incompatibilities with the default OpenGL environment. For X2GO, a workaround is in place, but in the web portal ("OOD", OpenOnDemand, the thing you log into by pointing your browser to <https://login.cluster.uni-hannover.de>), it may be necessary to manually loading Mesa/.21.1.7. Check for the module using `module --show-hidden spider Mesa` in case you are interested.

Gaussian jobs crash seemingly randomly, without producing reasonable error messages

Please add the following line to your job script:

```
#SBATCH -C CPU_MNF:intel
```

The Gaussian version we currently have does not run on AMD cpus.

A software package I use requires a newer libc (or another software component central to the running operating system) , but when I want to install it as root, I only get "you are not in sudoers". Please give me root acces/Please install that library!

→ No.

We have a highly modular software stack, and it is absolutely impossible to accommodate particular needs for particular versions of operating system libraries. To avoid everybody constantly changing the systems, you can only change things in your account, but not for the whole system. Installing a different version of a very common library would break quite a lot of things, so that is completely out of the question.

You usually can use a container to get your software to run. For details, please refer to [Modules & Application Software, section Apptainer Containers](#).

From:

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

Permanent link:

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

Last update: **2026/06/03 10:43**

