

# JupyterHub at NCAR

and related things

Davide Del Vento

# About the speaker

- Davide Del Vento, Ph.D.  
helping users of NCAR's HPC  
infrastructure since 2008
- It's fine to interrupt with questions in  
context of what I am discussing

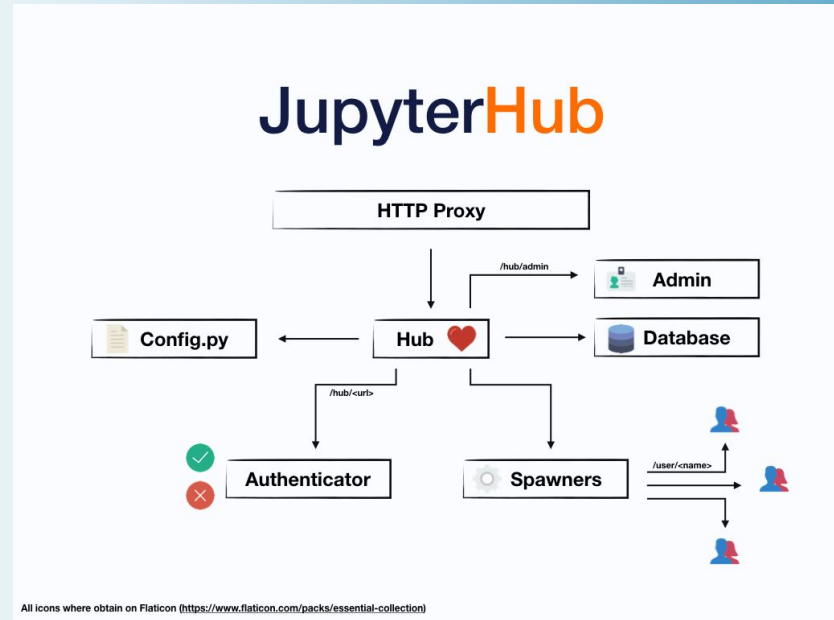


<https://staff.ucar.edu/users/ddvento>

<https://www.linkedin.com/in/delvento/>

# What is JupyterHub

- “The” way to serve Jupyter Notebooks to multiple users



# What is a Jupyter Notebook


- A single-user web app to create documents containing arbitrary, user-provided live code, equations, visualizations and narrative text
- Uses include: data cleaning and transformation, numerical simulation, statistical modeling, data visualization, machine learning, and much more

# How this differs from old one


- The Jupyter (formerly ipython) Notebook has been available at NCAR since at least early 2015... and it still is...
- But requires a clumsy setup with double-login, ssh-tunneling, etc.
- Yet, that is more robust and so still supported

Doc: <https://www2.cisl.ucar.edu/resources/jupyter-and-ipython>

### Number of Unique JupyterHub Users per Week



### Number of JupyterHub Sessions per Week



Computer  
CISL







# Useful documentation

Python - NCAR Package Library

2. Restore the alias by re-running the `ncar` module loader script.

### Creating your own clone of the NCAR Package Library

Creating a personal clone of the package library in your GLADE space is useful if you want to add or update packages, or develop Python code while using NCAR-provided packages. Simply specify the clone option (-c) to `ncar nvlib`, select a

JupyterHub at NCAR

UCAR NCAR

Closures/Emergencies Locations/Directions Find People

Support Resources About Us

NCAR is sponsored by National Science Foundation

STATUS NEWSROOM EVENTS

Google Custom Search Search

Conda environments on JupyterHub

www2.cisl.ucar.edu/resources/conda-environments

UCAR NCAR

Closures/Emergencies Locations/Directions Find People

Software Data Portals User Support Resources About Us

NCAR UCAR

CISL Computational & Information Systems Lab

NCAR is sponsored by National Science Foundation

DAILY BULLETIN RESOURCE STATUS NEWSROOM EVENTS

Google Custom Search Search

HOME » CONDA ENVIRONMENTS ON JUPYTERHUB

### CONDA ENVIRONMENTS ON JUPYTERHUB

Conda is not installed on the Cheyenne or Casper clusters, but you can install Miniconda in your own user space by following the instructions [here](#).

Miniconda is a lightweight version of the Anaconda distribution that contains the Conda package manager. After installing it, you can create new virtual Conda environments and install packages in them. See this [Conda documentation](#) for information on how to create environments and install packages.

Install the `ipykernel` package in Conda environments that you create. The [NCAR JupyterHub](#) will then detect and display your environments in the JupyterLab Launcher with names like "Python [conda:env myenv]."

To activate an environment and install the package, follow this example, in which "myenv" is the name of your Conda environment:

ncar\_pynv -tag myenv, you will see the clone has been detected.

HOME » SOFTWARE FOR HPC USERS » SOFTWARE LIBRARIES AND MODULES » JUPYTERHUB AT NCAR

### JUPYTERHUB AT NCAR

[Getting started](#) | [Python environments and kernels](#) | [Related documentation](#)

The JupyterHub deployment that CISL manages allows "push-button" access to NCAR's Cheyenne supercomputing resource and the Casper cluster of nodes used for data analysis and visualization, machine learning, and deep learning.

It gives users the ability to create, save, and share Jupyter Notebooks through the JupyterLab interface and to run interactive, web-based analysis, visualization and compute jobs on Cheyenne and Casper. JupyterHub is an alternative to X11 access for interacting with those resources to run jobs as well as for using web-based interactive shell functionality without the need to install or use software such as SSH or PuTTY.

**NOTE**

This JupyterHub deployment is a test system that is subject to more planned and unplanned downtimes than other resources in the NCAR/CISL high-performance computing ecosystem, such as login nodes. Responding to outages and providing updates to JupyterHub may be delayed while CISL staff support our full production HPC systems.

While JupyterHub provides added value and convenience, users are advised to be familiar with our [Jupyter](#) and [iPython](#) documentation as an alternative for continuing their work with HPC resources when JupyterHub is not available.

### Creating Jupyter kernels from NPL versions

If you wish to use the NCAR Package Library inside of other Jupyter instances on Cheyenne (for example, inside a personal NPL clone, conda environment, or a JupyterHub instance), you can run the following command to produce a user-space kernel from an available NPL version:

- Jupyter and iPython
- JupyterHub at NCAR
- Python - NCAR Package Library
- Math Kernel Library (MKL)
- NCAR Classic Libraries for Geophysics

### Getting started

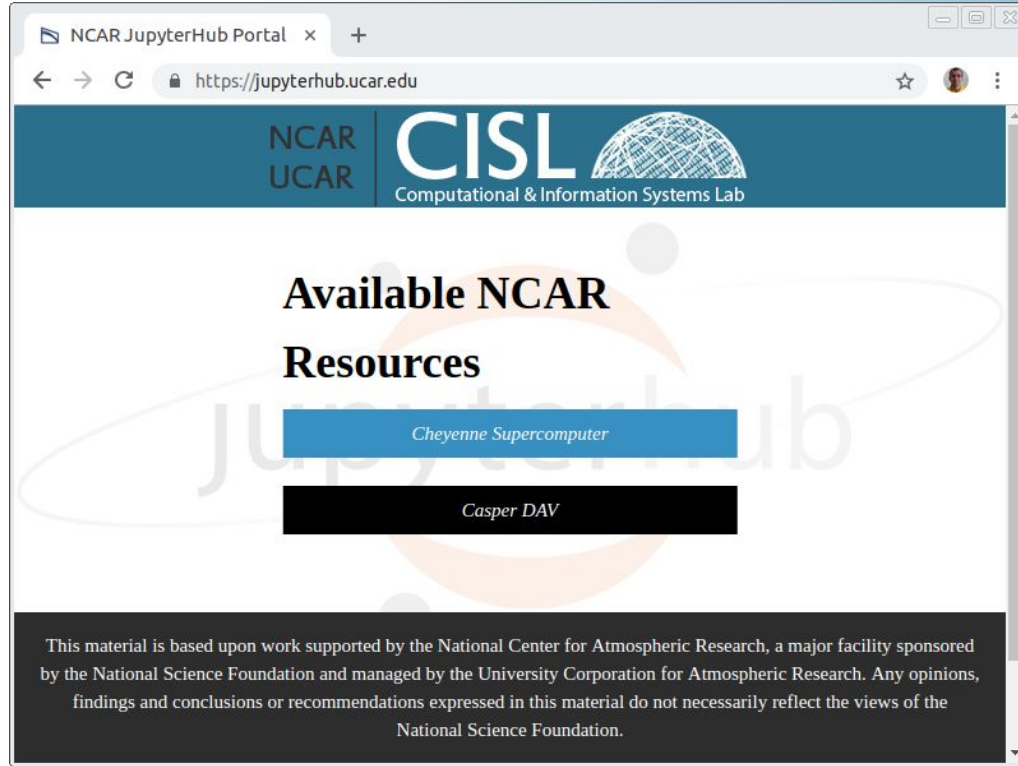
Use your web browser to go to [jupyterhub.ucar.edu](http://jupyterhub.ucar.edu). Chrome and Firefox are

# Useful documentation

- <https://www2.cisl.ucar.edu/resources/jupyterhub-ncar>
- <https://www2.cisl.ucar.edu/resources/python-%E2%80%93-ncar-package-library#clone>
- <https://www2.cisl.ucar.edu/resources/conda-environments>
- <https://www2.cisl.ucar.edu/resources/jupyter-and-ipython>
- Stay tuned: more to come....



# JupyterHub @ NCAR



NCAR UCAR

CISL  
Computational & Information Systems Lab

## Available NCAR Resources

*Cheyenne Supercomputer*

*Casper DAV*

This material is based upon work supported by the National Center for Atmospheric Research, a major facility sponsored by the National Science Foundation and managed by the University Corporation for Atmospheric Research. Any opinions, findings and conclusions or recommendations expressed in this material do not necessarily reflect the views of the National Science Foundation.

# Where to go from here

- Documentation provides additional details
- Pangeo project provides specific information e.g. for parallel analysis
- Future talk will cover kernel customization

# Future work

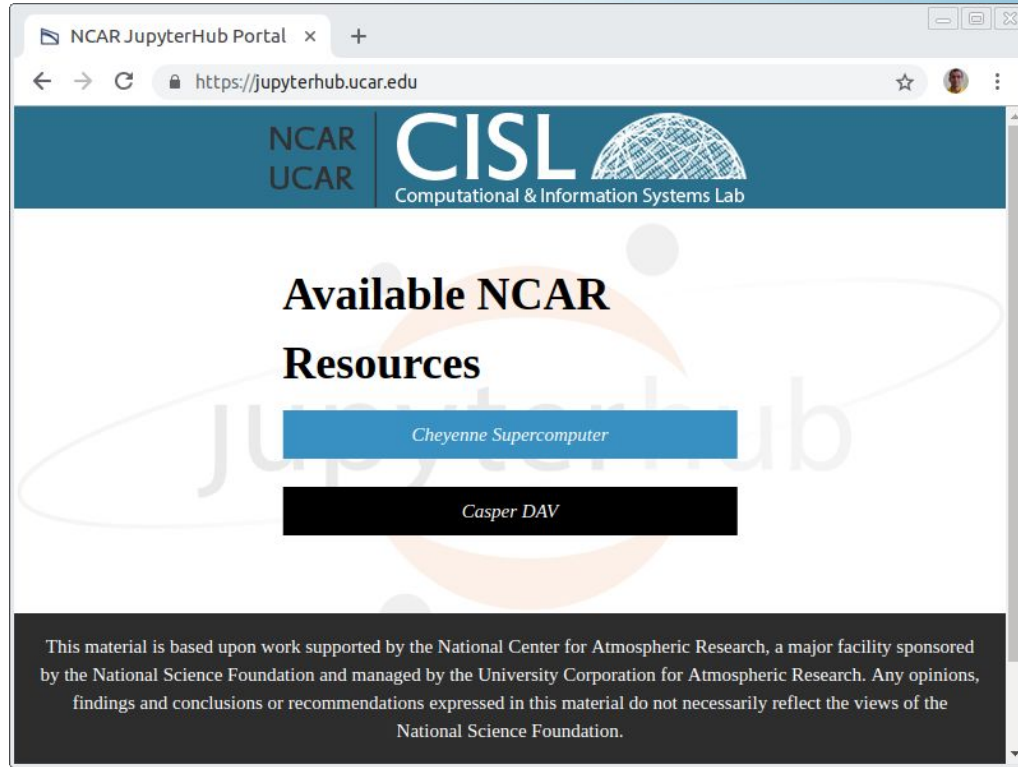
- Ability to load modules before kernel startup
- Improve user-friendliness of the login process (e.g. kernel logos overlays)
- More documentation and training (let us know what you need!!)

# The end

- Thank you for your attention!
- Please fill a quick survey which will be sent to you
- Do not hesitate to contact user support
- Yes, slides and recording will be provided
- Any more questions?

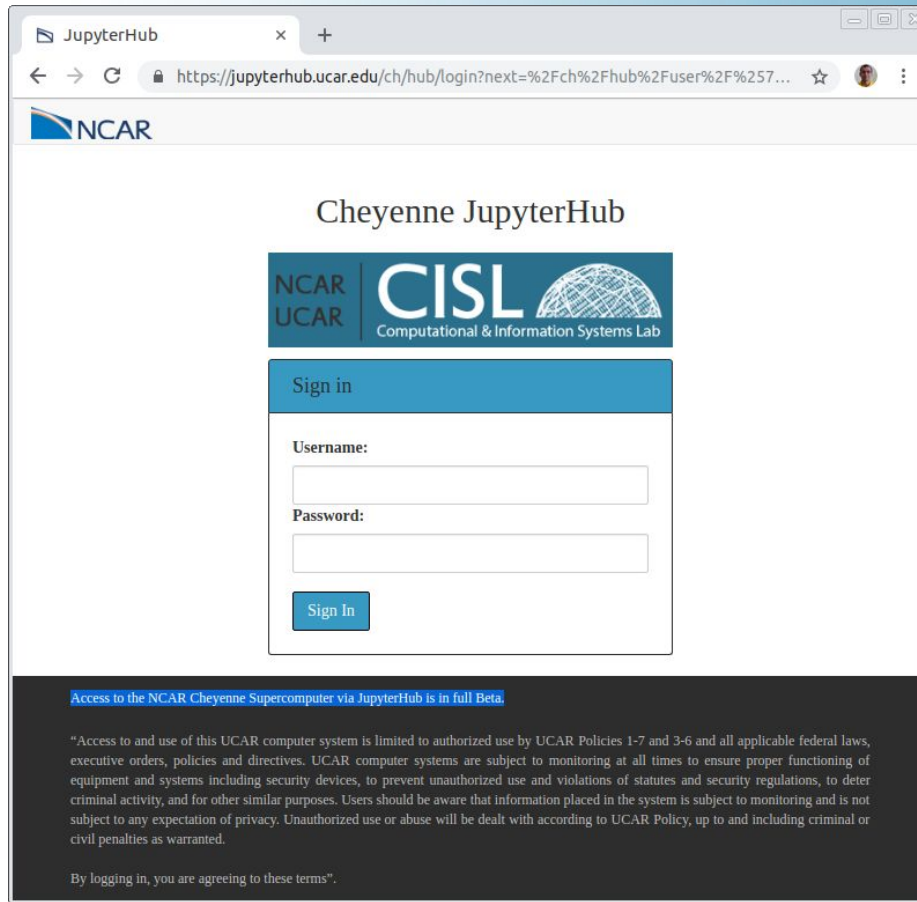
# Backup slides

# Logging in JupyterHub



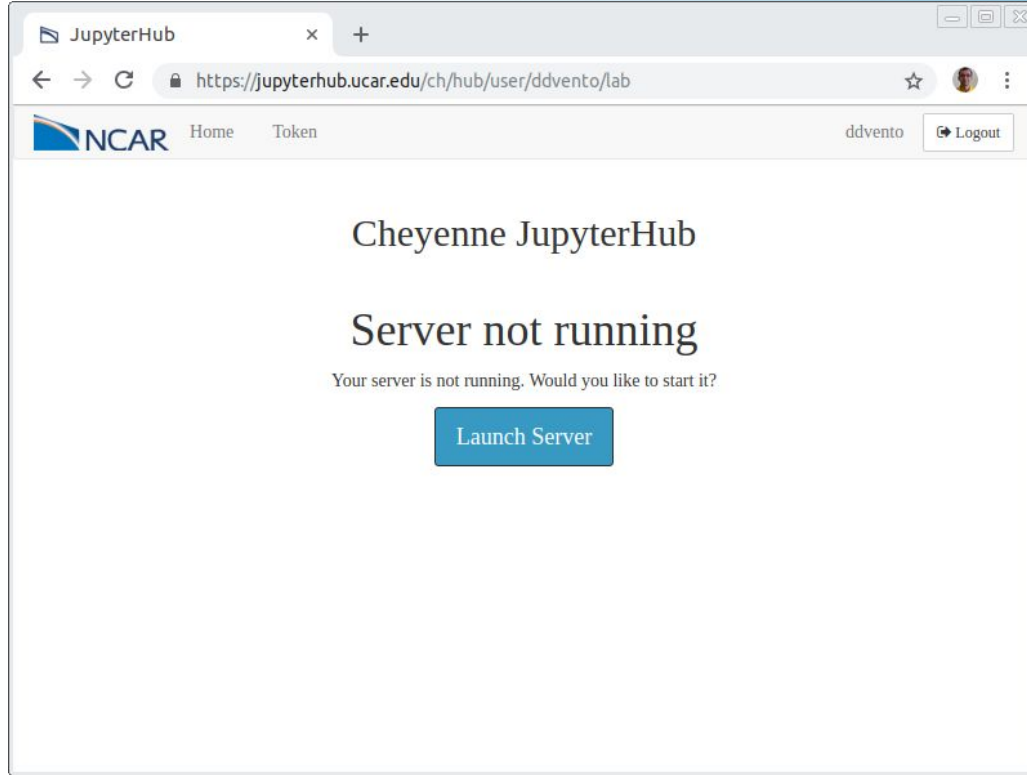


# Note the highlighted: this is beta



The screenshot shows a web browser window with the URL `https://jupyterhub.ucar.edu/ch/hub/login?next=%2Fch%2Fhub%2Fuser%2F%257...`. The page title is "Cheyenne JupyterHub". At the top left is the NCAR logo. Below it is a banner for "NCAR UCAR CISL Computational & Information Systems Lab". The main content is a "Sign in" form with fields for "Username:" and "Password:", and a "Sign In" button. At the bottom, a dark grey box contains a blue link: [Access to the NCAR Cheyenne Supercomputer via JupyterHub is in full Beta.](#) Below the link is a disclaimer: "Access to and use of this UCAR computer system is limited to authorized use by UCAR Policies 1-7 and 3-6 and all applicable federal laws, executive orders, policies and directives. UCAR computer systems are subject to monitoring at all times to ensure proper functioning of equipment and systems including security devices, to prevent unauthorized use and violations of statutes and security regulations, to deter criminal activity, and for other similar purposes. Users should be aware that information placed in the system is subject to monitoring and is not subject to any expectation of privacy. Unauthorized use or abuse will be dealt with according to UCAR Policy, up to and including criminal or civil penalties as warranted." At the very bottom of the dark box, it says "By logging in, you are agreeing to these terms".

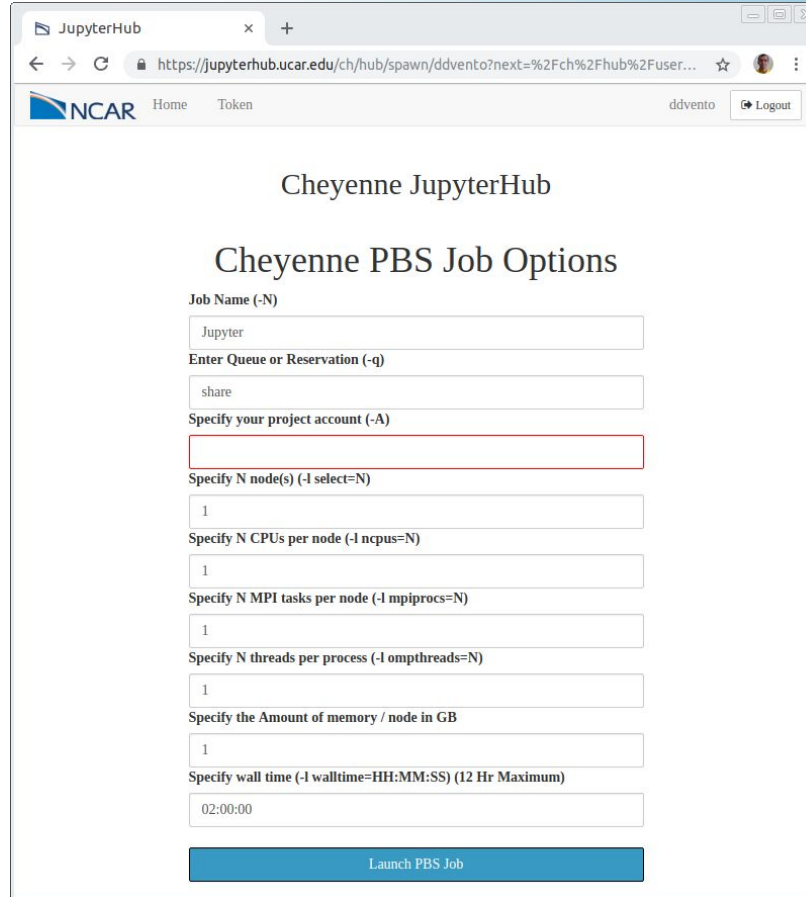
# Starting the Jupyter Server



# What is a Jupyter Server?

- A program responding to HTTP requests
- For a single user!!!
- The same program you can start yourself “the old way” or on your laptop
- Usually “inaccessible” from the outside
- Needs “something” to become accessible (e.g. ssh-tunnel)

# Submitting a PBS (or Slurm) job



The screenshot shows a web browser window with the URL `https://jupyterhub.ucar.edu/ch/hub/spawn/ddvento?next=%2Fch%2Fhub%2Fuser...`. The page title is "Cheyenne JupyterHub" and the main heading is "Cheyenne PBS Job Options". The form contains the following fields:

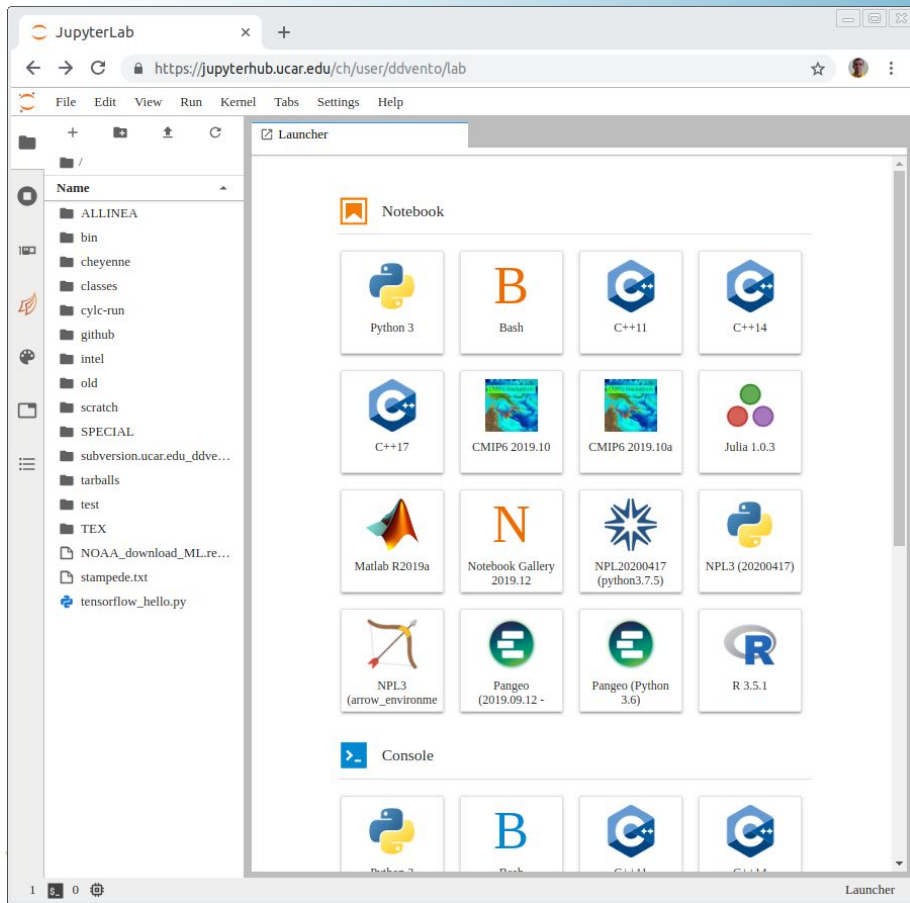
- Job Name (-N)**: Input field containing "Jupyter".
- Enter Queue or Reservation (-q)**: Input field containing "share".
- Specify your project account (-A)**: Empty input field.
- Specify N node(s) (-l select=N)**: Input field containing "1".
- Specify N CPUs per node (-l ncpus=N)**: Input field containing "1".
- Specify N MPI tasks per node (-l mpirprocs=N)**: Input field containing "1".
- Specify N threads per process (-l ompthreads=N)**: Input field containing "1".
- Specify the Amount of memory / node in GB**: Input field containing "1".
- Specify wall time (-l walltime=HH:MM:SS) (12 Hr Maximum)**: Input field containing "02:00:00".

A blue button labeled "Launch PBS Job" is located at the bottom of the form.

# Why submitting a job?

- Your node, your resources...
- i.e. others will not step on your toes (& vice versa)
- The memory is yours, if you run out, you have exhausted hw
- CPU is yours, if you run out...
- GPU (if you requested) is yours...

# Launch a kernel





# What is a kernel?

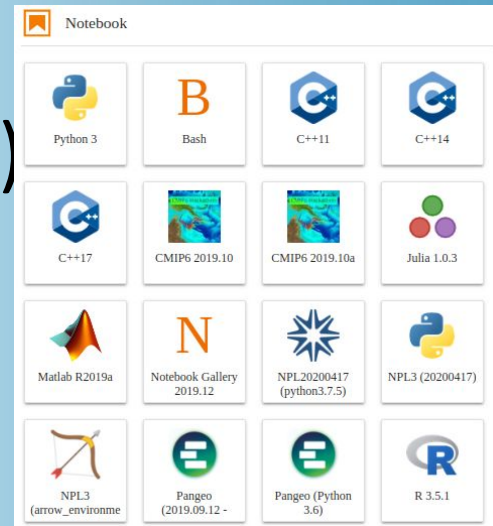
- The kernel is a program which interprets what you type, execute it in the language of choice, and keep the “session” in scope (so variables are not forgotten at each **RETURN**)

# What kernels can you use?

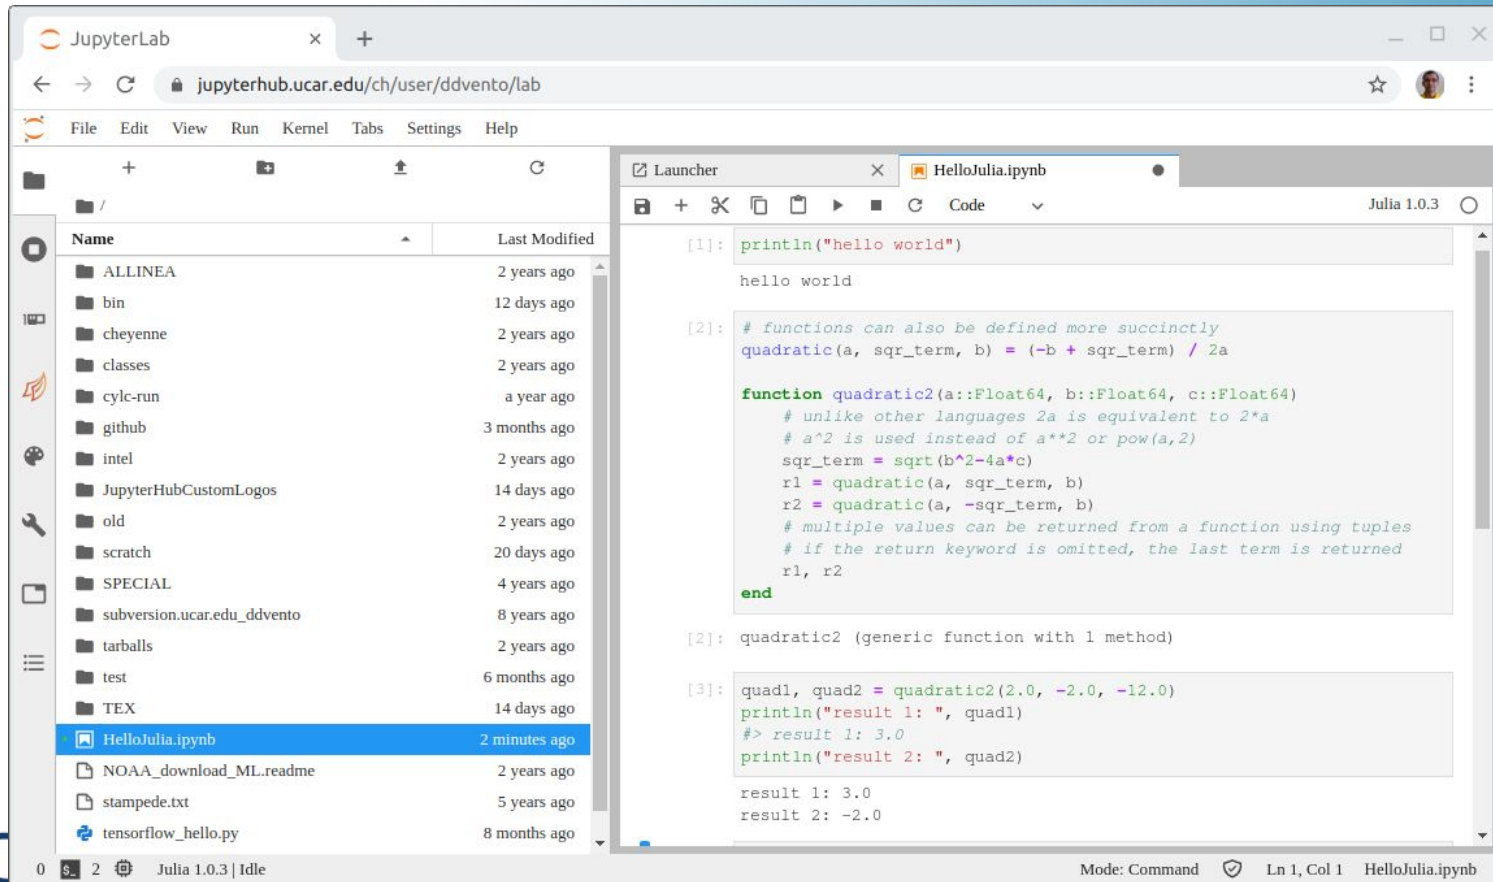
- Three varieties (that I know of):
  - Provided by NCAR/CISL  
(actually two kinds of these, next slide)
  - Provided by you as mods of CISL's ones
  - Provided by you as conda environments

# How to choose a kernel?

- What programming language would you like to use?
- Often Python, but many more are available (C, julia, matlab..)
- What libraries do you need?



# Working with a notebook



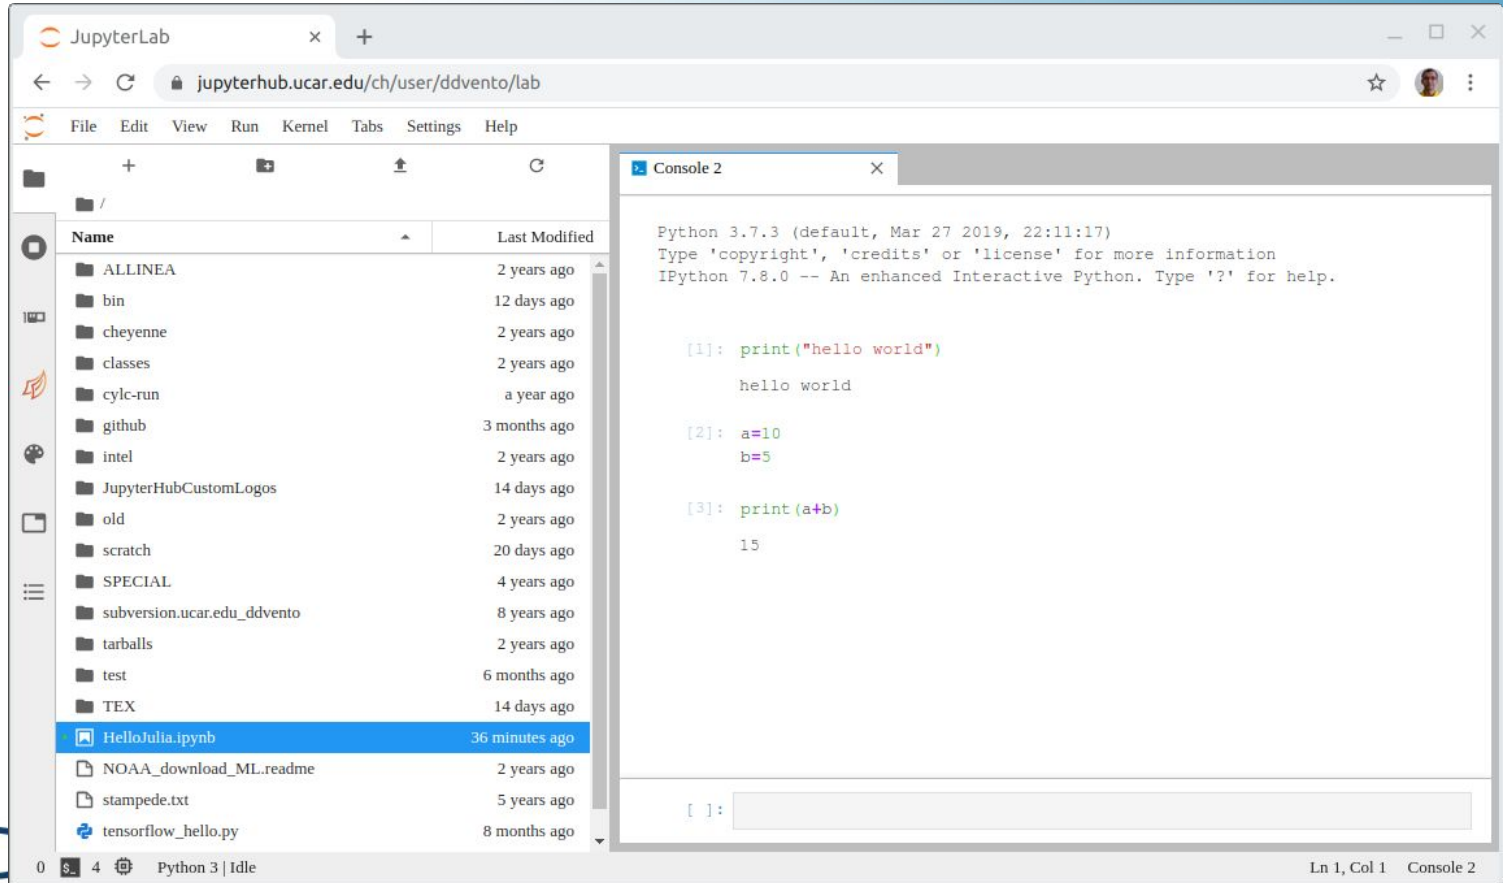
The screenshot shows a JupyterLab interface in a web browser. The browser address bar shows the URL `https://jupyterhub.ucar.edu/ch/user/ddvento/lab`. The interface includes a menu bar (File, Edit, View, Run, Kernel, Tabs, Settings, Help) and a file browser on the left. The file browser displays a list of files and folders with their last modified dates. The file `HelloJulia.ipynb` is selected and highlighted in blue.

The main area shows the notebook content for `HelloJulia.ipynb`. The code is as follows:

```
[1]: println("hello world")  
  
hello world  
  
[2]: # functions can also be defined more succinctly  
quadratic(a, sqr_term, b) = (-b + sqr_term) / 2a  
  
function quadratic2(a::Float64, b::Float64, c::Float64)  
    # unlike other languages 2a is equivalent to 2*a  
    # a^2 is used instead of a**2 or pow(a,2)  
    sqr_term = sqrt(b^2-4a*c)  
    r1 = quadratic(a, sqr_term, b)  
    r2 = quadratic(a, -sqr_term, b)  
    # multiple values can be returned from a function using tuples  
    # if the return keyword is omitted, the last term is returned  
    r1, r2  
  
end  
  
[2]: quadratic2 (generic function with 1 method)  
  
[3]: quad1, quad2 = quadratic2(2.0, -2.0, -12.0)  
println("result 1: ", quad1)  
#> result 1: 3.0  
println("result 2: ", quad2)  
  
result 1: 3.0  
result 2: -2.0
```

The status bar at the bottom indicates the current mode is `Command`, the cursor is at `Ln 1, Col 1`, and the active file is `HelloJulia.ipynb`.

# Working with the console



The screenshot displays the JupyterLab web interface. On the left, a file browser shows a directory structure with files like ALLINEA, bin, cheyenne, classes, cylc-run, github, intel, JupyterHubCustomLogos, old, scratch, SPECIAL, subversion.ucar.edu\_ddvento, tarballs, test, and TEX. The file 'HelloJulia.ipynb' is selected. On the right, a console window titled 'Console 2' shows the following text:

```
Python 3.7.3 (default, Mar 27 2019, 22:11:17)
Type 'copyright', 'credits' or 'license' for more information
IPython 7.8.0 -- An enhanced Interactive Python. Type '?' for help.

[1]: print("hello world")
      hello world

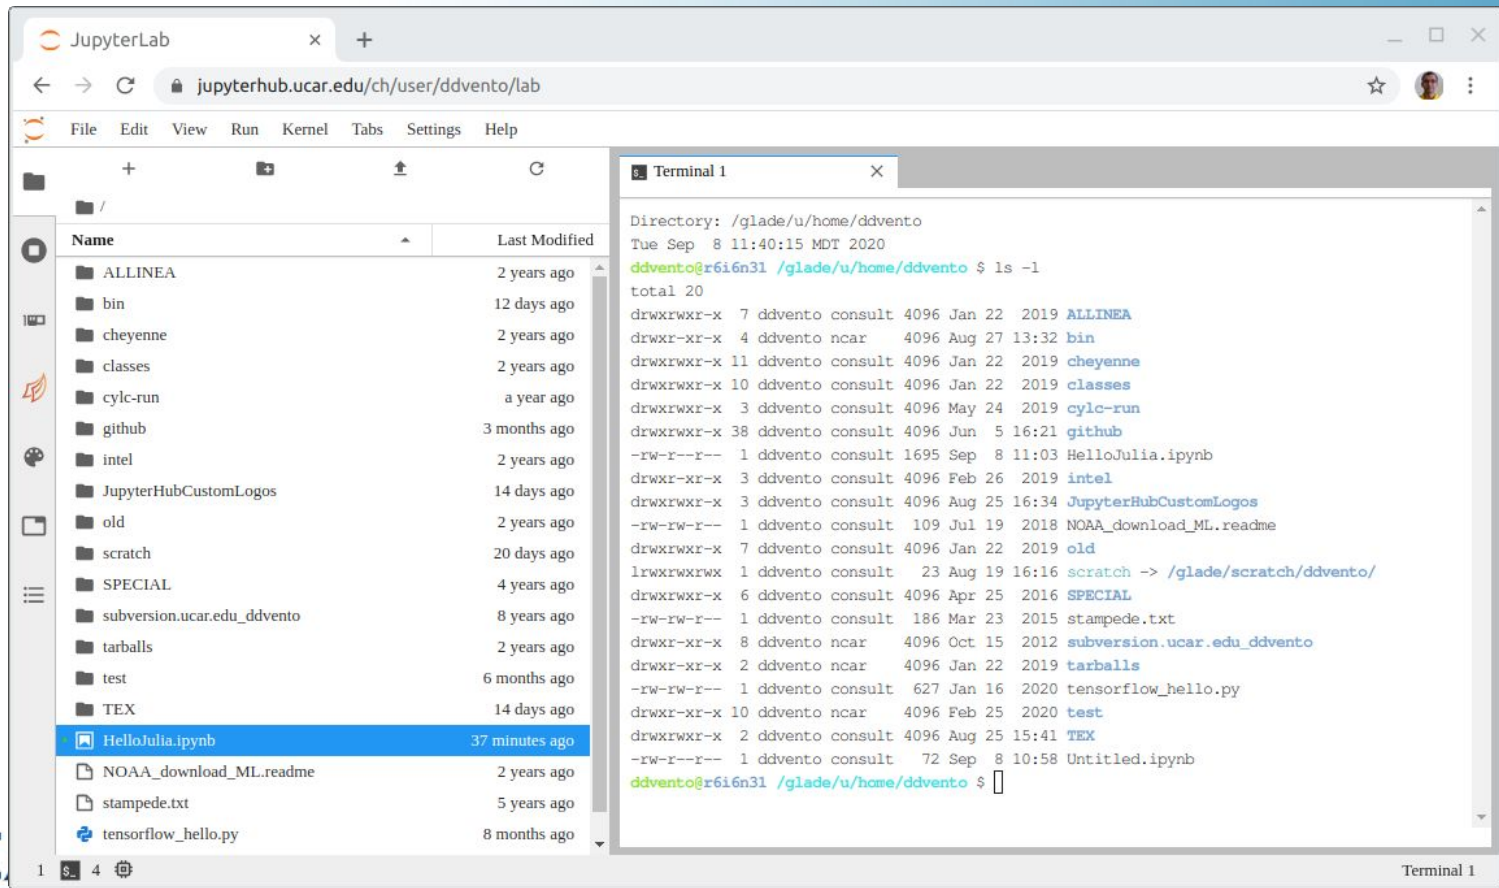
[2]: a=10
      b=5

[3]: print(a+b)
      15

[ ]:
```

The status bar at the bottom indicates 'Python 3 | Idle' and 'Ln 1, Col 1 Console 2'.

# Working with the shell



The screenshot shows the JupyterLab interface. On the left is a file browser showing a directory listing. On the right is a terminal window displaying the output of a `ls -l` command.

Name	Last Modified
ALLINEA	2 years ago
bin	12 days ago
cheyenne	2 years ago
classes	2 years ago
cylc-run	a year ago
github	3 months ago
intel	2 years ago
JupyterHubCustomLogos	14 days ago
old	2 years ago
scratch	20 days ago
SPECIAL	4 years ago
subversion.ucar.edu_ddvento	8 years ago
tarballs	2 years ago
test	6 months ago
TEX	14 days ago
HelloJulia.ipynb	37 minutes ago
NOAA_download_ML.readme	2 years ago
stampede.txt	5 years ago
tensorflow_hello.py	8 months ago

```
Terminal 1
Directory: /glade/u/home/ddvento
Tue Sep 8 11:40:15 MDT 2020
ddvento@r616n31 /glade/u/home/ddvento $ ls -l
total 20
drwxrwxr-x 7 ddvento consult 4096 Jan 22 2019 ALLINEA
drwxr-xr-x 4 ddvento ncar 4096 Aug 27 13:32 bin
drwxrwxr-x 11 ddvento consult 4096 Jan 22 2019 cheyenne
drwxrwxr-x 10 ddvento consult 4096 Jan 22 2019 classes
drwxrwxr-x 3 ddvento consult 4096 May 24 2019 cylc-run
drwxrwxr-x 38 ddvento consult 4096 Jun 5 16:21 github
-rw-r--r-- 1 ddvento consult 1695 Sep 8 11:03 HelloJulia.ipynb
drwxr-xr-x 3 ddvento consult 4096 Feb 26 2019 intel
drwxrwxr-x 3 ddvento consult 4096 Aug 25 16:34 JupyterHubCustomLogos
-rw-rw-r-- 1 ddvento consult 109 Jul 19 2018 NOAA_download_ML.readme
drwxrwxr-x 7 ddvento consult 4096 Jan 22 2019 old
lrwxrwxrwx 1 ddvento consult 23 Aug 19 16:16 scratch -> /glade/scratch/ddvento/
drwxrwxr-x 6 ddvento consult 4096 Apr 25 2016 SPECIAL
-rw-rw-r-- 1 ddvento consult 186 Mar 23 2015 stampede.txt
drwxr-xr-x 8 ddvento ncar 4096 Oct 15 2012 subversion.ucar.edu_ddvento
drwxr-xr-x 2 ddvento ncar 4096 Jan 22 2019 tarballs
-rw-rw-r-- 1 ddvento consult 627 Jan 16 2020 tensorflow_hello.py
drwxr-xr-x 10 ddvento ncar 4096 Feb 25 2020 test
drwxrwxr-x 2 ddvento consult 4096 Aug 25 15:41 TEX
-rw-r--r-- 1 ddvento consult 72 Sep 8 10:58 Untitled.ipynb
ddvento@r616n31 /glade/u/home/ddvento $
```



Back to [presentation](#)