Pangeo community platform for scientific data processing

OMP 2019/04/02, Guillaume Eynard-Bontemps, CNES/Pangeo

cnes
CENTRE NATIONAL D'ÉTUDES SPATIALES

Pangeo community goals and motivation

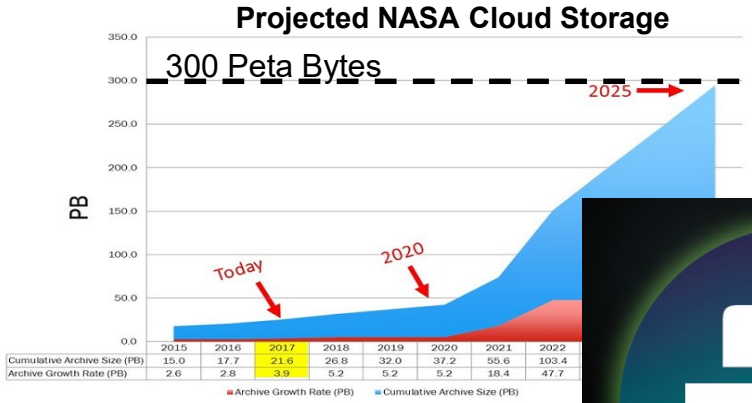

Problems

- Data volume crisis in (geo)sciences
- Software multiplication, non reproducibility
- Many copies of the same datasets
- Local vs HPC vs Cloud
- Technology gap: industry vs academia

Mission

To cultivate an ecosystem in wh
the next generation of open-sou
analysis tools for the geosciences
can be developed, distributed, and
sustained.

als/vision

oster collaboration around the open source Scientific Python ecosystem:

- open and collaborative development
- Welcoming and inclusive culture

- Support the development with domain-specific (geo)science and transverse packages
- Improve scalability of these tools to handle gigabytes to petabytescale datasets

EarthCube
TRANSFORMING GEOSCIENCES RESEARCH

NSF

Pangeo ecosystem

- Set of tools that will facilitate science at all scales
- Platform agnostic
- The core of the Pangeo ecosystem includes:
  - **Xarray** (data-model and toolkit for working with N-dimensional labeled arrays)
  - **Dask** (parallel computing)
  - **Jupyter** (interactive computing)
- Extensible: Series of 3rd party packages that build on top of core libraries
- Flexible: Individual components may be swapped in/out

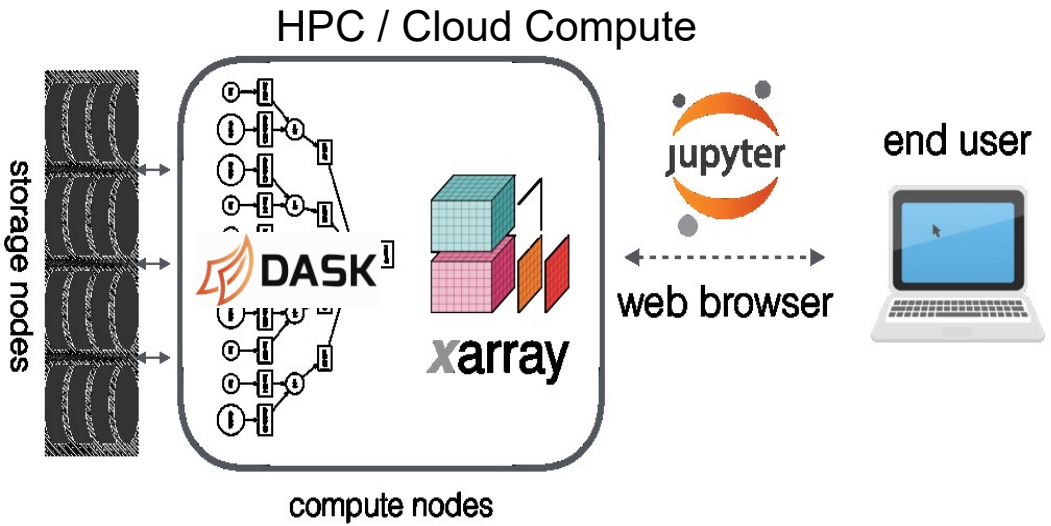

Examples of 3rd party packages in the Pangeo Ecosystem:

- Data discovery
- Regridding and GIS
- Vector calculus
- Signal processing
- Thermodynamics

Jupyter

- ***"The Jupyter Notebook is an open-source web application that allows you to create and share documents that contain live code, equations, visualizations and explanatory text. Uses include: data cleaning and transformation, numerical simulation, statistical modeling, machine learning and much more."***
- **Originally "Python-centric" but has been expanded to include over 40 popular programming languages (e.g. Julia and R)**
- **Check it out at: http://jupyter.org**

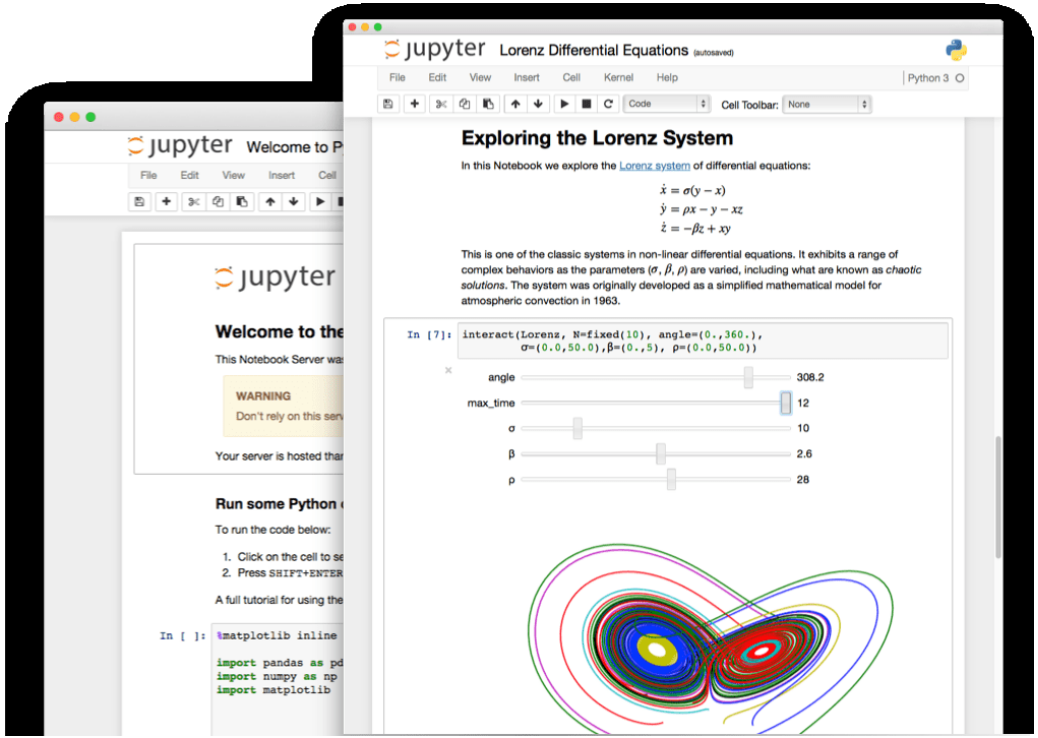

Dask

- **Dask is a flexible parallel computing library for analytic computing**
- **Parallel arrays allow us to seamlessly scale serial programs and workflows**
- **Dask is optimized for computation**
- **Can be utilized on a single machine or a cluster of machines**

Dask arrays coordinate many NumPy arrays arranged into a grid. These NumPy arrays may live on disk or on other machines.

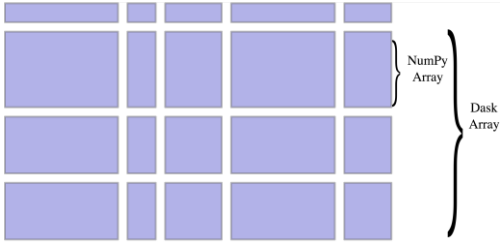

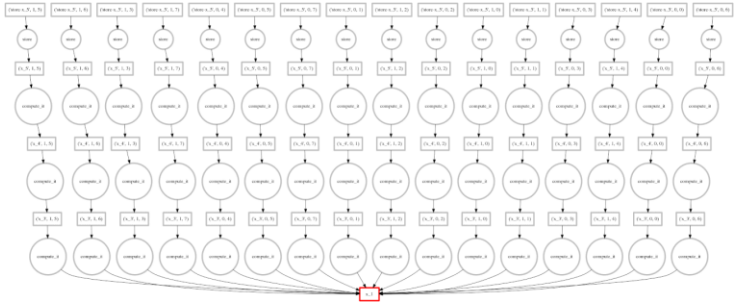

Example of a Dask task graph for a simple, embarrassingly parallel reduction operation.

Source: Dask documentation

Xarray

- **N-D labeled arrays and datasets in Python**
- **Data model emulates the Common Data Model (e.g. NetCDF)**
- **Key features:**
  - Label-based indexing
  - Interoperability with core scientific Python packages
  - Parallel computation using Dask
  - Wide range of input/output options
  - Robust data analysis and manipulation toolkit

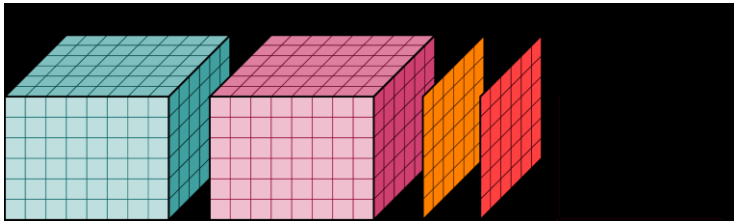

```
import xarray as xr

# Load a netCDF dataset
ds = xr.open_dataset('air_temperature.nc')
# Resample daily data to monthly means
ds = ds.resample('MS', dim='time', how='mean')
# Calculate a monthly climatology
climatology = ds.groupby('time.month').mean(dim='time')
# Calculate monthly anomalies
anomalies = ds.groupby('time.month') - climatology
# Plot an example monthly anomaly (June 2013)
anomalies.sel(time='2013-06')['air'].plot()
```

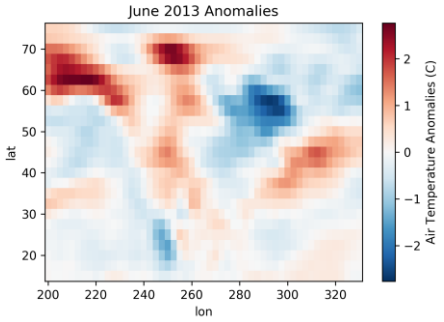

Hoyer, S. & Hamman, J., (2017). Xarray: N-D labeled Arrays and Datasets in Python. Journal of Open Research Software. 5(1), p.10. DOI: http://doi.org/10.5334/jors.148

Pangeo diversity

BUILD YOUR OWN PANGEO

Storage Formats	HDF	OPeNDAP	Cloud Optimized COG/Zarr/Parquet/etc.
ND-Arrays	NumPy	DASK	More coming…
Data Models	xarray	Iris	pandas $y_i t = \beta' x_{it} + \mu_i + \epsilon_{it}$
Processing Mode	jupyter Interactive	Batch 	Serverless
Compute Platform	HPC 	Cloud  Google Cloud Platform	Local

Pangeo HPC deployment

NCAR's Cheyenne Super Computer

- 145,152 processors
- 52.7 Pb of parallel disk storage
- InfiniBand high-speed interconnect

dask.distributed: parallel workers across many HPC nodes

Xarray for computational toolkit and I/O

Jupyter notebooks for interactive computing

New tools for deploying dask clusters on HPC

- e.g. **dask-jobqueue**[1]

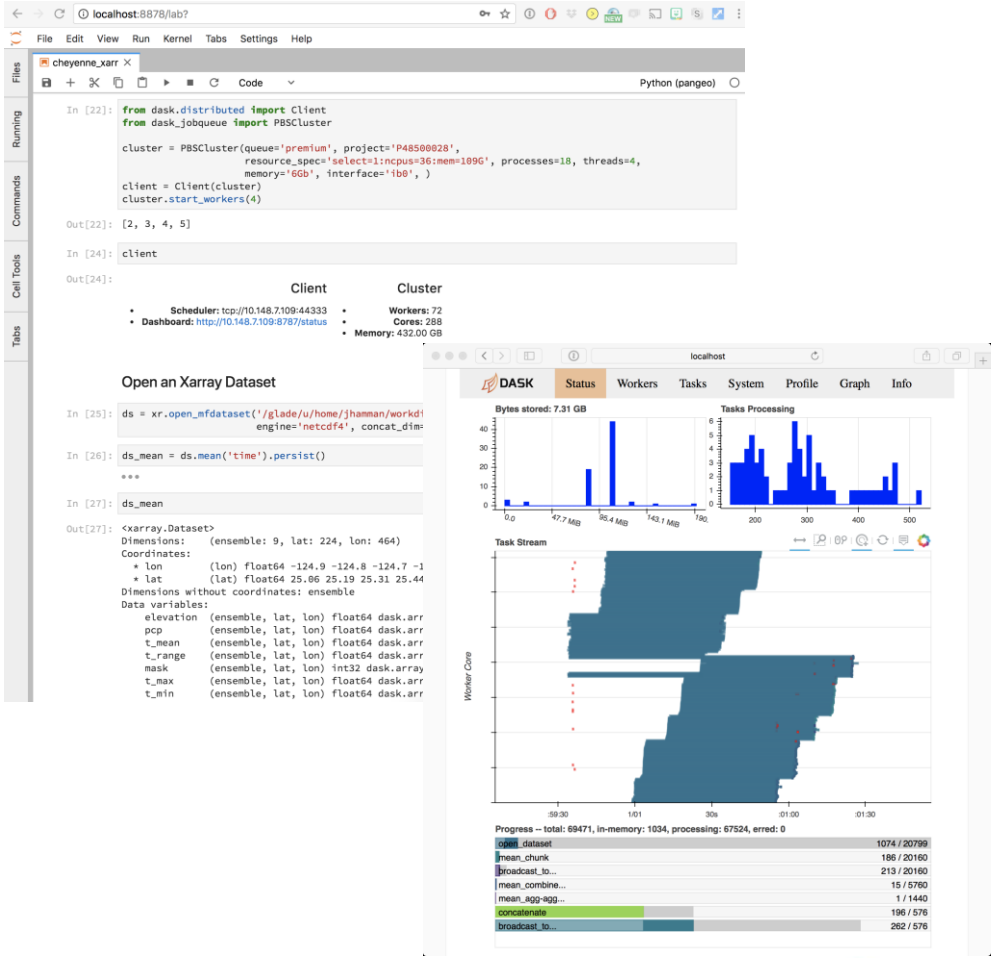

1: https://github.com/dask/dask-jobqueue

Pangeo public cloud deployment

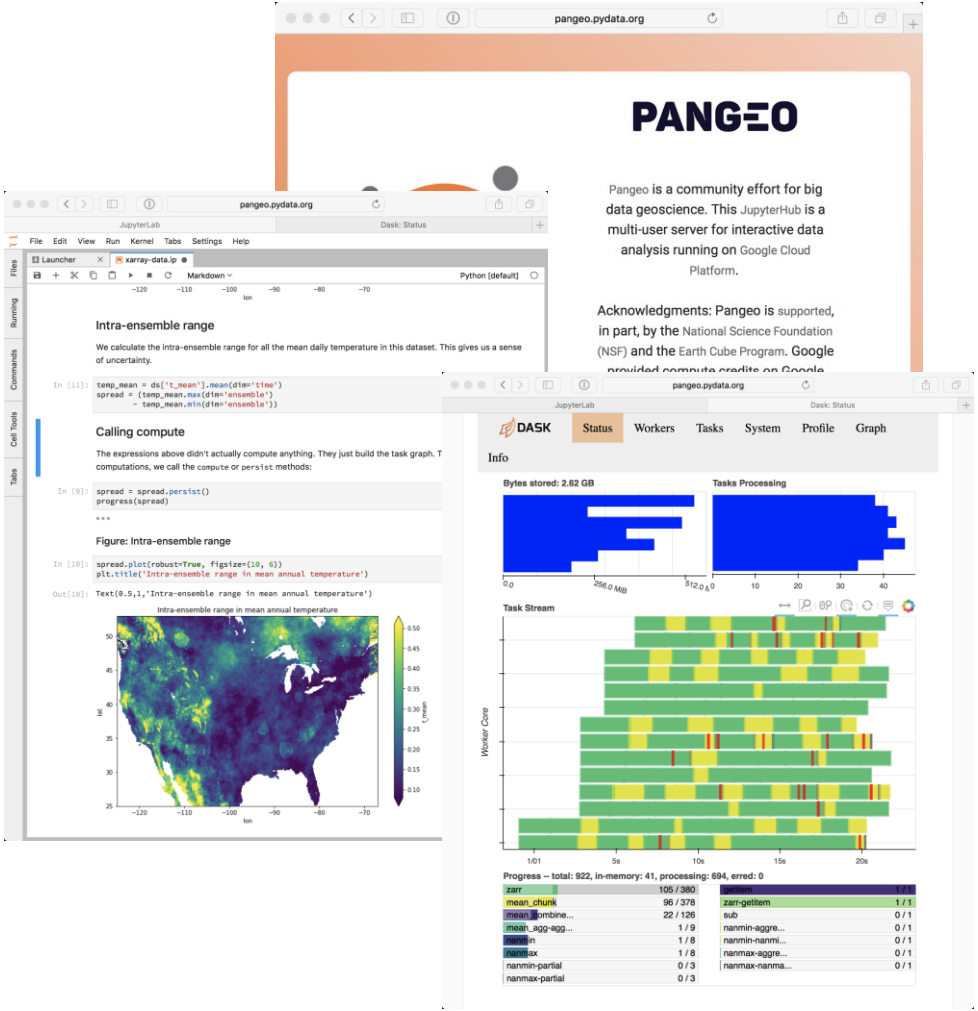

hub.pangeo.io

pangeo.binder.io

JupyterHub/BinderHub running on the Google Cloud

- Kubernetes for both Jupyter and Dask-distributed
  - Dask-kubernetes
- Exploring/evaluating:
  - Cloud storage
  - User environment customization
  - Data discovery
- Kubernetes Helm-chart (github.com/pangeo-data/helm-chart)
- CI/CD with Hubploy and CircleCI
- Deployments exist on AWS and Azure.

Pangeo vs state of the art

VS

Mature	Less Mature
Robust	Pretty strong
JVM/Python	Python only
Query optimized	Science optimized
Collections & Dataframes	Collections, DF, Arrays, Futures…
Python overhead	Python only
For big tabular data	For science data
Hadoop/Cloud/HPC	Hadoop/Cloud/HPC

VS

databases *(Rasdaman, SciDB…)*

Laptop to cluster

Serverless

NetCDF/TIFF no ingestion

Scales with Dask

Python only

Can build array db with Pangeo (Open data Cube)

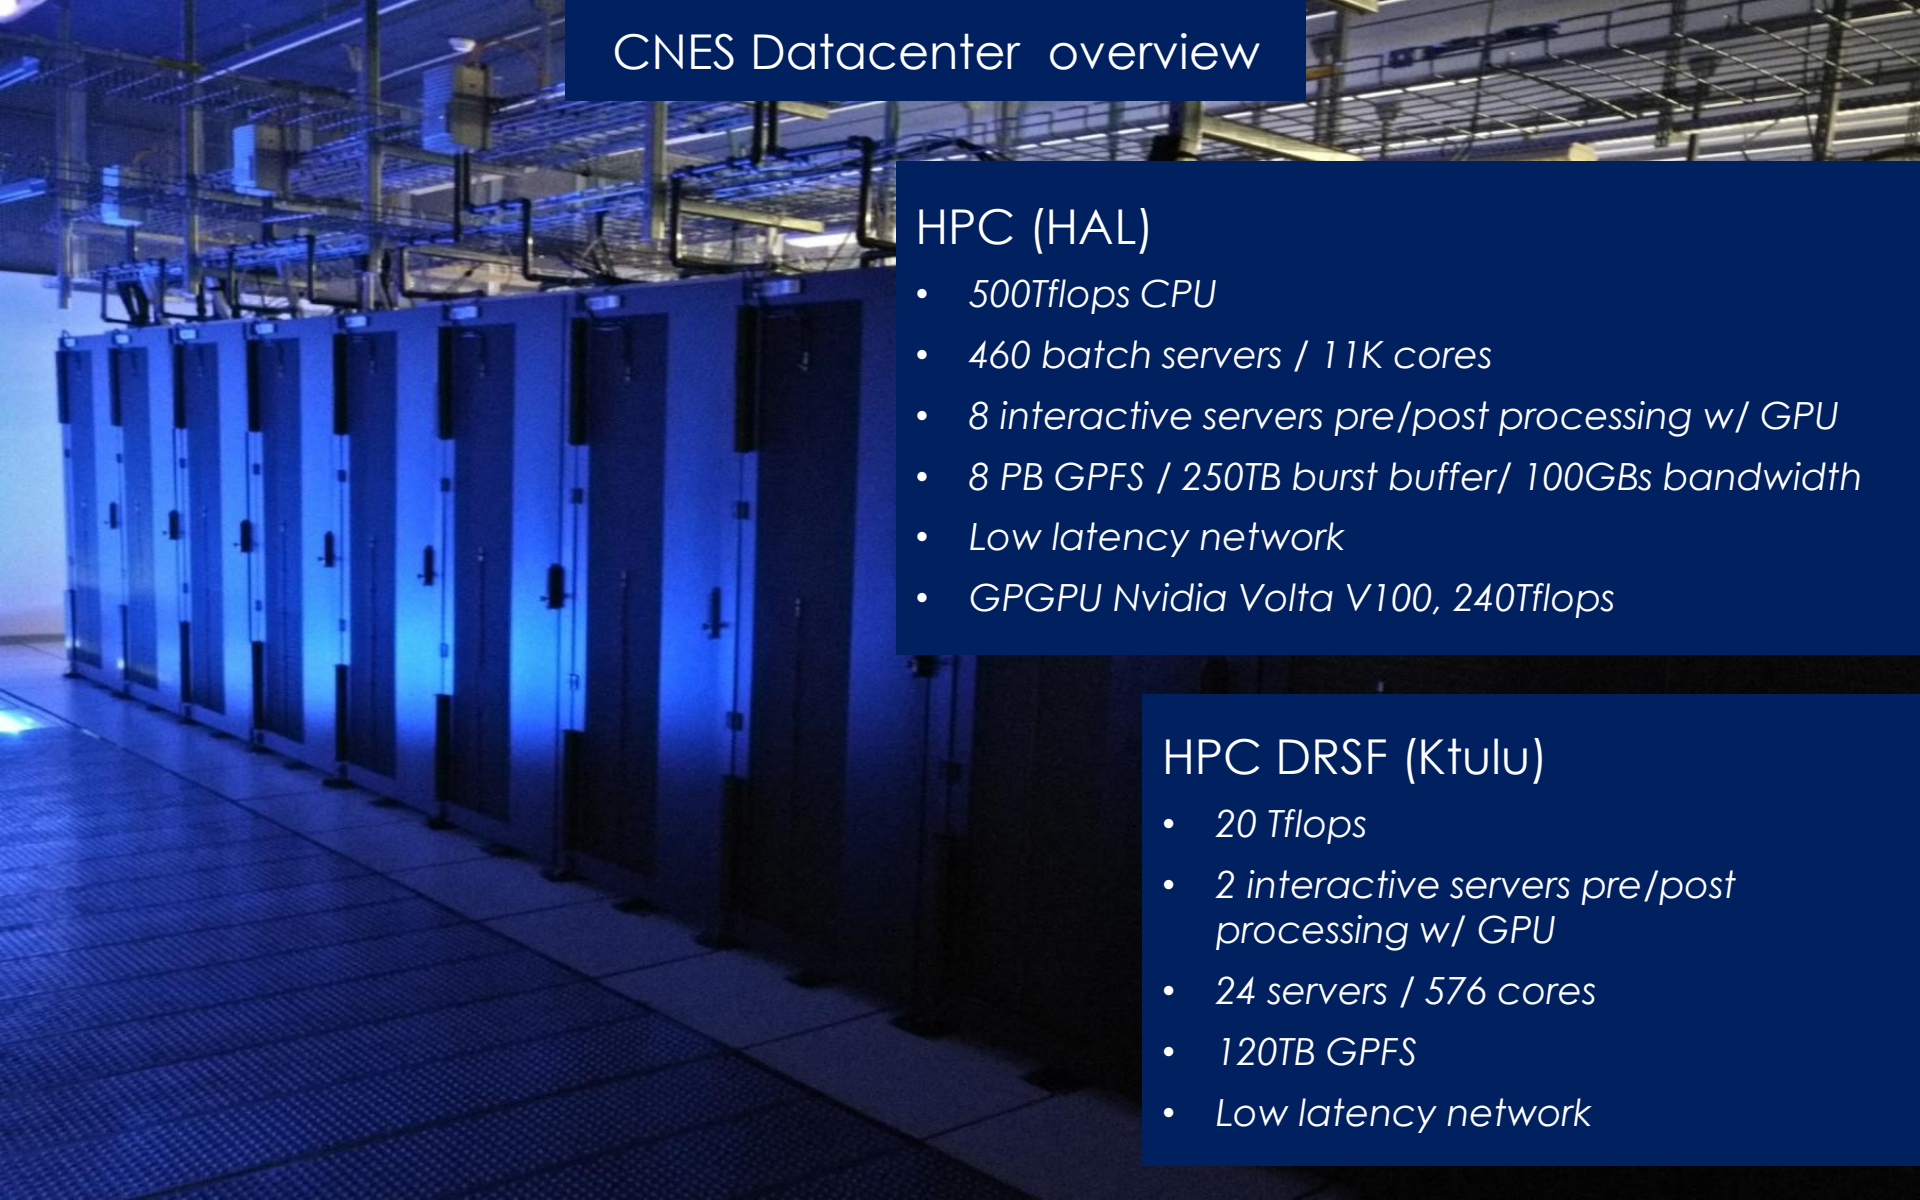

CNES Datacenter overview

HPC (HAL)

- *500Tflops CPU*
- *460 batch servers / 11K cores*
- *8 interactive servers pre/post processing w/ GPU*
- *8 PB GPFS / 250TB burst buffer/ 100GBs bandwidth*
- *Low latency network*
- *GPGPU Nvidia Volta V100, 240Tflops*

HPC DRSF (Ktulu)

- *20 Tflops*
- *2 interactive servers pre/post processing w/ GPU*
- *24 servers / 576 cores*
- *120TB GPFS*
- *Low latency network*

HPC usecases in CNES

Two main kinds of processing

Numerical simulation (HPC)

- Upstream phase, R&D
- Highly optimized technics
- Fine grain parallelism

Trends : multiscale, multiphysics

Data Processing (HTC)

- Downstream phase, operation
- Sensors data → scientific data
- Coarse grain parallelism

Trends : data volume explosion

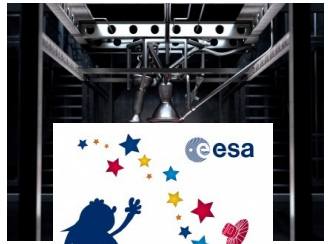

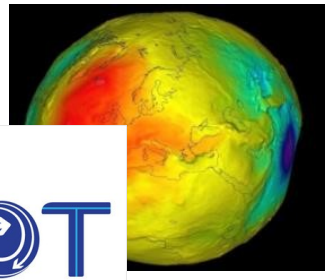

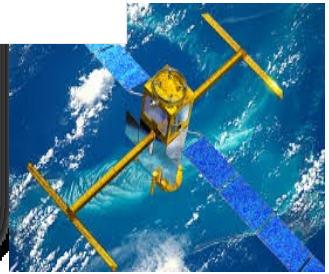

Pangeo at CNES

- JupyterHub and notebooks for interactive computing
  - Hub on a VM with qsub access
  - Batchspawner, Wrapspawner
- dask.distributed: parallel workers across many HPC nodes
- Xarray for computational toolkit and I/O
- New tool for deploying dask clusters on HPC: **dask-jobqueue**
  - Start a cluster from a notebook
  - Interactive (or not) distributed computing
  - Auto scaling capabilities

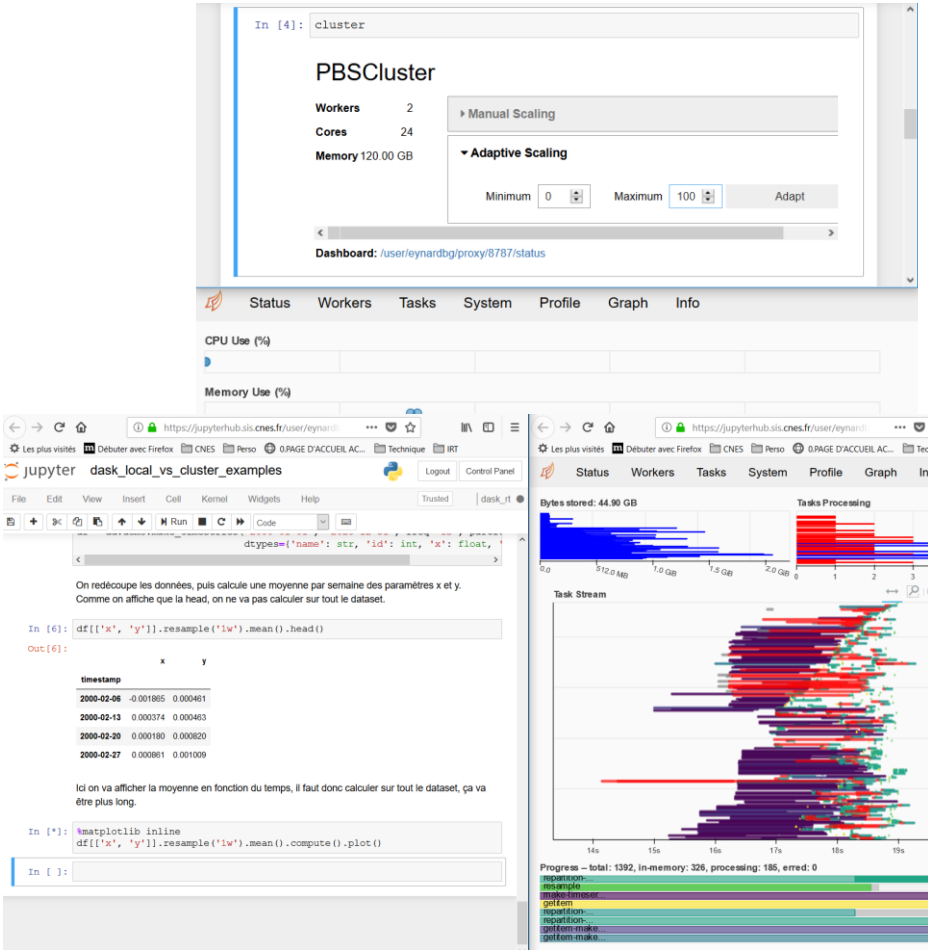

Demos @ CNES

Dask and dask-jobqueue basic example

Some realistic workload

Going deeper with Xarray

Image processing: NDVI

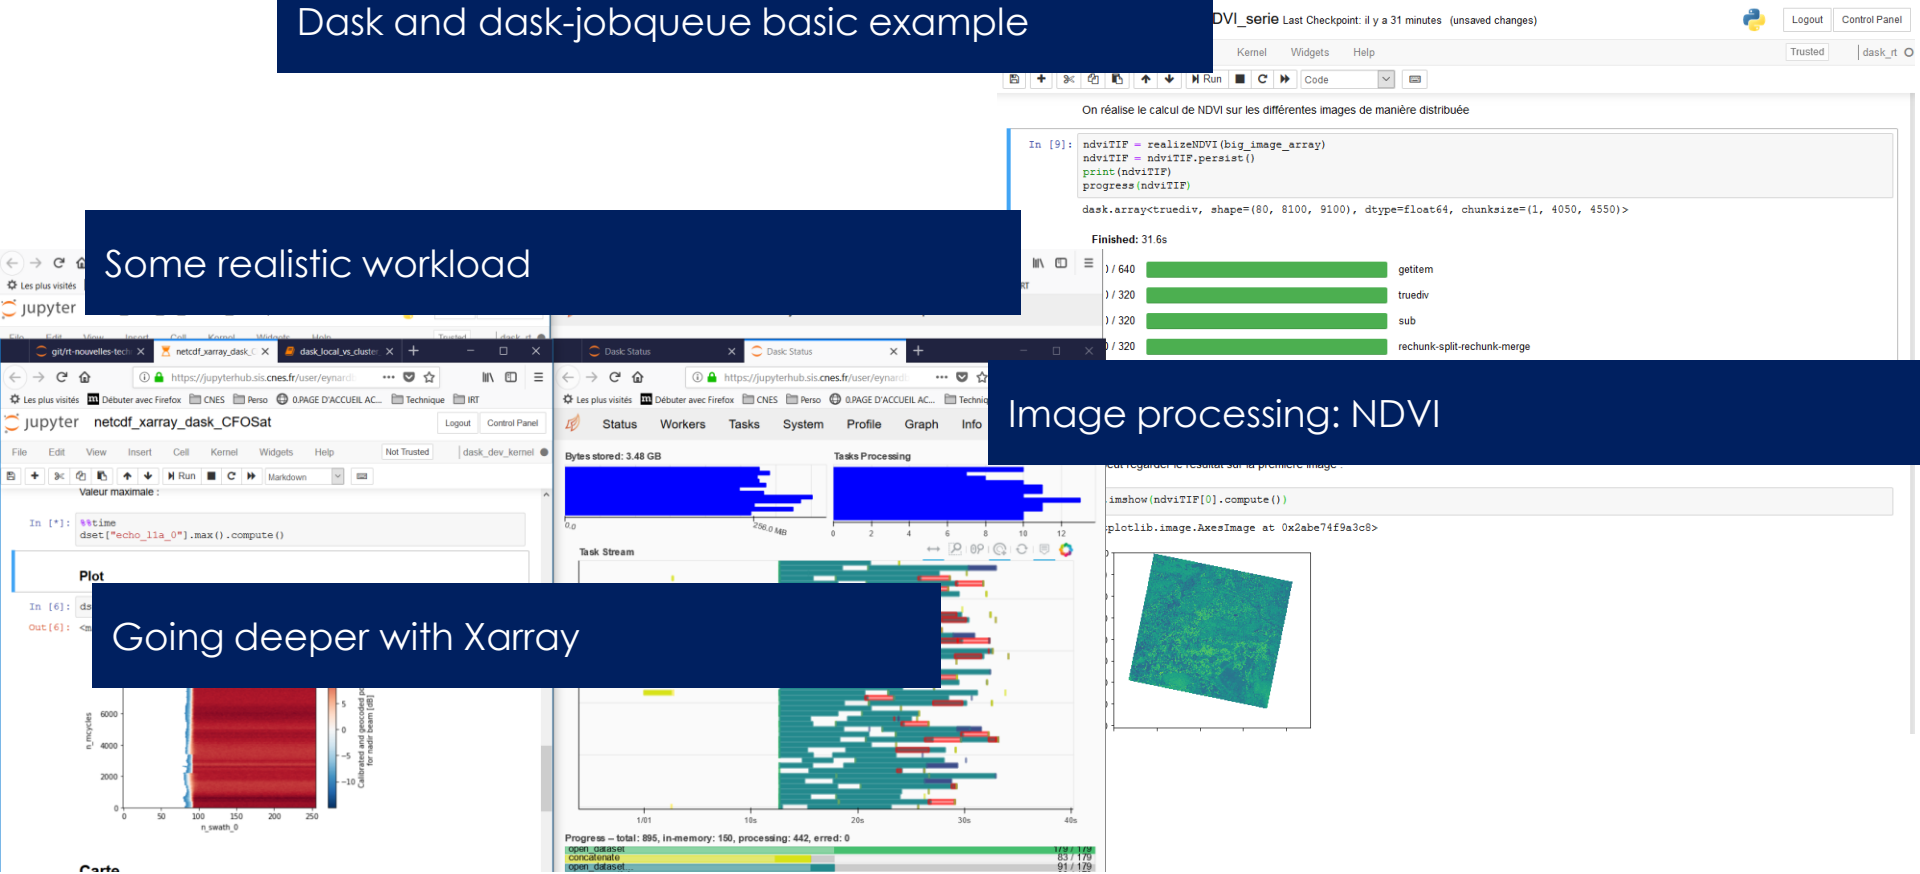

Pangeo Cloud use

Geoscience use cases:
http://pangeo.io/use_cases/index.html

Astronomy with GAIA catalog:

https://github.com/pangeo-data/pangeo/issues/255#issuecomment-427186915

Image processing and visualization

https://medium.com/pangeo/cloud-native-geoprocessing-of-earth-observation-satellite-data-with-pangeo-997692d91ca2

Conclusion, Useful links

Conclusions

- Pangeo ecosystem greatlty facilitates distributed computing and data analysis at scale
- It changes ways of doing it too
- Non monolithic platform built on top of existing Scientific Python stack and new related packages
- Community is always here to help
- Dask more versatile and easy to use than Spark.

Next steps

- Broaden users and use cases at CNES
- Encourage people to get in touch with Pangeo community
- Work in cooperation with others (Ongoing with Ifremer and CLS on SWOT aval data processing)
- Get involved!!
- Pangeo french meeting ont May 23rd

Pangeo wesite and discussions:

https://pangeo.io

https://github.com/pangeo-data/pangeo/issues

https://medium.com/pangeo

Pangeo Example + Binder:

https://github.com/pangeo-data/pangeo-example-notebooks

http://binder.pangeo.io/v2/gh/pangeo-data/pangeo-example-notebooks/master

Dask jobqueue:

https://github.com/dask/dask-jobqueue

Dask simple examples:

https://github.com/dask/dask-examples

My email

Guillaume.Eynard-Bontemps@cnes.fr

Tutorial

Github and binder link :
https://github.com/guillaumeeb/pangeo-tutorial-agu-2018

https://binder.pangeo.io/v2/gh/guillaumeeb/pangeo-tutorial-agu-2018/some_fixes

Dask only tutorial:

https://github.com/mrocklin/pydata-nyc-2018-tutorial

Other binder resources :

https://github.com/pangeo-data/pangeo-example-notebooks

https://github.com/pangeo-data/pangeo_ocean_examples

https://github.com/dask/dask-examples