Transitioning to NERSC

Overview

Resources at NERSC

- We have an allocation with…
  - ~10 M core hours (A Perlmutter core is ~2.25 more performant than a WC core - YMMV)
  - 450 NERSC GPU hours
  - 20 TB of storage for staging

- Big login nodes
- Jupyter notebook interface
- Lots of good documentation and support
- Perlmutter has been running very well – note monthly down day

Please fill out form if you want to use NERSC

- https://forms.office.com/g/91rYnwGcrs

- Allows us to understand needs
- We can help you with getting started

Wilson Cluster -> NERSC Transition

* Required

1. What is your name? *

Adam Lyon

2. What is your work e-mail address? *

lyon@fnal.gov

3. What is your ORCID id if you have one?

0000-0003-2563-6235

4. Please enter a short description of your project. If you are the PI, please list collaborators who will also be asking for a NERSC account. If you are one of those collaborators, then refer to the project title and PI. *

To take over the WORLD! Bwa ha ha ha!

I am the PI. I have no collaborators.

5. How many million Wilson Cluster CPU hours were you planning to use? *

1

6. Do you need to use GPUs for your project? *

- ( ) No GPUs
- (•) Single GPU
- ( ) Multiple GPUs

7. How much storage do you need at NERSC (TB)? *

0.5

8. How much data do you need to move to NERSC (remember that NERSC storage is meant for staging) (TB)? *

0.05

9. What is your data management plan for your data at NERSC and elsewhere? E.g. Where do you plan to stage in/out from? How long will data remain resident at NERSC? *

Output of my jobs will be sent to NERSC scratch storage and then moved 10 seconds later to my Microsoft OneDrive.

10. Do you already have a NERSC account? *

- (•) Yes
- ( ) No

11. If you do not already have a NERSC account, you may now proceed to https://docs.nersc.gov/getting-started/ and request one. If you do have a NERSC account, then make a request to join project m4599.

You agree to respond to NERSC and Fermilab requests for updates on the progress of your project and to abide by all computing policies at Fermilab and NERSC. *

- (•) I agree

Submit

NERSC Accounts

- We have a set of "group" accounts at NERSC with large allocations which are intended to be used by the HEP Community
  - Each has a "project" number associated with it.

- When applying for an account at NERSC you need to specify:
  - Project Number: m4599
  - Project P.I.: Andrew Norman (anorman)
  - This kicks off an approval process
    - It is normally fast (~ a few days), but has gotten stuck and taken weeks in some cases
  - When approved you are added to the project
    - We assign you a quota of CPU/GPU/Storage
      (so if you need more than the default amounts let me know)

Logging in to NERSC

- Logging in to NERSC is done via SSH w/ a 2-factor auth code
  - See instruction: (https://docs.nersc.gov/getting-started/)
  - You need to setup your MFA, and then from there get the sshproxy script from NERSC to ease repeated login

```
ssh <user>@perlmutter.nersc.gov
```

- Storage is located on the "community" file system (CFS) and is vissible in your login shell and via a Globus endpoint
- Submission is a standard slurm batch queue
  - You will want to specify your project number when submitting
  - There are separate queues for CPU and GPU resources
  - See documentation for more details on NERSC specific slurm options

Data Transfer

- NERSC uses "data transfer nodes" for large data ingest/egress
  - It is possible to transfer data directly from compute nodes, but it is slower and will bottleneck (i.e. it ends up routing through a DTN)
- For most data transfers use Globus.
  - https://www.globusonline.org/
  - You will want to install the globus personal endpoint on your local machine (if you want to transfer files from your laptop)
  - The perlmutter endpoint is:
    - "NERSC Perlmutter" and comes up from the search bar in Globus Online
    - Scheduling transfers is easy

Shared Resources

- There is project space for m4599
  - If you need to use a large mount of CFS project space let us know
  - We maintain a [poorly] organized directory structure for the different experiments

- CPU/GPU pool is shared across the project. If you plan to use a LARGE amount please let us know so we can adjust accordingly.