




# ENABLING TRANSFORMATIONAL SCIENCE

*PANGEO DEPLOYMENT AT NCI*

5<sup>th</sup> Feb 2021



# Special thanks to all supporters:



# Agenda

Time	topic
11:00-11:20	Overview of Pangeo Deployment at NCI ( <i>this presentation</i> )
11:20-11:40	Documentation and Live demo of setting up Pangeo ( <i>this presentation</i> )
11:40-12:00	Q&A and 10min break
12:00-12:10	Introduction to VDI and NCI-data-analysis environment
12:00-12:30	Overview of Xarray and Dask
12:30-12:45	Live demo
12:45-1:00	Q&A and wrap up the workshop



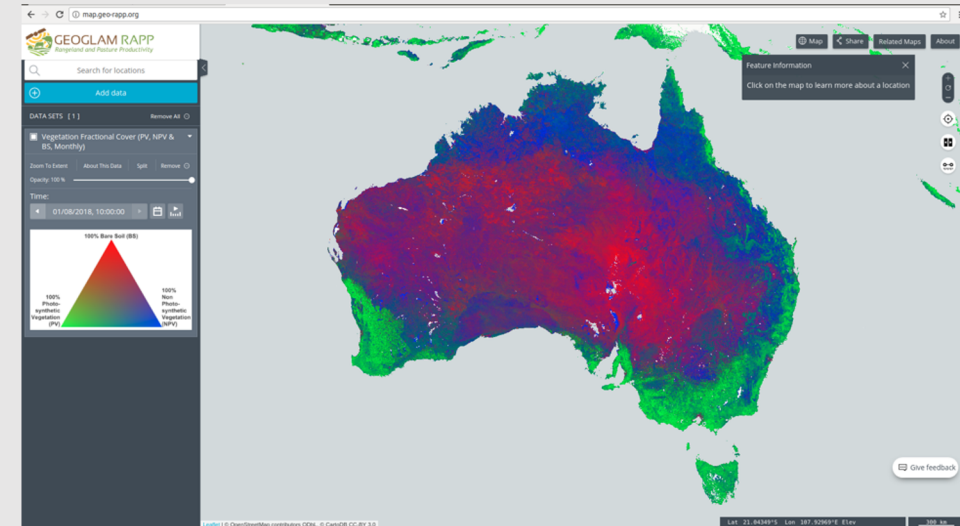


# WHAT DRIVES PROGRESS IN (GEO)SCIENCE?

$$\nabla \times \mathbf{e} = -\frac{\partial \mathbf{b}}{\partial t}$$
$$\nabla \times \mathbf{h} = \mathbf{j} + \frac{\partial \mathbf{d}}{\partial t}$$

Theory &  
Ideas

Observations  
/ Data

Simulations,  
Computation



(MODIS Fractional Cover loaded on the  
GEOGLAM RAPP map)

(Credit: Fernando Pérez)



# "THE GEARS OF THE ENGINE ARE STARTING TO GRIND"

Improved  
multiscale,  
nonlinear,  
noisy models

Exascale computing  
Cloud engineering  
skills  
Machine Learning  
pipelines

Theory &  
Ideas

Observations  
/ Data

Simulations,  
Computation


CMIP6 ~15-30 PB

Experimental Data  
rates of TB/day+

Heterogeneous,  
multimodal


(Credit: Fernando Pérez)

# ULTIMATE GOAL: REALLOCATE TIME!




(credit: Fernando Pérez)

## Traditional Way



## With data, software and compute available...



# PANGEO


Harnessing the power of cloud computing to study the whole Earth interactively




Interactivity



Distributed computing



Data models / numeric





# PANGEO DEPLOYMENT AT GADI: AN ACCELERATION TOOL



## Research use-cases

- Climate data analysis
- Weather data analysis
- Earth observation
- Geophysics
- Natural Hazards



## Tech developments

- Libraries and Tools
- Interactivity
- Cloud/HPC infrastructure
- Direct data access

# Pangeo – a big-data analysis platform at scale

Learning Goals for this presentation:

- Provide an overview of the Pangeo ecosystems
- Announce the data analysis environment
- Point to the documentation space
- Demonstrate data analysis examples

## What is Pangeo?

“Pangeo is first and foremost a *community* promoting open, reproducible, and scalable science.” (<http://pangeo.io>)



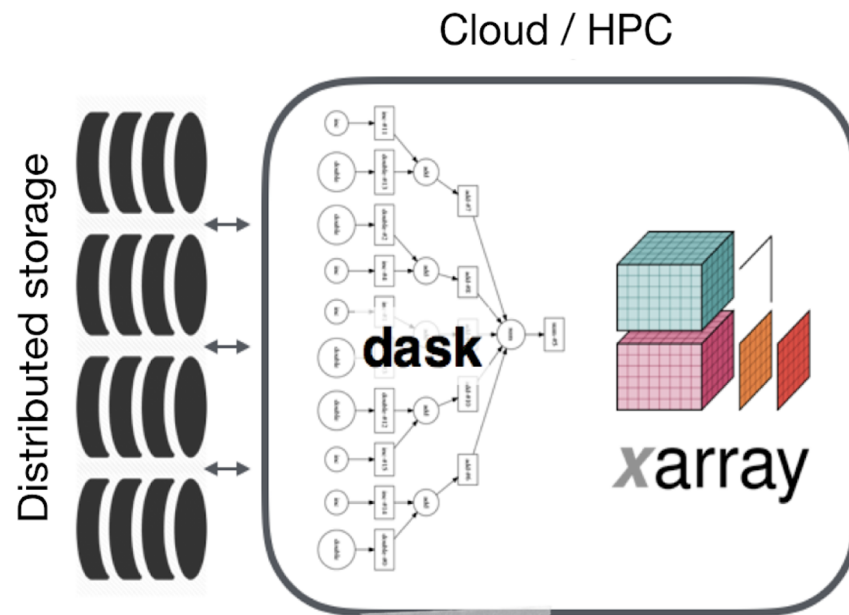
## Why do we need Pangeo?

Need better tools to support exploration of increasingly large data volumes



# PANGEO ARCHITECTURE

“Analysis Ready Data”  
stored on globally-available  
distributed storage.



Jupyter for interactive  
access remote systems

end user

web browser

Xarray provides data structures  
and intuitive interface for  
interacting with datasets

Parallel computing system allows  
users deploy clusters of compute  
nodes for data processing.

Dask tells the nodes what to do.


(<http://pangeo.io>)

## Step1: submit a job on Gadi to spin up a Jupyterlab server

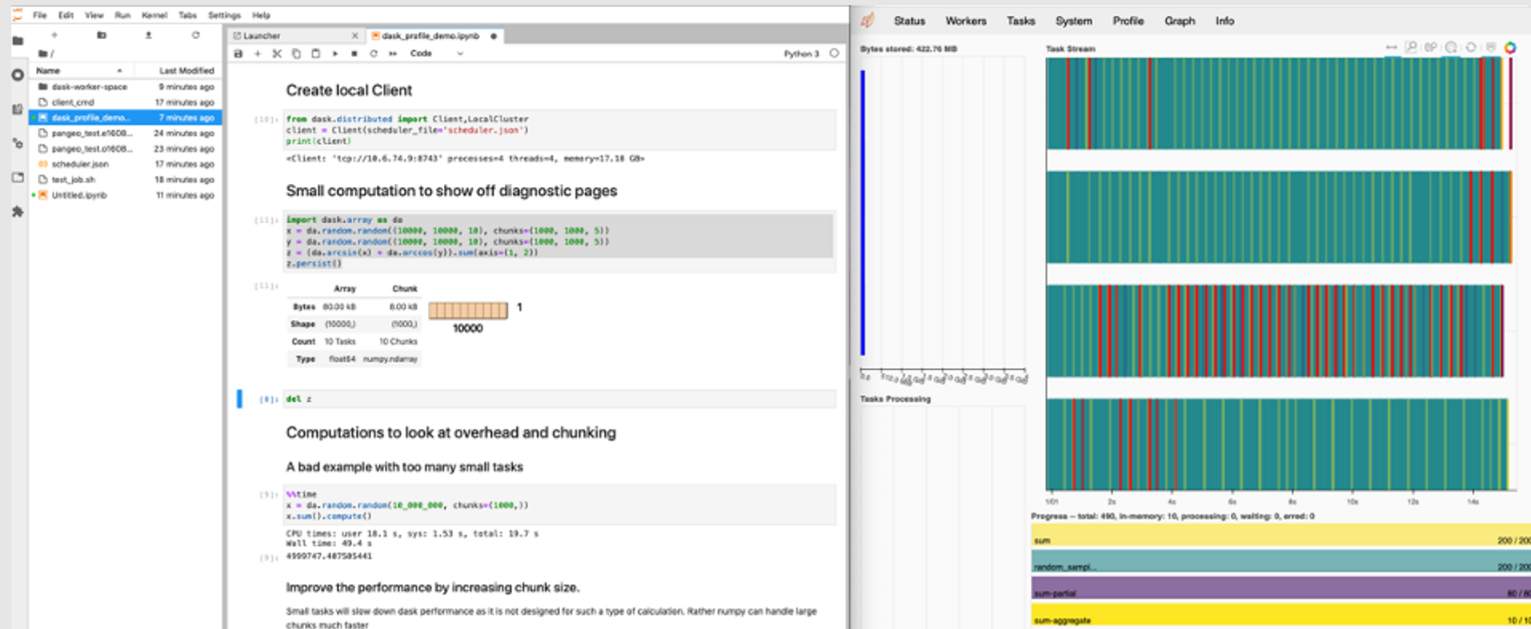
```
#!/bin/bash
#PBS -N pangeo_test
#PBS -P <project code>
#PBS -q normal
#PBS -l walltime=5:00:00
#PBS -l ncpus=96
#PBS -l mem=384GB
#PBS -l jobfs=100GB
#PBS -l storage=scratch/<project code>+gdata/<project code>
module use /g/data/dk92/apps/Modules/modulefiles
module load pangeo
module load NCI-data-analysis/2020.12
pangeo.ini.all.sh
sleep infinity
```

## Step2: ssh Gadi and open the ports on web browser (two tabs)

```
$ ssh -N -L 8388:gadi-cpu-clx-2224.gadi.nci.org.au:8388 abc123@gadi.nci.org.au &
```



## Step3: open the JupyterLab app and the dashboard



The image displays two side-by-side screenshots. The left screenshot shows a JupyterLab interface with a code editor containing Python code for creating a local client and performing computations. The code includes comments and function calls like `Client(LocalCluster)` and `dask.array`. The right screenshot shows the Dask dashboard, which includes a task stream visualization and progress bars for various tasks. The task stream shows a series of vertical bars representing tasks, and the progress bars show the status of different tasks, such as `sum`, `random_samp`, `sum-part`, and `sum-aggregate`.

## Scalable –

You can submit PBS jobs for the intensive compute parts of your workflow inside a notebook

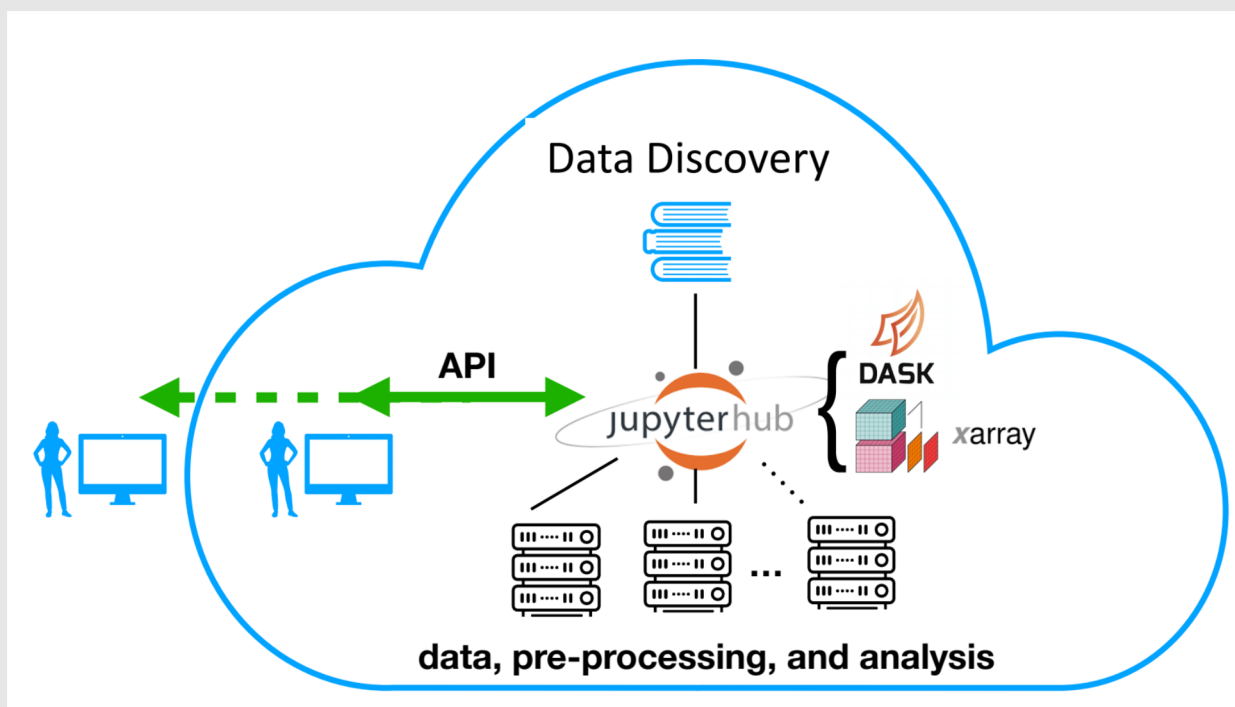
```
import dask.config
from dask.distributed import Client, LocalCluster
from dask_jobqueue import PBSCluster
walltime = '01:00:00'
cores = 96
memory = '160GB'

cluster = PBSCluster(walltime=str(walltime), cores=cores, memory=str(memory),
                    job_extra=['-P <project code>', '-l ncpus='+str(cores), '-l mem='+str(memory),
                               '-l storage=gdata/<project code>+gdata/<project code>+gdata/<project code>...'],
                    header_skip=["select"], python=os.environ["DASK_PBS_PYTHON"])
cluster.scale(jobs=2)
client = Client(cluster)
client
```



# Pangeo/Jupyter could be one of the essential interfaces to NCI's resources:

Move analysis to the data



**Instant access** to compute resources and data via PBS queueing system

**Democratize access** large computations accessed with **web browser**

**No downloading** data

**Scalable computational power** used and billed by time

**On-demand special resources (GPUs)**

**Reproducible workflows** thanks to network-accessible datasets and containerized software

## Things to be mindful of when using Pangeo

As the compute nodes do not allow internet access, libraries that download things on the fly won't work. For example, Cartopy might need to download a coastline database. In this case, you need to download the database onto Gadi, then point to the access point in your workflow.

Balance the resources requested for serial or parallel jobs. You can manage this by requesting minimum resource to get Pangeo started, then submit computationally intensive jobs inside your workflow to set up a Dask cluster as needed.

# Running notebooks on VDI vs using Pangeo on Gadi

Pangeo on Gadi can utilise Dask or xarray for parallel computing and data processing. However, this will likely require more resources such as CPU, memory and I/O throughput.

Note: Pangeo jobs should be submitted to the queue system and we recommend you request node-based resources.

## **Don't waste your SUs!**

The NCI-provided Virtual Desktop Infrastructure (VDI) is useful for:

- executing lightweight, serial/parallel Jupyter notebooks or Python scripts without consuming SUs
- developing scripts for the Pangeo environment

For more information on the VDI, see <https://opus.nci.org.au/display/Help/VDI+User+Guide>



# WHAT'S NEXT ?

- Deploy Pangeo on a Cloud infrastructure

# Live demo

- How to set up Pangeo on Gadi

<https://opus.nci.org.au/display/Help/5.1+Set+up+Pangeo>

## Step 1. Enabling Pangeo in your shell environment

To enable the Pangeo environment, you can use the following command within jobs, or within an interactive environment:

```
$ module load pangeo/<version>
```

```
Loading pangeo/2010.10
```

```
Loading requirement: intel-mkl/2019.3.199 python3/3.7.4 hdf5/1.10.5
```

```
netcdf/4.7.4
```

## Step 2. Configure your account on Gadi (once-only)

JupyterLab is bundled within the Pangeo environment and will be used to load notebooks and monitor jobs. To set up this environment, run the following two commands:

```
$ jupyter notebook --generate-config
```

```
$ jupyter notebook password
```

This is a stand-alone password that you will use later for accessing the JupyterLab server. It has nothing to do with your NCI login account and password. You can use this command to reset your password at any time.

## Step 3. Submitting a multi-node Pangeo job to Gadi

First create a directory where you will run the Jupyter notebook:

```
$ mkdir -p ~/pangeo/tutorial
```

Next you need to create a PBS shell script that will be used to launch a multi-node Pangeo job:

```
#!/bin/bash          ### see https://opus.nci.org.au/display/Help/PBS+Directives+Explained for explanation on PBS directives
#PBS -N pangeo_test
#PBS -P <project code>
#PBS -q normal      ### see https://opus.nci.org.au/display/Help/Queue+Limits for Gadi Queue Limits
#PBS -l walltime=5:00:00
#PBS -l ncpus=96
#PBS -l mem=384GB
#PBS -l jobfs=100GB
#PBS -l storage=scratch/<project code>+gdata/<project code>

module load pangeo
pangeo.ini.all.sh
sleep infinity
```

## Step 3. Submitting a multi-node Pangeo job to Gadi

```
#!/bin/bash
#PBS -N pangeo_test
#PBS -P <project code>
#PBS -q normal
#PBS -l walltime=5:00:00
#PBS -l ncpus=96
#PBS -l mem=384GB
#PBS -l jobfs=100GB
#PBS -l storage=scratch/<project code>+gdata/<project code>

module load pangeo
pangeo.ini.all.sh
sleep infinity
```

Note that for Gadi, the “#PBS -l storage=” flag is required which must include **all** the scratch and gdata project IDs that will be used in your codes.

Dask requires whole nodes so jobs beyond a single node must use multiples of 48 for their ncpus request.



## Step 3. Submitting a multi-node Pangeo job to Gadi

To submit the job to the queue:

```
$ qsub run_ipynb_job.sh
```

and take note of your `job_id` (which may look something like `1030967.gadi-pbs`).

Check to see when the job is running:

```
$ qstat <job_id>
```

## Step 3. Submitting a multi-node Pangeo job to Gadi

To ease the user experience of software installation, NCI provides a versioned virtual environment which includes a number of Python packages for data processing and visualisation. They are managed within project dk92. To use this virtual environment:

```
$ module use /g/data/dk92/apps/Modules/modulefiles
```

```
$ module load NCI-data-analysis/2020.12
```

A combination of Pangeo and NCI data-analysis virtual environments reduces the hassle of installing software and solving dependencies, and provides an interactive working environment for using JupyterLab. To use both Pangeo and this virtual environment, just add the above two lines to your PBS job submission script.

## Step 3. Submitting a multi-node Pangeo job to Gadi

```
#!/bin/bash
#PBS -N pangeo_test
#PBS -P <project code>
#PBS -q normal
#PBS -l walltime=5:00:00
#PBS -l ncpus=96
#PBS -l mem=384GB
#PBS -l jobfs=100GB
#PBS -l storage=scratch/<project code>+gdata/<project code>
```

```
module use /g/data/dk92/apps/Modules/modulefiles
module load pangeo
module load NCI-data-analysis/2020.12
pangeo.ini.all.sh
sleep infinity
```

## Step 4. Set up port forwarding to access from your desktop machine

Once the job is running on Gadi, there will be two files that are created in your current Gadi directory:

- `cliend_cmd`
- `scheduler.json`

The file `client_cmd` contains commands to forward network traffic from the defined port number of the worker node to an external client machine (i.e., your desktop) via the login node `gadi.nci.org.au`

## Step 4. Set up port forwarding to access from your desktop machine

Running:

```
$ more client_cmd
```

should give an output that looks something like:

```
ssh -N -L 8388:gadi-cpu-clx-2224.gadi.nci.org.au:8388 abc123@gadi.nci.org.au
```

```
ssh -N -L 8768:gadi-cpu-clx-2224.gadi.nci.org.au:8768 abc123@gadi.nci.org.au
```

For this example, JupyterLab uses port 8388 and Dask dashboard occupies port 8768 respectively from Gadi worker node gadi-cpu-clx-2224.

Note that both port numbers are randomly chosen for each job and will be different each time you run a new job.

## Step 4. Set up port forwarding to access from your desktop machine

Next, on your remote machine (e.g., Desktop), run each of the following commands separately if you have already set up SSH keys:

```
$ ssh -N -L 8388:gadi-cpu-clx-2224.gadi.nci.org.au:8388 abc123@gadi.nci.org.au &
```

```
$ ssh -N -L 8768:gadi-cpu-clx-2224.gadi.nci.org.au:8768 abc123@gadi.nci.org.au &
```

If you haven't set up SSH keys, instructions on how to do this can be found at <https://opus.nci.org.au/display/Help/Using+SSH+keys>

Alternatively, you should enter your Gadi password when prompted.




## Step 5. Connect to the remote JupyterLab server from your remote desktop computer

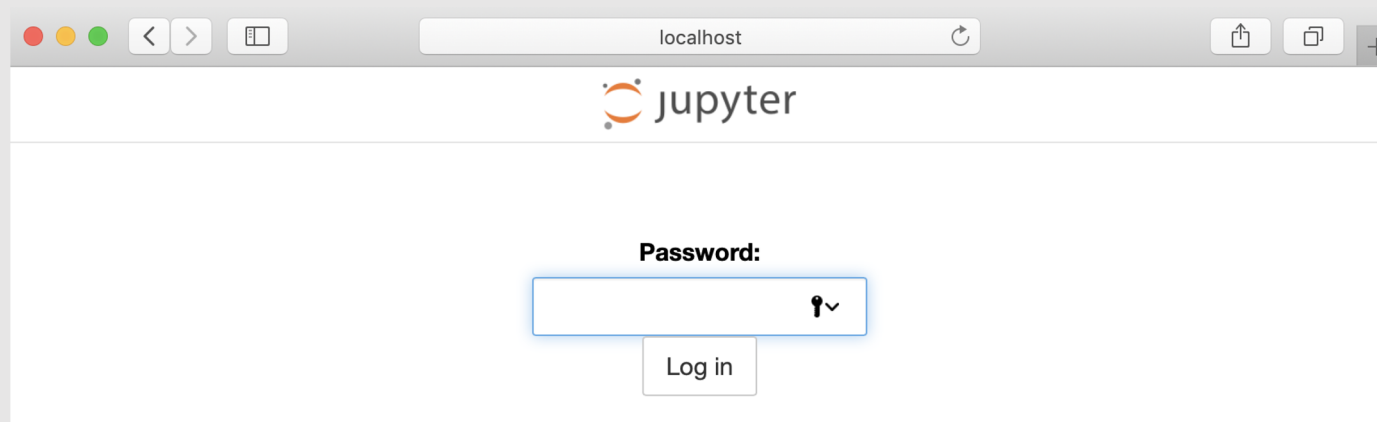
In a browser on your local desktop, type in the following URL:

“localhost:xxxx”

where xxxx is the port number from your client\_cmd file

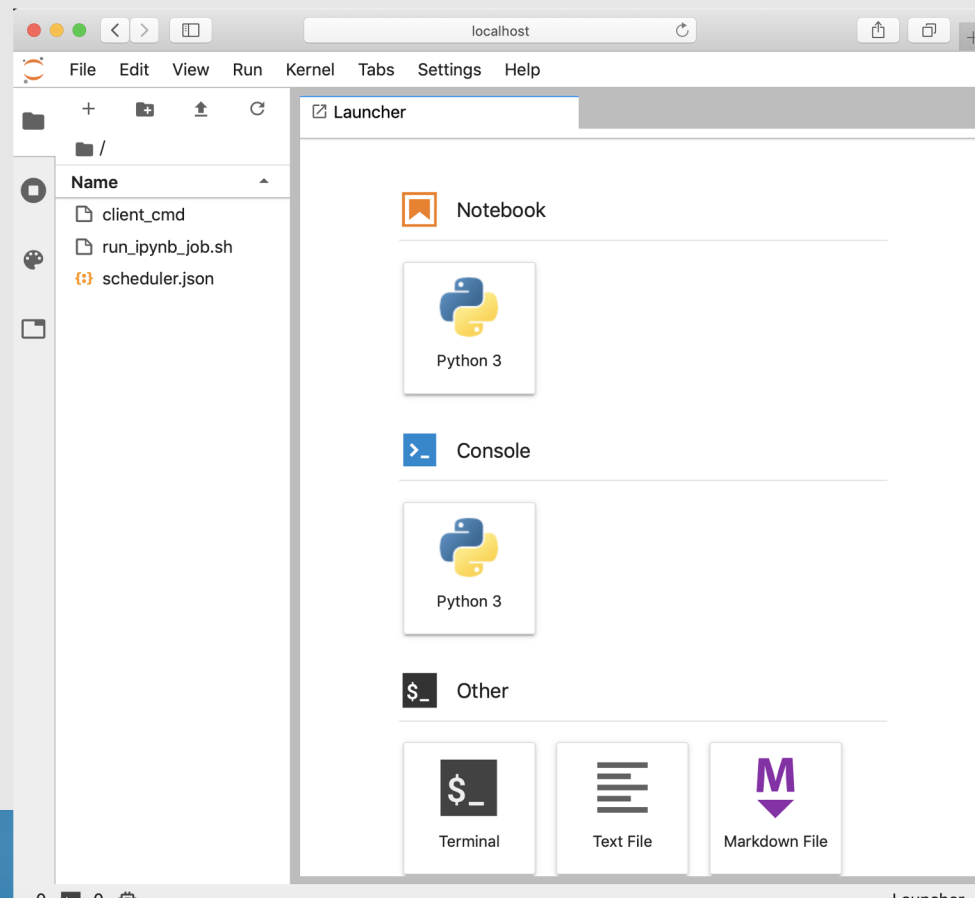


You will be prompted for your JupyterLab password that you created earlier:



## Step 5. Connect to the remote JupyterLab server from your remote desktop computer

Once your authentication has passed, a JupyterLab interface will be launched in a few seconds:



# Step 6. Importing a notebook example

You can drag and drop a notebook from your local machine into the JupyterLab. This file will also appear in your working directory on Gadi:

The screenshot illustrates the process of importing a notebook into JupyterLab. It is divided into three numbered sections:

- 1:** The JupyterLab interface on a local machine. The file browser on the left shows a directory containing several files, including `annual_ave.ipynb`, which is highlighted. A yellow callout box with a red arrow points to this file, stating: "Drag and drop from your local directory to JupyterLab".
- 2:** A file browser window on the local machine showing the `examples` directory. The file `annual_ave.ipynb` is circled in red, indicating it is the source of the notebook being imported.
- 3:** A terminal window on a remote Gadi machine. It shows the output of a `ls -l` command, which now includes `annual_ave.ipynb` in the list of files. A yellow callout box with a red arrow points to this entry, stating: "When the notebook is dragged and dropped into the JupyterLab directory, the file appears in your Gadi working directory immediately." Below the terminal output, a file browser window shows the `annual_ave.ipynb` file in the local working directory.

The central JupyterLab window displays the content of the `annual_ave.ipynb` notebook, which includes Python code for importing libraries, setting up a Dask client, and defining a function to correct time coordinates in DataArray files.

## Step 7. Setting up your Jupyter notebook to use the Dask server

To use the dask server established from the PBS job, it is necessary to add and run the following cell at the beginning of your notebook:

```
from dask.distributed import Client, LocalCluster
client = Client(scheduler_file='scheduler.json')
print(client)
```

The output will show the configuration of the client and Dask cluster. You can check that the number of cores matches what you requested in the job script. Now you can run your notebook as per usual.

To gracefully stop the PBS job:

```
!pango.end.sh
```

# Step 7. Setting up your Jupyter notebook to use the Dask server

View compute threads using Dask dashboard

From your local desktop, open a new tab in your web browser and type the second port in the client\_cmd file to open the Dask dashboard (e.g. localhost:8768). This will allow you to see the dynamic resources of the processing.



```
[ccc777@raijin4 ccc777]$ more run_ipynb_job.sh
#!/bin/bash
#PBS -N pangeo_test
#PBS -P c25
#PBS -q express
#PBS -l walltime=5:00:00
#PBS -l ncpus=32
#PBS -l mem=64GB
#PBS -l jobfs=100GB
module load pangeo/2019.10
pangeo.ini.all.sh
sleep infinity
```

- Make sure your project has enough kSU allocated to complete the job
- For the queue selection (-q), *normal* is recommended if the job is not urgent
- *Walltime* provides the limit of the job run, so make sure you specify enough walltime to run a computationally expensive job
- Watch for the new version of Pangeo in three months time



Add these lines into your notebook or python script!!

```
# start the dask client
from dask.distributed import
Client, LocalCluster
client =
Client(scheduler_file='scheduler.json')

... your work utilizing xarray&dask

# stop the pbs job.
! pangeo.end.sh
```

# Additional slides prepared for Q&A

DEPLOYING	Cloud	HPC
<b>Availability</b>	Resources “always” available	Resources heavily used (need to wait)
<b>Virtualization</b>	Containers and VMs everywhere!	Bare-metal
<b>Scalability</b>	Container-based (Kubernetes)	Job Queue Schedulers (PBS, LSF, SLURM, etc.)
<b>Ideal for...</b>	Interactive sessions with elastic scaling	Large tightly-coupled batch jobs
<b>Deployability</b>	Anyone can deploy anything, but requires some sysadmin knowledge	Only sysadmins can deploy anything, users can install and use approved software

# NCI's data analysis environments – system level

- Software applications are managed by sys admins
- Centralised installation with version control in /apps
- Search by "module avail package-name"
- Can request to install for users if widely used by the community

## Linux Environment Modules

VDI

<https://opus.nci.org.au/display/Help/3.+VDI+software>

Gadi

<https://opus.nci.org.au/display/Help/0.+Welcome+to+Gadi#id-0.WelcometoGadi-Applicationson/apps>

# NCI's data analysis environments – customer level

- User can install additional software in their own space
  - Short term use – scratch
  - Long term use - /g/data/compute-project
- Difference ways to customize personal data analysis environment using software manager like pip or conda

Customized Layer  
Export PYTHONPATH

## Linux Environment Modules

VDI  
<https://opus.nci.org.au/display/Help/3.+VDI+software>

Gadi  
<https://opus.nci.org.au/display/Help/0.+Welcome+to+Gadi#id-0.WelcometoGadi-Applicationson/apps>

# NCI's data analysis environments – dk92 python environment

- Provide a base Anaconda environment
  - Users can use it to create virtual environments in their own space
  - Or use NCI-data-analysis/2020.12 created and managed by NCI

## Supported Layer in dk92

```
“module use /g/data/dk92/apps/Modules/modulefiles  
module load NCI-data-analysis/2020.12”
```

## Linux Environment Modules

VDI

<https://opus.nci.org.au/display/Help/3.+VDI+software>

Gadi

<https://opus.nci.org.au/display/Help/0.+Welcome+to+Gadi#id-0.WelcometoGadi-Applicationson/apps>



# Customising your environment

You might need to customise your environment for reasons including:

- The packages are available on Gadi, but not included in the current Pangeo and virtual environment available in dk92
- The packages are not available on Gadi, and you need to install them
- You develop your own code

Users may need to load additional modules to run their own workflows. If a particular package is not available on Gadi, users can install it under their own working directory, or their computer project space if it is allowed. Please note, the Pangeo environment should always be loaded **before** adding other modules or installing new packages.

# Customising your environment

## Step 1: Enable Pangeo in your shell environment

To enable the Pangeo environment, you can use the following command within jobs, or within an interactive environment:

```
$ module load pangeo
```

```
Loading pangeo/2020.05
```

```
Loading requirement: intel-mkl/2019.3.199 python3/3.7.4 hdf5/1.10.5  
netcdf/4.7.3
```

## Step 2: Check if a module is available on Gadi

Let's see if **tensorflow** is available on Gadi:

```
[abc123@gadi-login-01 ~]$ module avail tensorflow
```

```
-----/apps/Modules/modulefiles-----
```

```
tensorflow/2.0.0 tensorflow/2.1.0 tensorflow/2.3.0
```

# Customising your environment

## Step 3a: Load modules if they are available on Gadi

If the module exists, it will list all the versions available. Pick the one that you would like to use:


```
[abc123@gadi-login-01 ~]$ module load tensorflow/2.0.0
```

Check if the module has loaded properly:

```
[abc123@gadi-login-01 ~]$ module list
```

Currently loaded Module files:

- |                         |                         |                          |
|-------------------------|-------------------------|--------------------------|
| 1) pbs                  | 5) netcdf/4.7.3         | 9) nccl/2.5.6-1+cuda10.1 |
| 2) intel-mkl/2019.3.199 | 6) pango/2020.05        | 10) openmpi/4.0.1        |
| 3) python3/3.7.4        | 7) cuda/10.1            | 11) tensorflow/2.0.0     |
| 4) hdf5/1.10.5          | 8) cudnn/7.6.5-cuda10.1 |                          |



# AND YES, *EXTENSIBLE* SOFTWARE! JUPYTERLAB - BEYOND NOTEBOOKS

The screenshot shows the JupyterLab interface. On the left is a file browser with a tree view of files and folders. The main area is split into three panes: a code editor on the left containing Python code for a polar plot, a console on the right showing the execution of the code, and a terminal at the bottom displaying system information like tasks, load average, and uptime.

This screenshot shows a more complex JupyterLab notebook. It features multiple code cells with Python code for data analysis, including histogram generation and MRI data processing. The interface includes a console, a terminal, and a file browser. Several plots are visible, including a histogram and a line plot of MRI intensity over time.

## Huge Team Effort!

C. Colbert, S. Corlay, A. Darian, B. Granger, J. Grout, P. Ivanov, I. Rose, S. Silvester, C. Willing, J. Zosa-Forde ...

