



**PANGEO**


A COMMUNITY-DRIVEN EFFORT FOR  
BIG DATA GEOSCIENCE




# WHAT DRIVES PROGRESS IN OCEANOGRAPHY?

**New Ideas**


$$E = \frac{\rho_0 |\mathbf{U}|}{\pi} \int_{|f|/|\mathbf{U}|}^{N/|\mathbf{U}|} P_{1D}(k) \sqrt{N^2 - |\mathbf{U}|^2 k^2} \sqrt{|\mathbf{U}|^2 k^2} dk$$



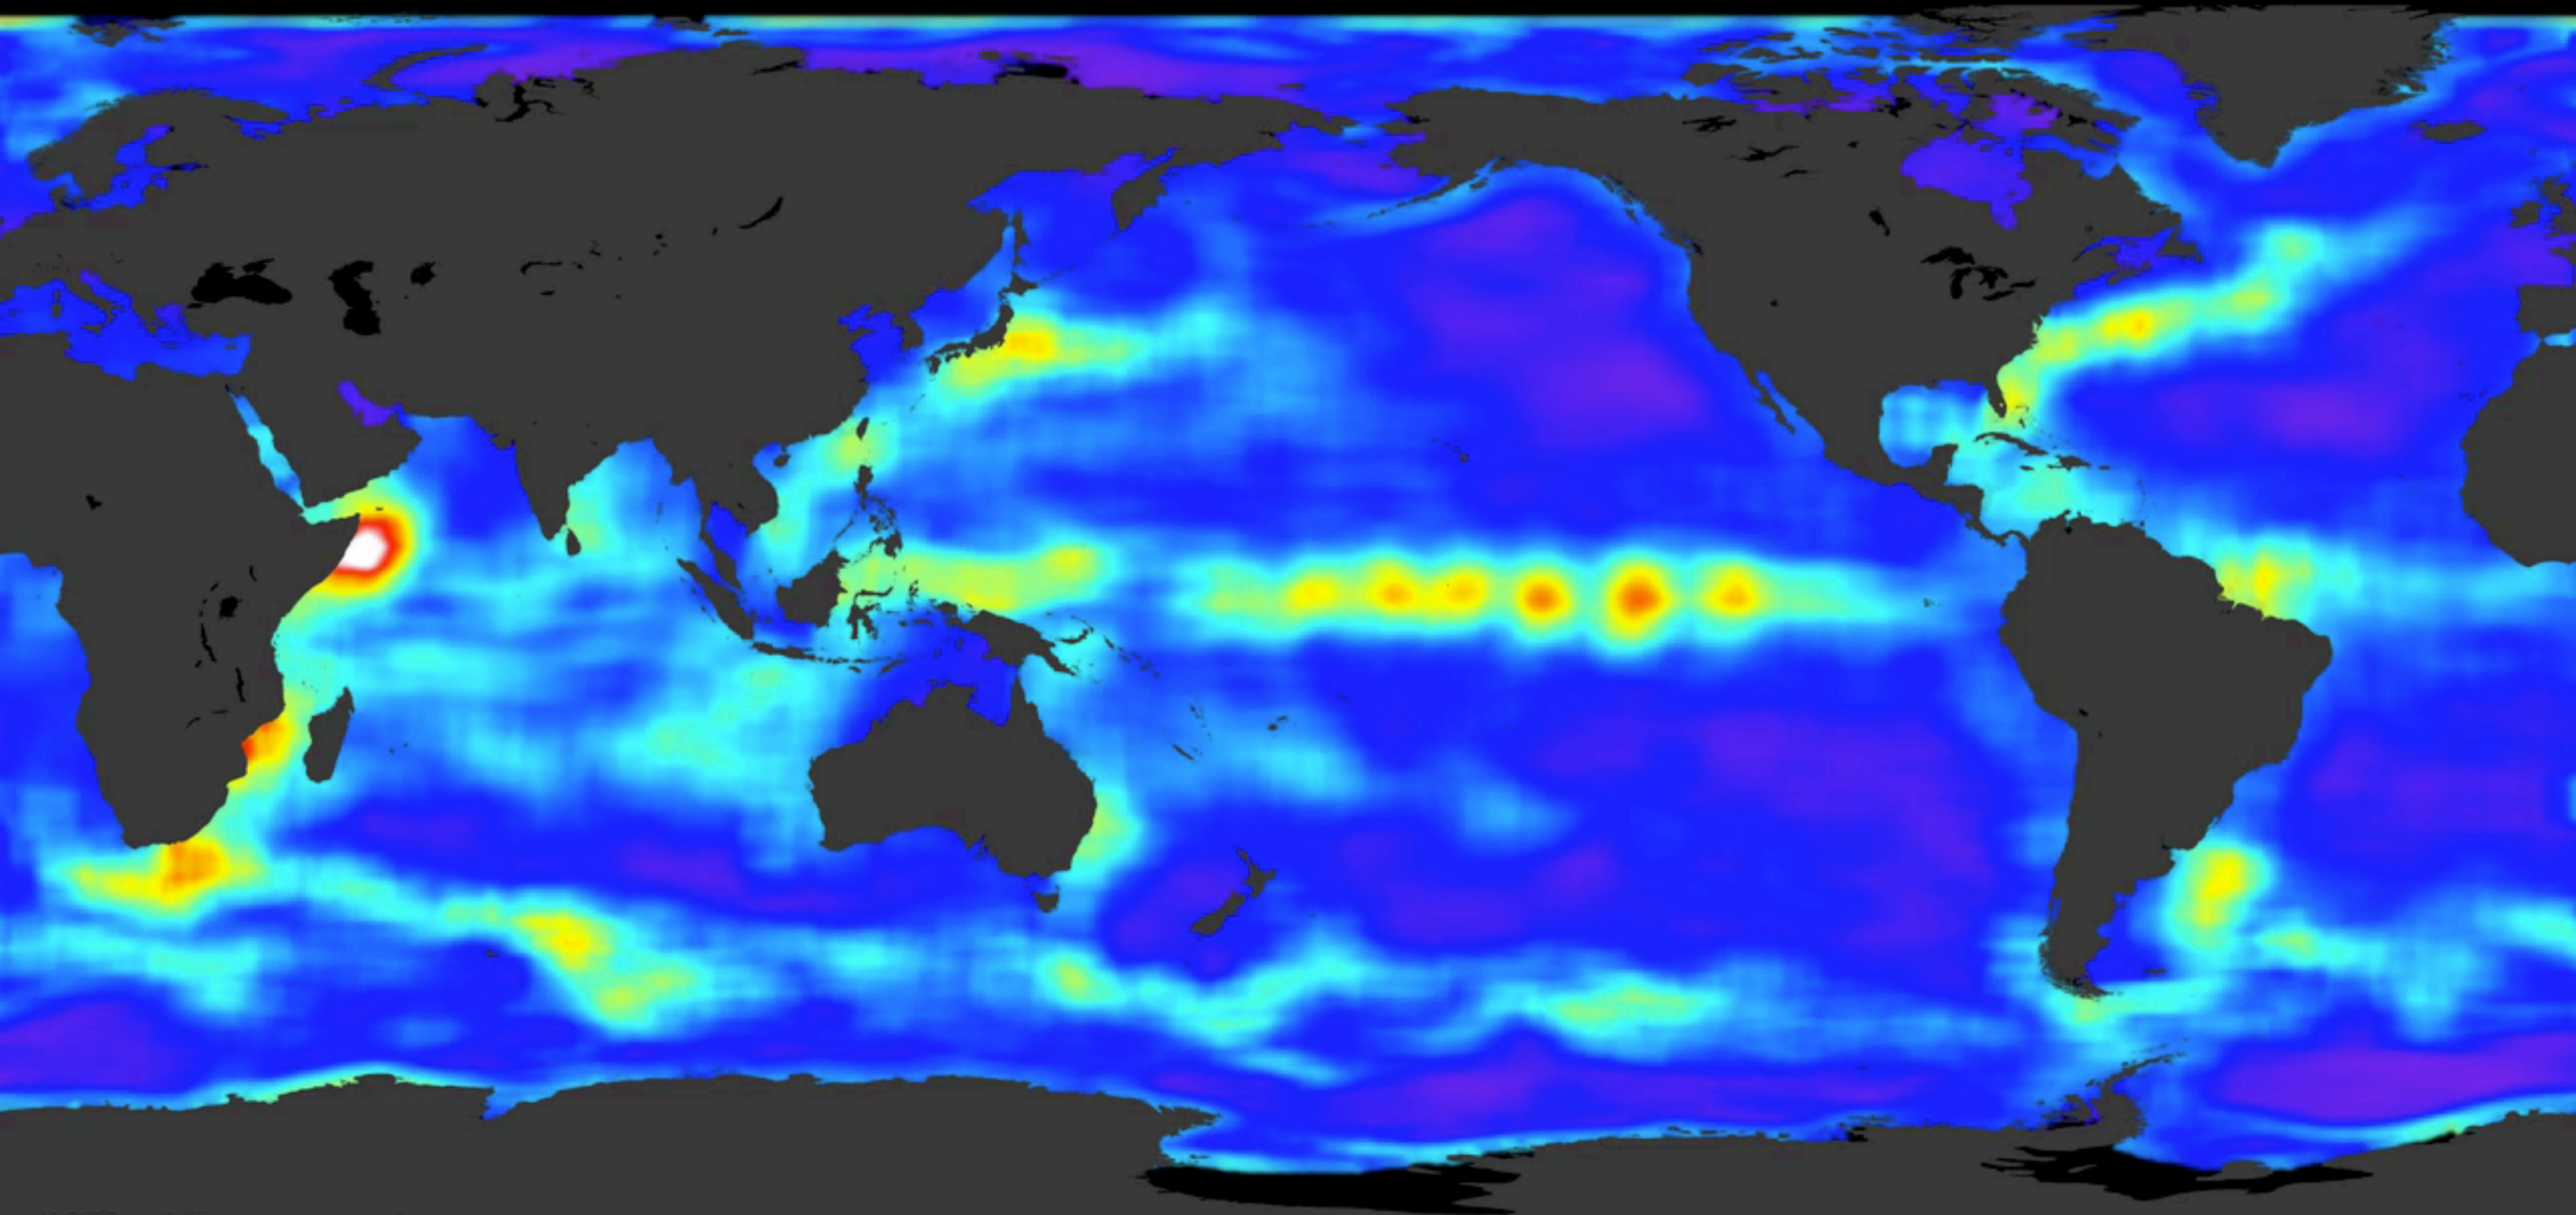
**New Simulations**



**New Observations**








1978 1.5 degree resolution  
Seasat

Credit: NASA SVS / PODAAC






Flourescence microscopy image of neurons by Jan Schmoranzer ([link](#))











*How should scientific data  
infrastructure be organized for the  
petabyte era?*



# WHAT IS PANGEO?

*“A community platform for Big Data geoscience”*

- Open Community
- Open Source Software
- Open Source Infrastructure



# PANGEO COMMUNITY




Lamont-Doherty Earth Observatory  
COLUMBIA UNIVERSITY | EARTH INSTITUTE



[HTTP://PANGEO.IO](http://PANGEO.IO)



# PANGEO TIMELINE






# SCIENTIFIC PYTHON FOR DATA SCIENCE


## Growth of major programming languages

Based on Stack Overflow question views in World Bank high-income countries



## Stack Overflow Traffic to Questions About Selected Python Packages

Based on visits to Stack Overflow questions from World Bank high-income countries



source: [stackoverflow.com](https://stackoverflow.com)



# PANGEO SOFTWARE




## SCIENTIFIC PYTHON FOR GEOSCIENCE



Credit: Stephan Hoyer, Jake Vanderplas (SciPy 2015)



# XARRAY DATASET: MULTIDIMENSIONAL VARIABLES WITH COORDINATES AND METADATA



*“netCDF meets pandas.DataFrame”*



Credit: Stephan Hoyer



# XARRAY MAKES SCIENCE EASY

```
import xarray as xr
ds = xr.open_dataset('NOAA_NCDC_ERSST_v3b_SST.nc')
ds
```

```
<xarray.Dataset>
```

```
Dimensions:  (lat: 89, lon: 180, time: 684)
```

```
Coordinates:
```

```
  * lat      (lat) float32 -88.0 -86.0 -84.0 -82.0 -80.0 -78.0 -76.0 -74.0 ...
```

```
  * lon      (lon) float32  0.0  2.0  4.0  6.0  8.0 10.0 12.0 14.0 16.0 18.0 20.0 ...
```

```
  * time     (time) datetime64[ns] 1960-01-15 1960-02-15 1960-03-15 ...
```

```
Data variables:
```

```
    sst      (time, lat, lon) float64 nan nan nan nan nan nan nan nan ...
```

```
Attributes:
```


```
    Conventions:  IRIDL
```

```
    source:       https://iridl.ldeo.columbia.edu/SOURCES/.NOAA/.NCDC/.ERSST/...
```



# XARRAY: LABEL-BASED SELECTION


```
# select and plot data from my birthday  
ds.sst.sel(time='1982-08-07', method='nearest').plot()
```





# XARRAY: LABEL-BASED OPERATIONS


```
# zonal and time mean temperature  
ds.sst.mean(dim=('time', 'lon')).plot()
```






# XARRAY: GROUPING AND AGGREGATION

```
sst_clim = sst.groupby('time.month').mean(dim='time')
sst_anom = sst.groupby('time.month') - sst_clim
nino34_index = (sst_anom.sel(lat=slice(-5, 5), lon=slice(190, 240))
               .mean(dim=('lon', 'lat'))
               .rolling(time=3).mean(dim='time'))
nino34_index.plot()
```



SST Anomaly in Nino 3.4 Region (5N-5S,120-170W)






# XARRAY

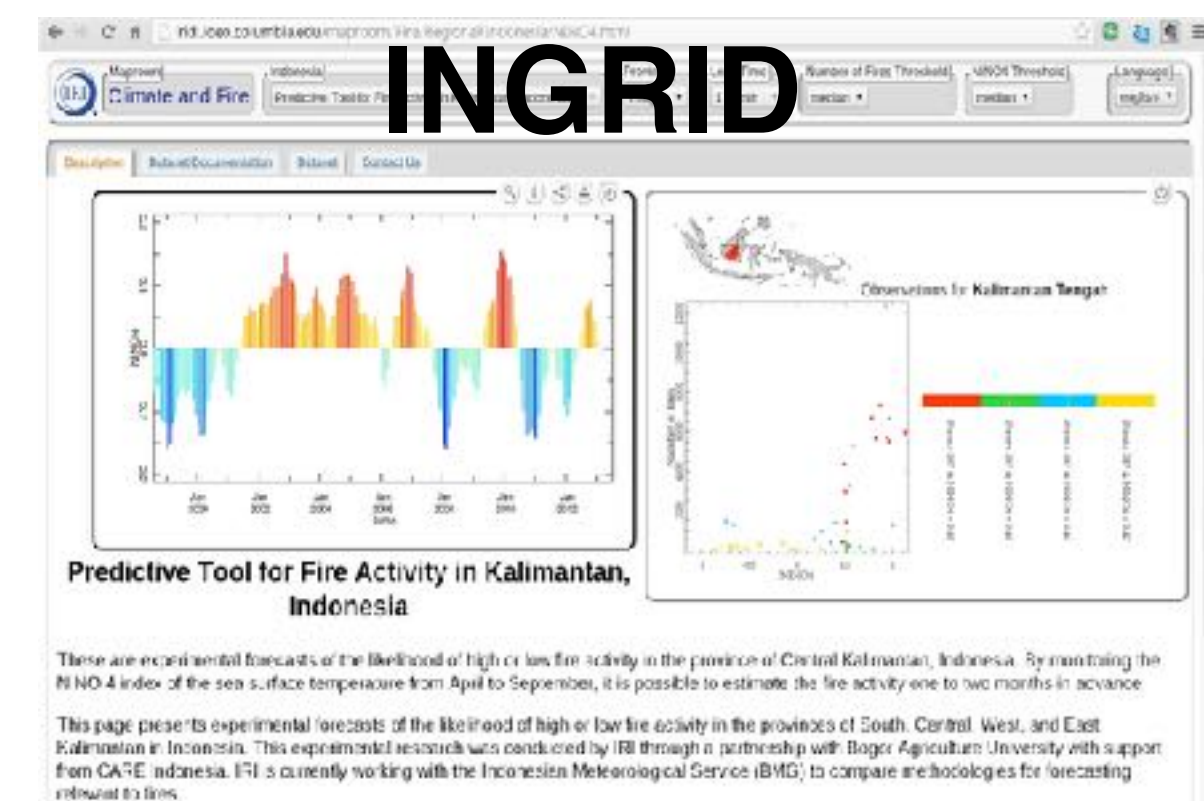
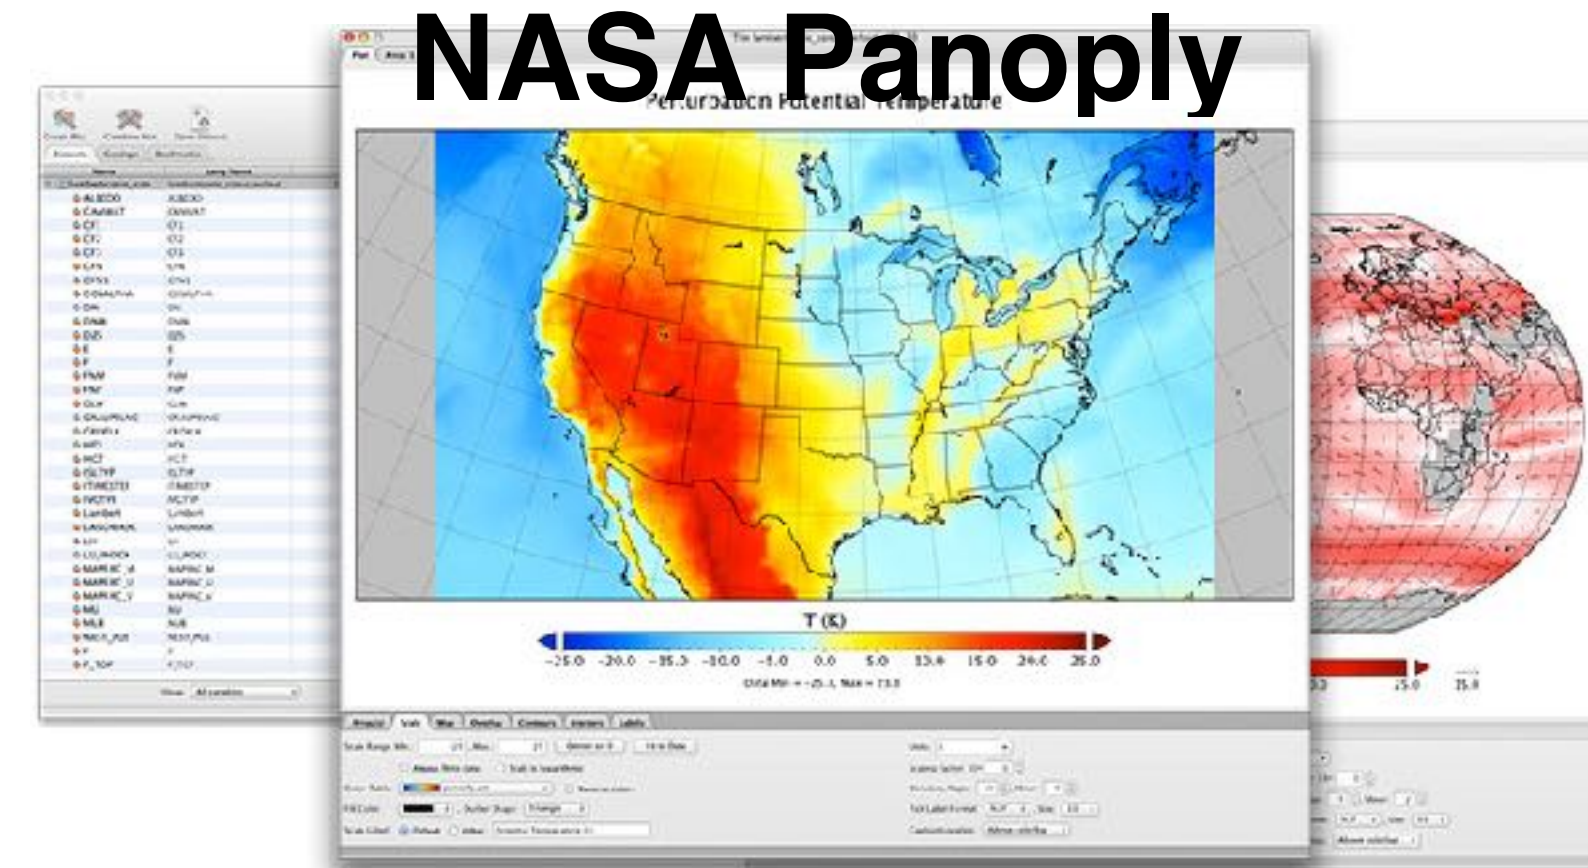

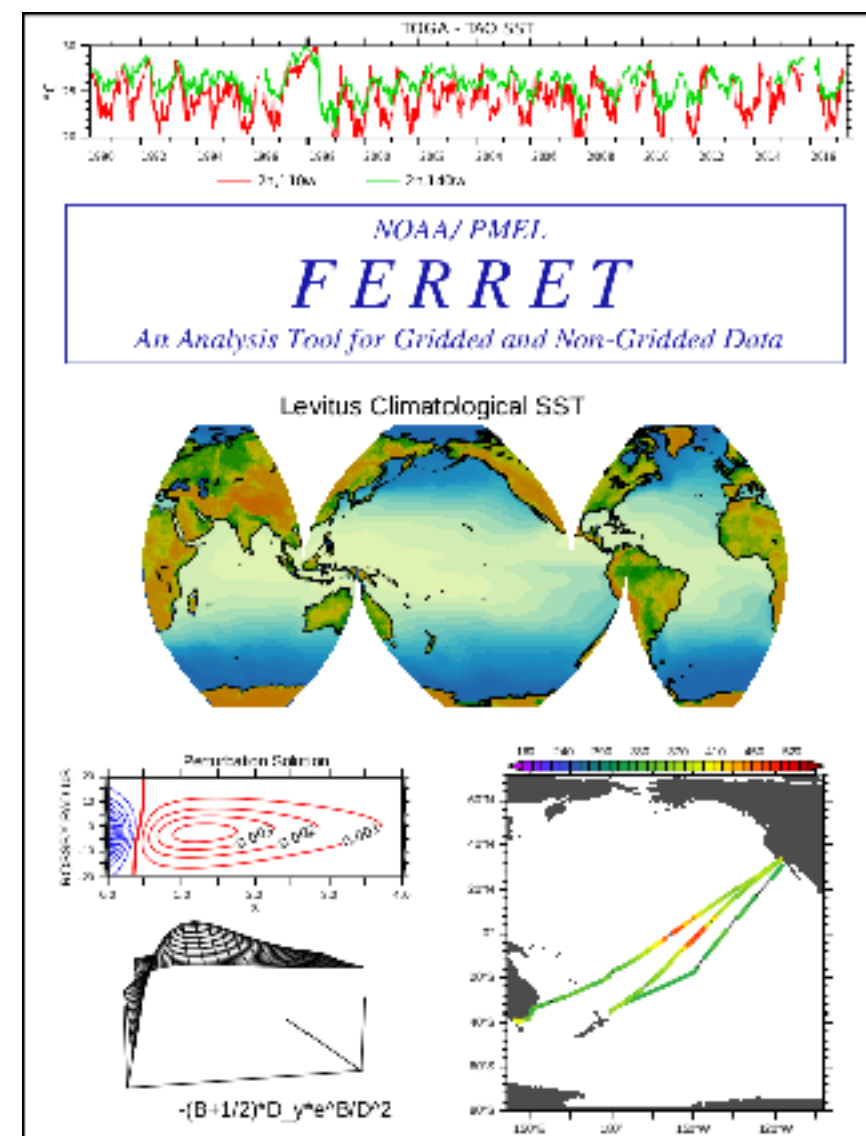
<https://github.com/pydata/xarray>

- label-based indexing and arithmetic
- interoperability with the core scientific Python packages (e.g., pandas, NumPy, Matplotlib)
- out-of-core computation on datasets that don't fit into memory (thanks dask!)
- wide range of input/output (I/O) options: netCDF, HDF, geoTIFF, zarr
- advanced multi-dimensional data manipulation tools such as group-by and resampling





# LEGACY SOFTWARE



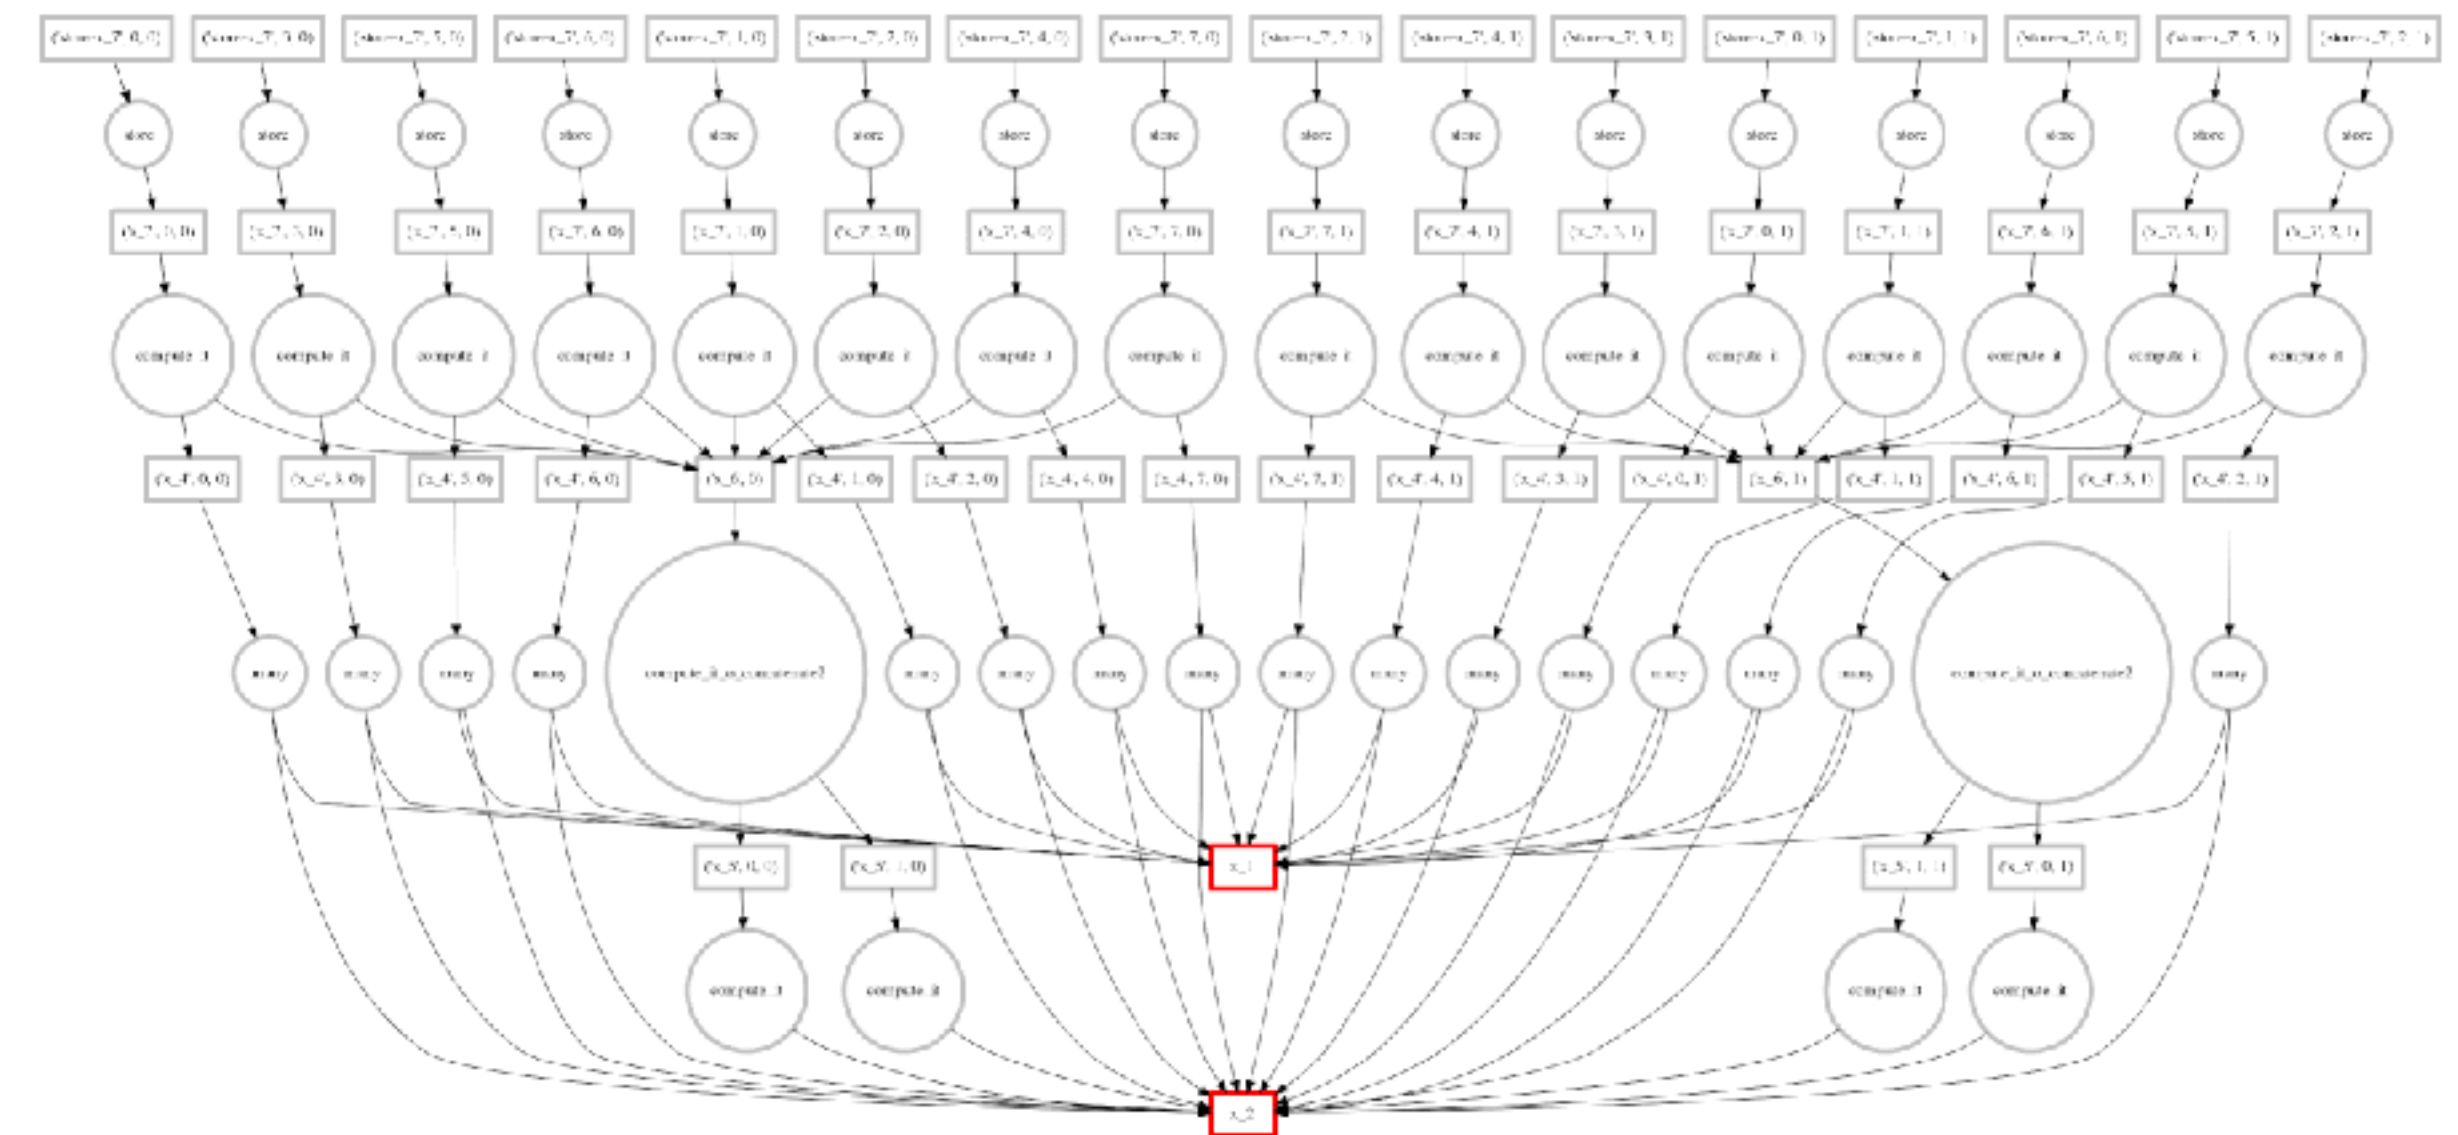


# DASK

<https://github.com/dask/dask/>

	8	8	8
5	('x', 0, 0)	('x', 0, 1)	('x', 0, 2)
5	('x', 1, 0)	('x', 1, 1)	('x', 1, 2)
5	('x', 2, 0)	('x', 2, 1)	('x', 2, 2)
5	('x', 3, 0)	('x', 3, 1)	('x', 3, 2)

ND-Arrays are split into chunks that comfortably fit in memory



Complex computations represented as a graph of individual tasks.

Scheduler optimizes execution of graph.

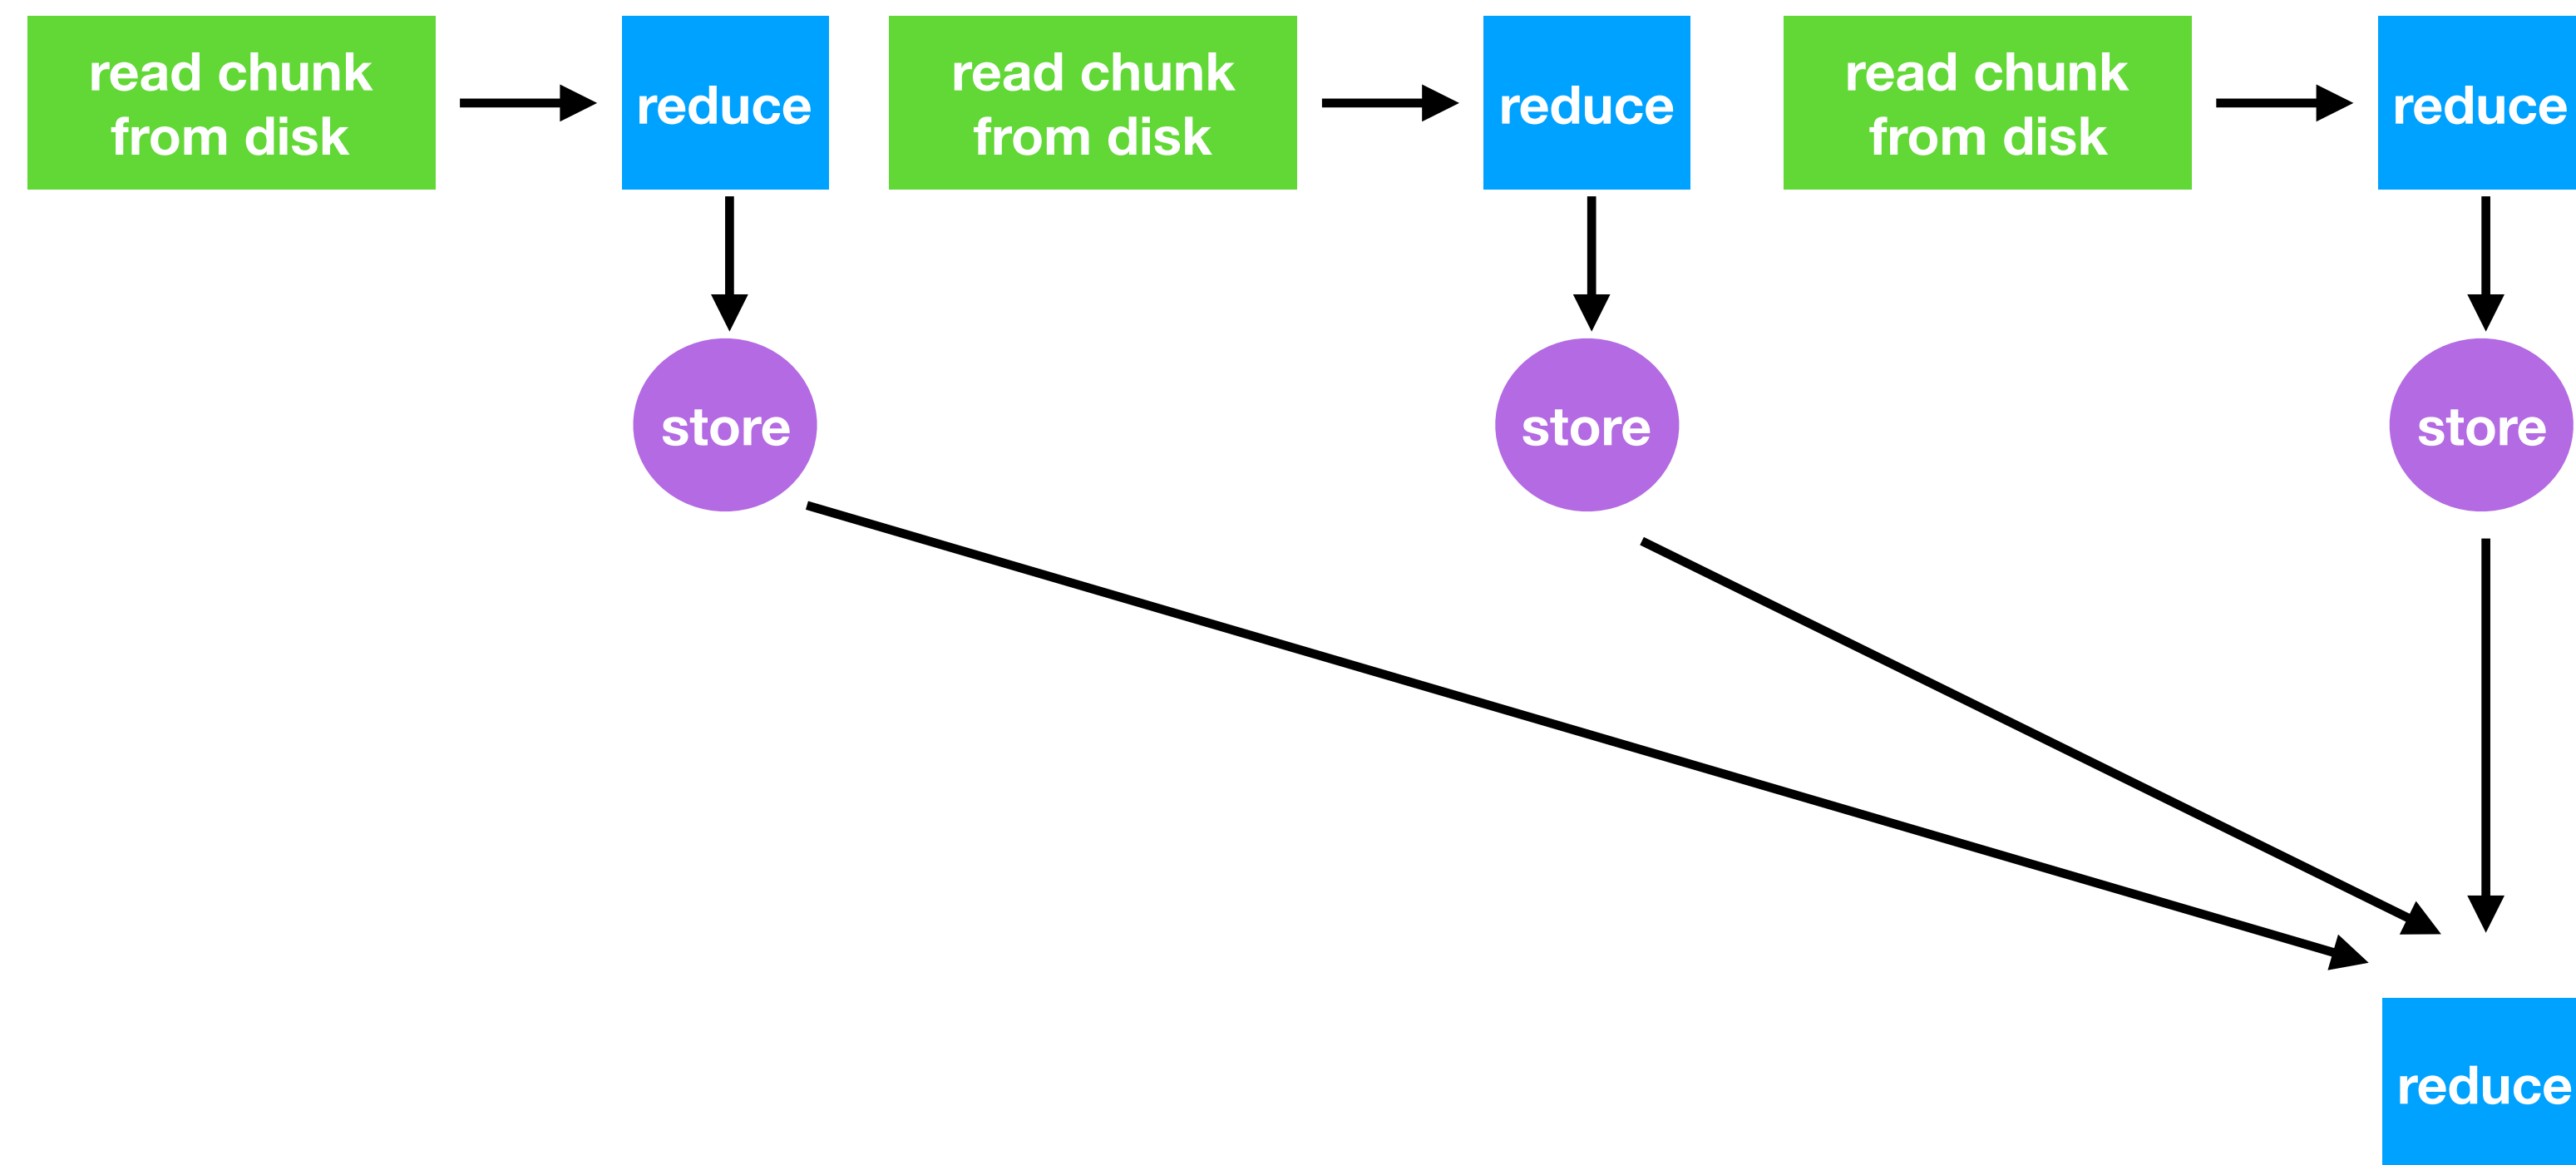


# EXAMPLE CALCULATION: TAKE THE MEAN!

**multidimensional  
array**

	8	8	8
5	('x', 0, 0)	('x', 0, 1)	('x', 0, 2)
5	('x', 1, 0)	('x', 1, 1)	('x', 1, 2)
5	('x', 2, 0)	('x', 2, 1)	('x', 2, 2)
5	('x', 3, 0)	('x', 3, 1)	('x', 3, 2)

**serial execution (a loop)**

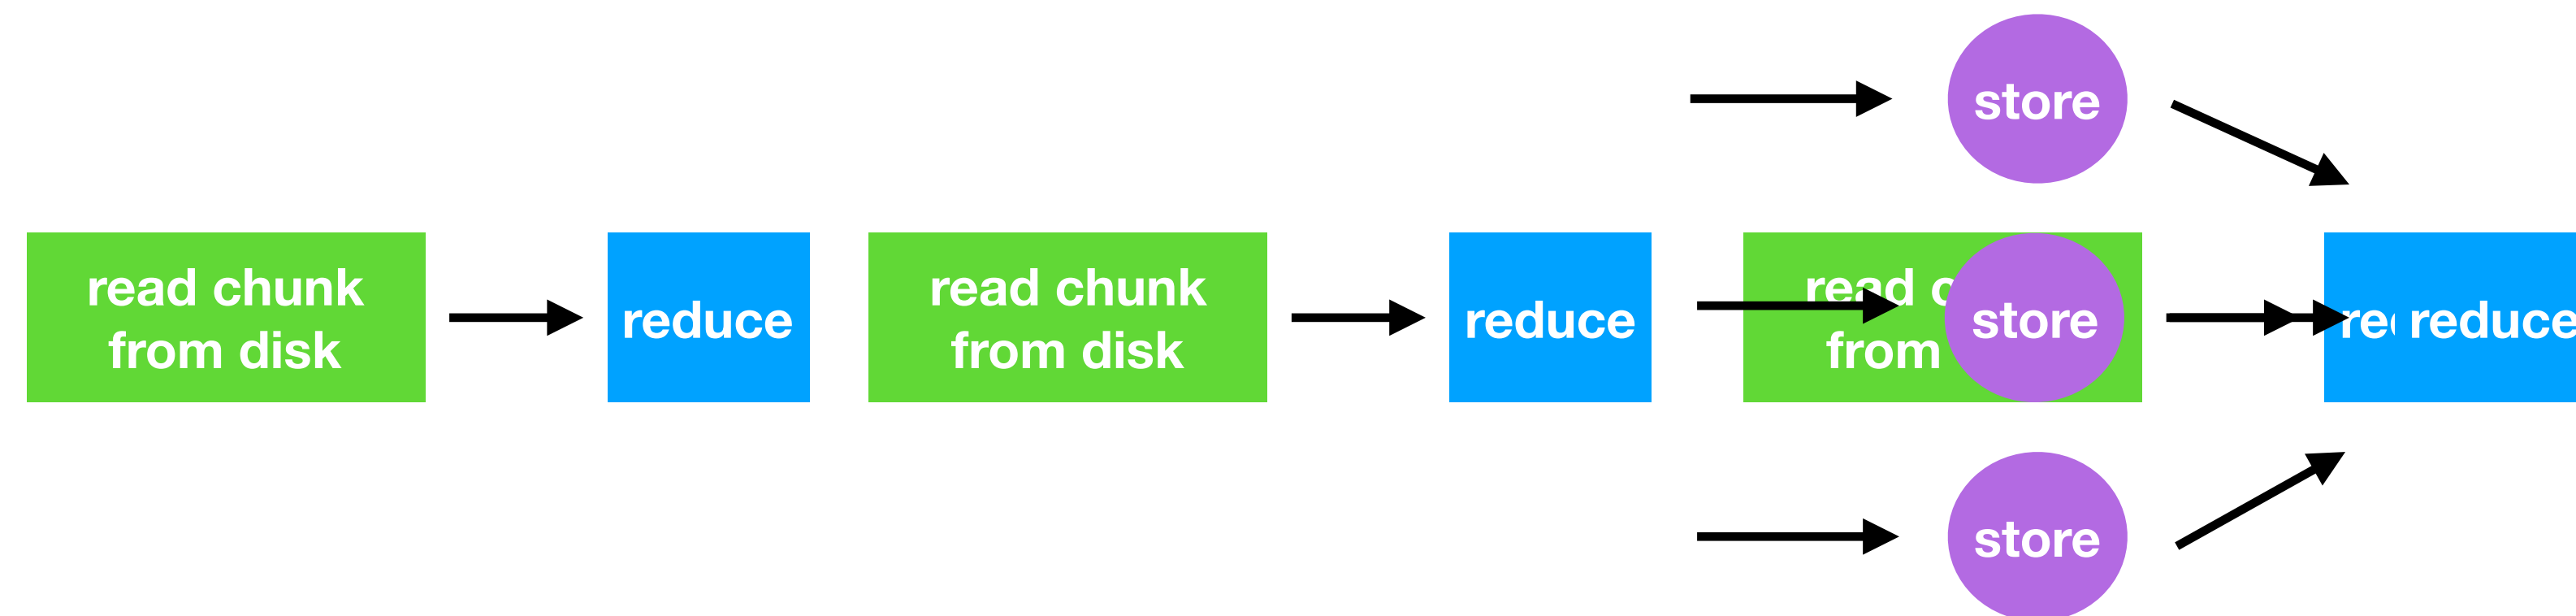




# EXAMPLE CALCULATION: TAKE THE MEAN!

**multidimensional  
array**

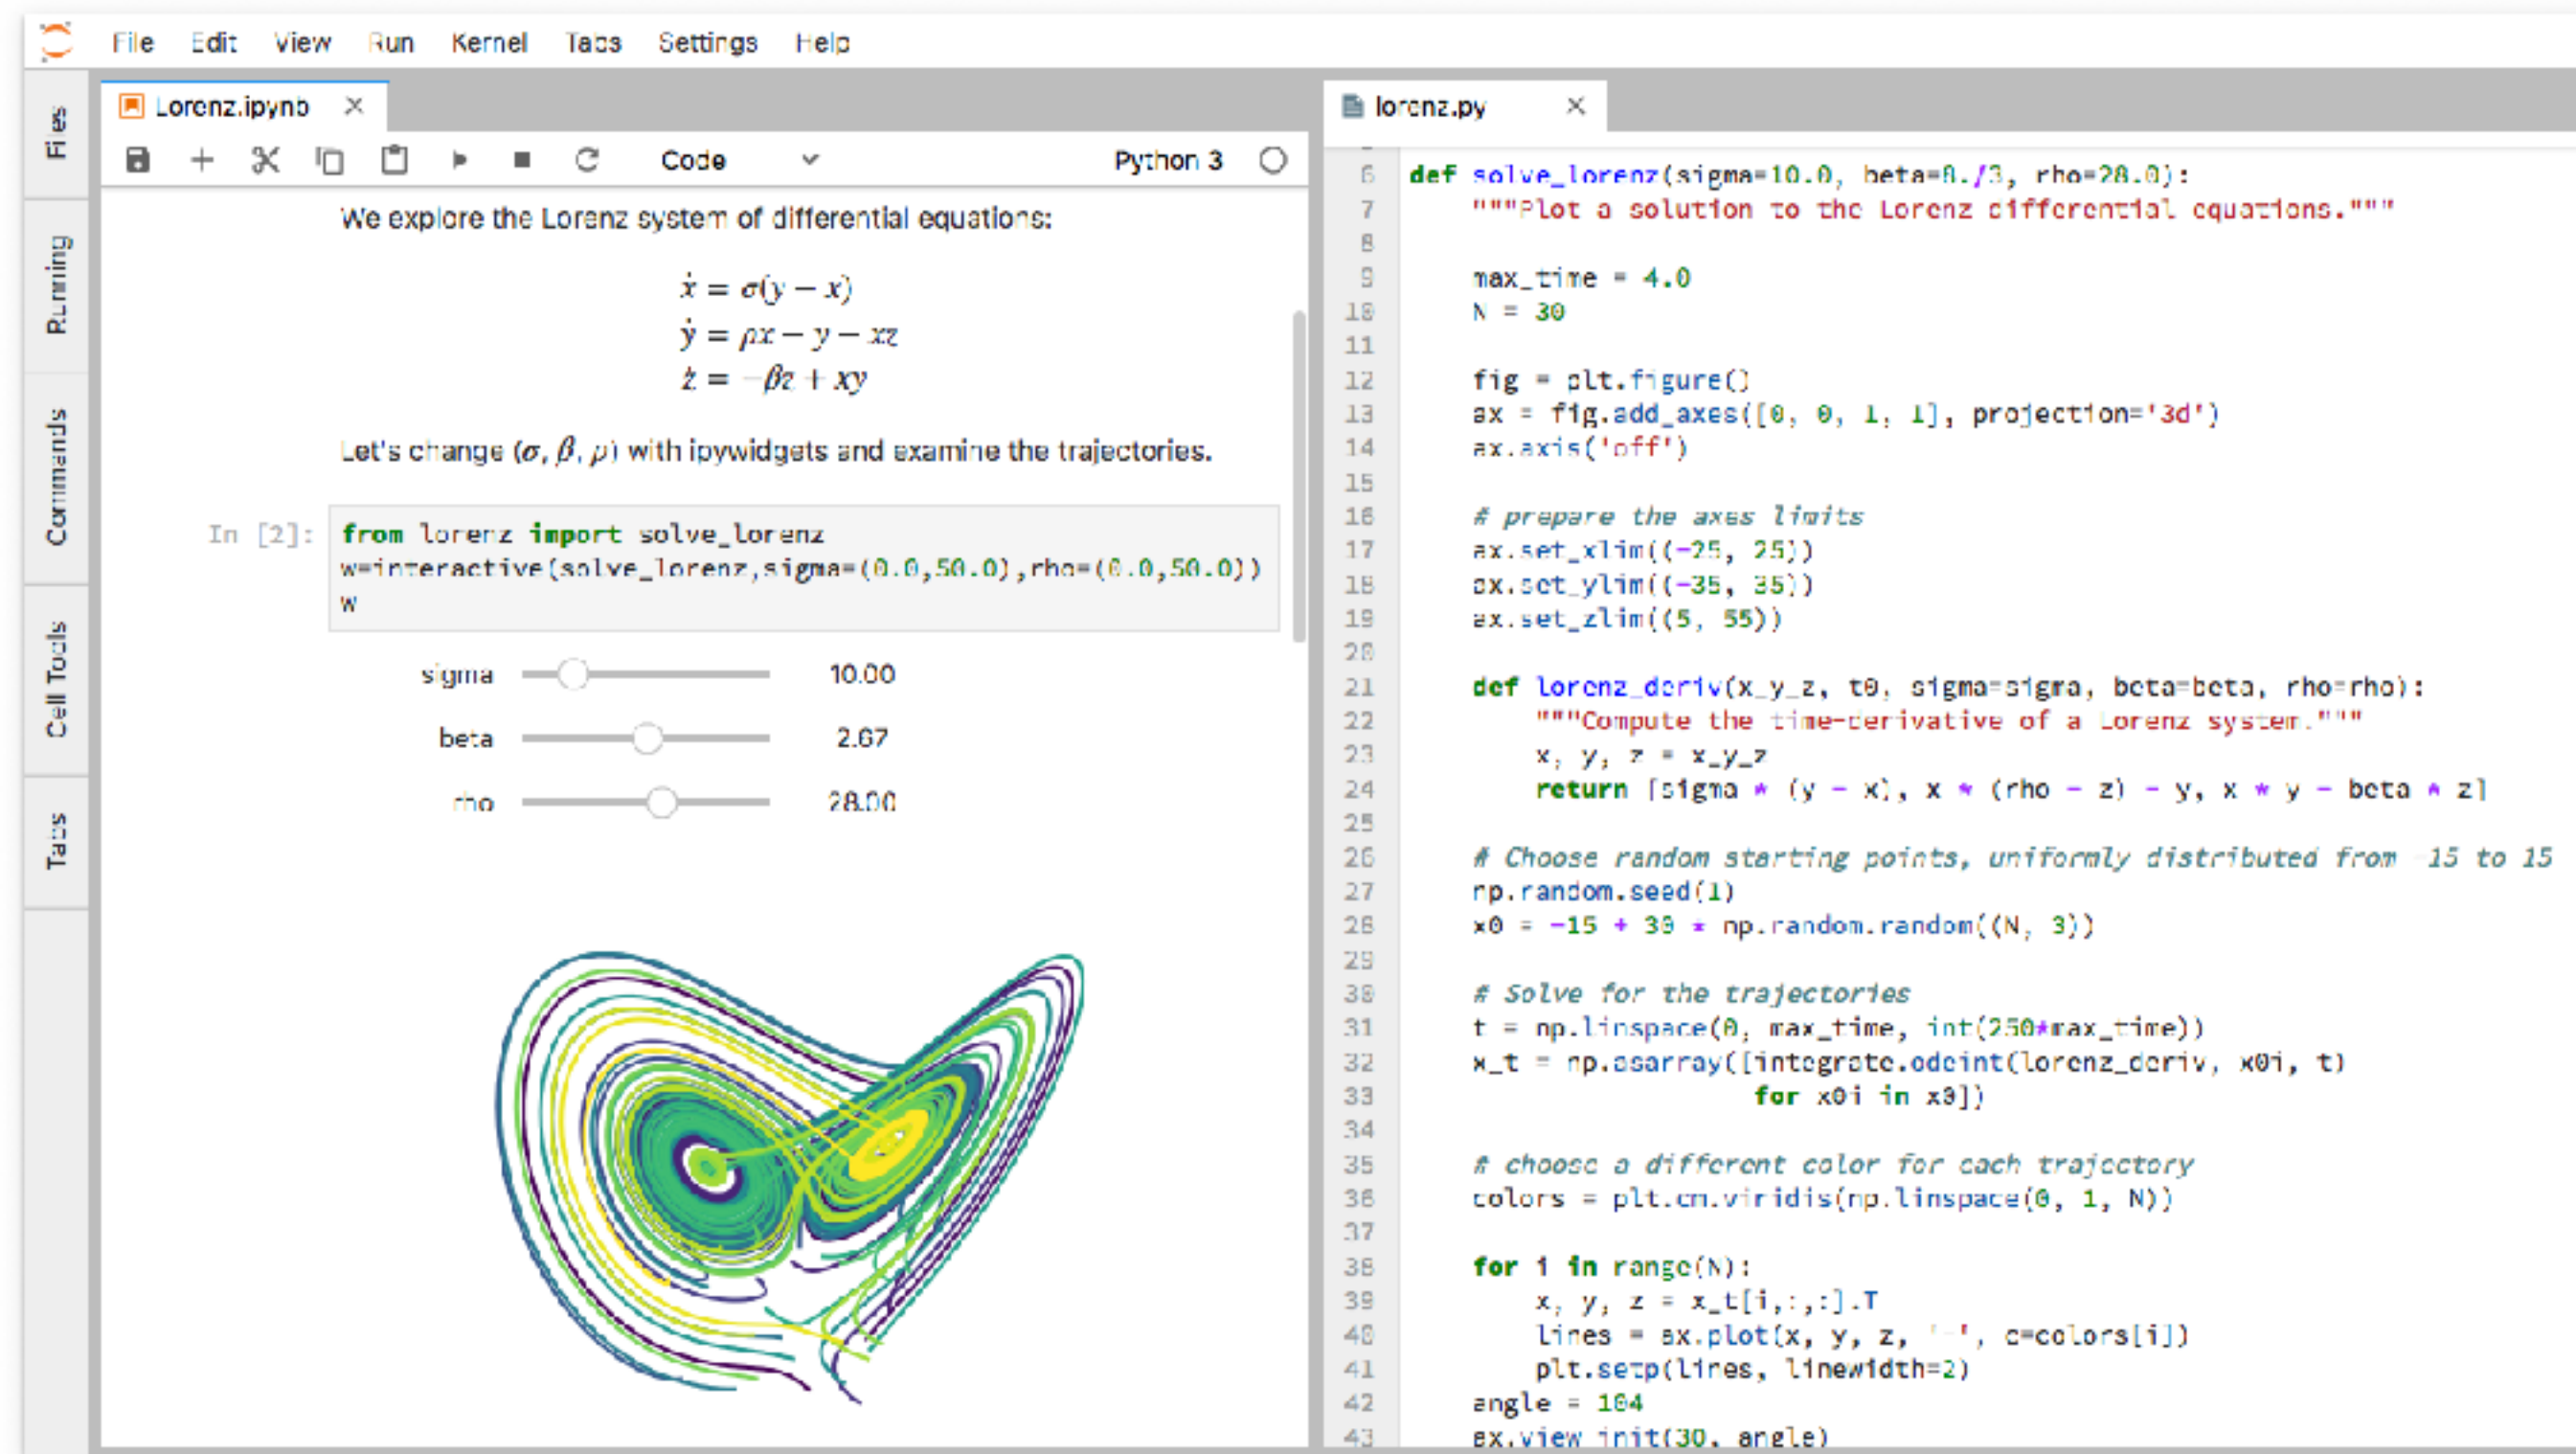

	8	8	8
5	('x', 0, 0)	('x', 0, 1)	('x', 0, 2)
5	('x', 1, 0)	('x', 1, 1)	('x', 1, 2)
5	('x', 2, 0)	('x', 2, 1)	('x', 2, 2)
5	('x', 3, 0)	('x', 3, 1)	('x', 3, 2)



**parallel execution (dask graph)**



# JUPYTER



*“Project Jupyter exists to develop open-source software, open-standards, and services for interactive computing across dozens of programming languages.”*




# PANGEO INFRASTRUCTURE




# FILE-BASED APPROACH

step 1: download



step 2: analyze



**Data provider's responsibilities**

**End user's responsibilities**



# SERVER-SIDE DATABASE

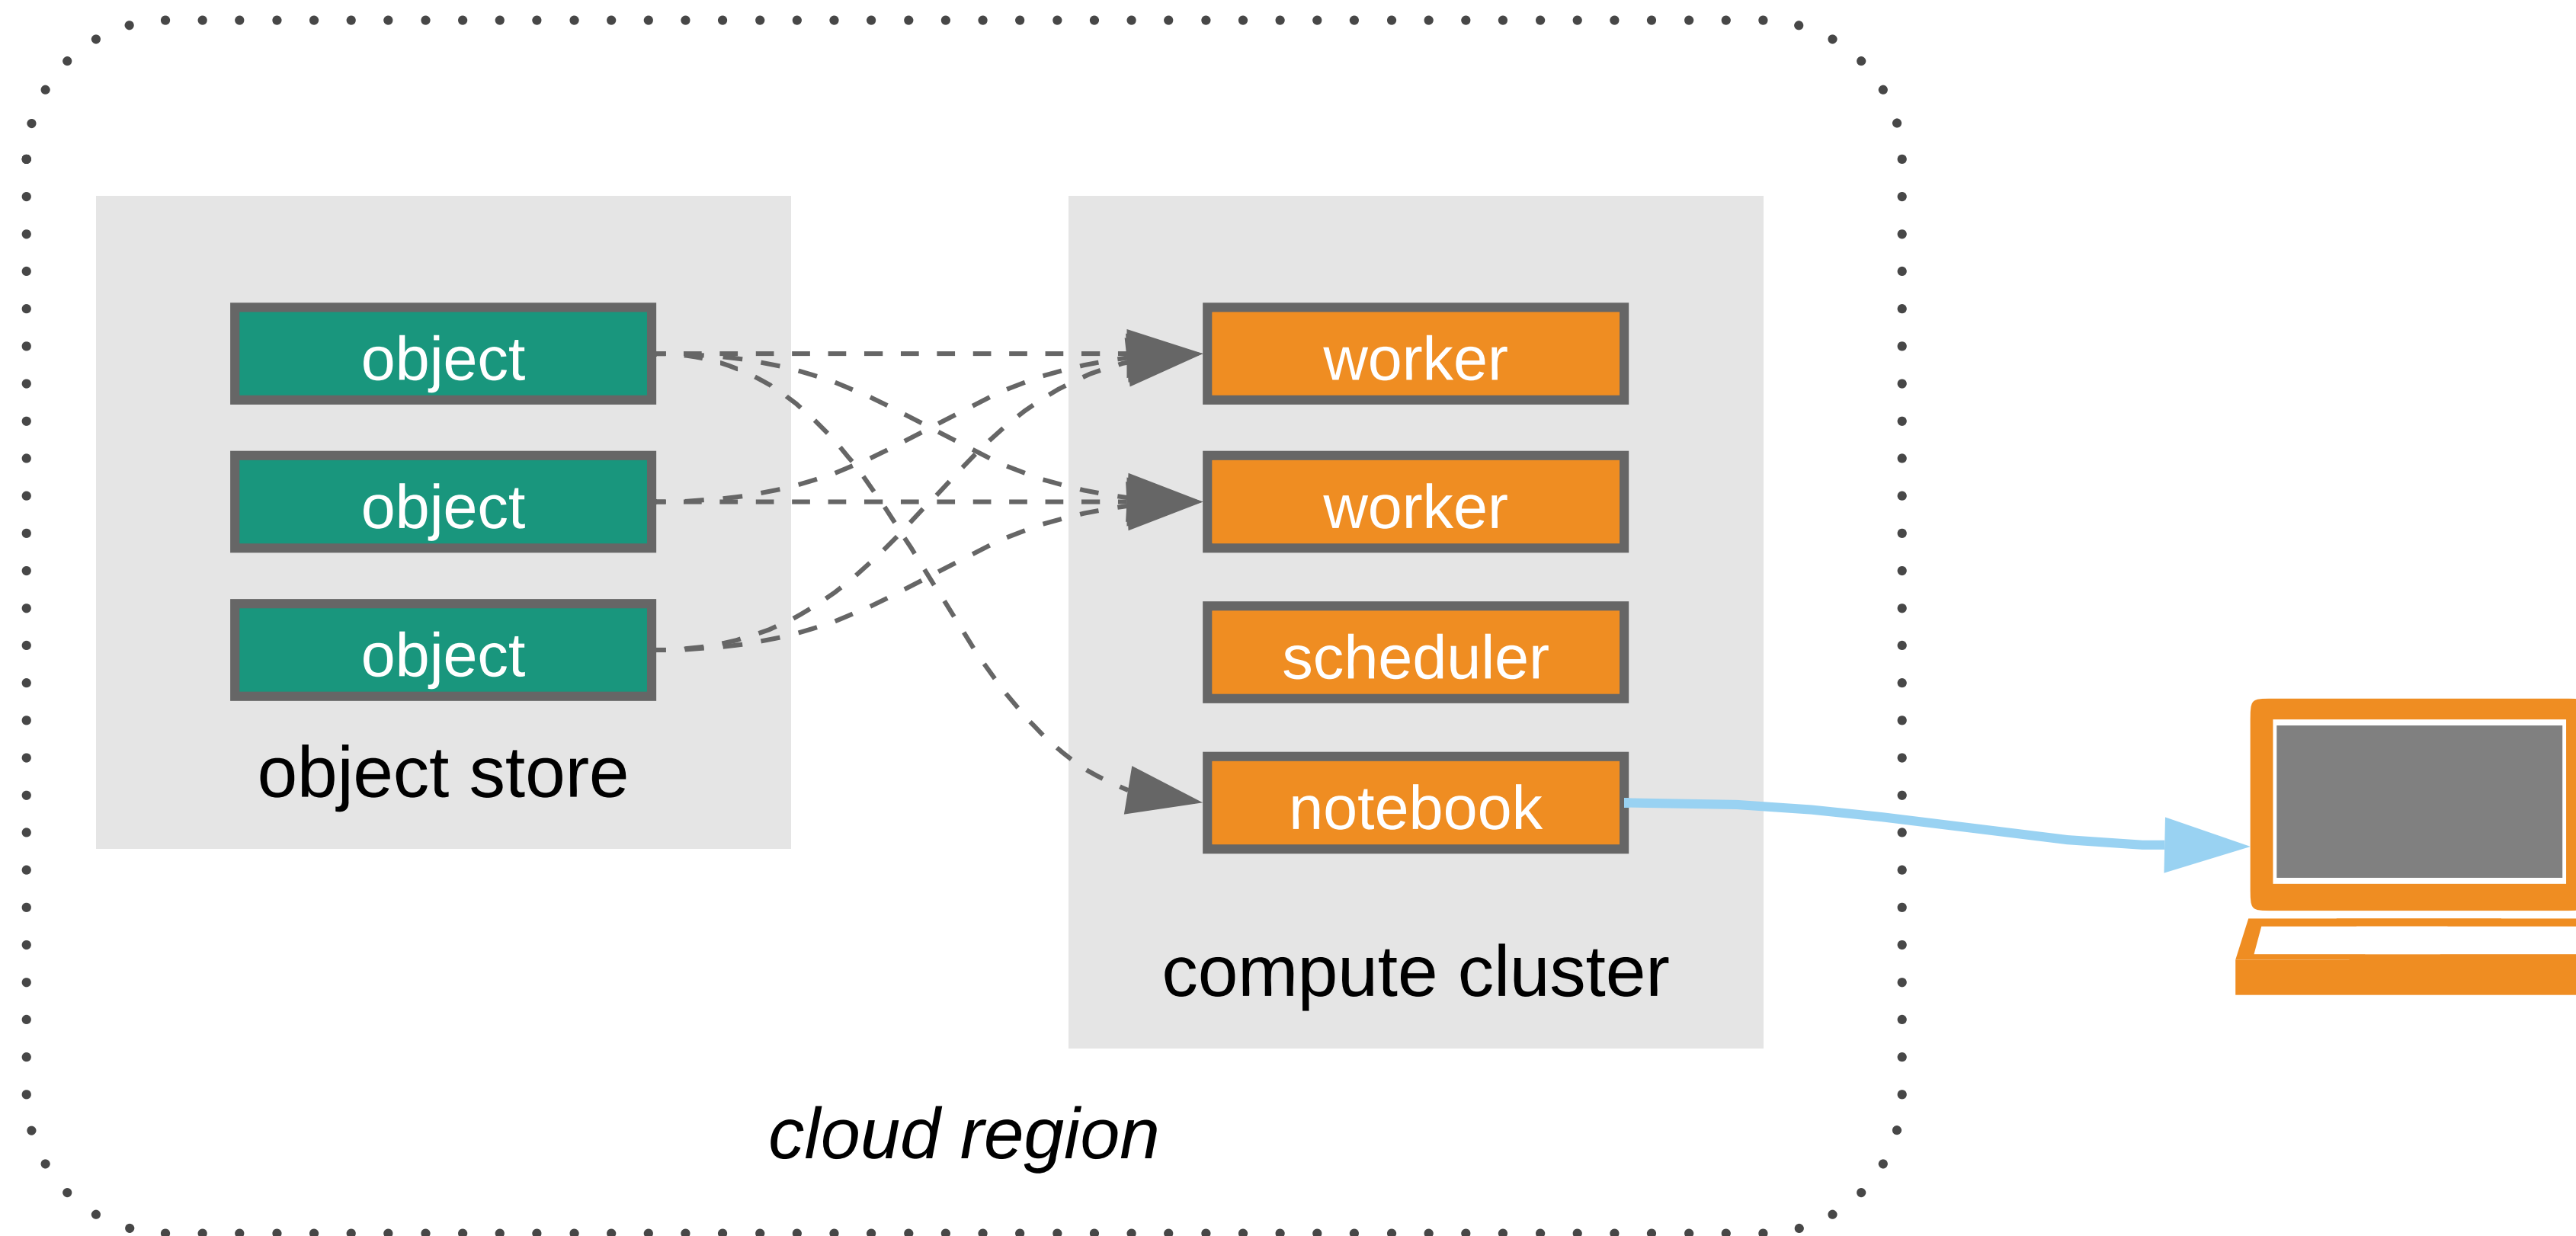


**Data provider's responsibilities**

**End user's responsibilities**



# CLOUD-NATIVE APPROACH



**Data provider's responsibilities**

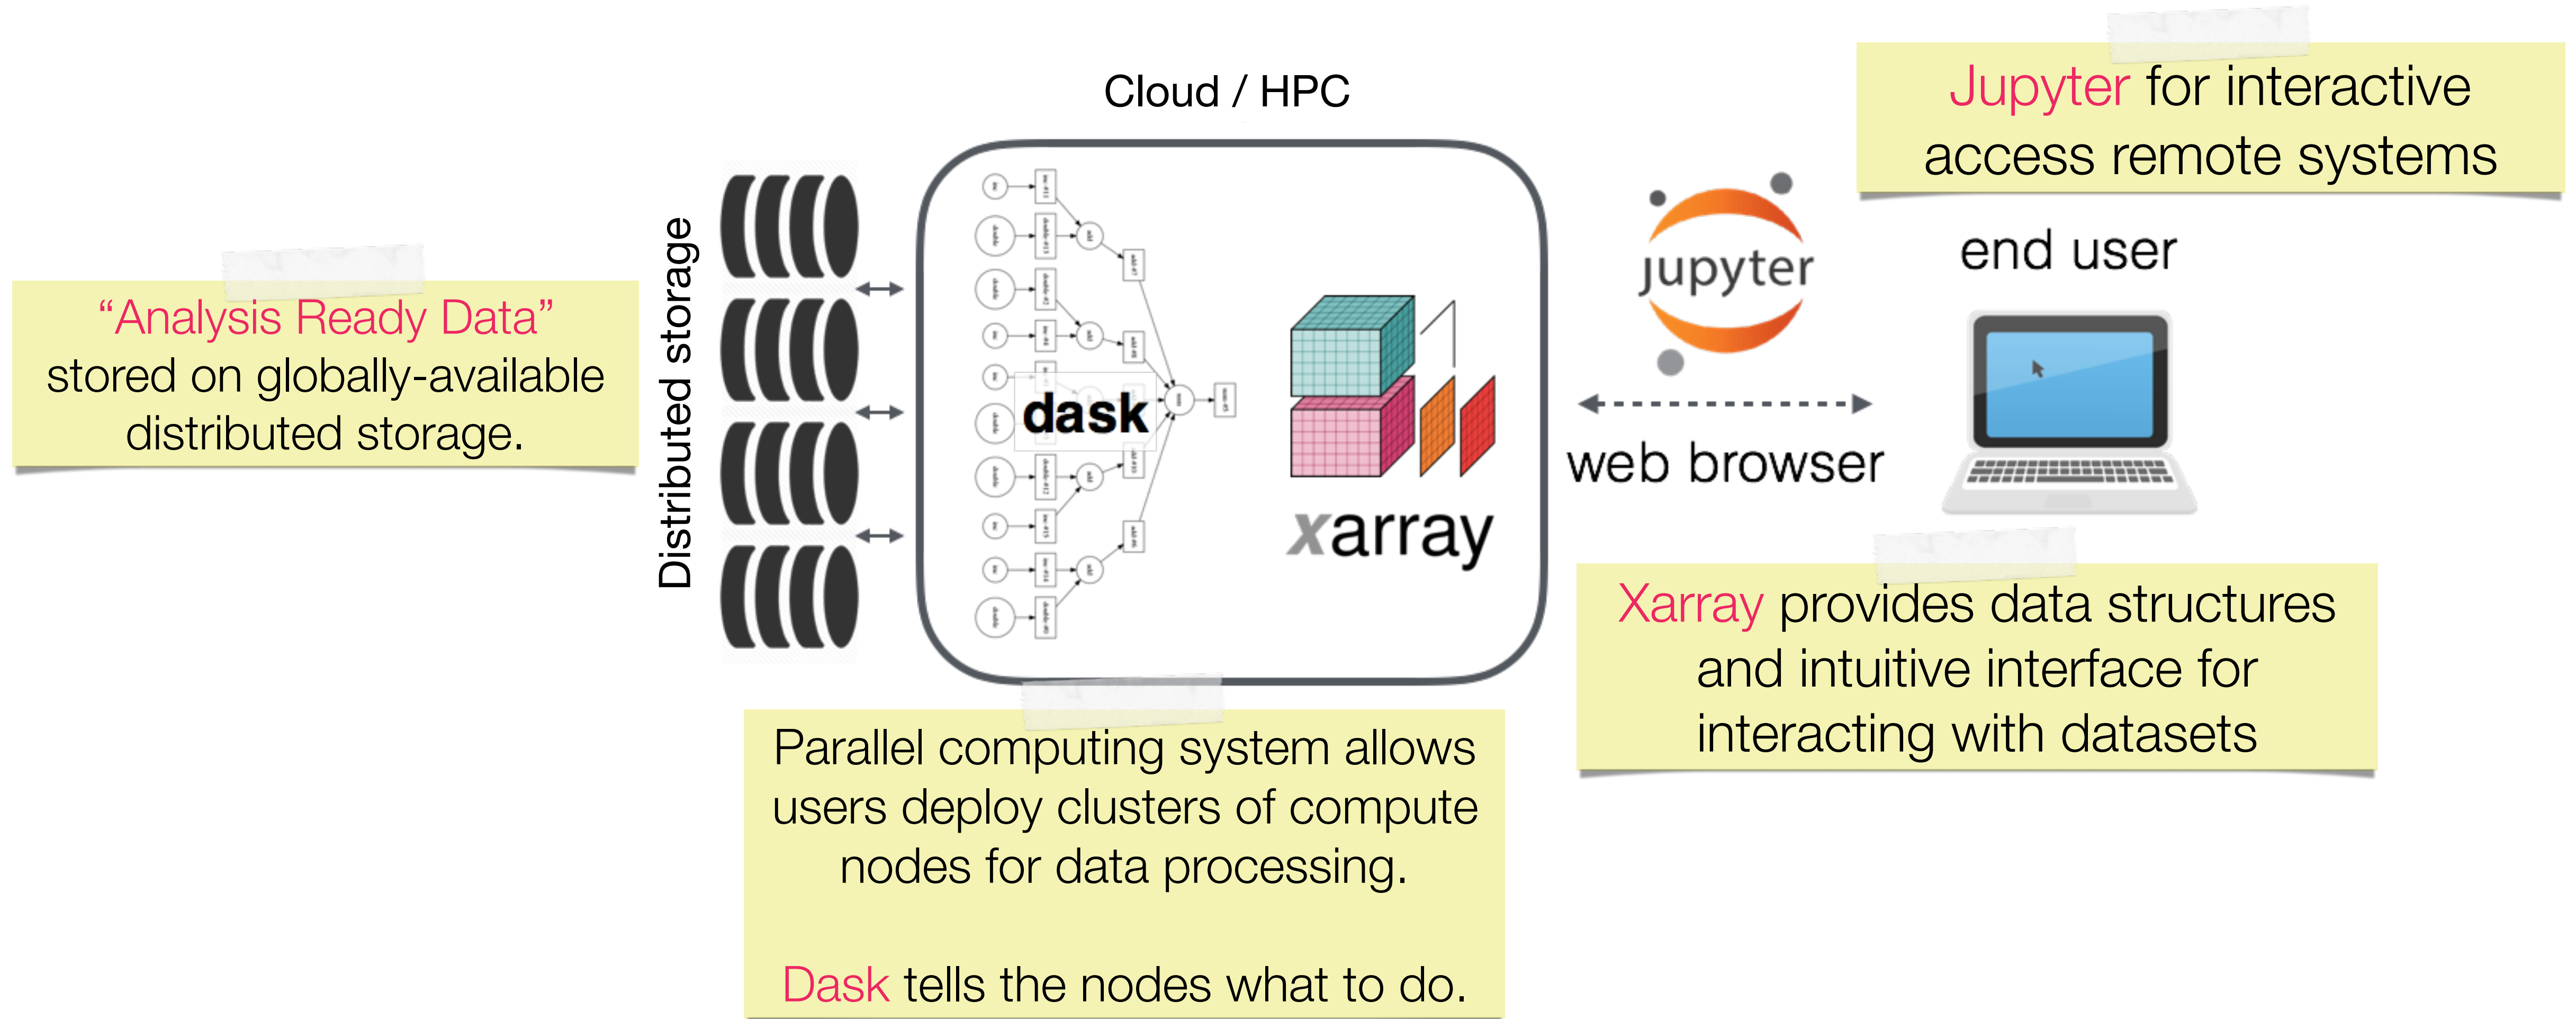
**End user's responsibilities**



# PANGEo PRINCIPLES FOR CLOUD-NATIVE SCIENCE INFRASTRUCTURE

- **Community-driven** - Our needs are no different from those of our peer institutions. By developing infrastructure collaboratively, we can accomplish much more than any one institution can alone.
- **Open source** - Because infrastructure is code, the code should be licensed in a way that enables the entire research community to reuse and build upon it.
- **Modular** - “all in one” solutions are impossible to maintain long term. Separation of concerns is a key principle of good software and systems engineering.
- **Vendor neutral** - Academic research infrastructure should use only vendor-neutral services APIs. If this principle is followed, it means we can redeploy our infrastructure anywhere.

# PANGEO ARCHITECTURE






# BUILD YOUR OWN PANGEO

Storage Formats			Cloud Optimized COG/Zarr/Parquet/etc.
ND-Arrays			More coming...
Data Models			
Processing Mode	 Interactive	Batch 	Serverless 
Compute Platform	HPC 	Cloud 	Local 

# PANGEO DEPLOYMENTS



NASA Pleiades



NCAR Cheyenne



kubernetes



Google Cloud Platform



[HTTP://PANGEO.IO/DEPLOYMENTS.HTML](http://pangeo.io/deployments.html)



# CONTINUOUS DEPLOYMENT

→ ↻ GitHub, Inc. [US] | <https://github.com/pangeo-data/pangeo-cloud-federation> ☆

README.md



This repository manages the continuous deployment of the [Pangeo](#) Cloud Federation JupyterHub Kubernetes clusters using [hubploy](#). It contains scripts to automatically redeploy when the image definition or chart parameters are changed.

Changing the image will typically take ~20 minutes, and changing a Helm config variable ~1 minute.

## Clusters

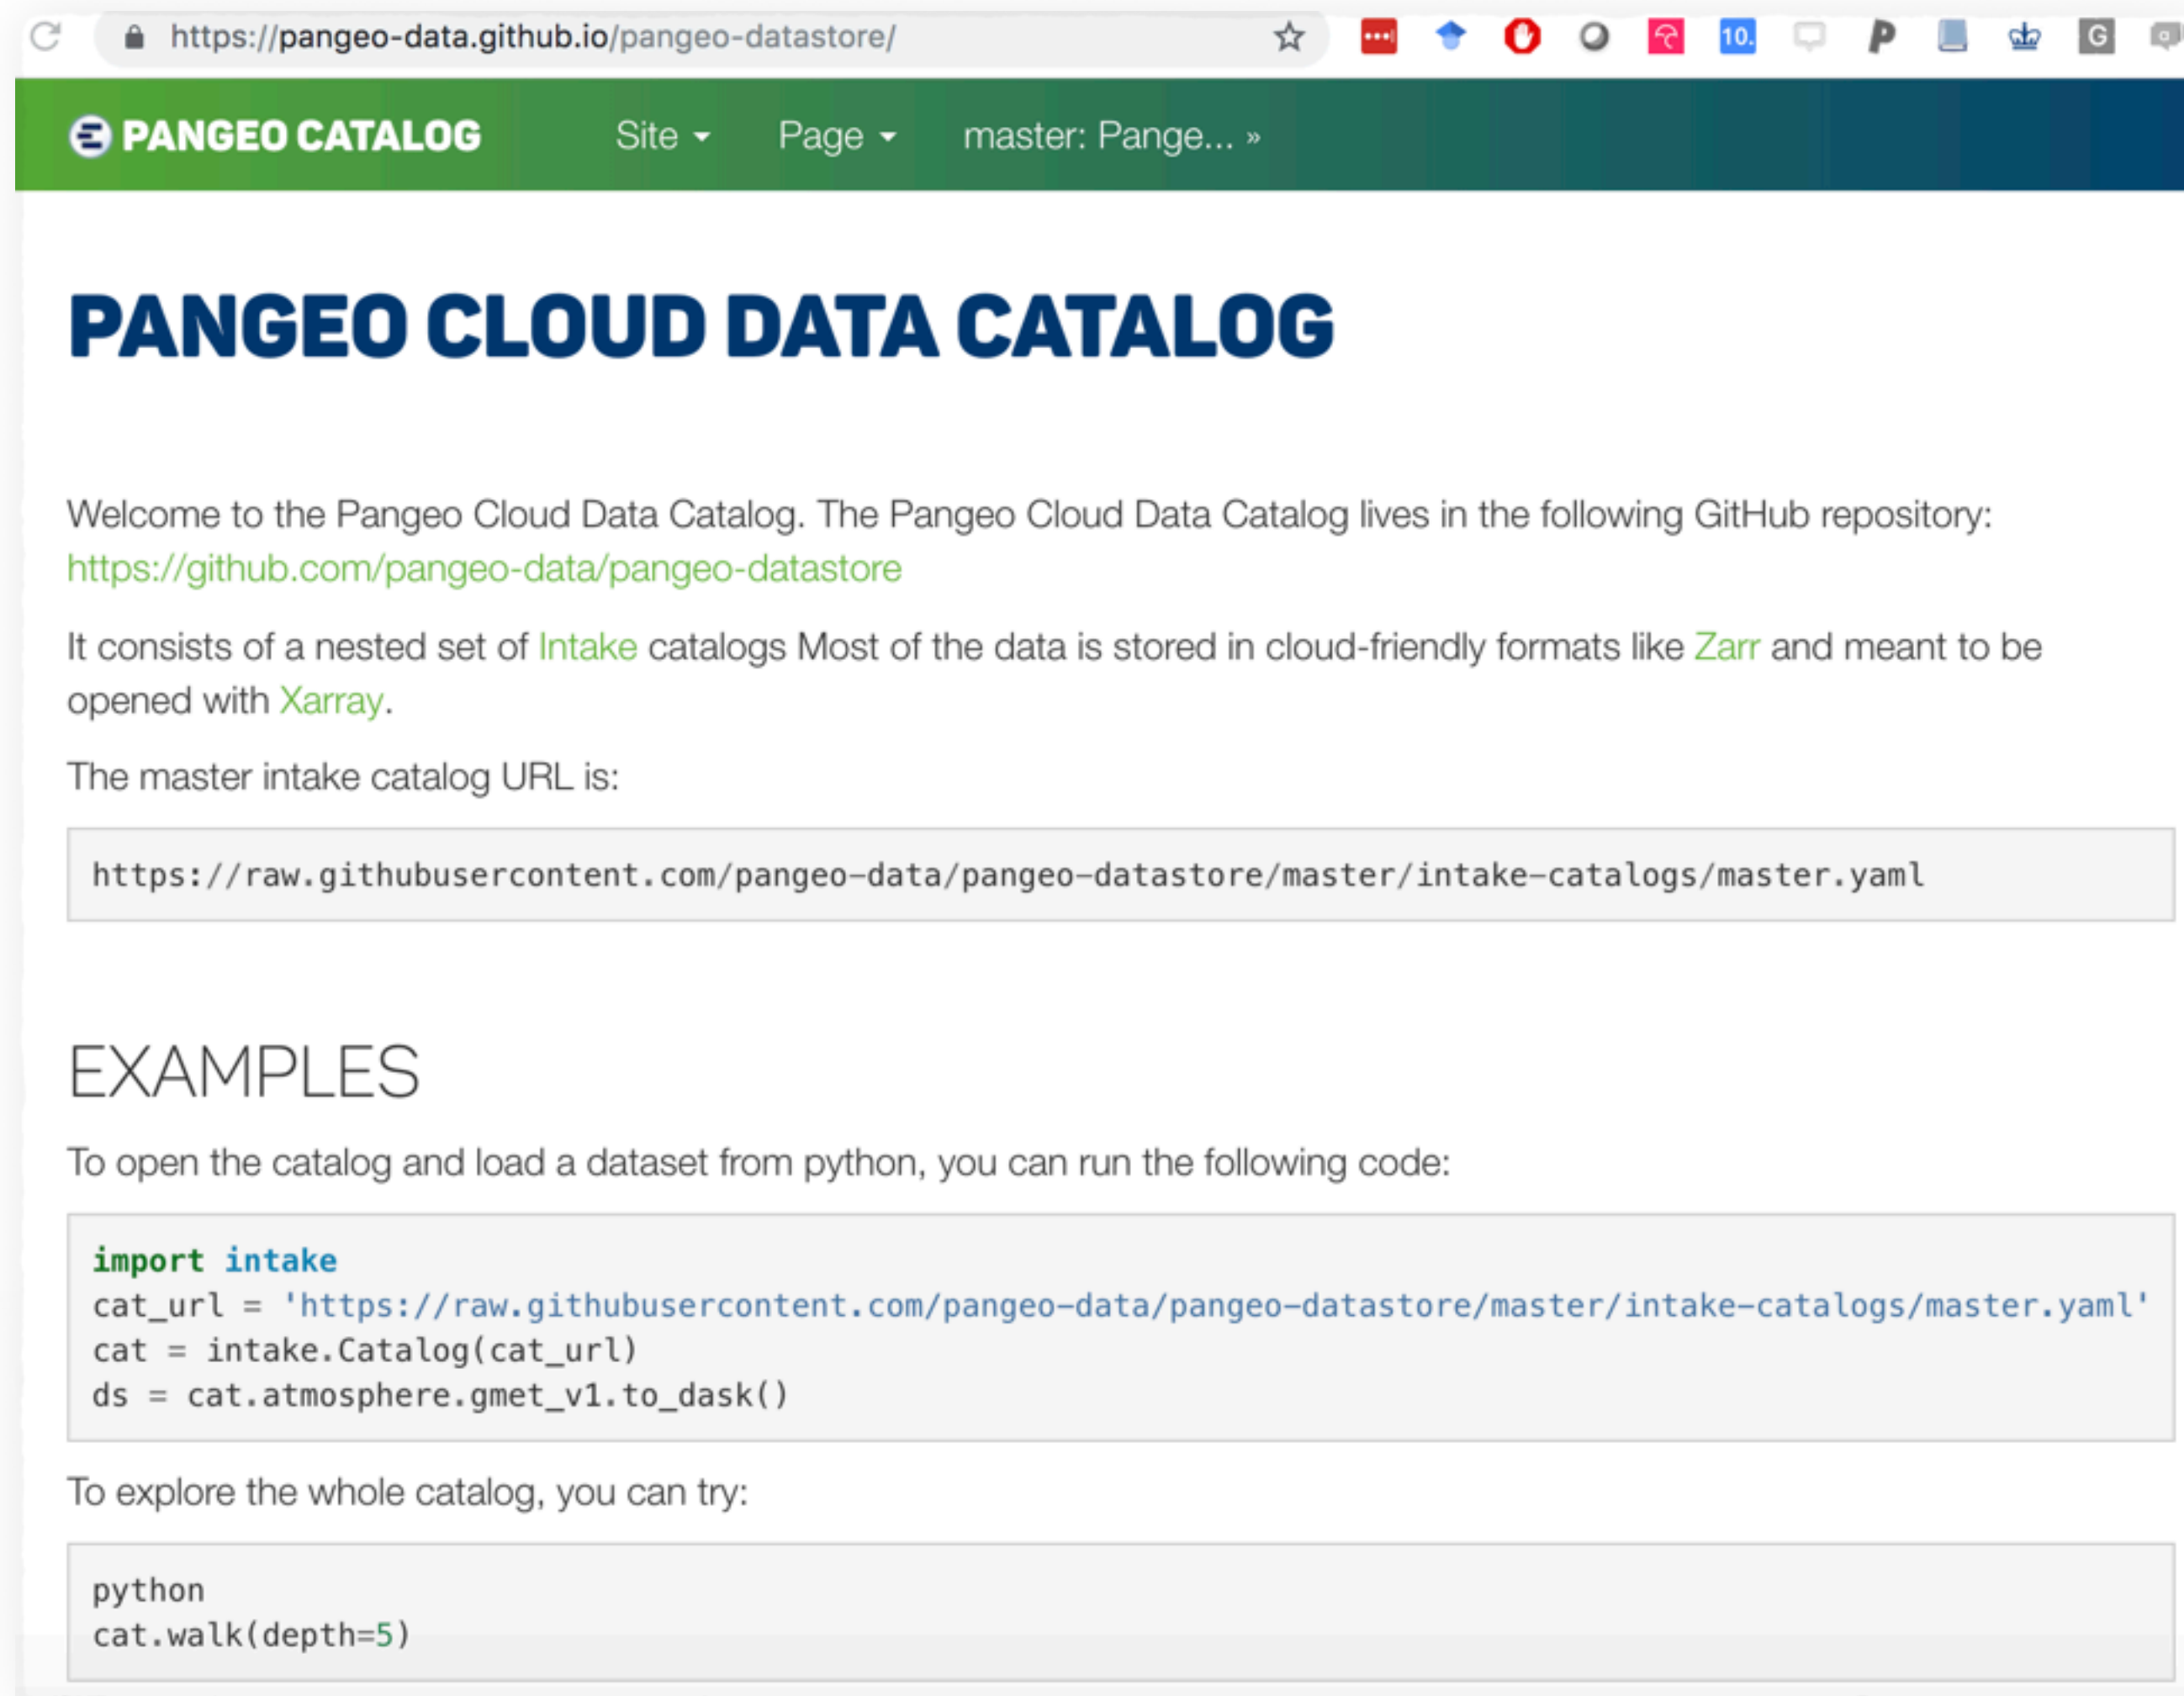
Name	Cloud: region	Staging URL	Production URL
dev	GCP: us-central1-b	<a href="https://staging.hub.pangeo.io">https://staging.hub.pangeo.io</a>	<a href="https://hub.pangeo.io">https://hub.pangeo.io</a>
ocean	GCP: us-central1-b	<a href="https://staging.ocean.pangeo.io">https://staging.ocean.pangeo.io</a>	<a href="https://ocean.pangeo.io">https://ocean.pangeo.io</a>
hydro	GCP: us-central1-b	<a href="https://staging.hydro.pangeo.io">https://staging.hydro.pangeo.io</a>	<a href="https://hydro.pangeo.io">https://hydro.pangeo.io</a>
nasa	AWS: us-east-1	<a href="https://staging.nasa.pangeo.io">https://staging.nasa.pangeo.io</a>	<a href="https://nasa.pangeo.io">https://nasa.pangeo.io</a>
icesat	AWS: us-west-2	<a href="https://staging.icesat.pangeo.io">https://staging.icesat.pangeo.io</a>	<a href="https://icesat.pangeo.io">https://icesat.pangeo.io</a>

## Build Status

Branch	Build
staging	 FAILED
prod	 PASSED

- <https://github.com/pangeo-data/pangeo-cloud-federation>
- Cloud-based clusters managed with helm /kubernetes
- Deployment is completely automated via GitHub / circleci
- Resources scale elastically with demand

# CLOUD DATA CATALOG



The screenshot shows a web browser window with the URL <https://pangeo-data.github.io/pangeo-datastore/>. The page has a green header with the PANGEO CATALOG logo and navigation links for Site, Page, and master: Pange... ». The main heading is **PANGEO CLOUD DATA CATALOG**. Below this, a welcome message states: "Welcome to the Pangeo Cloud Data Catalog. The Pangeo Cloud Data Catalog lives in the following GitHub repository: <https://github.com/pangeo-data/pangeo-datastore>". It then explains: "It consists of a nested set of **Intake** catalogs. Most of the data is stored in cloud-friendly formats like **Zarr** and meant to be opened with **Xarray**." The master intake catalog URL is provided in a code block: `https://raw.githubusercontent.com/pangeo-data/pangeo-datastore/master/intake-catalogs/master.yaml`. The **EXAMPLES** section follows, with instructions on how to open the catalog and load a dataset from python, accompanied by a code block: 

```
import intake
cat_url = 'https://raw.githubusercontent.com/pangeo-data/pangeo-datastore/master/intake-catalogs/master.yaml'
cat = intake.Catalog(cat_url)
ds = cat.atmosphere.gmet_v1.to_dask()
```



 Finally, it suggests exploring the whole catalog with the command: 

```
python
cat.walk(depth=5)
```

- <https://pangeo-data.github.io/pangeo-datastore/>
- Datasets stored in **zarr** format (cloud-native HDF-replacement)
- Cataloged using **intake**
- Automated testing of datasets



# CLOUD COSTS



# DEMO

<https://tinyurl.com/pangeo-ocean>





### Government HPC



### Commercial Cloud

Access	<div>✓ Available to all federally funded projects</div> <div>✗ Available <i>only</i> to federally funded projects</div>	<div>✓ Available globally to anyone with a credit card</div> <div>✗ Authentication is not integrated with existing research infrastructure</div>
Cost	<div>✓ Cost is hidden from researchers and billed by funding agencies</div> <div>✗ Allocations, quotas, limits</div>	<div>✗ Cost is borne by individual researchers and hidden from funding agencies</div> <div>✓ Economics of scale, unlimited resources</div>
Compute	<div>✓ Homogeneous, high performance nodes</div> <div>✗ Queues, batch scheduling, ssh access</div> <div>✗ Fixed-size compute</div>	<div>✓ Flexible hardware (big, small, GPU)</div> <div>✓ Instant provisioning of unlimited resources</div> <div>✓ Spot market: burstable, volatile</div>
Storage	<div>✓ Fast parallel filesystems (e.g. GPFS)</div>	<div>✓ Fast object storage</div>

# HOW TO GET INVOLVED

[HTTP://PANGEO.IO](http://pangeo.io)

- Use and contribute to xarray, dask, zarr, jupyterhub, etc.
- Access an existing Pangeo deployment on an HPC cluster, or cloud resources (<http://pangeo.io/deployments.html>)
- Adapt Pangeo elements to meet your projects needs (data portals, etc.) and give feedback via github: [github.com/pangeo-data/pangeo](https://github.com/pangeo-data/pangeo)
- Provide data in a cloud-optimized format



**The following slide is a backup in case live demo is not possible.**

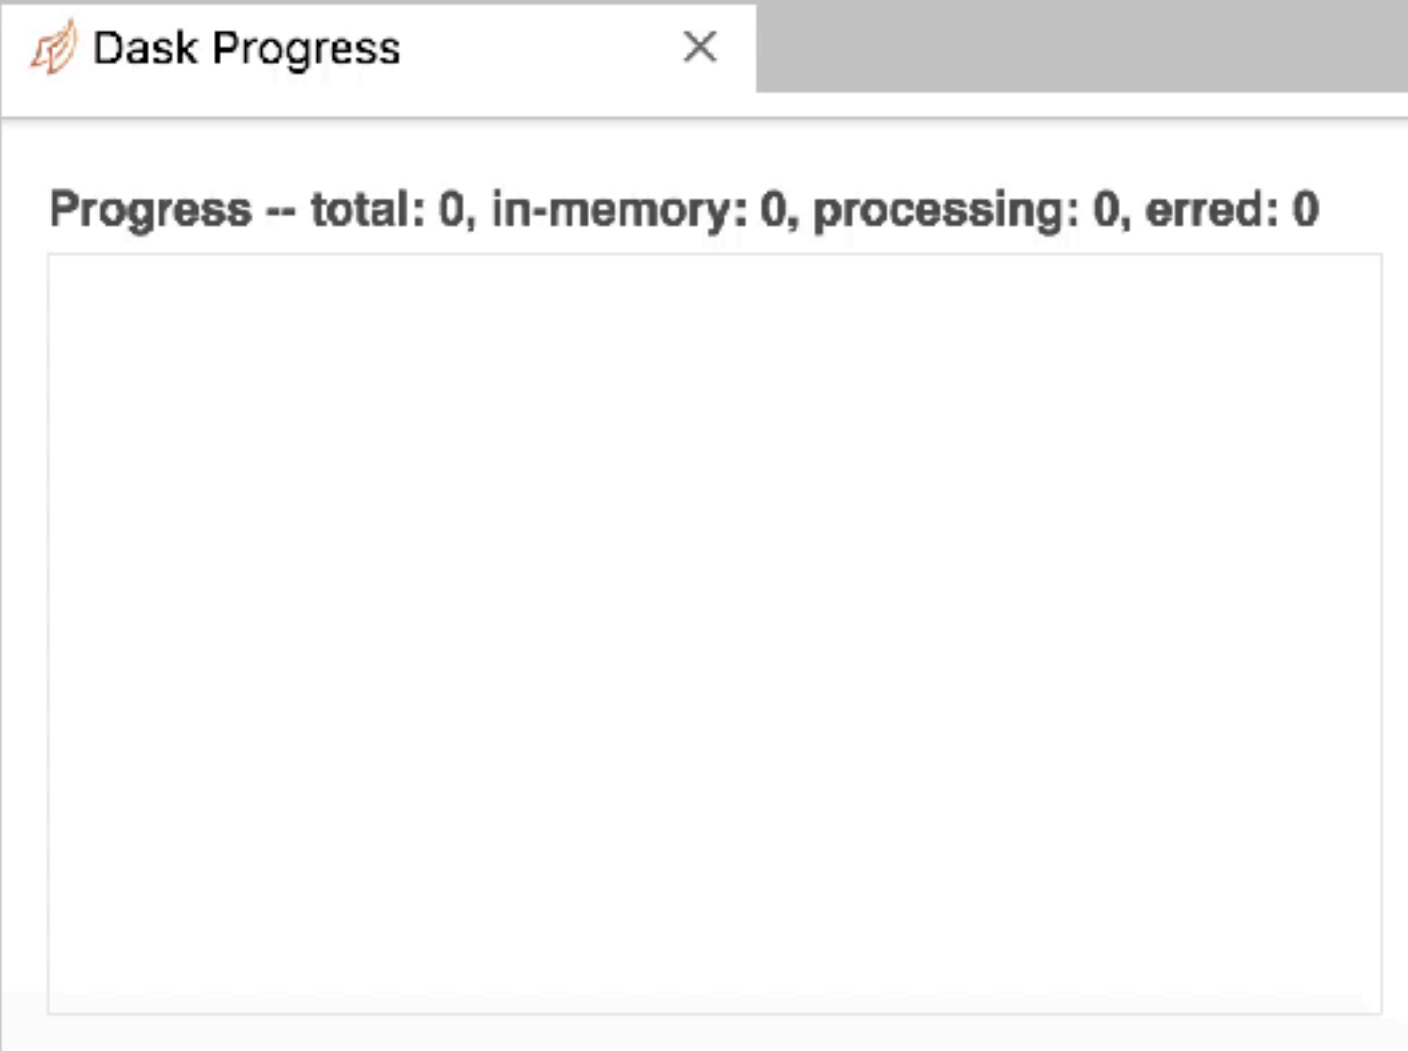
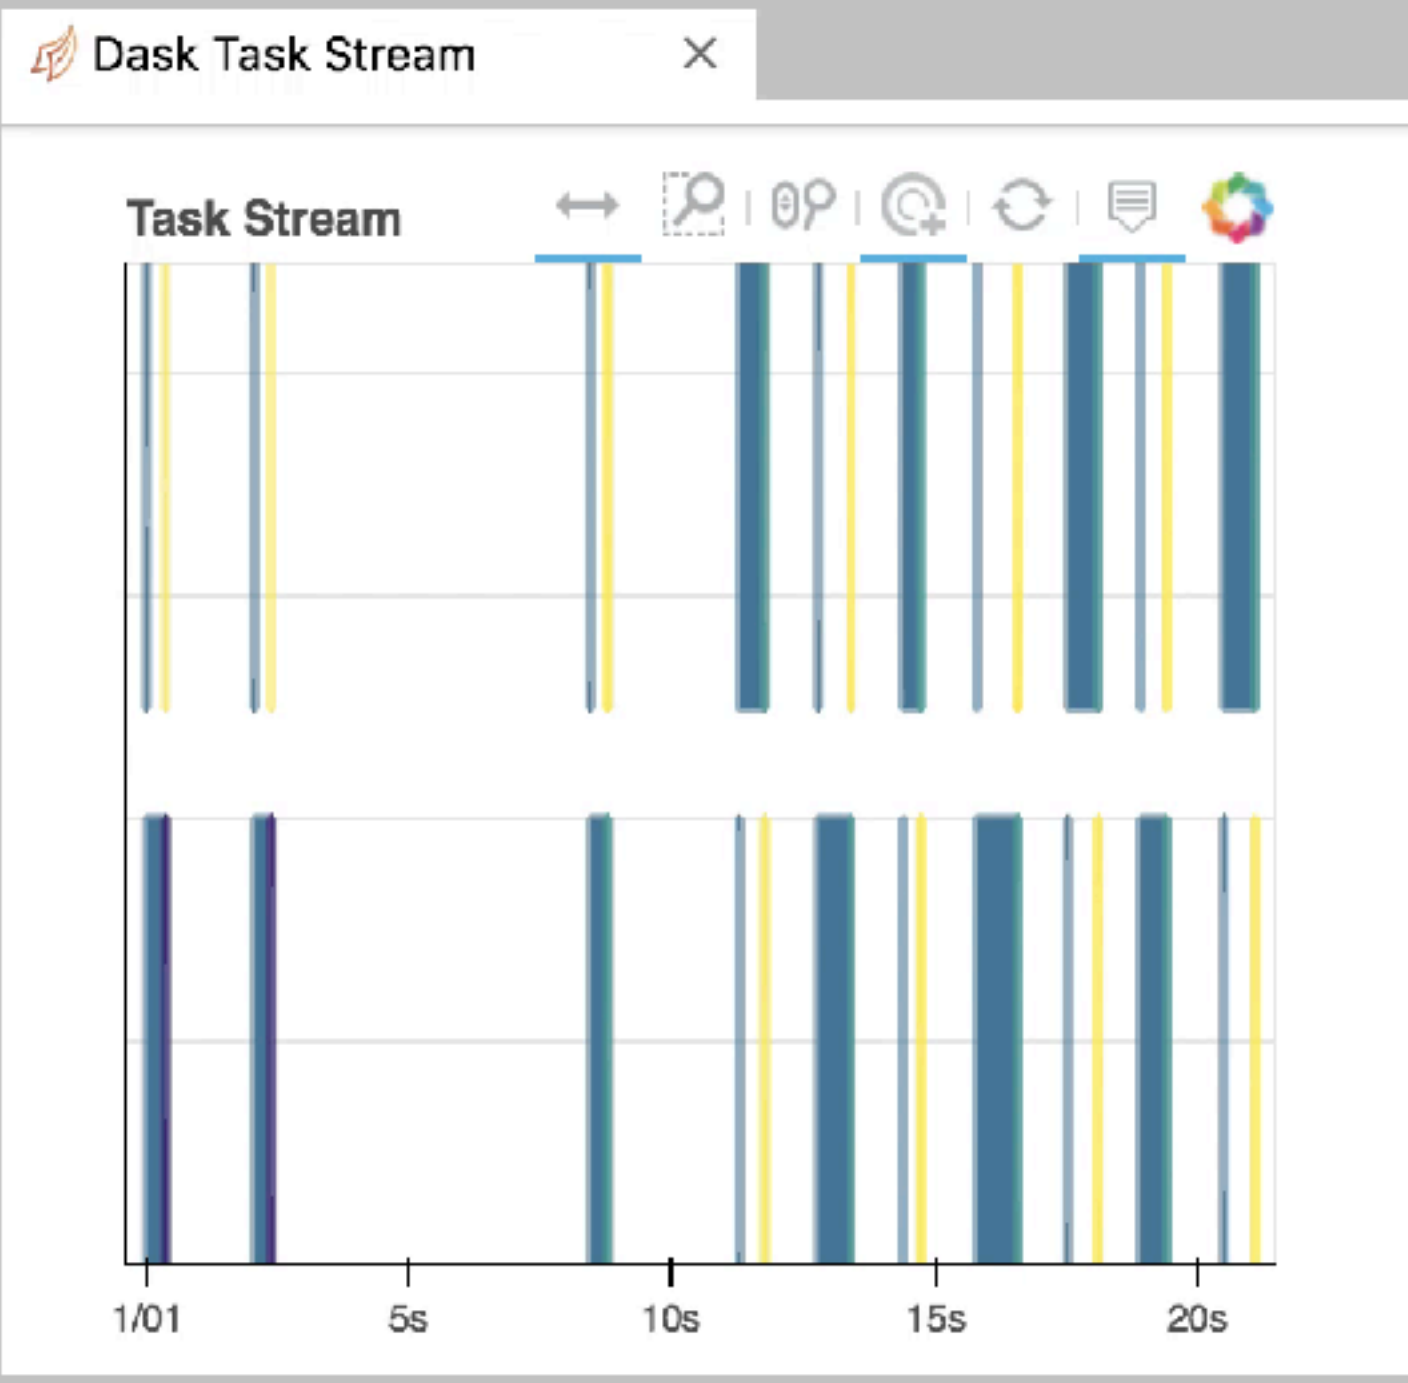
```
cesm-pop-highres-ocean
Python 3

[5]: <xarray.DataArray 'SST' (time: 14965, nlat: 2400, nlon: 3600)>
      dask.array<shape=(14965, 2400, 3600), dtype=float32, chunksize=(1, 2400, 3600)>
      Coordinates:
        * time      (time) float64 1.679e+04 1.679e+04 ... 3.175e+04 3.176e+04
        * nlon      (nlon) int64 0 1 2 3 4 5 6 7 ... 3593 3594 3595 3596 3597 3598 3599
        * nlat      (nlat) int64 0 1 2 3 4 5 6 7 ... 2393 2394 2395 2396 2397 2398 2399
          lon       (nlat, nlon) float64 dask.array<shape=(2400, 3600), chunksize=(2400, 3600)>
          lat       (nlat, nlon) float64 dask.array<shape=(2400, 3600), chunksize=(2400, 3600)>
      Attributes:
        cell_methods:  time: mean
        grid_loc:      2110
        long_name:     Surface Potential Temperature
        units:         degC

[ ]: %%output holomap='scrubber' fps=1

      sst_ds = hv.Dataset(sst, kdims=['time', 'nlon', 'nlat'])
      hv_sst = sst_ds.to(hv.Image, kdims=["nlon", "nlat"], dynamic=True)
      %opts Image [width=700 height=400] (cmap='magma')
      regrid(hv_sst, precompute=True)

[ ]:
```





- 12 GB / 24 hours -> must be a mistake