



JUPYTERLAB - SUPERCOMPUTING IN YOUR BROWSER

Training course "Introduction to the usage and programming of supercomputer resources in Jülich"

2022-11-21 | JENS H. GÖBBERT

(J.GOEBBERT@FZ-JUELICH.DE)

TIM KREUZER

(T.KREUZER@FZ-JUELICH.DE)

ALICE GROSCH

(A.GROSCH@FZ-JUELICH.DE)

MOTIVATION

your thinking, your reasoning, your insides, your ideas

“It is all about using and building a machinery **interface** **between** computational researchers and data, supercomputers, laptops, cloud **and** your thinking, your reasoning, your insides, your ideas about a problem.”

Fernando Perez, Berkely Institute for Data Science
Founder of Project Jupyter

<https://www.youtube.com/watch?v=xuNj5paMuow>

<https://jupyter.org>


Member of the Helmholtz Association

MOTIVATION

Rise of Jupyter's popularity

- In 2007, Fernando Pérez and Brian Granger announced „**ipython**: a system for interactive scientific computing“ [1]
- In 2014, Fernando Pérez announced a spin-off project from IPython called **Project Jupyter**.
 - IPython continued to exist as a Python shell and a kernel for Jupyter, while the Jupyter notebook moved under the Jupyter name.
- In 2015, GitHub and the Jupyter Project announced native rendering of Jupyter notebooks file format (.ipynb files) on the **GitHub**
- In 2017, the **first JupyterCon** was organized by O'Reilly in New York City. Fernando Pérez opened the conference with an inspiring talk. [2]
- In 2018, **JupyterLab** was announced as the next-generation web-based interface for Project Jupyter.
- In 2019, JupyterLab 1.0 ...
In 2020, JupyterLab 2.0 ...
In 2021, JupyterLab 3.0 ...
In 2022, JupyterLab 4.0 planned end of year

Counting how many Monthly Active Users (MAU) on GitHub are using Jupyter Notebooks



<https://www.benfrederickson.com/ranking-programming-languages-by-github-users/>
<https://github.com/benfred/github-analysis>

[1] Pérez F, Granger BE (2007) Ipython: a system for interactive scientific computing. Comput Sci Eng 9(3):21–29

[2] Pérez F, Project Jupyter: From interactive Python to open science -> <https://www.youtube.com/watch?v=xuNj5paMuow>

JUPYTER NOTEBOOK

creating reproducible computational narratives

Markdown Cells

Code Cells

Fourier transform

Fourier transforms are one of the universal tools in computational physics, which appear over and over again in different contexts. SciPy provides functions for accessing the classic [FFTPACK](#) library from NetLib, which is an efficient and well tested FFT library written in FORTRAN. The SciPy API has a few additional convenience functions, but overall the API is closely related to the original FORTRAN library.

To use the `fftpack` module in a python program, include it using:

```
[41]: from numpy.fft import fftfreq
      from scipy.fftpack import *
```

To demonstrate how to do a fast Fourier transform with SciPy, let's look at the FFT of the solution to the damped oscillator:

$$\frac{d^2x}{dt^2} + 2\zeta\omega_0 \frac{dx}{dt} + \omega_0^2 x = 0$$

where x is the position of the oscillator, ω_0 is the frequency, and ζ is the damping ratio. To write this second-order ODE on standard form we introduce $p = \frac{dx}{dt}$:


```
[42]: N = len(t)
      dt = t[1]-t[0]
      dt
```

```
[42]: 0.01001001001001001
```

```
[43]: # calculate the fast fourier transform
      # y2 is the solution to the under-damped oscillator from the previous section
      F = fft(y2[:,0])

      # calculate the frequencies for the components in F
      w = fftfreq(N, dt)
```

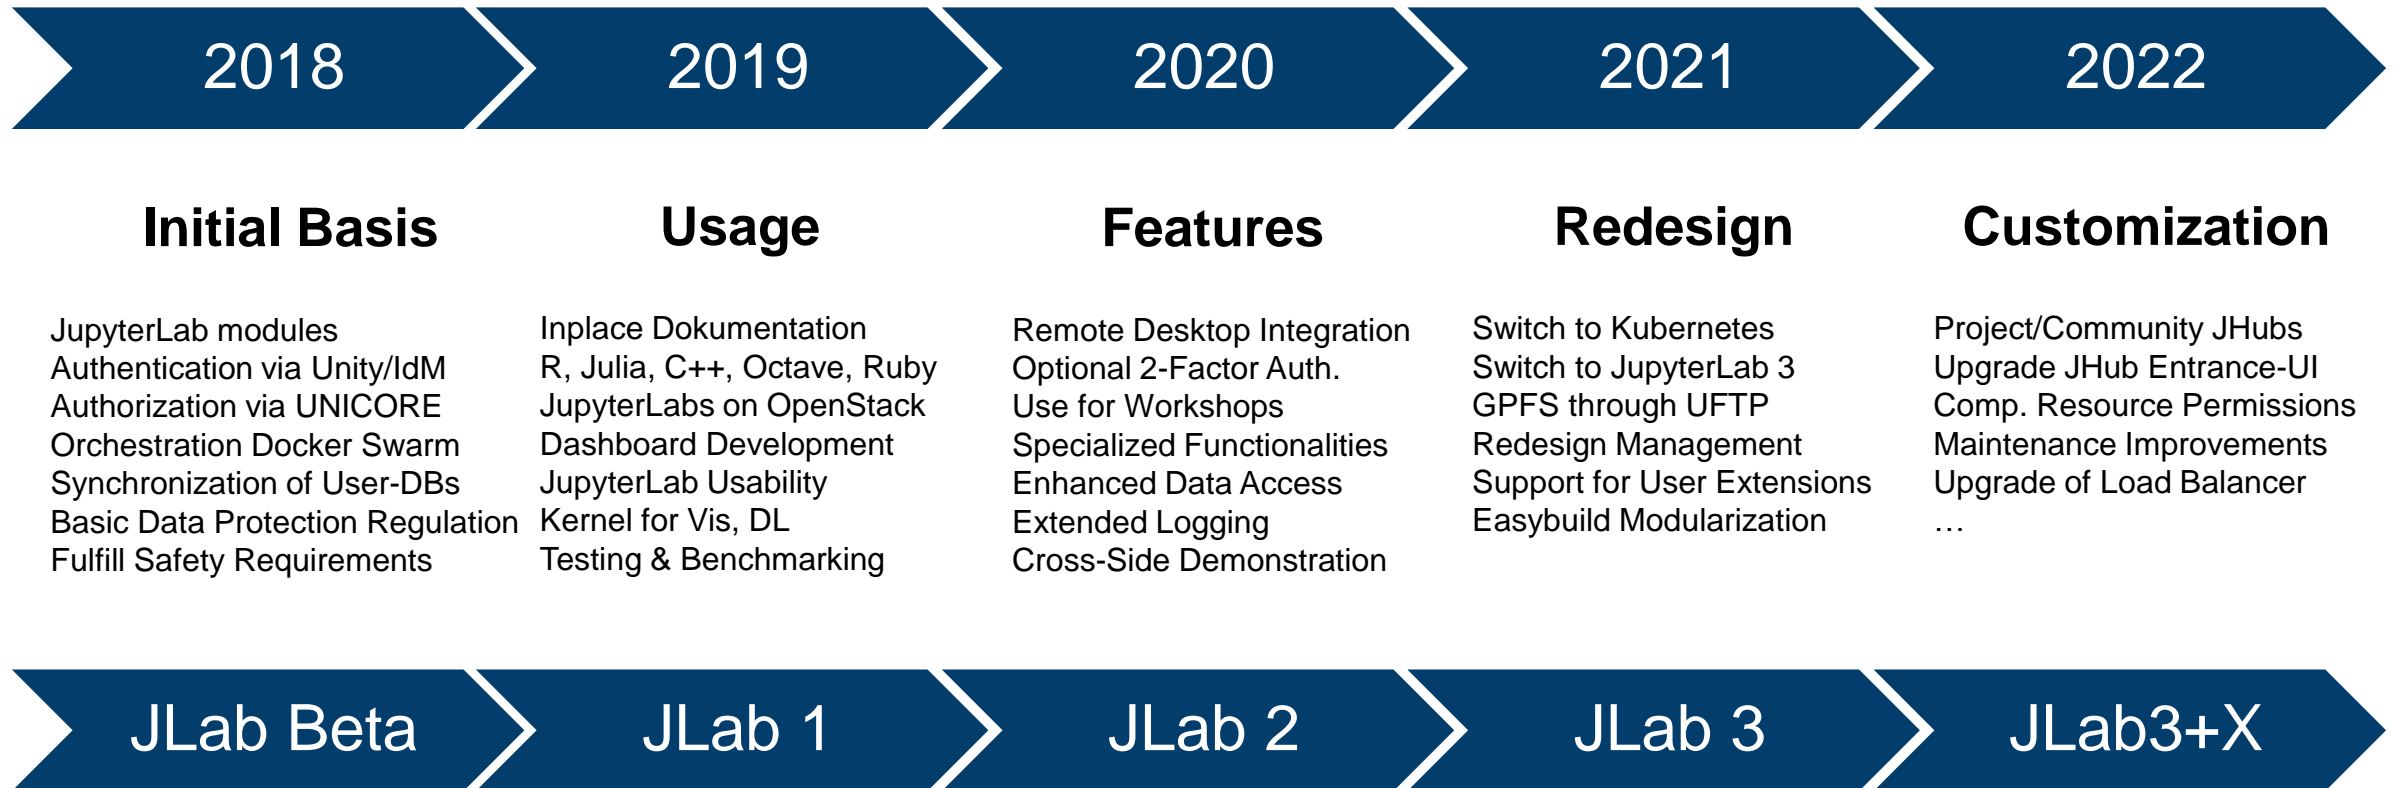
```
[44]: fig, ax = plt.subplots(figsize=(9,3))
      ax.plot(w, abs(F));
```



Output

Output

HISTORY OF JUPYTERLAB AT JSC



HISTORY OF JUPYTERLAB AT JSC

2018

Initial Basis

JupyterLab modules
Authentication via Unity/IdM
Authorization via UNICORE
Orchestration Docker Swarm
Synchronization of User-DBs
Basic Data Protection Regulation
Fulfill Safety Requirements

2022

Customization

Project/Community JHubs
Upgrade JHub Entrance-UI
Comp. Resource Permissions
Maintenance Improvements
Upgrade of Load Balancer
...


JLab Beta

JLab 1

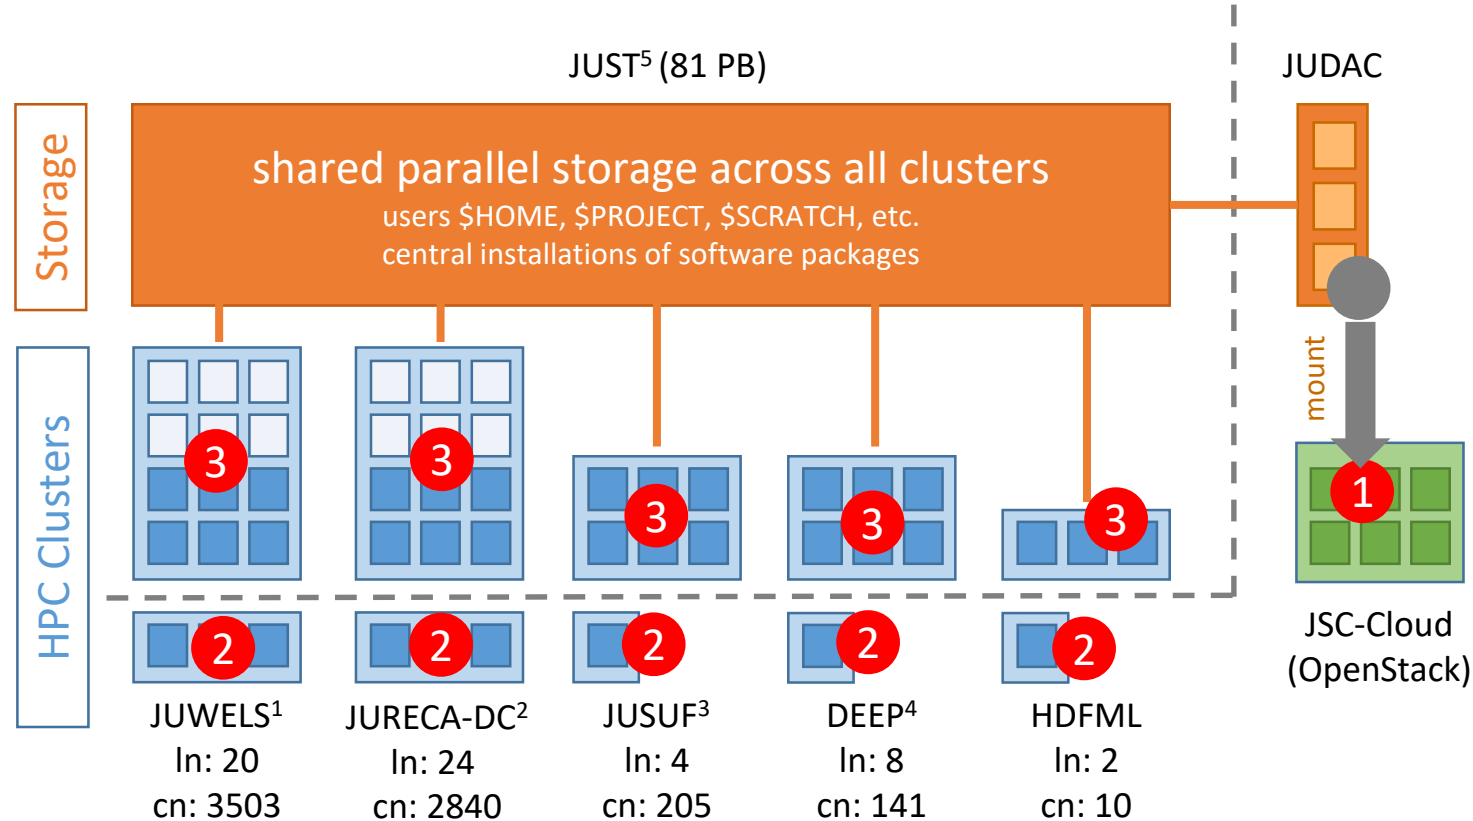
JLab 2

JLab 3

JLab3+X



JUPYTERLAB EVERYWHERE



no. login nodes = ln
no. compute nodes = cn

[1] <https://apps.fz-juelich.de/jsc/hps/juwels/configuration.html>

[2] <https://apps.fz-juelich.de/jsc/hps/jureca/configuration.html>

[3] <https://apps.fz-juelich.de/jsc/hps/jusuf/cluster/configuration.html>

[4] https://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/DEEP-EST/_node.html

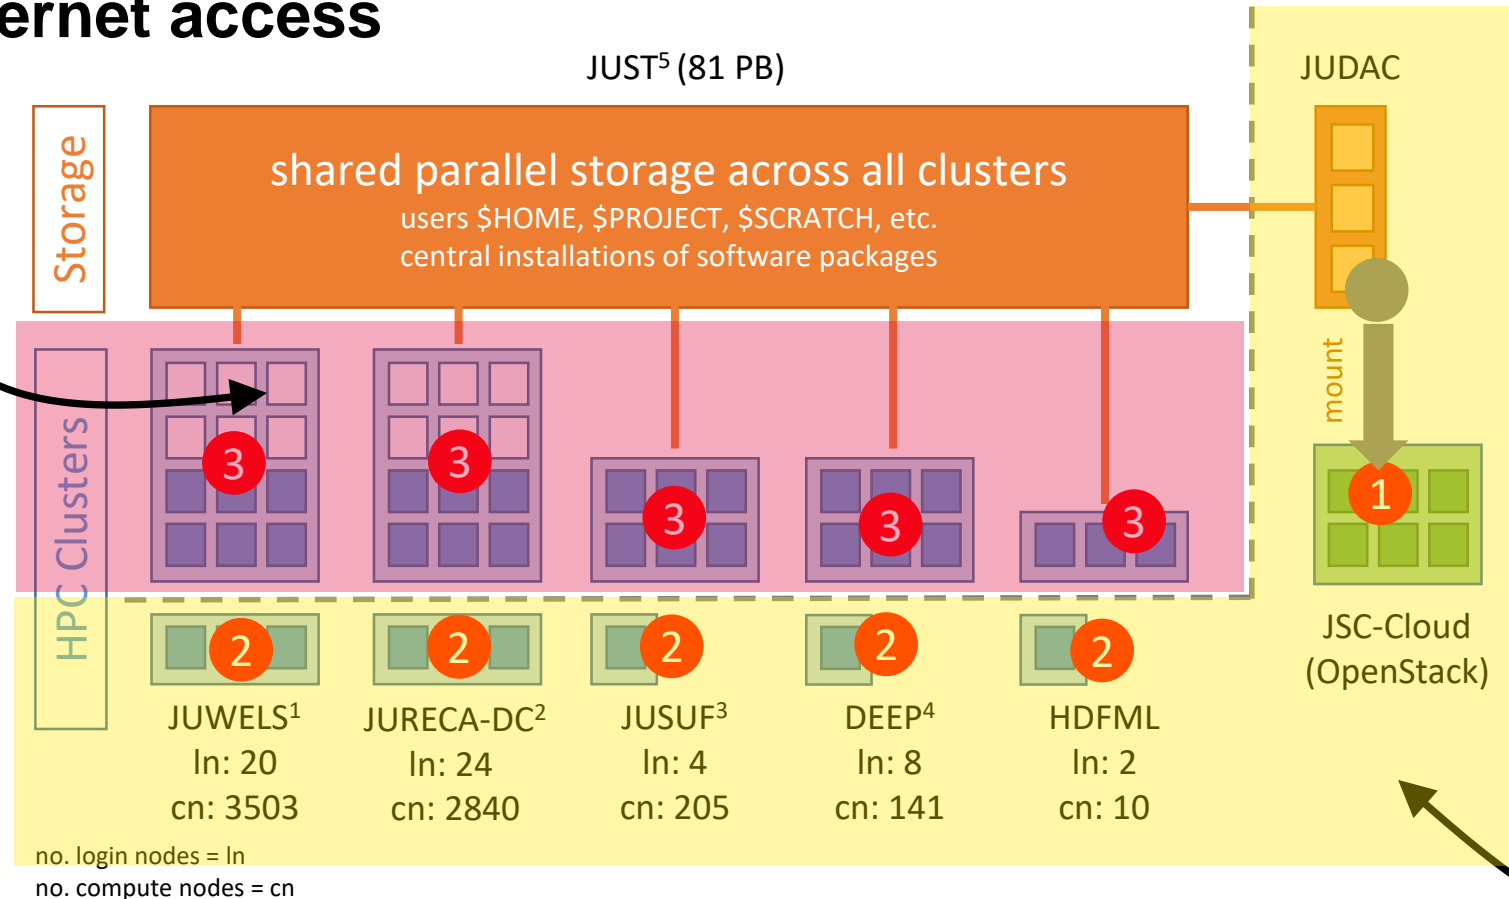
[5] https://www.fz-juelich.de/ias/jsc/EN/Expertise/Datamanagement/OnlineStorage/JUST/Configuration/Configuration_node.html

JupyterLab everywhere

- 1 JupyterLab on cloud
- 2 JupyterLab on login nodes
- 3 JupyterLab on compute nodes

JUPYTERLAB EVERYWHERE

NO internet access



JupyterLab everywhere

- 1 JupyterLab on cloud
- 2 JupyterLab on login nodes
- 3 JupyterLab on compute nodes

internet access

[1] <https://apps.fz-juelich.de/jsc/hps/juwels/configuration.html>
[2] <https://apps.fz-juelich.de/jsc/hps/jureca/configuration.html>
[3] <https://apps.fz-juelich.de/jsc/hps/jusuf/cluster/configuration.html>
[4] https://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/DEEP-EST/_node.html
[5] https://www.fz-juelich.de/ias/jsc/EN/Expertise/Datamanagement/OnlineStorage/JUST/Configuration/Configuration_node.html

TERMINOLOGY

What is JupyterLab

JupyterLab

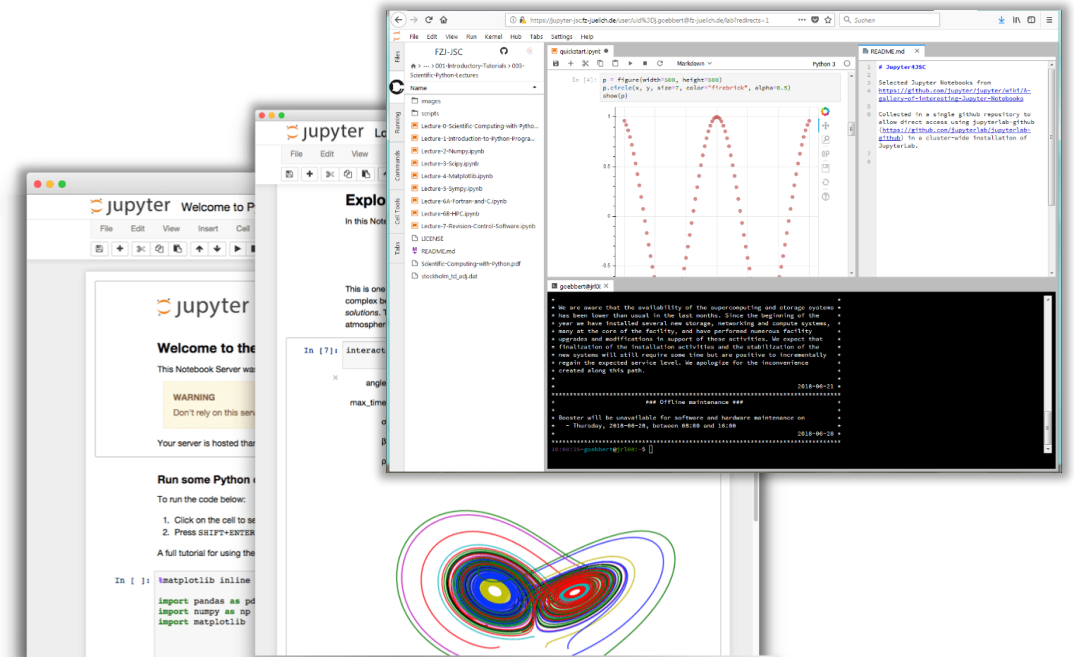
- **Interactive** working environment in the web browser
- For the creation of **reproducible** computer-aided narratives
- Very **popular** with researchers from all fields
- Jupyter = Julia + Python + R

Multi-purpose working environment

- Language agnostic
- Supports execution environments (“kernels”)
 - For dozens of languages: Python, R, Julia, C++, ...
- Extensible software design („extensions“)
 - many server/client plug-ins available
 - Eg. in-browser-terminal and file-browsing

Document-Centered Computing (“notebooks”)

- Combines code execution, rich text, math, plots and rich media.
- All-in-one document called Jupyter Notebook



<https://jupyterlab.readthedocs.io>

TERMINOLOGY

What is a Jupyter Notebook?

Jupyter Notebook

A notebook document (file extension **.ipynb**) is a document that can be rendered in a web browser

- It is a file, which stores your work in JSON format
- Based on a set of open standards for interactive computing
- Allows development of custom applications with embedded interactive computing.
- Can be extended by third parties
- Directly convertible to PDF, HTML, LaTeX ...
- Supported by many applications such as GitHub, GitLab, etc..

▼ 7:

```
cell_type:      "code"
execution_count: 4
metadata:       {}
outputs:        []
```

▼ source:

```
▼ 0:      "def calculate_pi(size_in_bytes, number_of_chunks):\n1:      "\n2:      "\n3:      "\n4:      "\n5:      "\n6:      "\n7:      "\n8:      "\n9:      "\n10:     "\n11:     "\n12:     "\n13:     "\n14:     "\n15:     "\n16:     "\n17:     "\n18:     "\n19:     "\n20:     "\n21:     "\n22:     "\n23:     "\n24:     "\n25:     "
```

```
"\n\n    Calculate pi using a Monte Carlo method.\n\n    an array of size size_in_bytes / 8 / 2, 2)\n\n    chunk_size = (int(array_shape[0] / number_of_chunks), 2)\n\n    # 2D random positions array\n    xy = data.uniform(low=0.0, high=1.0, size=array_shape[0] * chunk_size)\n\n    # specify chunk size, i.e. number of chunks\n\n\n    pi = 4 * xy_inside_circle.sum() / xy_inside_circle.size\n\n\n    # start Dask calculation\n\n    pi = pi.compute()\n\n\n    print(f"\nfrom {xy.nbytes / 1e9} GB randomly chosen positions")\n\n    print(f"    pi estimate: {pi}")\n\n    print(f"    pi error: {abs(pi - numpy.pi)}")\n\n    display(xy)\n\n\n    return pi
```

<https://jupyter-notebook.readthedocs.io/>
<https://github.com/jupyter/jupyter/wiki/A-gallery-of-interesting-Jupyter-Notebooks>

TERMINOLOGY

What is a Jupyter Kernel?

Jupyter Kernel

A “kernel” refers to the separate process which executes code cells within a Jupyter notebook.

Jupyter Kernel

- **run code** in different programming languages and environments.
- can be **connected to** a notebook (one at a time).
- **communicates** via ZeroMQ with the JupyterLab.
- Multiple **preinstalled** Jupyter Kernels can be found on our clusters
 - Python, R, Julia, Bash, C++, Ruby, JavaScript
 - Specialized kernels for deep learning, visualization, quantum computing
- You can easily **create your own kernel** which for example runs your specialized virtual Python environment.



<https://jupyter-notebook.readthedocs.io/>
<https://github.com/jupyter/jupyter/wiki/Jupyter-kernels>
<https://zeromq.org>

TERMINOLOGY

What is a JupyterLab Extension?

JupyterLab Extension

JupyterLab extensions can customize or enhance any part of JupyterLab.

JupyterLab Extensions


- provide new file viewers, editors, themes
 - provide renderers for rich outputs in notebooks
 - add items to the menu or command palette
 - add keyboard shortcuts
 - add settings in the settings system.
-
- Extensions can even provide an API for other extensions to use and can depend on other extensions.

The whole JupyterLab itself is simply a **collection of extensions** that are no more powerful or privileged than any custom extension.

With JupyterLab 3 **prebuild extensions** were introduced (in contrast to **source extensions**).

You can now (technically) extend a compiled JupyterLab 3+ with without the need to rebuild.

=> That allows a separation of a system-wide JupyterLab core installation and extensions installed by a user.



JUPYTERLAB - WHEREVER YOU PREFER

Local, Remote, Browser-only

Local installation:

- **JupyterLab** installed using conda, mamba, pip, pipenv or docker.
https://jupyterlab.readthedocs.io/en/stable/getting_started/installation.html

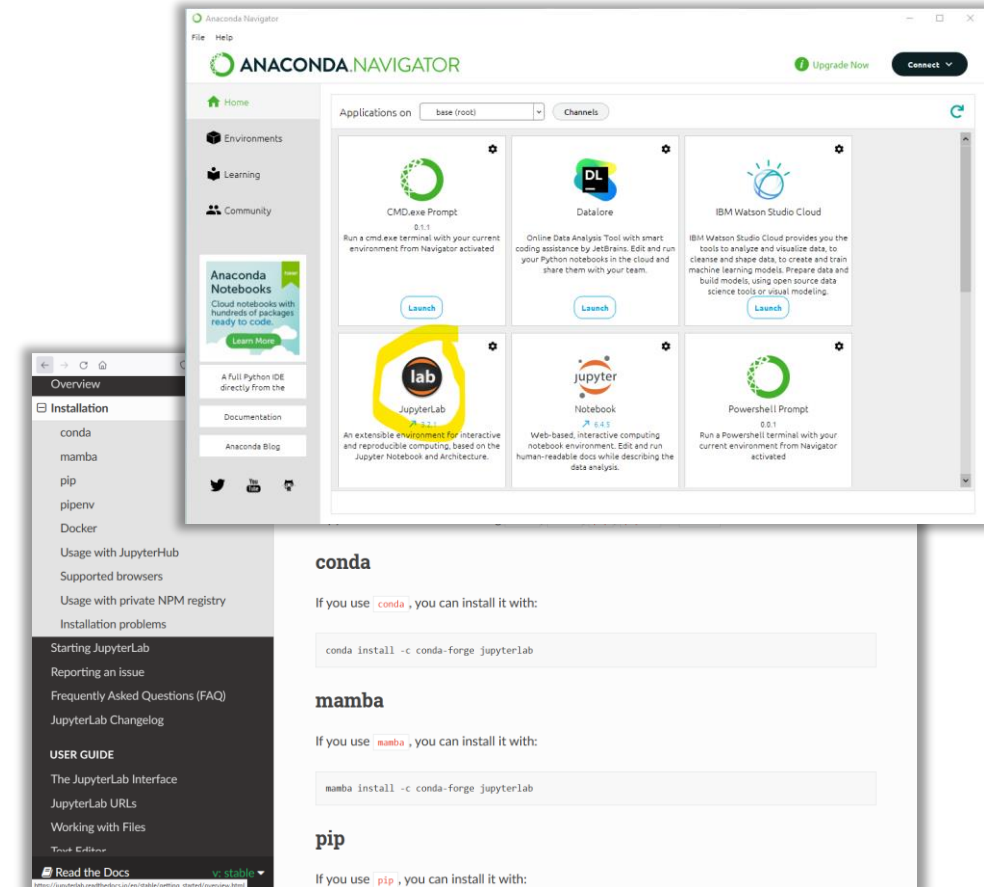


Remote (cluster) installation:

- **JupyterLab** installed on a remote server
 - in \$HOME (e.g. using pip or miniconda)
 - system-wide (e.g. with Easybuild, Spark) by the admins.

Browser-only installation (**experimental!**):

- **JupyterLab** local with server + client components in your browser
JupyterLite includes a browser-ready Python environment named Pyodide.
<https://blog.jupyter.org/jupyter-everywhere-f8151c2cc6e8>
<https://jupyter.org/try-jupyter/lab>



JUPYTERLAB - WHEREVER YOU PREFER

Local, Remote, Browser-only

Local installation:

- **JupyterLab** installed using conda, mamba, pip, pipenv or docker.
https://jupyterlab.readthedocs.io/en/stable/getting_started/installation.html

Remote (cluster) installation:

- **JupyterLab** installed on a remote server
 - in \$HOME (e.g. using pip or miniconda)
 - system-wide (e.g. with Easybuild, Spark) by the admins.

Browser-only installation (**experimental!**):

- **JupyterLab** local with server + client components in your browser
JupyterLite includes a browser-ready Python environment named Pyodide.
<https://blog.jupyter.org/jupyter-everywhere-f8151c2cc6e8>
<https://jupyter.org/try-jupyter/lab>



Tunnel the new JupyterLab to your local machine

Linux or Mac:
If your operating system is Linux or Mac use:

```
ssh -N -L <LOCAL_PORT>:<JLAB_NODE>:<JLAB_PORT> <USERID>@<LOGIN_NODE>.fz-juelich.de  
# example: ssh -N -L 8888:juwels04:8888 goebbert1@juwels01.fz-juelich.de  
  
# if you want to tunnel to juwels04 only, then you should set JLAB_NODE to "localhost"
```

Attention:

- LOGIN_NODE - Hostname of login node from the view of your local machine
- JLAB_NODE - Hostname of the node running JupyterLab from the view of LOGIN_NODE
- LOCAL_PORT - port on your local machine
- JLAB_PORT - port on the node running JupyterLab

Windows: In case your operating system is Windows, the setup of the tunnel depends on your ssh client. Here a short overview on how-to setup a tunnel with **PuTTY** is given.

It is assumed that PuTTY is already configured in a way that a general ssh connection to JUWELS is possible. That means that host name, user name and the private ssh key (using PuTTY's Pageant) are correctly set. You already made a first connection to JUWELS using PuTTY.

To establish the ssh tunnel start PuTTY and enter the "SSH-->tunnels" tab in the PuTTY configuration window before connecting to JUWELS. You have to enter the source port (eg. <LOCAL_PORT> = 8888) and the destination (eg. juwels01.fz-juelich.de:8888) and then press add. After pressing add, the tunnel should appear in the list of forwarded ports and you can establish the tunnel by pressing the open button.

JUPYTERLAB - WHEREVER YOU PREFER

Local, Remote, Browser-only

Local installation:

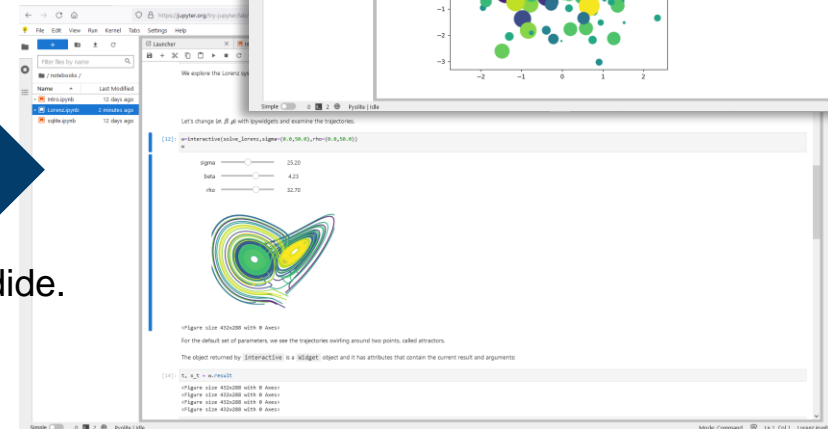
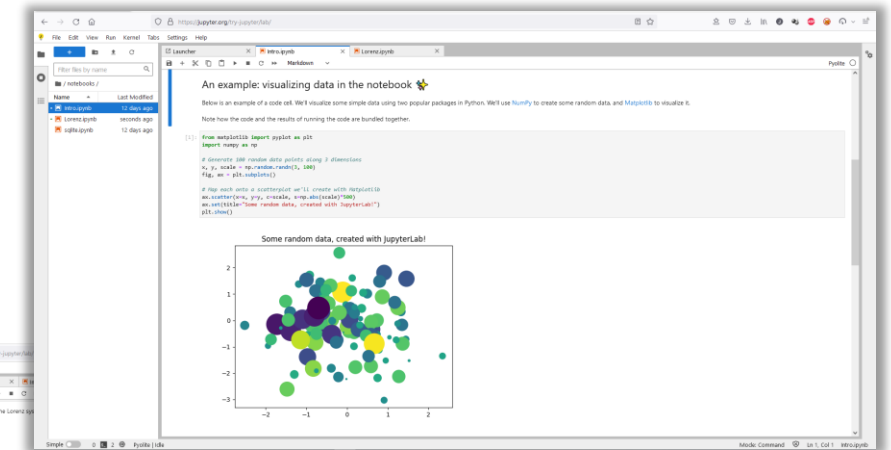
- **JupyterLab** installed using conda, mamba, pip, pipenv or docker.
https://jupyterlab.readthedocs.io/en/stable/getting_started/installation.html

Remote (cluster) installation:

- **JupyterLab** installed on a remote server
 - in \$HOME (e.g. using pip or miniconda)
 - system-wide (e.g. with Easybuild, Spark) by the admins.

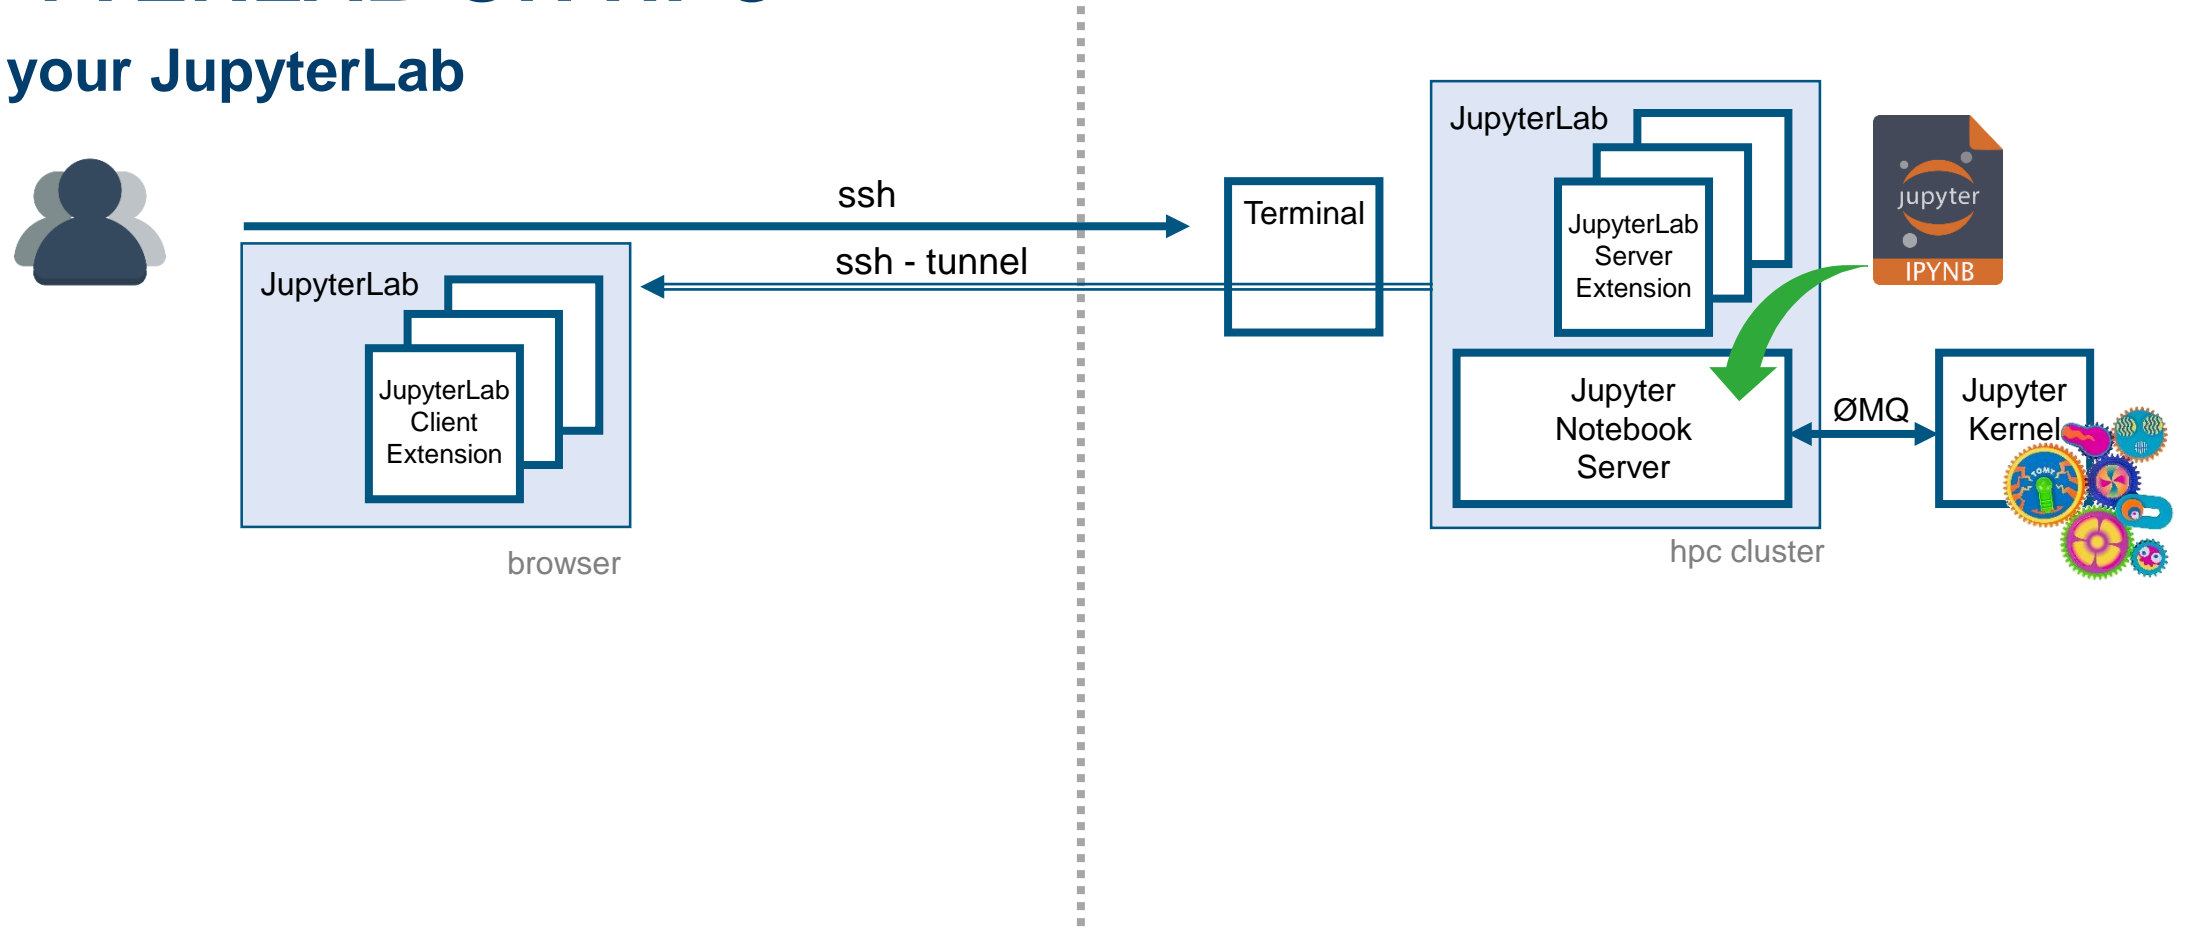
Browser-only installation (**experimental!**):

- **JupyterLab** local with server + client components in your browser
JupyterLab includes a browser-ready Python environment named Pyodide.
<https://blog.jupyter.org/jupyter-everywhere-f8151c2cc6e8>
<https://jupyter.org/try-jupyter/lab>



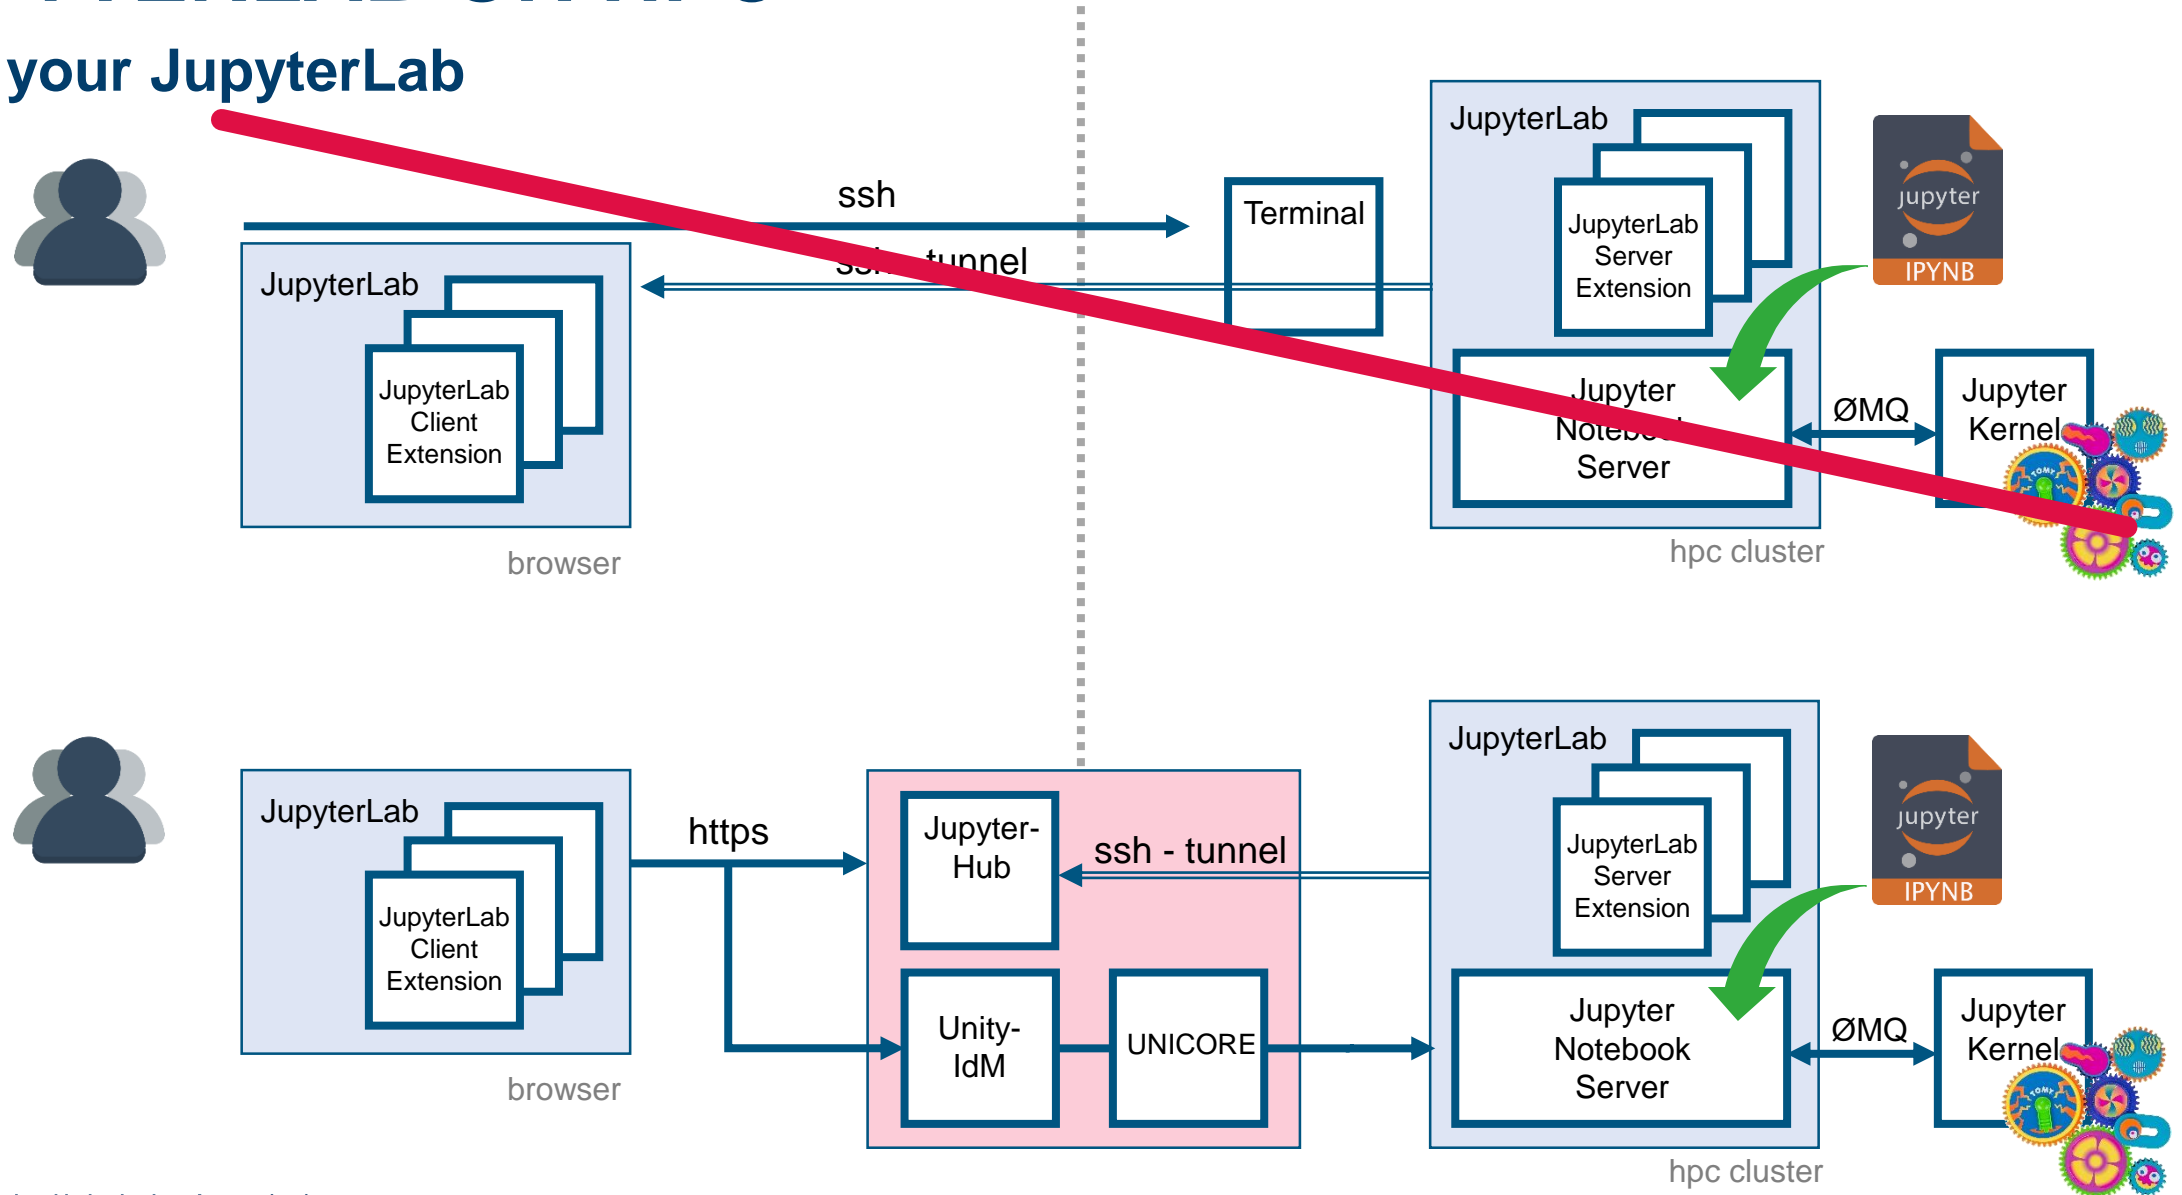
JUPYTERLAB ON HPC

Start your JupyterLab



JUPYTERLAB ON HPC

Start your JupyterLab



JUPYTER-JSC WEBSERVICE

Prepare

1) Register & Login

- ✓ <https://judoor.fz-juelich.de>

2) Join a project

- ✓ Wait to get joined by the project PI

3) Sign usage agreement

- ✓ Wait for creation of HPC accounts

4) Check Connected Services:

- ✓ jupyter-jsc

The screenshot displays the Jülich Supercomputing Centre (JSC) user interface. At the top, there is a navigation bar with links for 'Your account', 'Mentoring', 'Search', and 'Detailed Statistics'. The main content area is divided into several sections:

- Account:** A sidebar on the left lists fields for 'Salutation', 'E-mail address', 'Telephone', and 'Address'. The 'E-mail address' field has a blue checkmark icon.
- Mentored projects:** A section with a button labeled 'Mentored projects'.
- Systems:** A table listing available systems:
 - judac:** Labeled 'training2109' with a green checkmark. A link 'Manage SSH-keys' is present. A note states 'Usage agreement confirmed on 18.04.2021' with a green checkmark.
 - jureca:** Labeled 'JURECA-DC_GPU: training 2211' with a red X. A note states 'You need to sign the usage agreement to access this system' with a red X.
- Projects:** A section with a button 'Join a project' and a project entry: 'Interactive High-Performance Computing with Jupyter @ JSC' with a green checkmark and the label 'training2211'.
- Software:** A section with a button 'Join a project'.
- Connected Services:** A section showing a list of services: 'trac', 'llview', 'jards', 'gitlab', and 'jupyter-jsc'. The 'jupyter-jsc' service has a green checkmark.

JUPYTER-JSC WEBSERVICE

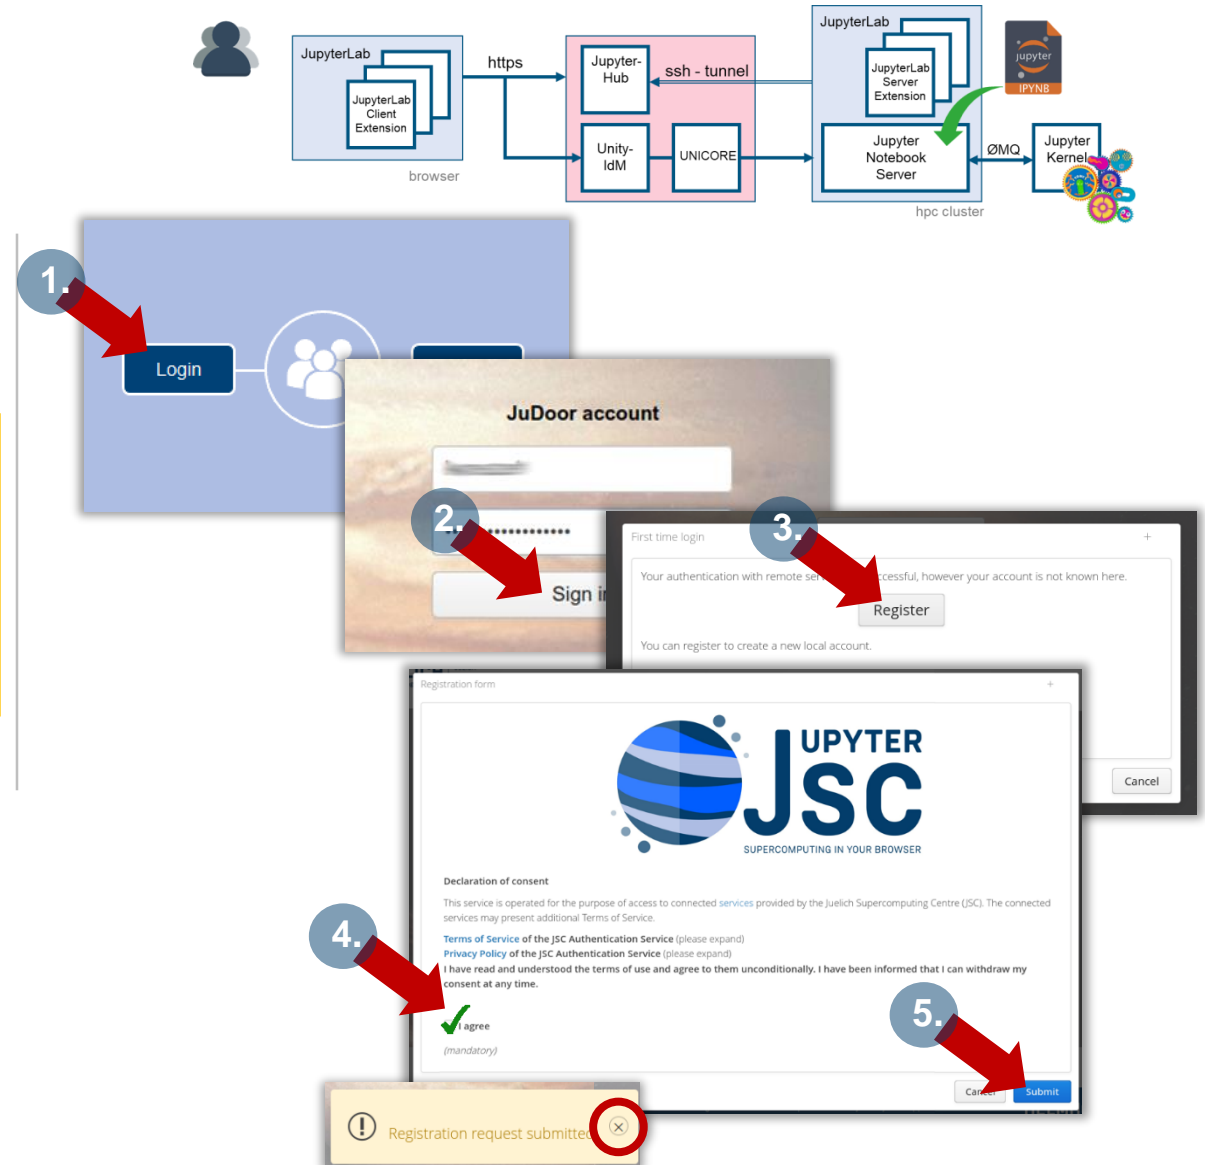
First time login

=> <https://jupyter-jsc.fz-juelich.de>

Jupyter-JSC first time login

- Requirements:
 - Registered at judoor.fz-juelich.de
 - (check "Connected Services" = jupyter-jsc)
 - Project membership + signed systems usage agreement
 - Waited ~10 minutes

1. Login at <https://jupyter-jsc.fz-juelich.de>
2. Sign in with your JSC account
3. Register to Jupyter-JSC
4. Accept usage agreement
5. Submit the registration
6. Wait for email and confirm your email address



JUPYTER-JSC WEBSERVICE

Start your JupyterLab

JUPYTER-JSC WEBSERVICE

Your server is starting up...

You will be redirected automatically when it's ready for you.

Name	System	Partition	Project	Status	Actions
juwels_cluster	JUWELS	devel	ccsys	70%	Cancel

JupyterLab + New

You can configure your existing JupyterLab's by expanding the corresponding table row.

Name	System	Partition	Project	Status	Actions
hdfcloud_3.3	HDF-Cloud	N/A	N/A		Start
juwelsbooster_login	JUWELS	LoginNodeBooster	ccsys		Start
juwels_cluster	JUWELS	devel	ccsys	30%	Open Cancel

Supercomputing in Your Browser

We are pleased to bring "Supercomputing in your browser". Jupyter-JSC is designed to provide you with the rich high performance computing (HPC) ecosystem to the world's most popular software: web browsers. JupyterLab is a web-based interactive development environment for Jupyter notebooks, code, and data. JupyterLab is flexible to support a wide range of workflows in data science, scientific computing, and machine learning. Read more.

Please use your JSC account to log in or register if you have not already done so. It's also possible to log in via Helmholtz AAI.

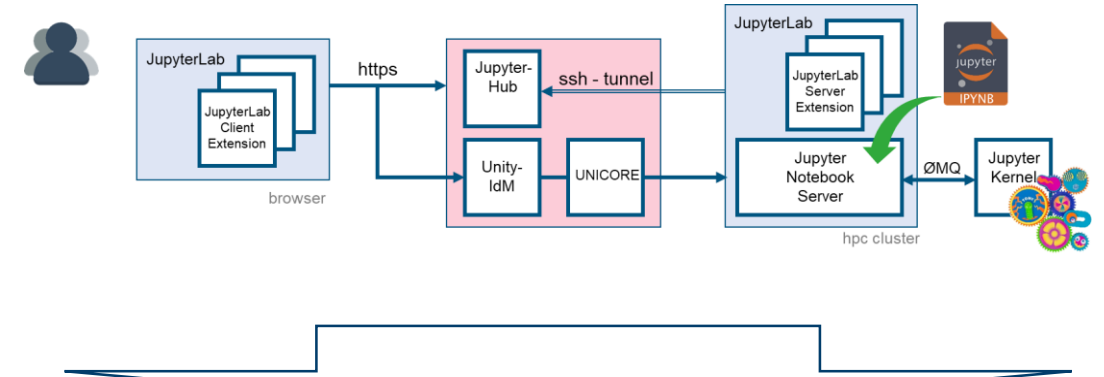
[Login](#) [Register](#)

JUPYTER-JSC

JUWELS **JURECA** **JUSUF** **DEEP** **HDFML** **HDF-Cloud**

© Forschungszentrum Jülich | [Imprint](#) | [Privacy Policy](#) | [Support](#) | [Terms of Service](#)

HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES



GPU DASHBOARDS

- GPU Utilization
- GPU Memory
- GPU Resources
- PCIe Throughput
- WLink Throughput
- WLink Timeline
- Matching Resources

```
1 import math
2 import numpy as np
3 from numba import cuda
4 import matplotlib.pyplot as plt
5 matplotlib.interactive = True

[1]: len(cuda.gpus)

[2]: 4

[3]: len(cuda.gpus[0].name)

[4]: len(cuda.gpus[0].name)

[5]: b'Tesla V100-SXM2-16GB'
```

GPU Memory

GPU Memory: 332.40 MB

Variables

- math: module
- np: module
- cuda: module
- plt: module
- matplotlib: module
- numba.cuda.compiler.Dispatcher: module
- size: 400
- iterations: 100
- my_numpy_array: array

Breakpoints

- amp/jupyterlab_30146/402220956.py: 4

JUPYTER-JSC WEBSERVICE

Control Panel

A. Jupyter-JSC – Add new JupyterLab

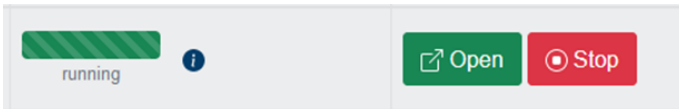


B. Configuration Dialog


- set Name, Type, System, Account, Project, Partition

C. Jupyter-JSC – Actions

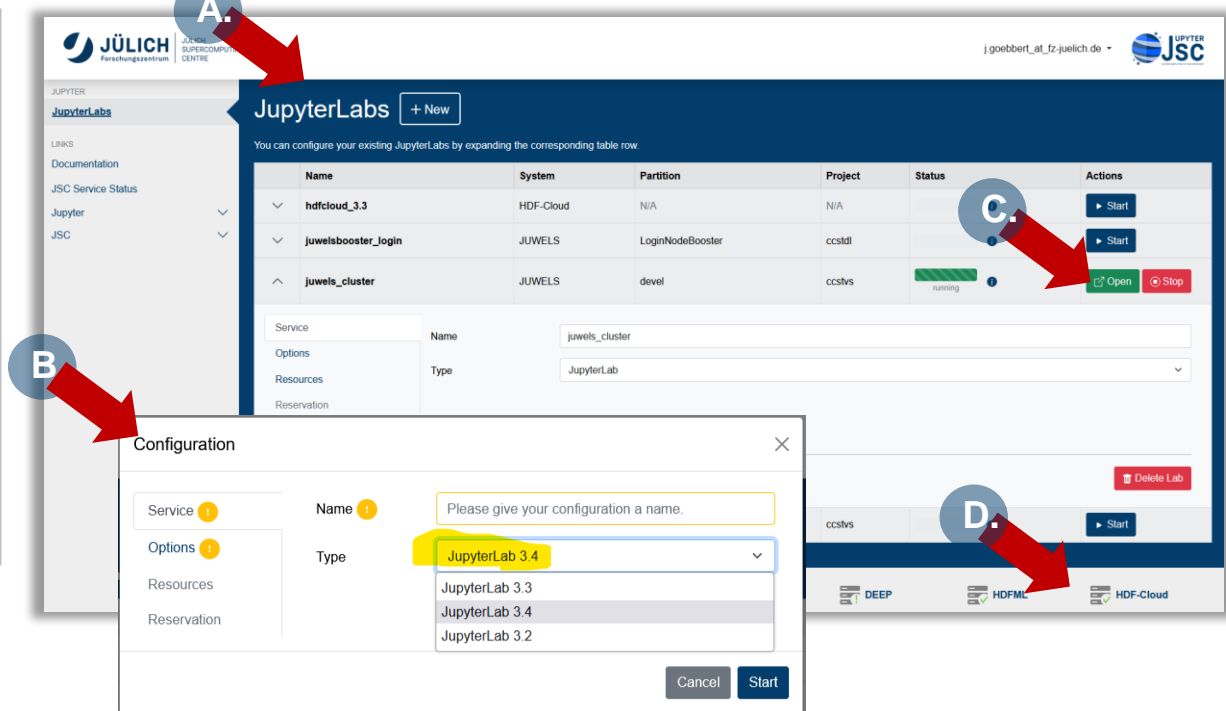


- Open/Stop a running JupyterLab
- Change/Delete **configuration**



D. Jupyter-JSC -- Statusbar




- Upcoming maintenance (mouse hover for details)
- System offline



E. Jupyter-JSC – Logout

Logout will ask what you want to do with the running JupyterLabs – be careful what you answer!



JUPYTER-JSC WEBSERVICE

JupyterLab Configuration

Jupyter-JSC – Configuration

Available options **depend on**


- user account settings visible in judoor.fz-juelich.de
- system specific usage agreement on JuDoor is signed (!!!)
- currently available systems in all of your projects

Basic options

- Type:
multiple versions of JupyterLab are installed
- System:
JUWELS, JURECA, JUSUF, DEEP, HDFML, HDF-Cloud
- Account:
In general users only have a single account
- Project:
project which have access to the selected system
- Partition:
partition which are accessible by the project
(this includes the decision for LoginNode and ComputeNode)

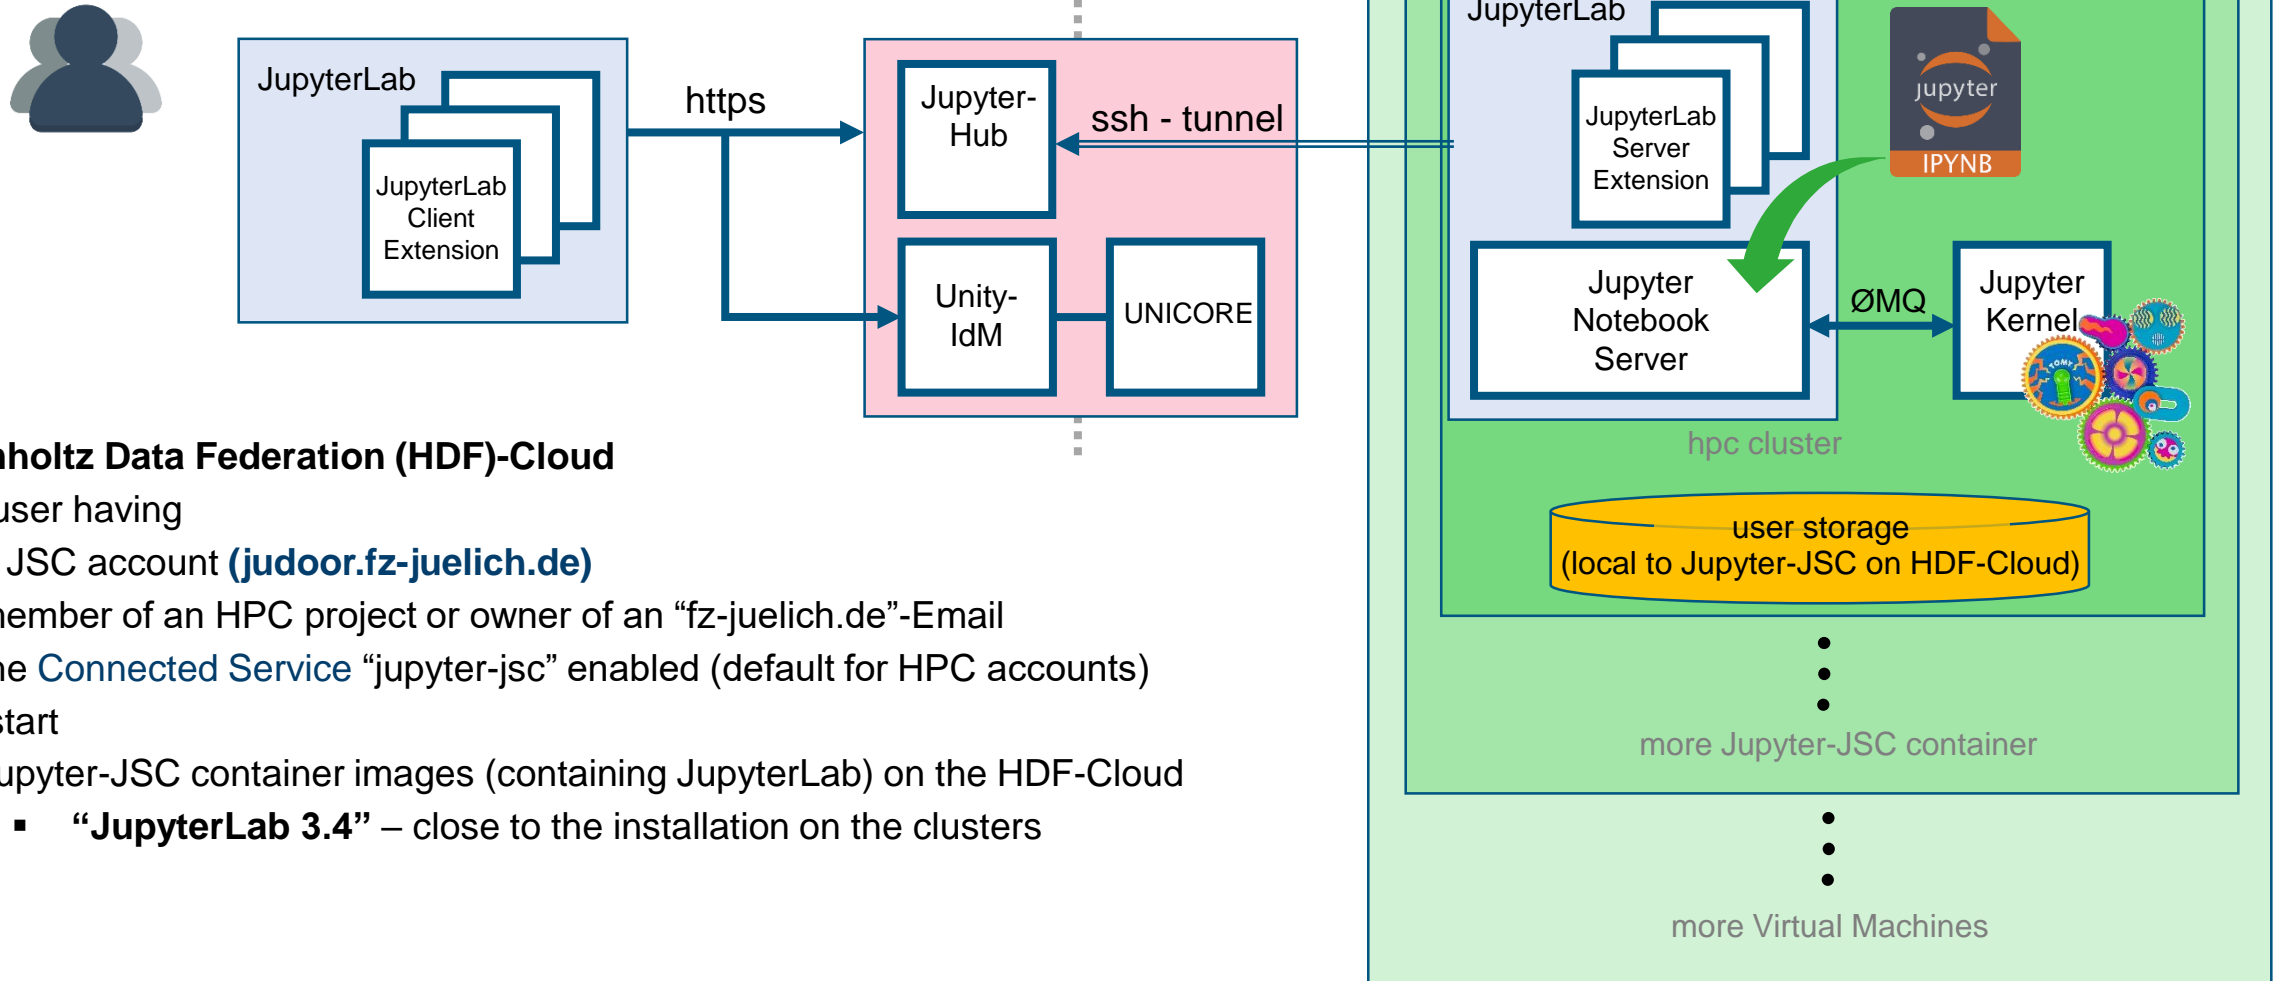
Extra options

- Partition == compute Nodes, Runtime, GPUs, ...



JUPYTER-JSC WEBSERVICE

System: HDF-Cloud

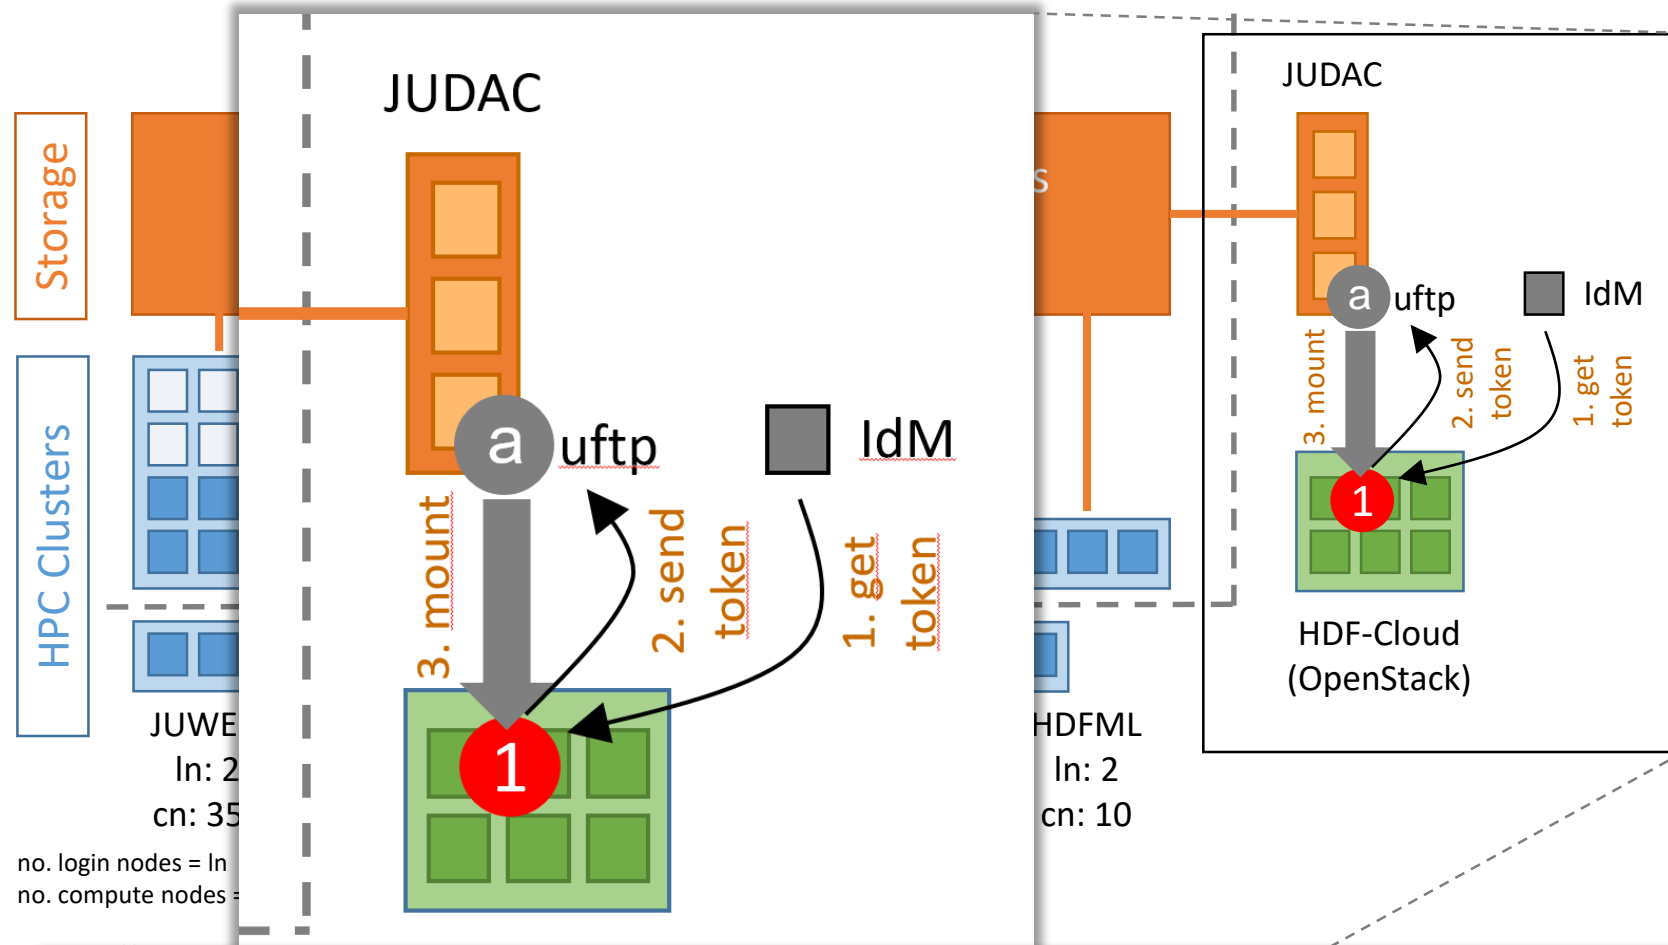


Helmholtz Data Federation (HDF)-Cloud

Any user having

- a JSC account (judoor.fz-juelich.de)
 - member of an HPC project or owner of an "fz-juelich.de"-Email
 - the [Connected Service](#) "jupyter-jsc" enabled (default for HPC accounts)
- can start
- Jupyter-JSC container images (containing JupyterLab) on the HDF-Cloud
 - **"JupyterLab 3.4"** – close to the installation on the clusters

HOW TO MOUNT GPFS ON HDF-CLOUD



[1] https://gitlab.jsc.fz-juelich.de/jupyter4jsc/prace-2022.04-jupyter4hpc/-/blob/main/day_2/2_hpc-environment/1-hdf-cloud_mount-hpc-storage.ipynb

[2]

[3]

[4]


[5] https://www.jsc.fz-juelich.de/infrastructure/configuration_nodeadm

JUPYTER-JSC SECRETS

Very important to know



Secret 1: Support button

- Let us know, if something does not work.
We can only fix it, if we know it.



Secret 2: Reload on connection loss


- “Server Not Running”
means, that your browser just lost connection
=> **Just hit “Dismiss” !!!**
(as soon as you are online again)
- “File Save Error for <...>”
means, that your browser just lost connection
=> **Just hit “Dismiss” !!!**
(as soon as you are online again)



You can **always** safely hit the “Reload” button of your browser, if the connection to JupyterLab ever gets lost.
(it will just restart JupyterLab on the browser-site)

JUPYTER-JSC SECRETS

For experts only 😊




Secret 3: Jupyter-JSC logs

- Jupyter-Lab gets started by UNICORE on our HPC systems
- On startup UNICORE created the directory `$SCRATCH_<project>/unicore-jobs/<random-hash>/`
 - In the terminal of a running JupyterLab, this directory is `$JUPYTER_LOG_DIR`
- In this directory you find
 - `stdout` -> terminal output of jupyterlab messages
 - `stderr` -> terminal output of jupyterlab error messages
 - `.start` -> details how your JupyterLab got started

Secret 4: change to a different JupyterLab version

- In `.start` you can see, that
 - `$HOME/.jupyter/start_jupyter-jsc.sh` is used to prepare the environment for JupyterLab. This script must ensure that the command `jupyter` is available in `$PATH`.



```
#!/bin/bash
```


```
module purge  
module load Stages/2022  
module load GCCcore/.11.2.0  
module load JupyterCollection/2022.3.4
```

Switch to a customized JupyterLab with
`$HOME/.jupyter/start_jupyter-jsc.sh`

It enables you to switch to an older/newer/other version of JupyterLab, if the default one gives you trouble or is missing features.

JUPYTER-JSC WEBSERVICE

Some comments about the UI



Annotations on the JupyterLab interface:

- open filebrowser**: Points to the file browser icon in the sidebar.
- open launcher**: Points to the launcher icon in the sidebar.
- tutorials & examples**: Points to the tutorials and examples icon in the sidebar.
- sidebar with core and extensions features**: Points to the sidebar area.
- indicates active notebook cell**: Points to the active notebook cell in the main area.
- type of active notebook cell**: Points to the cell type dropdown menu.
- no close, but go back to Jupyter-JSC's control panel**: Points to the 'Control Panel' button in the top right.
- memory consumption (keep an eye on that!)**: Points to the memory usage indicator in the top right.
- Type of Jupyter kernel this notebook is connected to (click to change)**: Points to the kernel name dropdown menu.
- notebook cell**: Points to a code cell in the main area.

Code cell content:

```
[*]: # INPUT NEEDED:
KERNEL_NAME=${USER}_kernel

export KERNEL_NAME=$(echo "${KERNEL_NAME}" | awk '{print tolower($0)}')
echo ${KERNEL_NAME} # double check

• change if you like

• List directories where JupyterLab will search for kernels

JUPYTER_SEARCH_PATH (for kernels-directory)
Jupyter search paths for kernels-directories"
if [ -d $JUPYTER_PATH ]; then
echo "$HOME/.local/share/uvotter"
```


JUPYTERLAB EXTENSIONS

JUPYTERLAB EXTENSIONS


Installed by default at Jupyter-JSC

IPyVolume

3d plotting for Python in the Jupyter notebook based on IPython widgets using WebGL




<https://github.com/maartenbreddels/ipyvolume>



JupyterLab-Git

JupyterLab extension for version control using Git



<https://github.com/jupyterlab/jupyterlab-git>

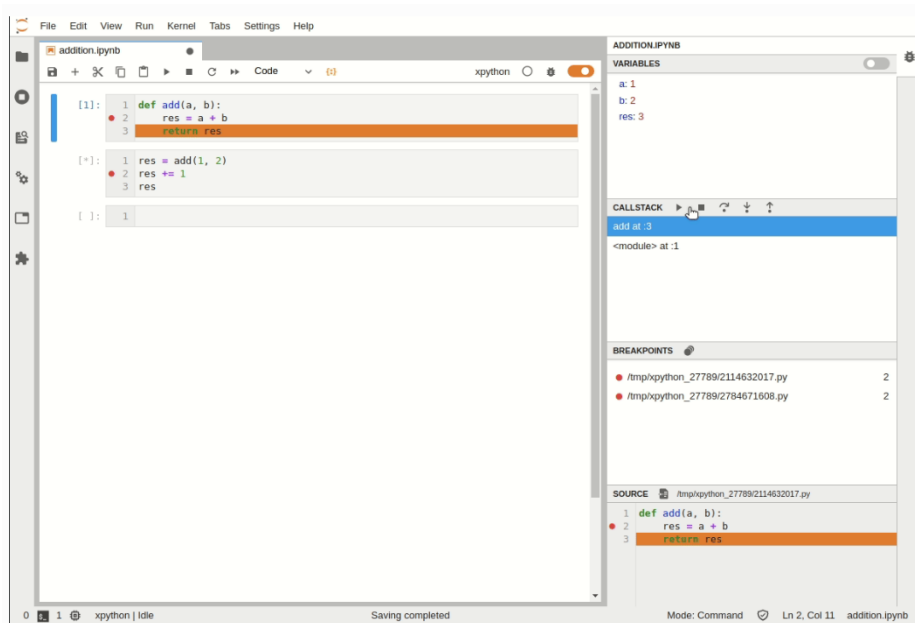
JUPYTERLAB EXTENSIONS

Installed by default at Jupyter-JSC


JupyterLab - Visual Debugger

JupyterLab >= 3 ships with a Debugger front-end by default.

This means that notebooks, code consoles and files can now be debugged from JupyterLab directly! For the debugger to be enabled and visible, a kernel with support for debugging is required.

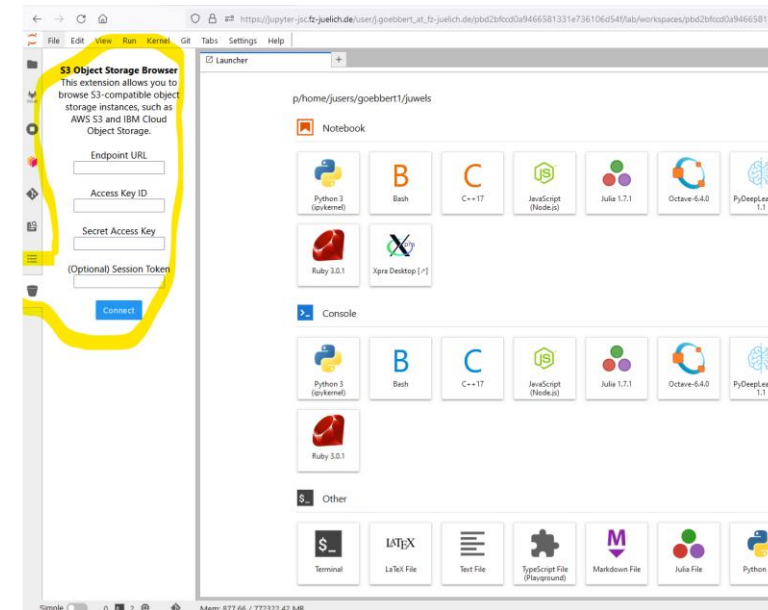


<https://jupyterlab.readthedocs.io/en/stable/user/debugger.html>



JupyterLab-S3-browser

A JupyterLab extension for browsing S3-compatible object storage



<https://github.com/IBM/jupyterlab-s3-browser>

JUPYTERLAB EXTENSIONS

Installed by default at Jupyter-JSC

PyThreeJS

A Python / ThreeJS bridge utilizing the Jupyter widget infrastructure.
<https://threejs.org> - lightweight, 3D library with a default WebGL renderer.

```
In [9]: f = """
function f(origu,origv) {
  // scale u and v to the ranges I want: [0, 2*pi]
  var u = 2*Math.PI*origu;
  var v = 2*Math.PI*origv;

  var x = Math.sin(u);
  var y = Math.cos(v);
  var z = Math.cos(u*v);

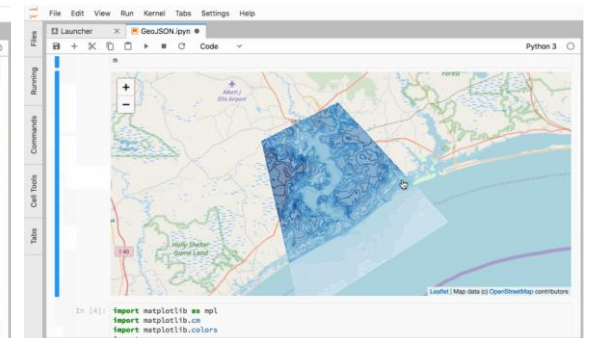
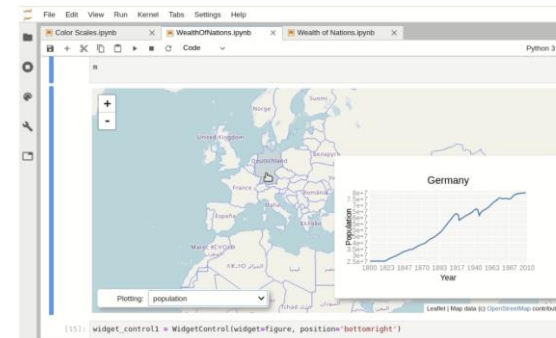
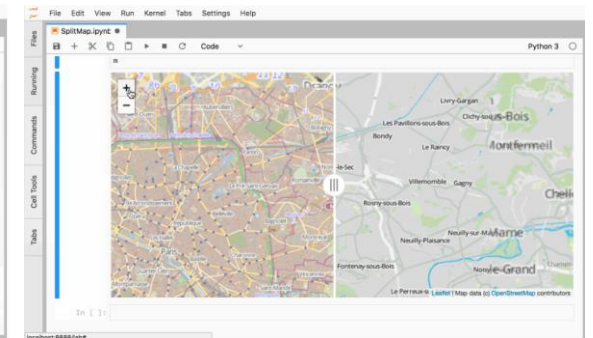
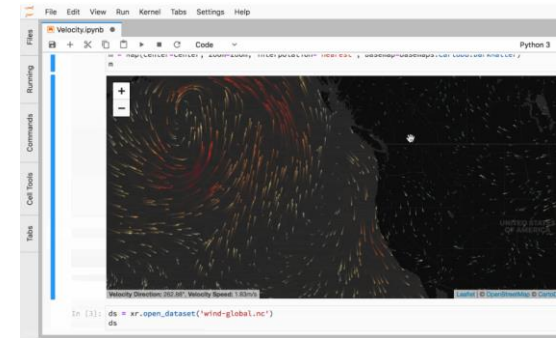
  return new THREE.Vector3(x,y,z)
}
"""
surf_g = ParametricGeometry(func=f);
surf = Mesh(geometry=surf_g, material=LambertMaterial(color='green', side='FrontSide'))
surf2 = Mesh(geometry=surf_g, material=LambertMaterial(color='yellow', side='BackSide'))
scene = Scene(children=[surf, surf2, AmbientLight(color='#777777')])
c = PerspectiveCamera(position=[5, 5, 3], up=[0, 0, 1],
                      children=[DirectionalLight(color='white',
                                                  position=[3, 5, 1],
                                                  intensity=0.6)])
renderer = Renderer(camera=c, scene=scene, controls=[OrbitControls(controlling=c)])
display(renderer)
```



<https://github.com/jupyter-widgets/pythreejs>

IPyLeaflet

A Jupyter / Leaflet bridge enabling interactive maps in the Jupyter notebook.



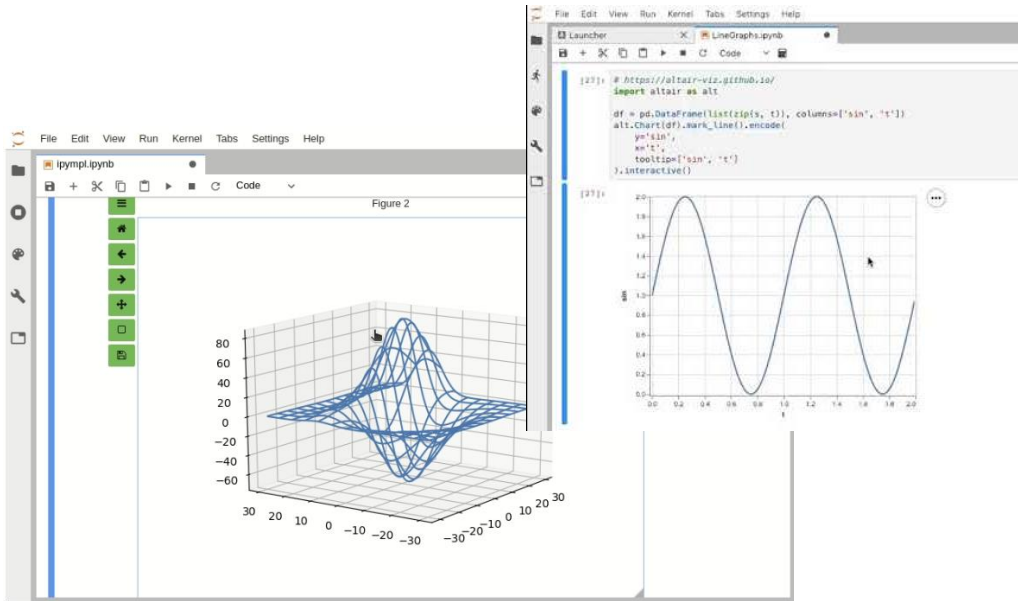
<https://github.com/jupyter-widgets/ipyleaflet>

JUPYTERLAB EXTENSIONS

Installed by default at Jupyter-JSC

IPyMPL - matplotlib

Leveraging the Jupyter interactive widgets framework, ipympl enables the interactive features of matplotlib in the Jupyter notebook and in JupyterLab.

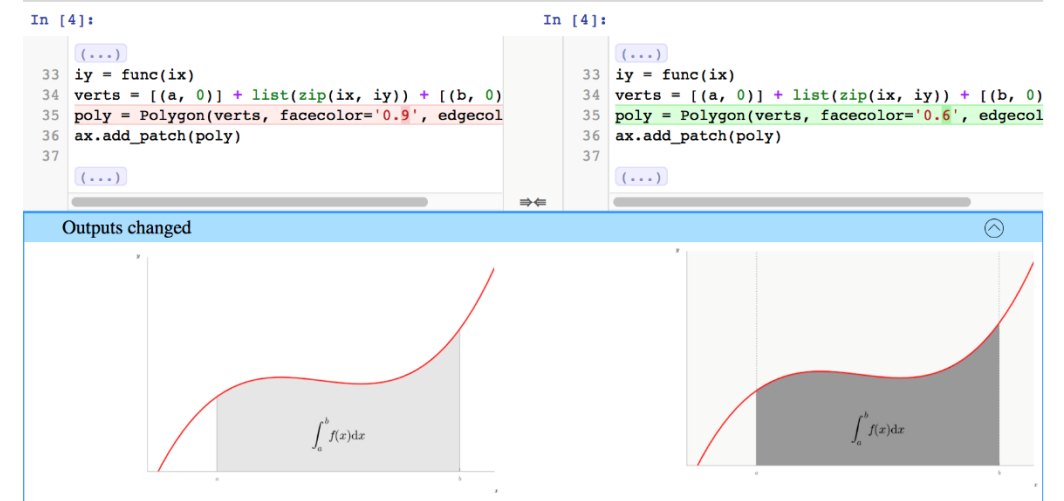


<https://github.com/matplotlib/ipympl>

Member of the Helmholtz Association

NBDime

Tools for diffing and merging of Jupyter notebooks.



<https://github.com/jupyter/nbdime>

JUPYTERLAB EXTENSIONS

Installed by default at Jupyter-JSC

Plotly

JupyterLab extension for the interactive and browser-based graphing library Plotly.

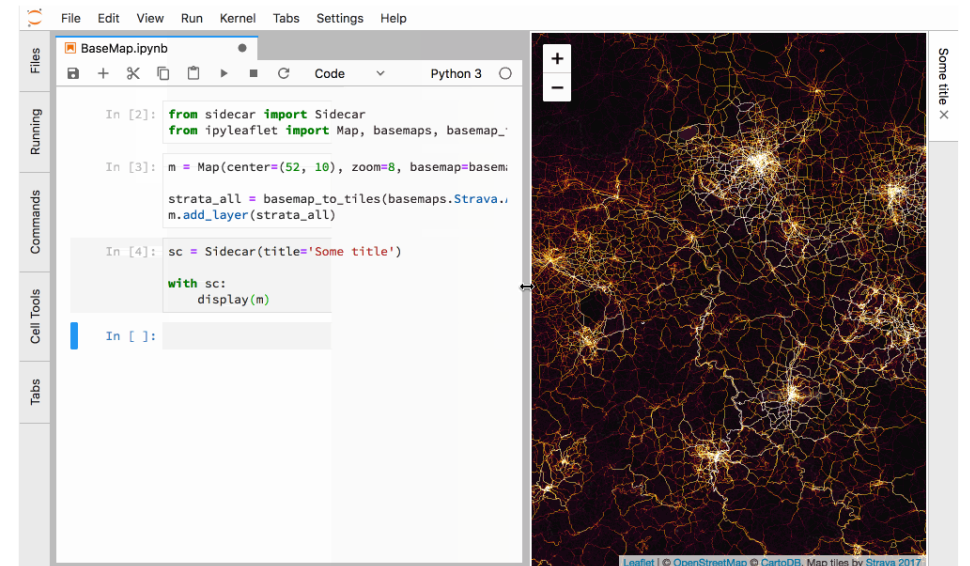
<https://plotly.com/python/>



<https://github.com/plotly/plotly.py>

JupyterLab-Sidecar

A sidecar output widget for JupyterLab.



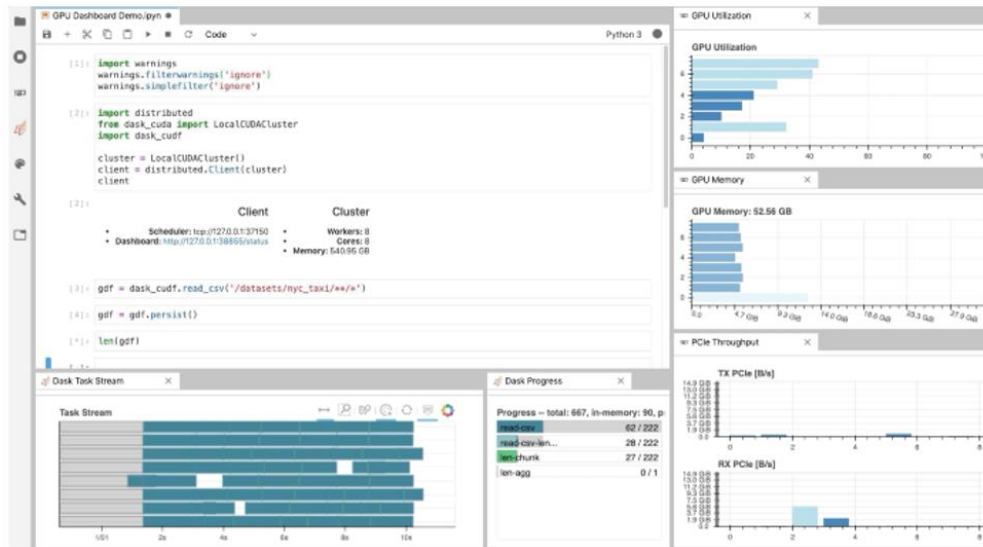
<https://github.com/jupyter-widgets/jupyterlab-sidecar>

JUPYTERLAB EXTENSIONS


Installed by default at Jupyter-JSC

NVDashboard

NVDashboard is an open-source package for the real-time visualization of NVIDIA GPU metrics in interactive Jupyter Lab environments.




<https://github.com/rapidsai/jupyterlab-nvdashboard>
<https://developer.nvidia.com/blog/gpu-dashboards-in-jupyter-lab/>



Voilà

Voilà turns Jupyter notebooks into standalone web applications.




<https://github.com/voila-dashboards/voila>

JUPYTERLAB EXTENSIONS

Installed by default at Jupyter-JSC

Extensions	old version	new version	type
Core			
@jupyterlab/server-proxy	v2.1.0	v3.1.0	prebuild
@jupyter-widgets/jupyterlab-manager	v2.0.0	v3.0.1	prebuild
jupyterlab-datawidgets	v6.3.0	v7.0.0	source
UI Enhancements			
@jlab-enhanced/recents		v3.0.1	prebuild
@jlab-enhanced/favorites	v2.0.0	v3.0.0	prebuild
jupyterlab-topbar-extension	v0.5.0	v0.6.1	
jupyterlab-system-monitor	v0.6.0	v0.8.0	prebuild
@jupyter-server/resource-usage		v0.6.0	n/a
jupyterlab-theme-toggle	v0.5.0	v0.6.1	source
jupyterlab-controlbtn	jupyterlab-control	v0.5.0	n/a
@jupyterlab/toc	v4.0.0	integrated into JupyterLab 3	
Developer Tools			
@jupyterlab/git	v0.23.3	v0.32.4	prebuild
jupyterlab-gitlab	v2.0.0	v3.0.0	prebuild
@krassowski/jupyterlab-lsp	v2.1.3	v3.9.0	prebuild
nbdime-jupyterlab	v2.1.0	v3.1.0	prebuild
@ryantam626/jupyterlab_code_formatter	v1.3.8	v1.4.10	prebuild
@jimbarr/jupyterlab_spellchecker	v0.2.0	v0.7.2	prebuild
jupyterlab-nvddashboard		v0.6.0	prebuild



Data Visualization

jupyter-matplotlib	v0.7.4	v0.9.0	prebuild
@bokeh/jupyter_bokeh	v2.0.4	v3.0.4	prebuild
jupyterlab-plotly	v4.14.3	v5.3.1	
bqplot	v0.5.22	v0.5.32	prebuild
@pyviz/jupyterlab_pyviz	v1.0.4	v2.1.0	prebuild
jupyter-leaflet	v0.13.3	v0.14.0	prebuild
ipyvolume	v0.6.0-alpha.5	v0.6.0-alpha.8	prebuild
jupyter-threejs	v2.2.0	v2.3.0	prebuild
@jupyter-widgets/jupyterlab-sidecar	v0.5.0	v0.6.1	prebuild

Framework Integrations

dask-labextension	v3.0.0	v5.1.0	prebuild
@jupyterlab/latex	v2.0.1	v3.1.0	prebuild
jupyter-webrtc	v0.5.0	v0.6.0	prebuild

Dashboard Development

jupyter-vue	v1.5.0	v1.6.1	
jupyter-vuetify	v1.6.1	v1.8.1	
@voila-dashboards/jupyterlab-preview	v1.1.0	v2.1.0-alpha.2	prebuild
jupyterlab-dash	v0.4.0	v0.4.0	prebuild

Welcome

jupyterlab_iframe	v0.3.0	v0.4.0	source
jupyterlab-tour		v3.1.3	prebuild

JUPYTER KERNEL

JUPYTER KERNEL

How to create your own Jupyter Kernel


Jupyter Kernel

A “kernel” refers to the separate process

with

Ju

-
-
-
-



You can easily **create your own kernel** which for example runs your specialized virtual Python environment.

<https://github.com/jupyter/jupyter/wiki/Jupyter-kernels>

JUPYTER KERNEL

How to create your own Jupyter Kernel



Jupyter Kernel

A “kernel” refers to the separate process which executes code cells within a Jupyter notebook.

Jupyter Kernel

- run code in different programming languages **and environments**.
- can be connected to a notebook (one at a time).
- communicates via ZeroMQ with the JupyterLab.
- Multiple **preinstalled** Jupyter Kernels can be found on our clusters
 - Python, R, Julia, Bash, C++, Ruby, JavaScript
 - Specialized kernels for visualization, quantum computing

You can easily **create your own kernel** which for example runs your specialized virtual Python environment.



<https://github.com/jupyter/jupyter/wiki/Jupyter-kernels>

JUPYTER KERNEL

How to create your own Jupyter Kernel


Jupyter Kernel

A “kernel” refers to the separate process which executes code cells within a Jupyter notebook.

Jupyter Kernel

- run code in different programming languages **and environments**.
- can be connected to a notebook (one at a time).
- communicates via ZeroMQ with the JupyterLab.
- Multiple **preinstalled** Jupyter Kernels can be found on our clusters
 - Python, R, Julia, Bash, C++, Ruby, JavaScript
 - Specialized kernels for visualization, quantum computing

You can easily **create your own kernel** which for example runs your specialized virtual Python environment.



Building your own Jupyter kernel is a three step process

1. Create/Pimp new **virtual Python environment**
`venv`
2. Create/Edit **launch script** for the Jupyter kernel
`kernel.sh`
3. Create/Edit Jupyter **kernel configuration**
`kernel.json`

<https://github.com/jupyter/jupyter/wiki/Jupyter-kernels>

JUPYTER KERNEL

How to create your own Jupyter Kernel

Jupyter Kernel

A “kernel” refers to the separate process which executes code cells within a Jupyter notebook.

Jupyter Kernel


- run code in different programming languages **and environments**.

https://gitlab.version.fz-juelich.de/jupyter4jsc/j4j_notebooks/-/blob/master/001-Jupyter/Create_JupyterKernel_general.ipynb

clusters

- Python, R, Julia, Bash, C++, Ruby, JavaScript
- Specialized kernels for visualization, quantum computing

You can easily **create your own kernel** which for example runs your specialized virtual Python environment.



Building your own Jupyter kernel is a three step process

1. Create/Pimp new **virtual Python environment**
venv

<https://github.com/jupyter/jupyter/wiki/Jupyter-kernels>

JUPYTER KERNEL

Run your Jupyter kernel configuration

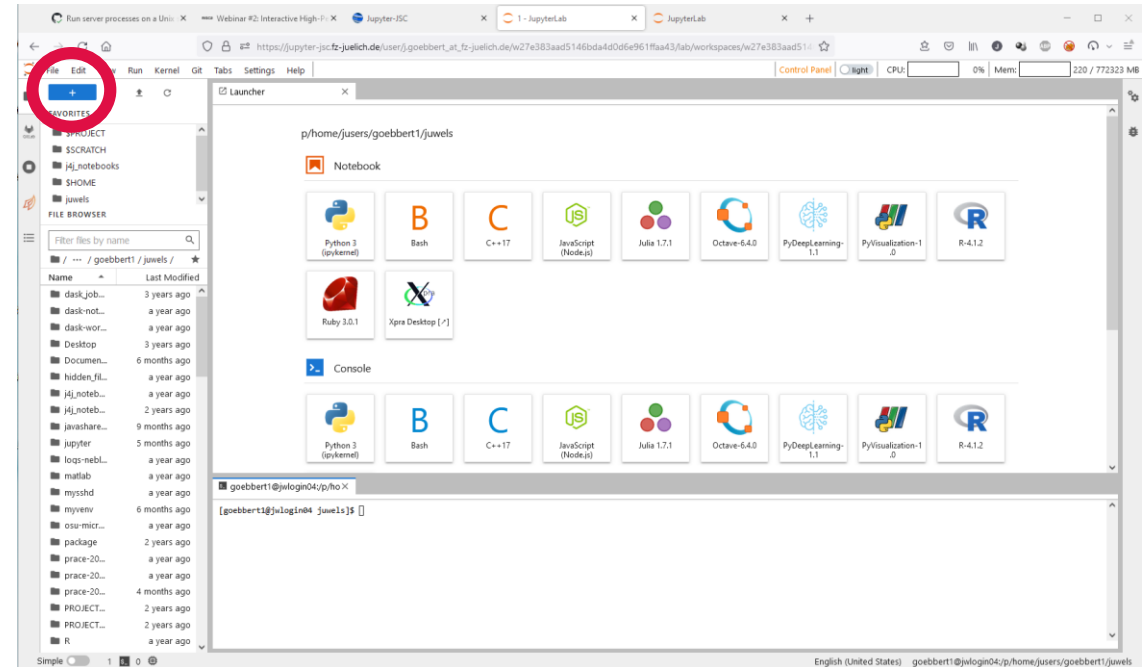

Run your Jupyter Kernel

1. <https://jupyter-jsc.fz-juelich.de>
2. Choose system where your Jupyter kernel is installed in `~/.local/share/jupyter/kernels`
3. Select your kernel in the launch pad or click the kernel name.

Conda

How to base your Jupyter Kernel on a Conda environment:

https://gitlab.version.fz-juelich.de/jupyter4jsc/j4j_notebooks/-/blob/master/001-Jupyter/Create_JupyterKernel_conda.ipynb



https://gitlab.version.fz-juelich.de/jupyter4jsc/j4j_notebooks/-/blob/master/001-Jupyter/Create_JupyterKernel_general.ipynb

Member of the Helmholtz Association

JUPYTERLAB – REMOTE DESKTOP

Run your X11-Applications in the browser

Jupyter-JSC gives you easy access to a remote desktop

1. <https://jupyter-jsc.fz-juelich.de>
2. Click on “Xpra”

Xpra - X Persistent Remote Applications

is a tool which runs X clients on a remote host and directs their display to the local machine.

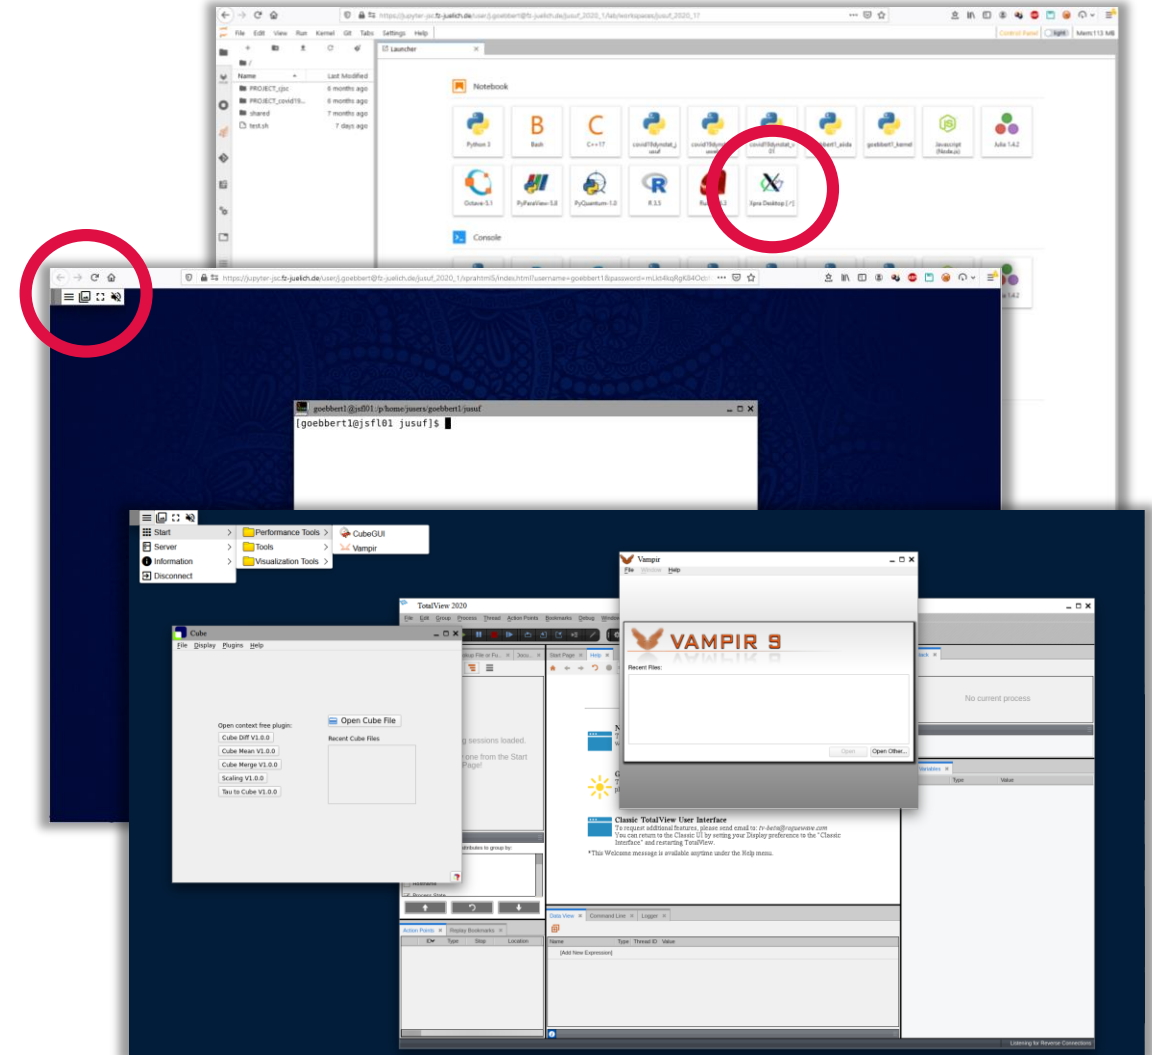
- Runs in a browser
- allows dis-/reconnection without disrupting the forwarded application
- <https://xpra.org>

The remote desktop will run on the same node as your JupyterLab does (this includes compute nodes).

It gets killed, when you stop your JupyterLab session.

Hint:

- CTRL + C -> CTRL + Insert
- CTRL + V -> SHIFT + Insert



JUPYTERLAB – REMOTE DESKTOP

Run your X11-Applications in the browser

Jupyter-JSC gives you easy access to a remote desktop

1. <https://jupyter-jsc.fz-juelich.de>
2. Click on “Xpra”

Xpra - X Persistent Remote Applications

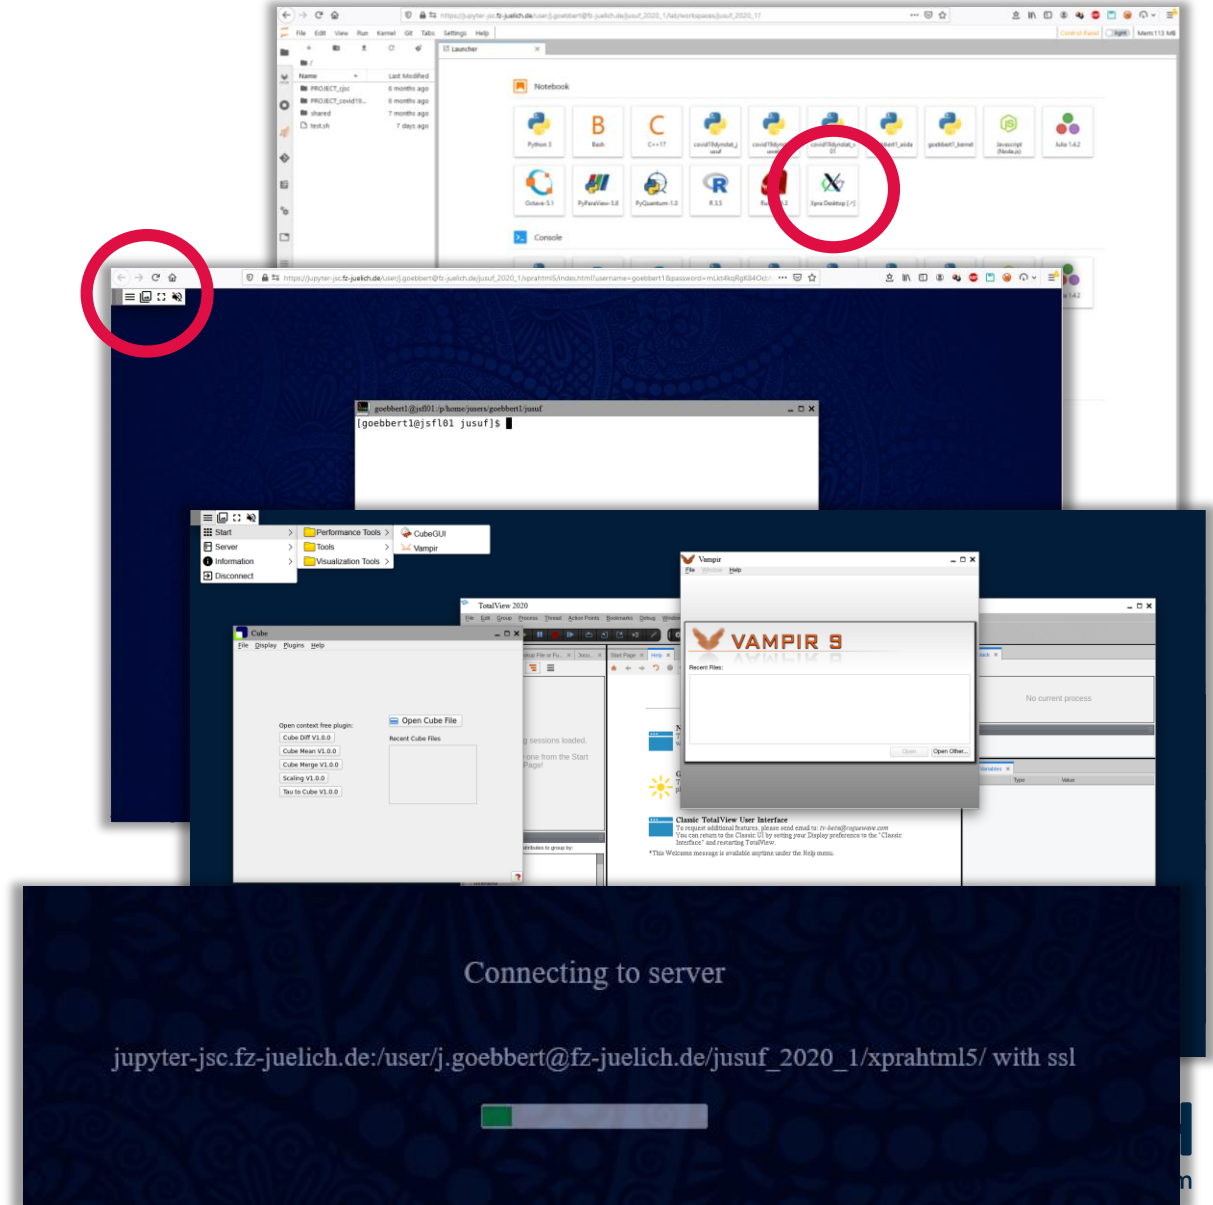
is a tool which runs X clients on a remote host and directs their display to the local machine.

- Runs in a browser
- allows dis-/reconnection without disrupting the forwarded application
- <https://xpra.org>

If the connection got lost at some point,
just hit the “reload” button of your browser.

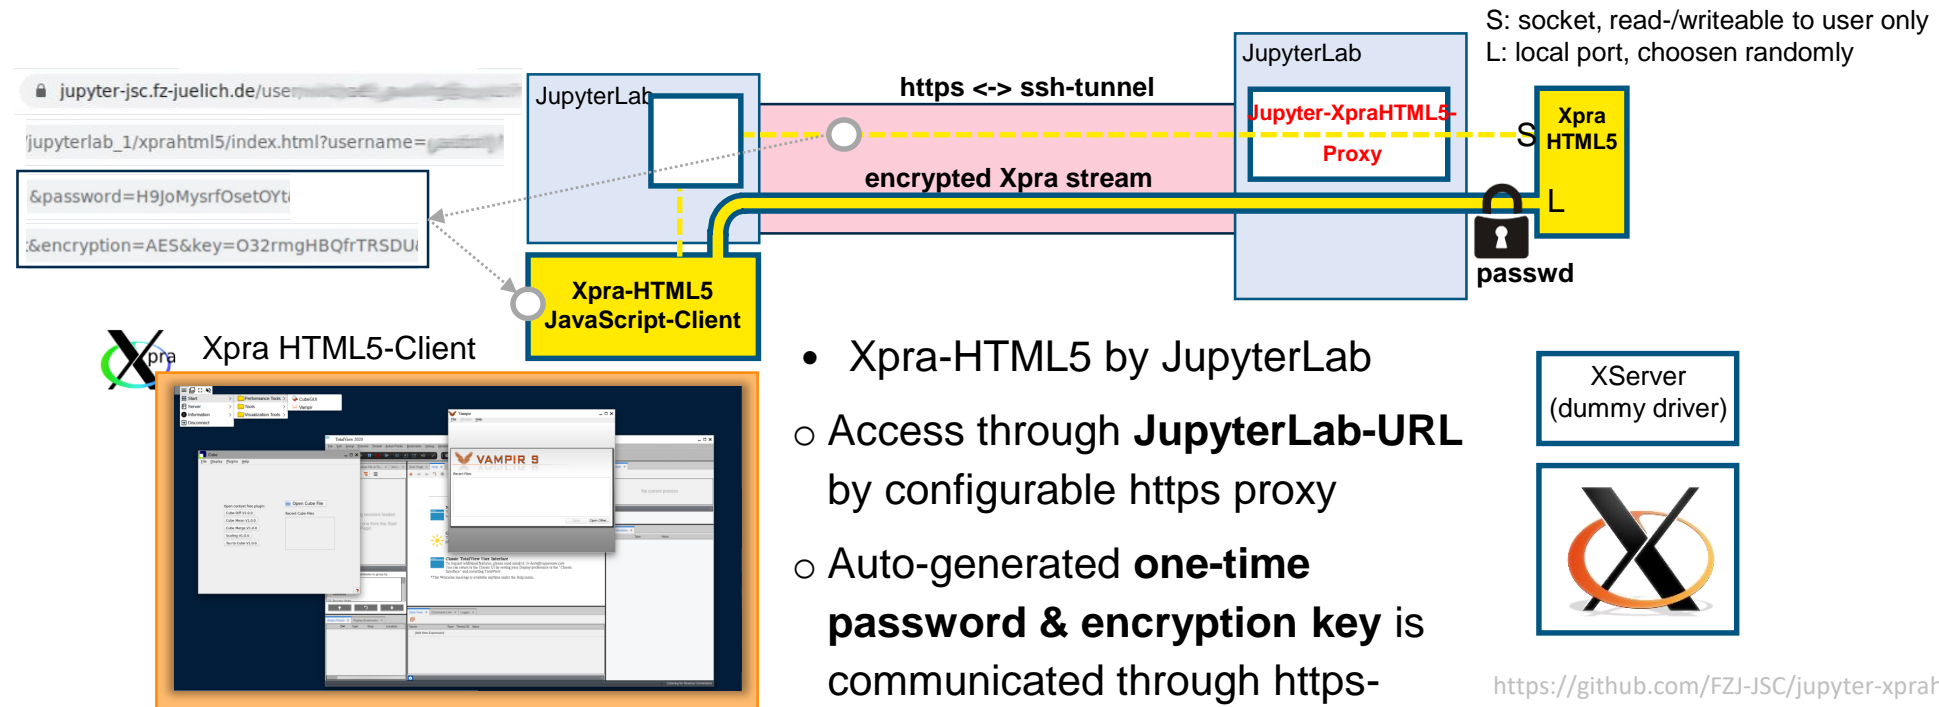
Hint:

- CTRL + C -> CTRL + Insert
- CTRL + V -> SHIFT + Insert



JUPYTERLAB – REMOTE DESKTOP

Run your X11-Applications in the browser



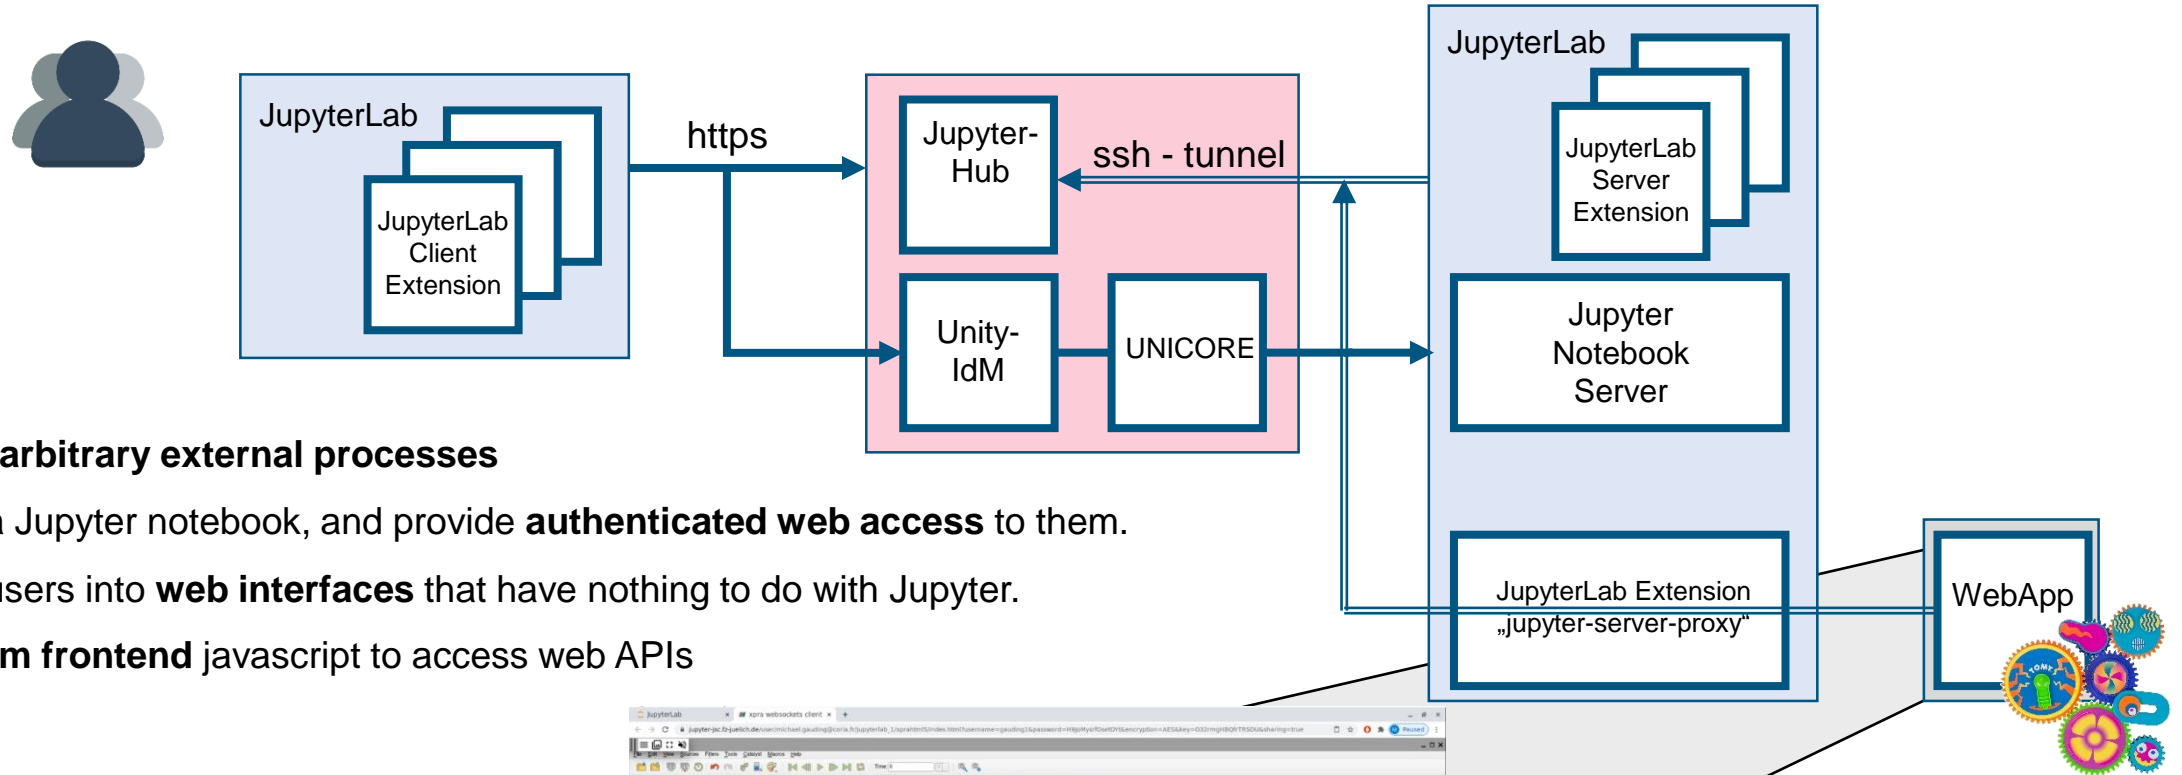
- Xpra-HTML5 by JupyterLab
 - Access through **JupyterLab-URL** by configurable https proxy
 - Auto-generated **one-time password & encryption key** is communicated through https-proxy

<https://github.com/FZJ-JSC/jupyter-xprahtml5-proxy>

JUPYTER CAN DO MORE


JUPYTERLAB – WEBSERVICE PROXY

Extension: jupyter-server-proxy



Allows to run **arbitrary external processes**

- alongside a Jupyter notebook, and provide **authenticated web access** to them.
- launching users into **web interfaces** that have nothing to do with Jupyter.
- **access from frontend javascript** to access web APIs



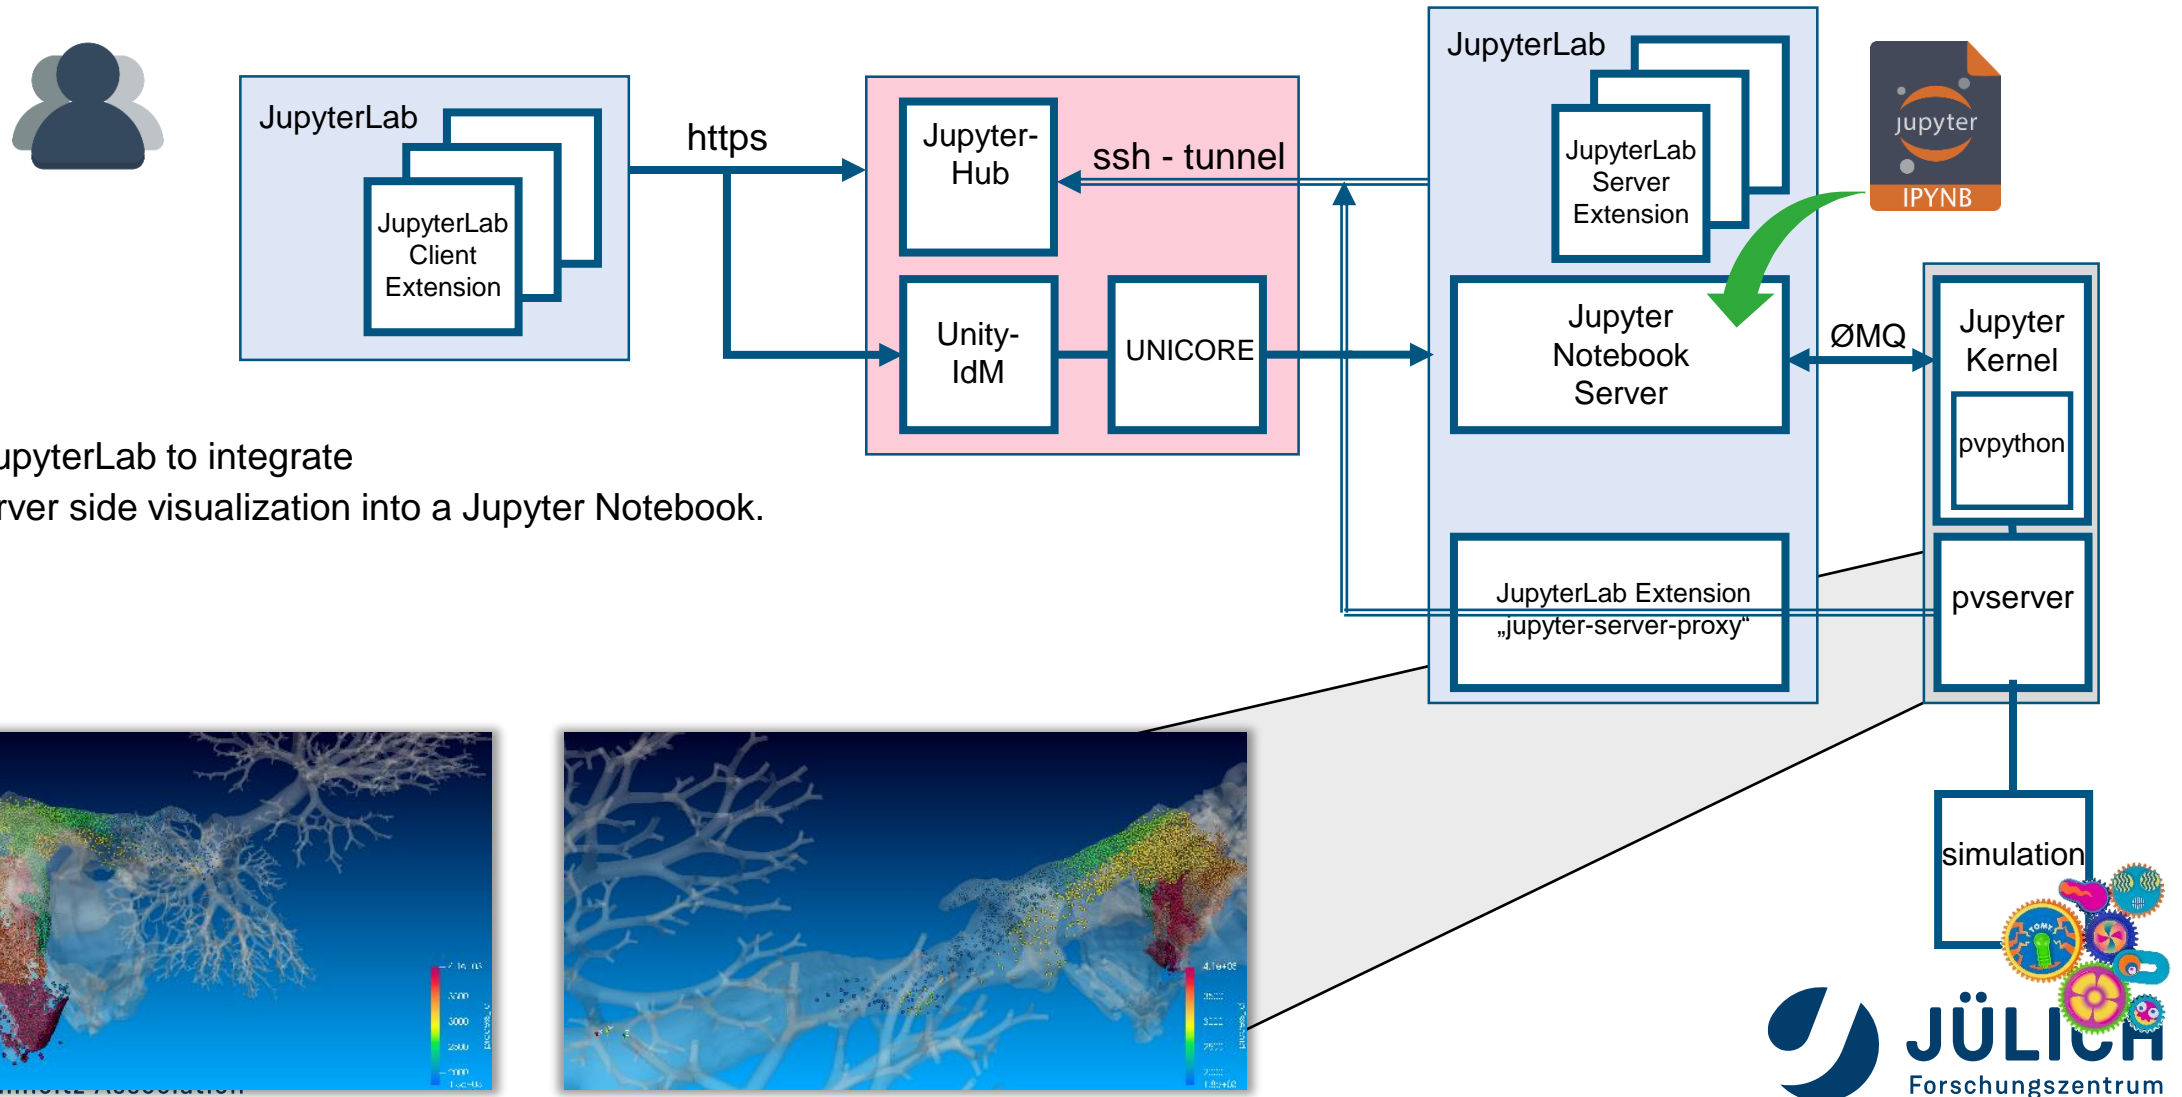
Turbulent mixing with variable density, subset of 1939x600x3584 grid points, Michael Gauding, CORIA

<https://github.com/jupyterhub/jupyter-server-proxy>

Member of the Helmholtz Association

JUPYTERLAB – WEBSERVICE PROXY

Extension: jupyter-server-proxy



How to use JupyterLab to integrate interactive server side visualization into a Jupyter Notebook.

JUPYTERLAB – WEBSERVICE PROXY

Extension: jupyter-server-proxy

Accessing Arbitrary Ports or Hosts

If you have a web-server running on the server listening on <port>, you can access it through the notebook at

<notebook-base>/proxy/<port>

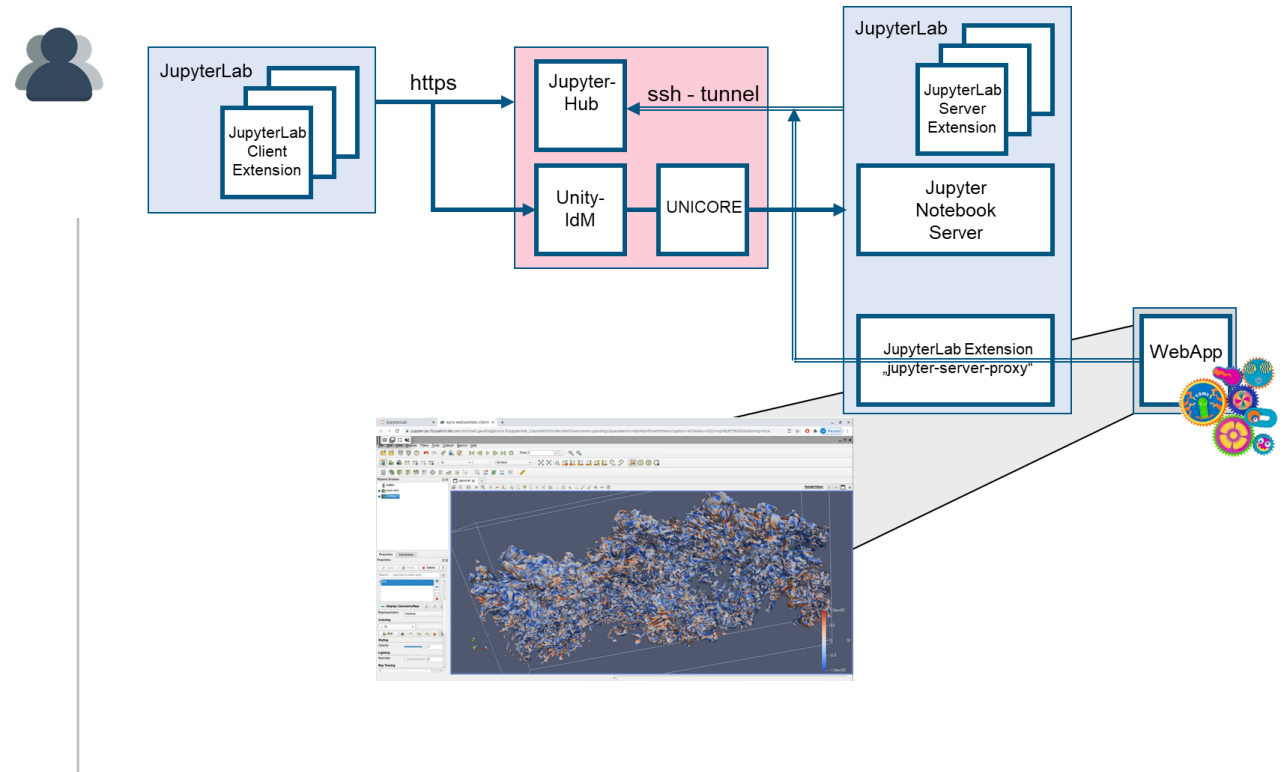
The URL will be rewritten to remove the above prefix.

You can disable URL rewriting by using

<notebook-base>/proxy/absolute/<port>

so your server will receive the full URL in the request.

This works for all ports listening on the local machine.




Example:


https://jupyter-jsc.fz-juelich.de/user/j.goebbert@fz-juelich.de/juwels_login/proxy/12345

DASHBOARDS WITH JUPYTER/VOILA

Voilà turns Jupyter notebooks into standalone web applications



- **Rendering** of live Jupyter notebooks with interactive widgets with the look-and-feel of a stand-alone web app.
- Voilà disallows execute requests from the front-end, **preventing** execution of arbitrary code.
- **Enables** HPC users to develop easily web applications from their Jupyter notebooks.



BENEFITS

Why Jupyter is so popular among Data Scientists

Some of the reasons ...

- Jupyter allows to **view the results of the code in-line** without the dependency of other parts of the code.
- Jupyter mixes easy for users who extend their code **line-by-line with feedback** attached all along the way
- Jupyter Notebooks support visualization and include rendering data in **live-graphics and charts**.
- Jupyter is maintaining the **state of execution of each cell** automatically.
- Supports IPyWidget packages, which provide **standard user interface** for exploring code and data interactively.
- Jupyter is platform and language **independent** with a well known file format in JSON.
- JupyterLab is build to be extentended, has a hudge community and a long history and is therefore future-proof.

TUTORIALS

Get started with Jupyter/JupyterLab – published 2018 but still a great starting point

Possible start to enter the world of interactive computing with IPython in Jupyter/JupyterLab:

- Leverage the Jupyter Notebook for interactive data science and visualization
- High-performance computing and visualization for data analysis and scientific modeling
- A comprehensive coverage of scientific computing through many hands-on, example-driven recipes with detailed, step-by-step explanations



<https://ipython-books.github.io>

<https://github.com/ipython-books/cookbook-2nd>

QUESTIONS?

