

Jupyter Hub on SciNet

Ramses van Zon

SciNet HPC Consortium

January 11, 2017

Interactive Python on SciNet from your Browser

or . . .

Interactive R on SciNet from your Browser



Interactive Python

- Python has always (1991-) been interactive.
- IPython has always (2001-) offered an improved interactive environment.
- The IPython Notebook has been (2013-2015) a web-based, notebook-style interface.
- Jupyter is the spin-off of the *notebook* part of IPython.
- Jupyter can support support other languages as well, notably R (but, apparently, also Julia, Haskell, and Ruby).
- Jupyter hub is a server app that allows multiple users to spawn notebooks from their browser.

Setup on SciNet

Disclaimer: this is an experimental setup running on good, but out-of-warranty hardware.

- Two jupyterhub servers, each with 128 GB of memory and 16 cores.
- Access using an ssh tunnel via login.scinet.utoronto.ca:

```
$ ssh USER@login.scinet.utoronto.ca -L8888:jupyterhub:8000 -N -f
```

- This will select (round-robin) one of the two jupyterhub servers.
- Point your browser to 'localhost:8888' and log in with your SciNet account.
- The browser should now show the files in your \$HOME on SciNet.
(if not, try reloading the page, it may have timed out).
- You can open or create Python 2, Python 3, and R notebooks.



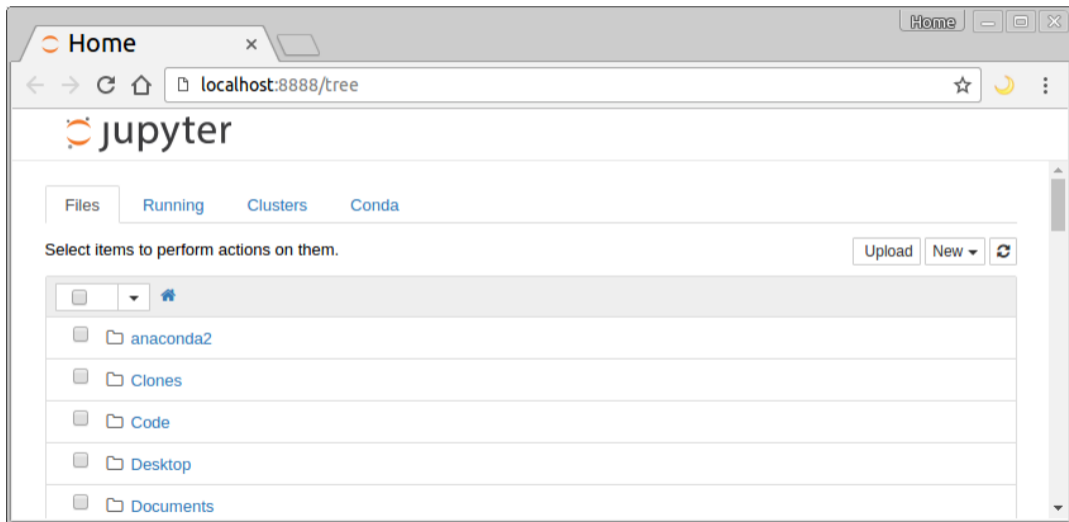
Jupyter Notebook



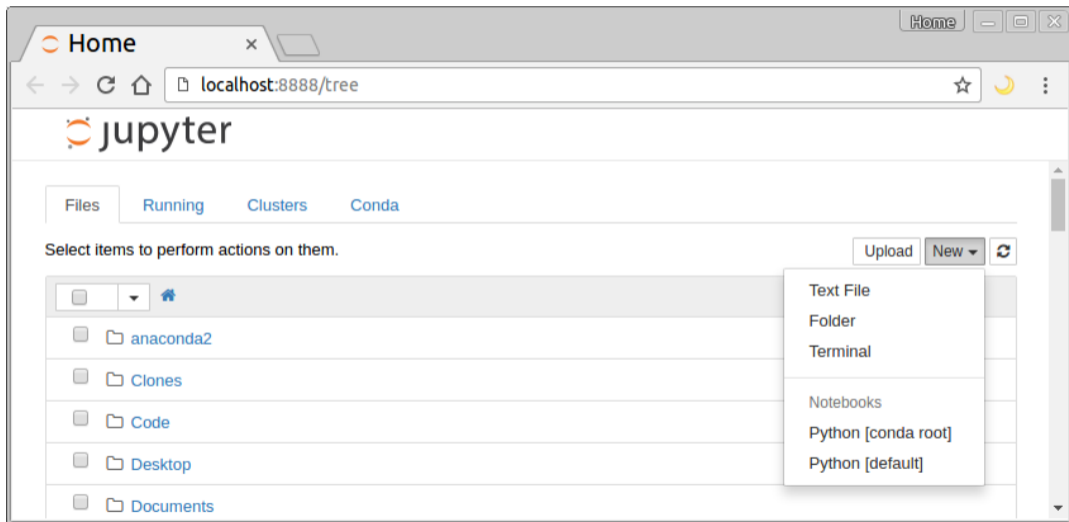
- Runs an interface in your web browser for python
- In that interface, **notebooks** contain python code and **rich text** or **graphics** outputs interspersed.
- You can save, load, modify notebooks.
- Can use it on your own computer too. If you've installed anaconda, you should have this. E.g.

```
$ [INSTALLDIR]/jupyter notebook
```

Jupyter Notebook - Screenshot



Jupyter Notebook - Screenshot



The screenshot displays the Jupyter Notebook web interface in a browser window. The address bar shows the URL `localhost:8888/tree`. The page title is "jupyter". Below the title, there are navigation tabs for "Files", "Running", "Clusters", and "Conda". The "Files" tab is active, showing a file browser interface. The text "Select items to perform actions on them." is displayed above the file list. The file list includes a home icon, a dropdown arrow, and a list of folders: "anaconda2", "Clones", "Code", "Desktop", and "Documents". A "New" button is visible, and its dropdown menu is open, showing options: "Text File", "Folder", "Terminal", "Notebooks", "Python [conda root]", and "Python [default]".

Jupyter Notebook - Screenshot

The screenshot displays the Jupyter Notebook web interface in a browser window. The address bar shows the URL `localhost:8888/tree`. The page title is "jupyter". Below the title, there are navigation tabs for "Files", "Running", "Clusters", and "Conda". The "Files" tab is active, showing a file browser interface. The interface includes a "Select items to perform actions on them." prompt, an "Upload" button, a "New" dropdown menu, and a refresh button. The file tree shows the following folders:

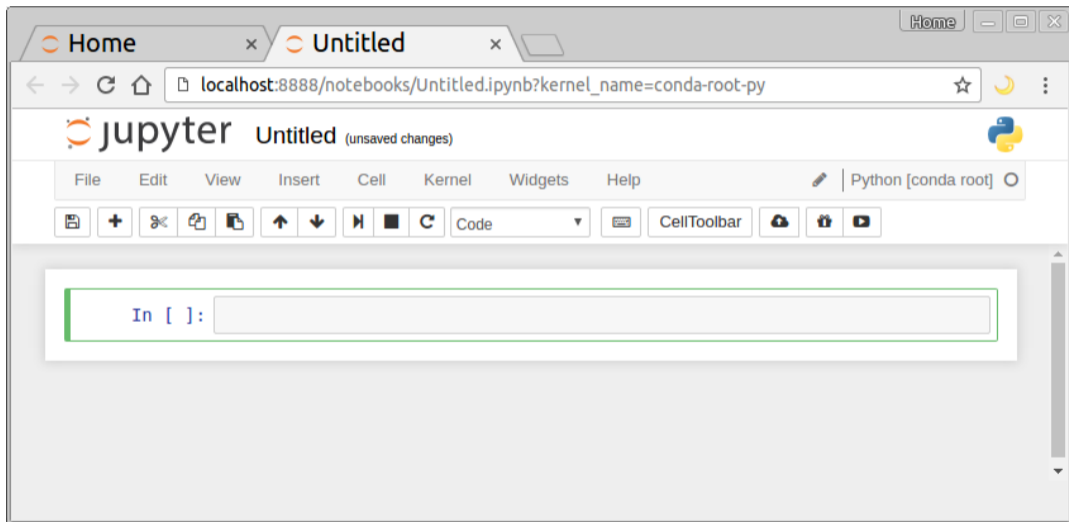
- Home
- anaconda2
- Clones
- Code
- Desktop

The "New" dropdown menu is open, showing the following options:

- Text File
- Folder
- Terminal
- Notebooks
- Python [conda root]
- Python [default]

The address bar at the bottom of the browser window shows `localhost:8888/tree#`.

Jupyter Notebook - Screenshot



Jupyter Notebook - Screenshot

Home x Untitled x

localhost:8888/notebooks/Untitled.ipynb?kernel_name=conda-root-py

jupyter Untitled (autosaved)

File Edit View Insert Cell Kernel Widgets Help Python [conda root]

Code CellToolbar

```
In [1]: print "Hello, world"
Hello, world
```

```
In [ ]:
```

Jupyter Notebook - Tips to get started

- Jupyter can also browse your (SciNet) files and edit them.
- Use the 'new' button to create a new python notebook.
- Give your notebooks reasonable names.
- To execute a python input line, press Shift-Enter.
- Save your work periodically (even though there is autosave).
- To work similarly to `ipython --pylab`, do:

```
In [1]: from pylab import *  
        %matplotlib notebook
```

Drawbacks of a Notebook Environment

- Notebook files (.ipynb) are not scripts.
- Does not (always) work well with version control.
- Designed to run in browser.
- Graphics is inline, which is great for quick exploration but make tweaking a plot harder (IPython+X works better for this).
- You can jump around in the notebook, and execute different parts: hard to keep track of what you did.

Advantages of a Notebook environment

- You can jump around in the notebook, and execute different parts: Easier exploration, experimentation and debugging.
- Auto-save
- You can rerun parts of your code (while, e.g., keeping large data in memory)
- You can add text portions, making your notebook more like an article.
- Which in turn can be useful for sharing, demos, teaching, . . .
- You can still export as a script.
- Also has a terminal.

Advantages of using SciNet's JupyterHub

SciNet's JupyterHub allows different users to run notebooks, using your SciNet credentials.

- Runs on SciNet, where it can see your files.
- When using graphics: no need to setup X or vnc.
- Dedicated, large memory servers (shared by all users though).
- No need to set up your own JupyterHub on Scinet.
- Set up for Python 2, Python 3, and R.
- Large number of packages preinstalled.

Demo

...

Running Jupyter Notebook on a Compute Node

- 1 (Once) Create symbolic links from your home directory to your scratch:

```
$ mkdir -p $SCRATCH/admin $HOME/.local/share
$ mv $HOME/.local/share/jupyter $SCRATCH/admin/jupyter
$ mv $HOME/.ipython $SCRATCH/admin/ipython
$ ln -s $SCRATCH/admin/jupyter $HOME/.local/share/jupyter
$ ln -s $SCRATCH/admin/ipython $HOME/.ipython
```

- 2 Grab your compute node, e.g.

```
$ qsub -I -l nodes=1:ppn=8,walltime=2:00:00
```

- 3 At the interactive prompt, type

```
$ hostname # prints the HOSTNAME; write it down or copy it.
$ module load gcc/4.4.6 ffmpeg/3.1 anaconda3/hub
$ cd $SCRATCH; jupyter notebook --ip=* --no-browser
```

- 4 In a different terminal on your own machine, open an ssh tunnel

```
$ ssh USERNAME@login.scinet.utoronto.ca -L8889:HOSTNAME:8888 -N -f
```

- 5 Now point your browser at localhost:8889



Questions

...