

Cloud-based Cluster Computing for PyData

Olivier Grisel - SUCCES, Paris 2016



Outline

Demo of distributed ML in Python

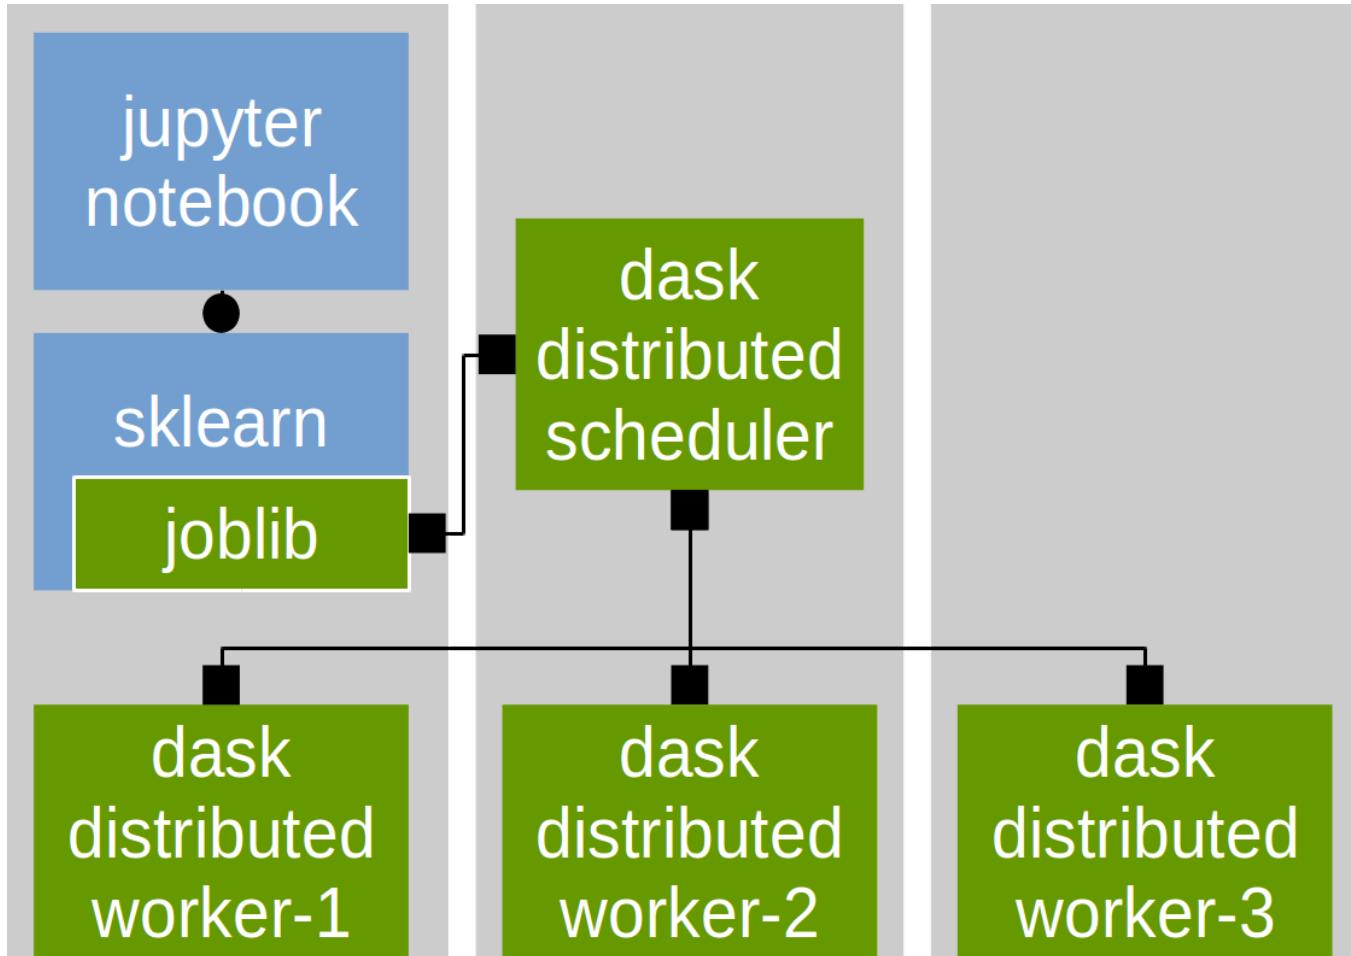
Some PyData tools for interactive distributed computing

Cluster orchestration with Docker and Kubernetes



Distributed ML with joblib
and dask/distributed

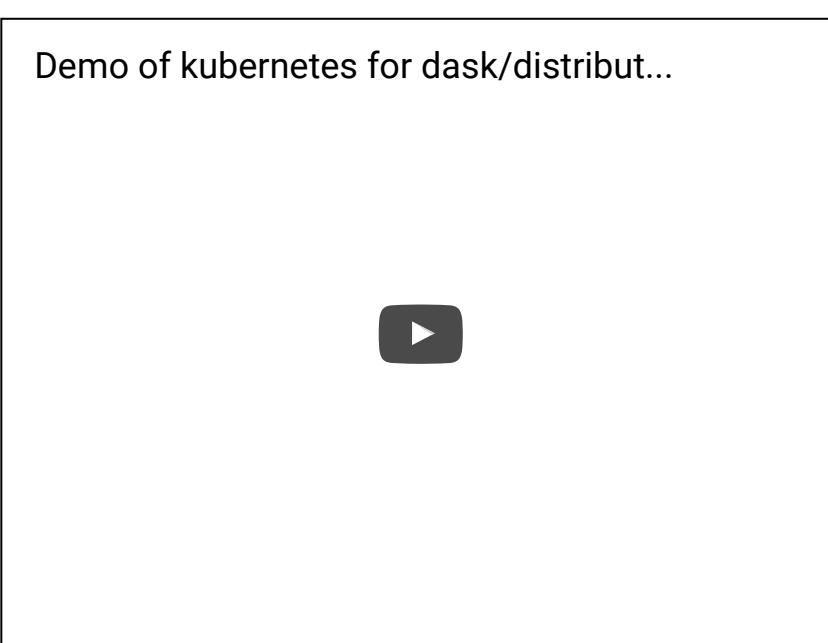
A cluster for sklearn



Demo

<https://github.com/oqrisel/docker-distributed>

Have a look at the notebooks in the examples folder.



<https://youtu.be/6mKNSEQ0FIQ>

Jupyter notebook

The screenshot shows a Jupyter Notebook interface with two panes. The left pane is a 'Welcome to the Jupyter Notebook' page, and the right pane is a notebook cell titled 'Exploring the Lorenz System'.

Exploring the Lorenz System

In this Notebook we explore the [Lorenz system](#) of differential equations:

$$\begin{aligned}\dot{x} &= \sigma(y - x) \\ \dot{y} &= \rho x - y - xz \\ \dot{z} &= -\beta z + xy\end{aligned}$$

This is one of the classic systems in non-linear differential equations. It exhibits a range of complex behaviors as the parameters (σ , β , ρ) are varied, including what are known as *chaotic solutions*. The system was originally developed as a simplified mathematical model for atmospheric convection in 1963.

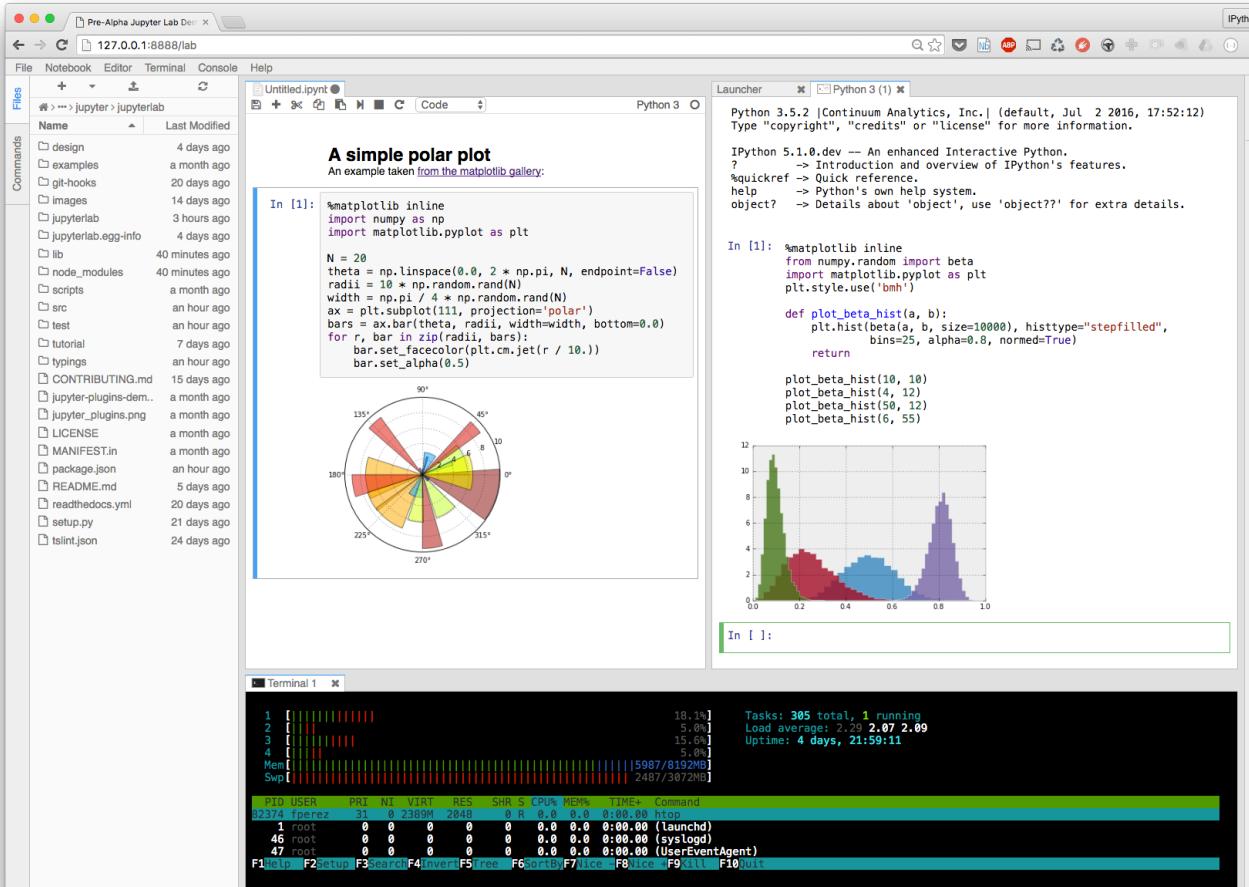
In [7]:

```
interact(Lorenz, N=fixed(10), angle=(0.,360.),
          σ=(0.0,50.0),β=(0.,5), ρ=(0.0,50.0));
```

angle: 308.2
max_time: 12
 σ : 10
 β : 2.6
 ρ : 28

Figure: A 3D plot of the Lorenz attractor, showing a complex, chaotic trajectory forming a butterfly shape with multiple colored loops (red, green, blue, yellow).

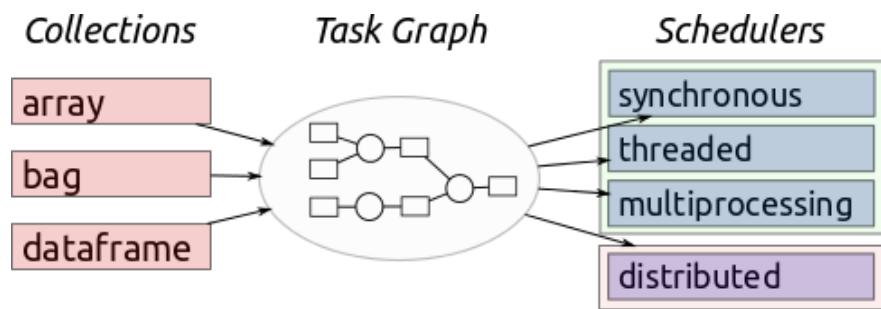
JupyterLab



dask & distributed

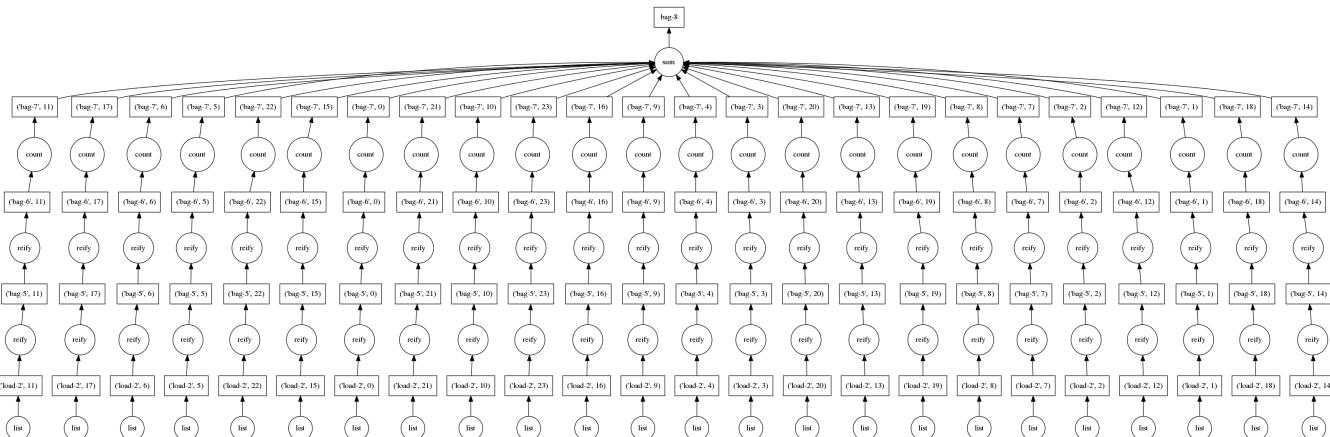
- Collection API similar to NumPy array and Pandas DataFrame objects
- Custom workloads via tasks scheduling API
- Pure python, easily integrated in PyData environment
- Low overhead (scheduling latency: ~1ms per tasks)
- Scales up: runs on clusters with 1000s of cores
- Scales down: runs on a laptop in a single process

dask architecture



dask bags & compute graphs

```
>>> import dask.bag as db
>>> b = db.from_s3('githubarchive-data', '2015-01-01-*json.gz')
    .map(json.loads)
    .map(lambda d: d['type'] == 'PushEvent')
    .count()
```



dask arrays & compute graphs

```
>>> import dask.array as da  
>>> x = da.ones((5000, 1000), chunks=(1000, 1000))  
>>> u, s, v = da.linalg.svd(x)
```

dask collection-oriented API

Dask DataFrame mimics Pandas

```
import pandas as pd
df = pd.read_csv('2015-01-01.csv')
df.groupby(df.user_id).value.mean()

import dask.dataframe as dd
df = dd.read_csv('2015-*.*.csv')
df.groupby(df.user_id).value.mean().compute()
```

Dask Array mimics NumPy

```
import numpy as np
f = h5py.File('myfile.hdf5')
x = np.array(f['/small-data'])

x = x.mean(axis=1)

import dask.array as da
f = h5py.File('myfile.hdf5')
x = da.from_array(f['/big-data'],
                  chunks=(1000, 1000))
x = x.mean(axis=1).compute()
```

Dask Bag mimics iterators, Toolz, and PySpark

```
import dask.bag as db
b = db.read_text('2015-*.*.json.gz').map(json.loads)
b.pluck('name').frequencies().topk(10, lambda pair: pair[1]).compute()
```

dask task-oriented API

Dask Delayed mimics for loops and wraps custom code

```
from dask import delayed
L = []
for fn in filenames:                      # Use for loops to build up computation
    data = delayed(load)(fn)              # Delay execution of function
    L.append(delayed(process)(data))     # Build connections between variables

result = delayed(summarize)(L)
result.compute()
```

The `concurrent.futures` interface provides general submission of custom tasks:

```
from dask.distributed import Client
client = Client('scheduler:port')

futures = []
for fn in filenames:
    future = client.submit(load, fn)
    futures.append(future)

summary = client.submit(summarize, futures)
summary.result()
```

Container-based orchestration with docker & kubernetes

Docker is:

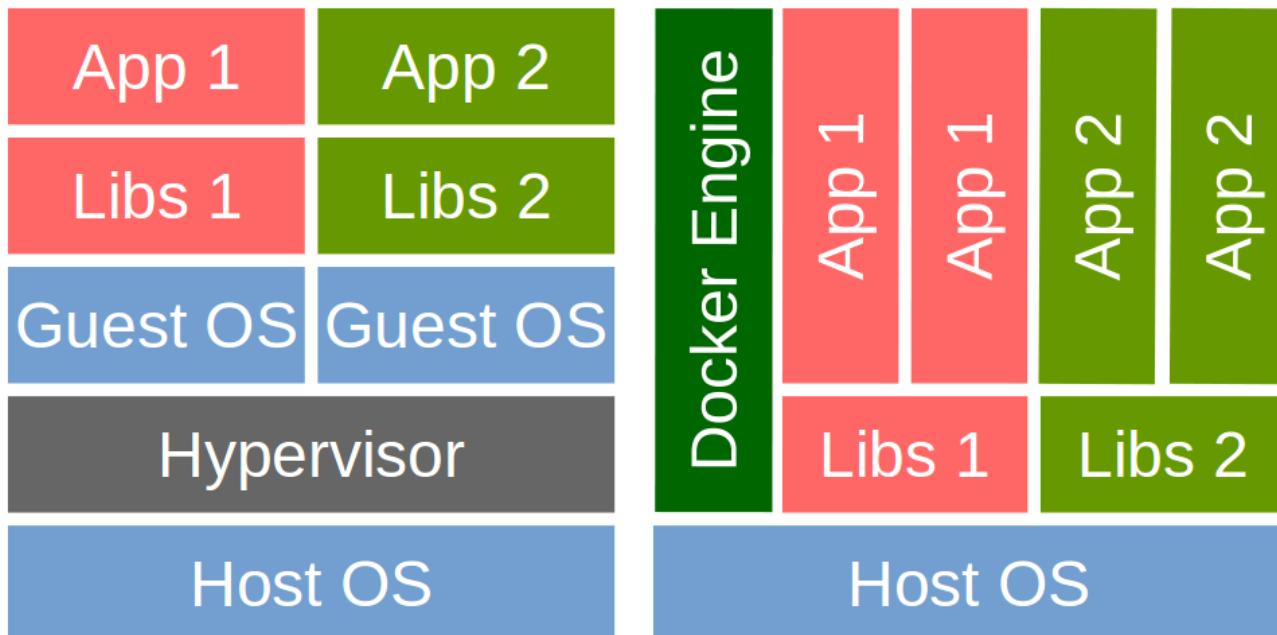
Not a virtual machine system

Linux container + Layered file-system +
Abstract network

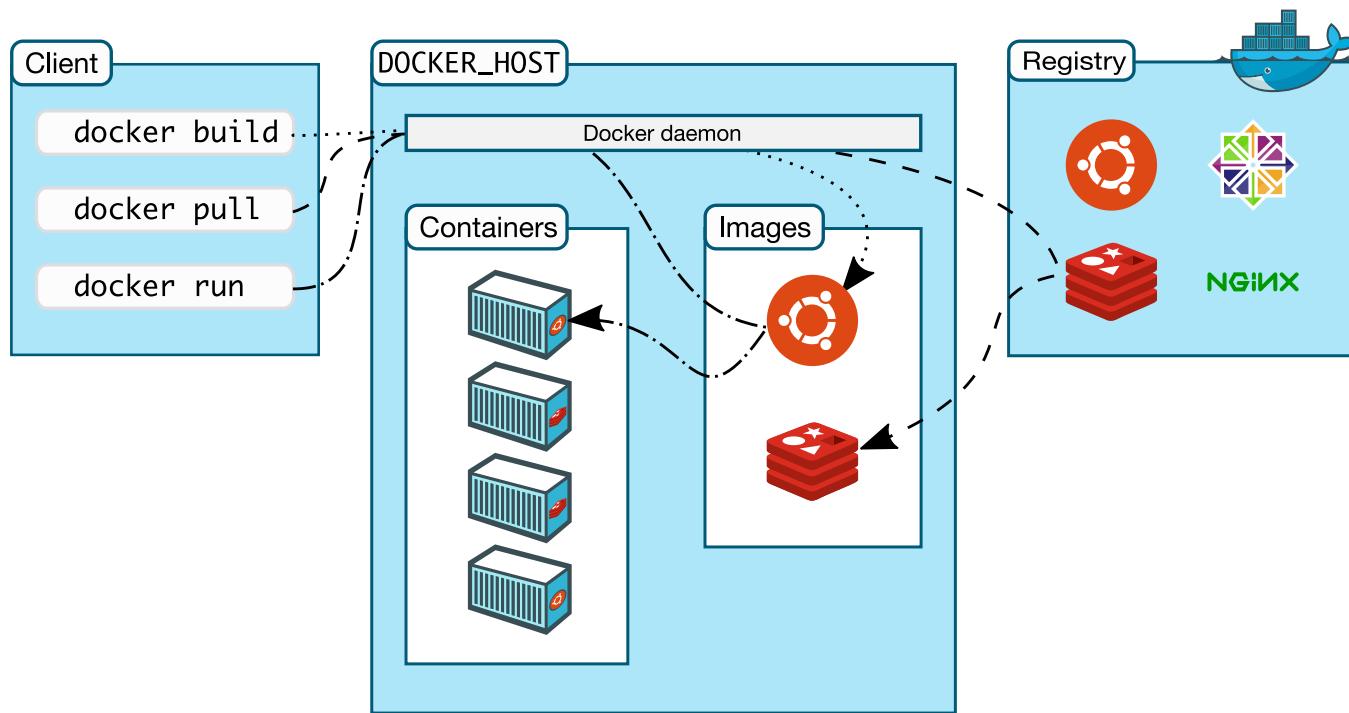
User-friendly command line interface

It can run in a VM (or not)

VM vs Container



Docker architecture





Building an image

```
$ git clone https://github.com/ogrisel/docker-distributed
$ ls docker-distributed
docker-compose.yml
Dockerfile
examples
kubernetes
README.md
requirements.txt
```

```
$ git clone https://github.com/ogrisel/docker-distributed
$ ls docker-distributed
docker-compose.yml
Dockerfile
examples
kubernetes
README.md
requirements.txt
```

```
$ git clone https://github.com/ogrisel/docker-distributed
$ ls docker-distributed
docker-compose.yml
Dockerfile
examples
kubernetes
README.md
requirements.txt
```

```
FROM debian:jessie
MAINTAINER Olivier Grisel <olivier.grisel@ensta.org>

RUN apt-get update -yqq && apt-get install -yqq wget bzip2 git \
&& rm -rf /var/lib/apt/lists/*

# Configure environment
ENV LC_ALL=C.UTF-8 LANG=C.UTF-8

RUN mkdir /work
WORKDIR /work

# Install Python 3 from miniconda
RUN wget -O miniconda.sh \
https://repo.continuum.io/miniconda/Miniconda3-latest-Linux-x86_64.sh \
&& bash miniconda.sh -b -p /work/miniconda \
&& rm miniconda.sh

ENV PATH="/work/bin:/work/miniconda/bin:$PATH"

RUN conda install -y \
pip \
notebook \
pandas \
scikit-learn \
&& conda clean -tipsy

# Install the master branch of distributed and dask
COPY requirements.txt .
RUN pip install -r requirements.txt && rm -rf ~/.cache/pip/

# Add the example notebooks
COPY examples .
```

docker build

```
$ docker build -t ogrisel/distributed .
Sending build context to Docker daemon 178.2 kB
Step 1 : FROM debian:jessie
--> 1b01529cc499
Step 2 : MAINTAINER Olivier Grisel <olivier.grisel@ensta.org>
--> Using cache
--> 37887ee139f1
Step 3 : RUN apt-get update -yqq && apt-get install -yqq wget [...]
--> Using cache
--> 3c2b8caccb80
[...]
```

docker build & docker push

```
$ docker build -t ogrisel/distributed .
Sending build context to Docker daemon 178.2 kB
Step 1 : FROM debian:jessie
--> 1b01529cc499
Step 2 : MAINTAINER Olivier Grisel <olivier.grisel@ensta.org>
--> Using cache
--> 37887ee139f1
Step 3 : RUN apt-get update -yqq && apt-get install -yqq wget [...]
--> Using cache
--> 3c2b8caccb80
[...]
$ docker run --rm ogrisel/distributed \
    python -c "import sklearn; print(sklearn.__version__)"
0.18.1
$ docker push ogrisel/distributed
```



Orchestration

Orchestration tools



Docker Swarm and
Docker Compose



Kubernetes



DC/OS and Mesos

Kubernetes and GKE

```
$ gcloud config set compute/zone europe-west1-d
$ gcloud container clusters create cluster-1 \
    --num-nodes 3 \
    --machine-type n1-highcpu-32 \
    --scopes bigquery,storage-rw \
    --preemptible \
    --no-async
```

Creating cluster cluster-1...done.

Created [https://container.googleapis.com/v1/.../cluster-1].

kubeconfig entry generated for cluster-1.

NAME	ZONE	MACHINE_TYPE	NUM_NODES	STATUS
cluster-1	europe-west1-d	n1-highcpu-32	3	RUNNING

Kubernetes and GKE

```
$ gcloud config set compute/zone europe-west1-d
$ gcloud container clusters create cluster-1 \
    --num-nodes 3 \
    --machine-type n1-highcpu-32 \
    --scopes bigquery,storage-rw \
    --preemptible \
    --no-async
```

Creating cluster cluster-1...done.

Created [https://container.googleapis.com/v1/.../cluster-1].

kubeconfig entry generated for cluster-1.

NAME	ZONE	MACHINE_TYPE	NUM_NODES	STATUS
cluster-1	europe-west1-d	n1-highcpu-32	3	RUNNING

```
$ gcloud container clusters get-credentials cluster-1
```

Fetching cluster endpoint and auth data.

kubeconfig entry generated for cluster-1.

Kubernetes and GKE

```
$ git clone https://github.com/ogrisel/docker-distributed
$ cd docker-distributed
$ kubectl create -f kubernetes/
service "dscheduler" created
service "dscheduler-status" created
replicationcontroller "dscheduler" created
replicationcontroller "dworker" created
service "jupyter-notebook" created
replicationcontroller "jupyter-notebook" created
```

Kubernetes and GKE

```
$ git clone https://github.com/ogrisel/docker-distributed
$ cd docker-distributed
$ kubectl create -f kubernetes/
service "dscheduler" created
service "dscheduler-status" created
replicationcontroller "dscheduler" created
replicationcontroller "dworker" created
service "jupyter-notebook" created
replicationcontroller "jupyter-notebook" created

$ kubectl get services
NAME           CLUSTER-IP      EXTERNAL-IP    PORT(S)
dscheduler     10.115.249.189  <none>        8786/TCP,9786/TCP
dscheduler-status 10.115.244.201  130.211.50.206  8787/TCP
jupyter-notebook 10.115.254.255  146.148.114.90  80/TCP
kubernetes     10.115.240.1    <none>        443/TCP
```

distributed-worker.yml

```
apiVersion: v1
kind: ReplicationController
metadata:
  name: dworker-controller
spec:
  replicas: 3
  selector:
    name: dworker
  template:
    metadata:
      labels:
        name: dworker
    spec:
      containers:
        - name: dworker
          image: ogrisel/distributed:latest
          args: ["dask-worker", "dscheduler:8786"]
```

Conclusion

Jupyter + Dask a nice interactive distributed environment for PyData

Containers make cluster setup reproducible and cloud agnostic

Docker Swarm vs Kubernetes vs Mesos: learn any one and understand others

Thank you



Rackspace for free cloud resources to support SciPy projects and for helping the Python community in general.



Google for GCP credits to test distributed Python with Kubernetes.



Inria for supporting my work on scikit-learn and related projects.

Thank you!

Sample configuration:

- github.com/oqrisel/docker-distributed

Those slides:

- oqrisel.github.io/decks/2016_succes_paris_pydata_kubernetes

Image credits

- "Supermicro Storage" by Robert: <https://flic.kr/p/e17qr3>
- "smiiile!" by mooglet: <https://flic.kr/p/am9zQe>
- "Docker Architecture" <https://docs.docker.com/engine/understanding-docker/>
- "Container terminal" by Martin Abegglen: <https://flic.kr/p/b7rw5D>
- "Symphonic Orchestra and Choir of the WYD" by HazteOir.org:
<https://flic.kr/p/c5B6Cd>