

How to Submit a Good Help Ticket


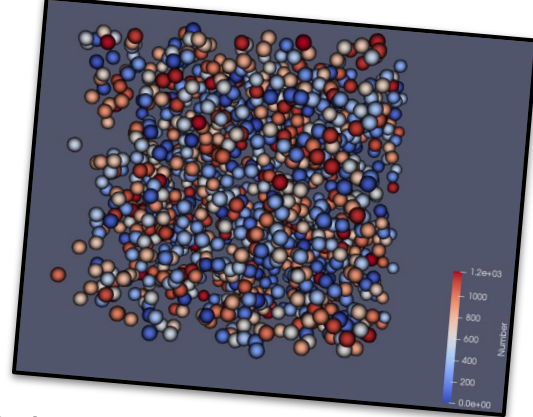


Erik Palmer
Software Integration Engineer
Programming Environments and Models
April 18, 2024

Who am I?



- Consultant since June 2022
- First came to Berkeley Lab in 2017 as a summer intern with NSF-MSGI – worked with Center for Computational Science and Engineering on AMReX applications
- Grew up in Berkeley, but now work remote from Michigan



What we will cover: all the questions about tickets you were too afraid to ask!

And...

- Ticket MYTHS vs FACTS
- Common problems and troubleshooting tips
- DOs and DON'Ts for filing a good ticket

For example:

Do: ask questions during this presentation


Don't: assume you are the only person who needs help!

But First! (Not so) Trivial Trivia!



About how many consultants does NERSC have?

- A. 10
- B. 20
- C. 40
- D. 80



About how many consultants does NERSC have?


A. 10

B. 20

C. 40

D. 80

Answer: C. About 40 (~37). NERSC as a whole has about 135 staff.



How long should it take a consultant to reply to your question?

- A. 10 minutes
- B. 2 hours
- C. 4 hours
- D. 8 hours



How long should it take a consultant to reply to your question?

- A. 10 minutes
- B. 2 hours
- C. 4 hours**
- D. 8 hours

Answer: C. Note these are business hours, not regular hours. Business hours are 8am - 5pm Pacific Time!



Ticket MYTHS vs. FACTS

MYTH: NERSC consultants do NOT like
answering tickets

MYTH: NERSC consultants do NOT like answering tickets

FACT: Many of us work at NERSC because we like helping people. It can very satisfying to help someone solve a technical problem!

MYTH: NERSC consultants do NOT like
answering tickets

FACT: Many of us work at NERSC because we
like helping people. It can very satisfying to help
someone solve a technical problem!

Your Problems > My Problems

MYTH: I should NOT submit a ticket because my questions/problems are “stupid”. Only experts should submit tickets.

MYTH: I should NOT submit a ticket because my questions/problems are “stupid”. Only experts should submit tickets.

FACT: All users can and should submit tickets.
Beginner questions are very welcome!

MYTH: I should NOT submit a ticket because my questions/problems are “stupid”. Only experts should submit tickets.

FACT: All users can and should submit tickets. Beginner questions are very welcome!

“My job is too easy” –No One Ever

MYTH: I need to spend days stuck on a problem before I file a ticket.

MYTH: I need to spend days stuck on a problem before I file a ticket

FACT: Sometimes the answer might be easy or something we've seen before. Go ahead and submit your question after you've done some troubleshooting, NOT after you've been stuck for days.

MYTH: I need to attach all relevant files to the ticket in ServiceNow.

MYTH: I need to attach all relevant files to the ticket in ServiceNow.

FACT: All NERSC consultants have the power to impersonate NERSC users on our system. If you provide the full paths, we can find your files to look at if needed.

Troubleshooting Tips

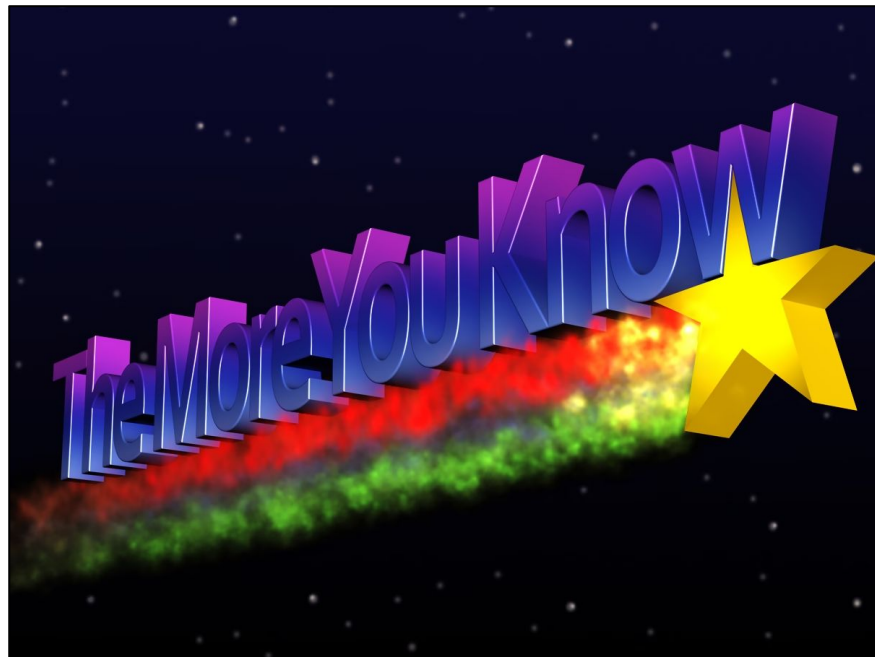
When you submit a ticket, you want to resolve the problem ASAP so you can get back to your work.

We also want to help you quickly and efficiently!

If you can do a little work up front to explain your problem and provide the relevant info, most likely we can solve it faster.

Before you file a ticket, there a few basic things you can check:

- [NERSC Live Status](#) shows current system status (down, degraded, etc)
- If you're on the NUG Slack, this is a good place to check to see if others are reporting similar issues
- [showquota](#)- are you over quota?



Before you file a ticket, there a few basic things you can check:

- [Iris](#)- is your account current? Do you still have compute hours?
- Dotfiles - is there anything in your `.bashrc`, `.bash_profile` files that you might have forgotten about or that needs to be updated? You'd be surprised how often these settings cause problems for users.
- Try the search bar in our [docs](#)
- If none of these help, it's time to submit a ticket!



DOs and DON'Ts for submitting a good ticket

DO: be specific!

Hard to troubleshoot

- My code is slow
- My job won't start
- Perlmutter is broken



It's always helpful to include the actual error message, even if it's long!

Better to troubleshoot

- JobID 123456 was 3x slower than JobID 234566.
- The jobscript located at `$HOME/submit_job.sh` works on Cori but not on Perlmutter. This was the error message:...
- Running python `$SCRATCH/test_cori.py` crashes with this error message.

DON'T: just paste your error message without any context!

- We are pretty good at computers, but without any context, it can be really hard to look at an error message and understand what went wrong.
- We'll have to ask you to provide more information, which just slows the process down
- **Just like school, show your work!**



DO: show your work!

- Imagine you are looking at your problem for the first time
- What are all the steps a NERSC consultant needs to reproduce the error from a fresh login?
- What troubleshooting steps have you tried?
- If you are using a software package, where is the repo/documentation?
- If a job failed, what was the jobid?
- [More tips in our docs](#)




DON'T: send screenshots!

As consultants we like to be able to copy and paste the paths in your error message to look at your files, to search stackoverflow, etc.

If you give us a screenshot we can't copy/paste anything, which just slows us down
This also applies to pictures of your screen you took with your phone

DO: Please paste your error message in plain text, even if it seems really long!



Now you try!





Update & Exit

Save & Stay

Resolve Incident

Responses



* Title

"Perlmutter" is broken



Notes

Event Details

Related Records

Staff work notes
(NERSC private)

Staff work notes (NERSC private)



Additional
comments
(Customer visible)

Additional comments (Customer visible)

Post

Guest
Comments

Hi, I can't login to perlmutter! It is broken. Please fix it.



Update & Exit

Save & Stay

Resolve Incident

Responses



* Title "Perlmutter" is broken



Notes

Event Details

Related Records

Staff work notes
(NERSC private)

Staff work notes (NERSC private)



Additional
comments
(Customer visible)

Additional comments (Customer visible)

Post

Guest
Comments

Hi, I can't login to perlmutter! It is broken. Please fix it.

What information might you need to troubleshoot this? What could this user have given us to make this an easier ticket to solve?

If the User had given this instead....

Guest
Comments

Hi, I tried to login to Perlmutter using Jupyter but I am seeing a “disk quota full” error. How can I address this so I can login? Thank you!

The consultant could have noticed the problem quicker!

Consulting
Comments

Thanks for your question and for providing the error message!

This means that your \$HOME directory is likely full. You
-Consultant

Laurie's Top Tip: Check your dotfiles regularly and keep them updated!





* Title

Cupy doesn't work on Perlmutter

Guest
Comments

Hello, I am trying to run a python program and I get an error when I import cupy. I think cupy is broken on Perlmutter - can you please reinstall it. Thank you!

Error message:

```
stephey@nid004603:~> python -c "import cupy as cp;cp.arange(10)"
```

```
Traceback (most recent call last):
```

```
File "<string>", line 1, in <module>
```

```
File "/global/common/software/das/stephey/conda/conda_envs/cupy/lib/python3.10/site-packages/cupy/_creation/ranges.py", line 58, in arange
```

```
    ret = cupy.empty((size,), dtype=dtype)
```

```
File "/global/common/software/das/stephey/conda/conda_envs/cupy/lib/python3.10/site-packages/cupy/_creation/basic.py", line 22, in empty
```

```
    return cupy.ndarray(shape, dtype, order=order)
```

```
File "cupy/_core/core.pyx", line 136, in cupy._core.core.ndarray.__new__
```

```
File "cupy/_core/core.pyx", line 224, in cupy._core.core._ndarray_base._init
```

```
File "cupy/cuda/memory.pyx", line 742, in cupy.cuda.memory.alloc
```

```
File "cupy/cuda/memory.pyx", line 1419, in cupy.cuda.memory.MemoryPool.malloc
```

```
File "cupy/cuda/memory.pyx", line 1439, in cupy.cuda.memory.MemoryPool.malloc
```

```
File "cupy/cuda/device.pyx", line 47, in cupy.cuda.device.get_device_id
```

```
File "cupy_backends/cuda/api/runtime.pyx", line 178, in cupy_backends.cuda.api.runtime.getDevice
```

```
File "cupy_backends/cuda/api/runtime.pyx", line 143, in cupy_backends.cuda.api.runtime.check_status
```

```
cupy_backends.cuda.api.runtime.CUDARuntimeError: cudaErrorNoDevice: no CUDA-capable device is detected
```

```
stephey@nid004603:~>
```



* Title

Cupy doesn't work on Perlmutter

Guest
Comments

Hello, I am trying to run a python program and I get an error when I import cupy. I think cupy is broken on Perlmutter - can you please reinstall it. Thank you!

Error message:

```
stephey@nid004603:~> python -c "import cupy as cp;cp.arange(10)"
```

```
Traceback (most recent call last):
```

```
File "<string>", line 1, in <module>
```

```
File "/global/common/software/das/stephey/conda/conda_envs/cupy/lib/python3.10/site-packages/cupy/_creation/ranges.py", line 58, in arange
```

```
    ret = cupy.empty((size,), dtype=dtype)
```

```
File "/global/common/software/das/stephey/conda/conda_envs/cupy/lib/python3.10/site-packages/cupy/_creation/basic.py", line 22, in empty
```

```
    return cupy.ndarray(shape, dtype, order=order)
```

```
File "cupy/_core/core.pyx", line 136, in cupy._core.core.ndarray.__new__
```

```
File "cupy/_core/core.pyx", line 224, in cupy._core.core._ndarray_base._init
```

```
File "cupy/cuda/memory.pyx", line 742, in cupy.cuda.memory.alloc
```

```
File "cupy/cuda/memory.pyx", line 1419, in cupy.cuda.memory.MemoryPool.malloc
```

```
File "cupy/cuda/memory.pyx", line 1439, in cupy.cuda.memory.MemoryPool.malloc
```

```
File "cupy/cuda/device.pyx", line 47, in cupy.cuda.device.get_device_id
```

```
File "cupy_backends/cuda/api/runtime.pyx", line 178, in cupy_backends.cuda.api.runtime.getDevice
```

```
File "cupy_backends/cuda/api/runtime.pyx", line 143, in cupy_backends.cuda.api.runtime.check_status
```

```
cupy_backends.cuda.api.runtime.CUDARuntimeError: cudaErrorNoDevice: no CUDA-capable device is detected
```

```
stephey@nid004603:~>
```

What else could this user have included in this ticket to make it quicker to troubleshoot?

If the User had given this instead....

Guest
Comments

Hello, I am trying to run a python program and I get an error when I import cupy. I think cupy is broken on Perlmutter - can you please reinstall it. Thank you!

Error message:

...

Slurm Command:

```
salloc -N 1 -t 30 -C cpu -A m000 -q interactive
```

The consultant could have noticed the problem quicker!

Consulting
Comments

Thanks for your question and for providing your slurm command! Your command requests a cpu node, which does not have CUDA devices for cupy to talk to.

Please use:

```
salloc -N 1 -t 30 -C gpu -A nstaff -q interactive
```

This will get you a gpu-accelerated node, with CUDA devices.

Let me know if you have more questions!

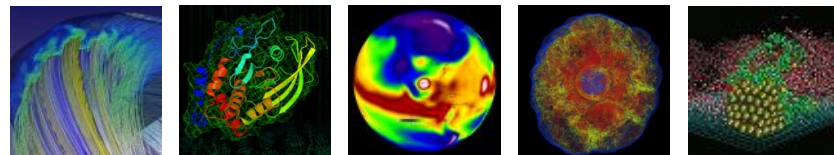
-Consultant

TL;DR

- ❖ Our job is to help you be productive on NERSC systems!
- ❖ If you get stuck and basic troubleshooting doesn't help, please submit a ticket
 - Beginner questions welcome
 - Don't spend days stuck on something- it might be an easy fix
- ❖ When you submit a ticket:
 - Be specific about the problem
 - Provide the error message and give context
 - Show your work and provide all steps to reproduce (ie. package information, slurm commands etc.)
 - Provide plain text instead of screenshots
- ❖ If you can give us all the information up front, we can most likely solve your problem more quickly!

2023-2024 NERSC EARLY CAREER WINNER TALK:

Building the Largest 3d Map of the Universe with DESI and NERSC



Anthony Kremin, PhD
Project Scientist
Lawrence Berkeley National
Laboratory (LBNL)

Early Career Award Winner - Innovative Use of High Performance Computing



- Dr. Anthony Kremin is a Project Scientist at Lawrence Berkeley National Laboratory working on the Dark Energy Spectroscopic Instrument.
- Before arriving at Berkeley Lab as a Postdoctoral Researcher, he received his Ph.D. in Physics from the University of Michigan, Ann Arbor.
- For the past four years, he has been working with DESI to develop and run the spectroscopic pipeline used to manage, process, and aggregate the raw spectroscopic data into meaningful cosmological catalogs using NERSC computing resources.
- Honored for developing and implementing a pioneering approach to enable data processing from the Dark Energy Spectroscopic Instrument (DESI), helping fulfill DESI's mission to construct the world's largest 3D map of the universe and allow fundamental tests of cosmological physics.