

Frequently Asked Questions

- [How do I apply for an account?](#)
- [Login issues](#)
 - [Error: "Access denied because you are not enrolled"](#)
 - [Connection closed after Duo success](#)
- [I want to use <insert application name> on HAL! Can you install it?](#)
- [How do I share my files and collaborate with other HAL users?](#)
- [How much data can I have on HAL?](#)
- [My job is not running!](#)
 - [Reason: \(Priority\)](#)
 - [Reason: \(ReqNodeNotAvail\)](#)
 - [Reason: \(Resources\)](#)
 - [Reason: \(AssocGrpGRES\)](#)
 - [Reason: \(Reservation\)](#)
- [I want to install Tensorflow/PyTorch/Caffe but I can't install one of its dependencies!](#)
- [I'm a new user and I am trying to access the HAL OnDemand with the link on the Wiki but am getting a "home directory not found" error!](#)
- [I'm a new user and I'm getting "sbatch: error: Batch job submission failed: Invalid account or account/partition combination specified" from s\(w\)batch/s\(w\)run/HAL OnDemand](#)

How do I apply for an account?

First, please fill out this [application form](#). The form asks why you need an account and what resources you will be using.

Next, you will receive an email with the link that you must follow to request the actual user account on HAL. You must follow this link to create an NCSA user ID (if you do not have one) and to request membership in the corresponding LDAP group. Check your emails and follow the instructions.

Login issues

Error: "Access denied because you are not enrolled"

As of May 2020, two-factor authentication via Duo is now a requirement for access to HAL. You are receiving this error message because you are not enrolled in NCSA Duo. Go to <https://go.ncsa.illinois.edu/2fa> and follow the instructions. **Note that the campus Duo is a separate system and you will not be able to use it to access NCSA systems.**

Connection closed after Duo success

This indicates that you have Duo setup but are not currently authorized to use HAL. If you have never used HAL, please follow the instruction in the "How do I apply for an account?" section of this page to request access.

However, if you are reading this, it's more likely that you had access at some point in the past, but were revoked later due to inactivity or some other reasons. To restore access, contact us at help+isl@ncsa.illinois.edu and we will look at your case.

I want to use <insert application name> on HAL! Can you install it?

Firstly, please check if the application you want supports ppc64le architecture. HAL uses IBM's POWER9 architecture in order to achieve improved multi-GPU performance, but this comes at the cost that common x86 software may not work on HAL. If an application states it supports ppc64le, it still may not work on HAL, because the older POWER8 architecture uses the same ppc64le identifier but is not 100% compatible. We are happy to help you test the application if this is the case.

Once you identify a version of the application that supports POWER9, see the following guidelines for installation:

- Is the application free and open source?
 - Closed-source applications may have license terms that apply to research institutions. Note that "classroom" or "student" licenses are typically invalid for multi-user clusters like HAL.
 - If you have a license for your personal use, you can install it in your home directory. We will not approve a request to install such an application system-wide.
 - Closed-source applications also need to have stated official support for IBM POWER9. Adding support for a new architecture is a complex project that can take more than a year, especially without the support of an open-source community.
 - Not all open-source applications can run on all architectures. If an open-source application doesn't have official support for IBM POWER9, check with the developers to see if it has any dependencies that don't work on IBM POWER9. Sometimes, the application itself is architecture-independent, but some of its dependencies are not, so it still won't work (for example, some machine learning framework that uses Intel-specific machine code to accelerate computation). You can try to install the application in your home directory and ask for help by submitting a ticket to help+isl@ncsa.illinois.edu.
- Do you think it will be useful for all users?
 - If it's an application with limited scope that is specifically required for just your project, consider installing it in your home directory. If it's a Python package, you can clone one of the system-wide Anaconda environments and install the package in the cloned environment.
 - If you think the application can be utilized by all users, submit a ticket to help+isl@ncsa.illinois.edu, and we will review the request. This usually takes one to a few business days and we may deny the request if we decide it should not be installed system-wide.

How do I share my files and collaborate with other HAL users?

Contact us through Slack or help+isl@ncsa.illinois.edu and provide a list of people you want to share with (with their usernames on HAL, preferably), and a name you want to call the shared folder. We will create a folder under /projects for your use.

How much data can I have on HAL?

Data should not be kept on HAL for extended periods of time. HAL is not backed up in any way and is supported on an as-available basis. Each user is limited to 2TB and 20,000,000 files/directories (whichever limit is reached first).

My job is not running!

Use the following command to get a list of your jobs (replace user_name with your username):

```
squeue -u user_name
```

The right-most column will contain a reason for each of the pending jobs. Refer to the list below for detailed explanations.

Reason: (Priority)

There is at least one pending job with a higher priority than this job. The priority for a job depends on a couple of factors, the biggest of which is recent usage. Most likely you are seeing this reason after running some combination of a large number of jobs, jobs using a large amount of resources, or jobs that run for a long time. The recent usage factor slowly decays in a two week period, which means any usage prior to two weeks before the job was submitted will not impact the priority of the job. You can check your recent usage here: <https://go.illinois.edu/halfairshare>

Jobs that are pending for this reason may remain pending for a long time if the recent usage factor has reduced your priority below most of the active users. If there is a sufficient difference between someone's recent usage and that of yours, and the difference in the recent usage factor is large enough to exceed the waiting time factor, their job may receive a higher priority and therefore run before your job, even if it is submitted after your job.

Reason: (ReqNodeNotAvail)

Some of the nodes specifically requested by the job is not available, which can mean the node is running jobs with a higher priority, reserved in a reservation, manually drained by an administrator for maintenance, or unavailable due to some issues. This job will run when all the requested nodes become available.

Reason: (Resources)

This job is at the front of the queue, but there are not enough resources for it to start running. This job will start running as soon as enough resources become available. The priority calculation favors large jobs, so when resources gradually become available, smaller jobs with similar recent usage factor won't run before this job and take away the available resource. Note that if someone has much lower recent usage than you do, their jobs can still run before your job, because the bonus from their recent usage factor can exceed the bonus from your job size factor.

Reason: (AssocGrpGRES)

This means you have reached the limit of resources that can be allocated to one user at any given time. There are three limits in place: a maximum of 5 running jobs; a maximum of 5 nodes running jobs; and a maximum of 16 GPUs running jobs. This job will run as soon as some of your running jobs finish and free up the resources.

Reason: (Reservation)

This job is submitted to an inactive reservation. If the reservation is in the future, it will run when the reservation starts. If the reservation has ended, it will be stuck in the queue forever until it's deleted.

I want to install Tensorflow/PyTorch/Caffe but I can't install one of its dependencies!

STOP HERE. HAL uses the specialized IBM POWER9 architecture, which means it sacrifices compatibility with traditional x86_64 software. Many things need to be re-compiled for it and this is usually a very tedious process. We have common Machine Learning frameworks already installed and ready to use, refer to [Getting started with WMLCE \(former PowerAI\)](#) for more details. If you need a newer version than we are currently providing, we will usually wait until the next version of WMLCE is released. We can help you compile a newer version if there is a very specific need for it. In that case, contact us via Slack or send a ticket to help+isl@ncsa.illinois.edu and someone will respond to your request.

I'm a new user and I am trying to access the HAL OnDemand with the link on the Wiki but am getting a "home directory not found" error!

You need to access HAL system via "ssh hal.ncsa.illinois.edu" first to initialize your home folder. After your home folder created, you can access HAL-OnDemand without any problem.

I'm a new user and I'm getting "sbatch: error: Batch job submission failed: Invalid account or account/partition combination specified" from s(w)batch/s(w)run/HAL OnDemand

Your account has not been properly initialized. Try logging in and out of hal-login2.ncsa.illinois.edu via SSH a few times. If it isn't working, contact an admin on Slack or send a ticket to help+isl@ncsa.illinois.edu.