


Dask and Jupyter at NERSC



Rollin Thomas

**NERSC Data & Analytics Services
NUG Monthly Teleconference
2019-12-12**

What Problem Does Dask Try to Solve?

The NERSC logo is located in the top right corner. It consists of the letters "NERSC" in a bold, white, sans-serif font, set against a dark blue background with a bright light flare effect emanating from behind the text.

Python is now a dominant language for data analytics and general programming, major platform for ML/DL

Growth fueled by many factors, including:

- Computational libraries: NumPy, SciPy, Pandas, Scikit-Learn...
- Libraries for visualization: Matplotlib, Seaborn, Bokeh...
- Tools for interactivity & sharing: Jupyter notebooks & ecosystem...

These things were not designed to scale beyond a single machine.

⇒ Dask: Scalable Analytics in Python

- Developed to scale up these Python tools
- Multi-core machines, distributed clusters, HPC



Big Picture of How Dask Works



With Dask, you start a **cluster** of some processes, kind of like Spark:

- 1 scheduler
- N workers
- 1 dashboard (optional, highly recommended)


(Many ways to start clusters.
We'll focus on *one* leveraging
Cori compute nodes today.)

Connect to the scheduler with a **client**:

- From a script
- From a notebook

Submit work to the cluster through client:

- **Tasks** are scheduled dynamically
- Lazy evaluation, futures/promises
- Has handy “big data” collections
- TCP, IB, UCX (new)...



Actually Dask has handy primitives too, but we won't talk about those today.

Adapted from <https://docs.dask.org/en/latest/>

Why Talk about Dask at NERSC Now?

The NERSC logo is located in the top right corner. It consists of the letters "NERSC" in a bold, white, sans-serif font, set against a dark blue background with a bright light flare effect emanating from behind the text.

More and more users have been asking about being able to use it. In particular they want to use it from Jupyter.

Dask ecosystem has matured to support HPC better.

What are the best practices to use Dask on Cori compute?

- [dask-mpi](#) → *uses MPI to launch cluster*
- [dask-jobqueue](#) → *starts workers in batch jobs*
- [Containers](#) or at least [/global/common/software](#) to scale launch
- <https://docs.dask.org/en/latest/setup/hpc.html>

We had to do some work:

- Network between Jupyter and compute nodes
- Contribute to jupyter-server-proxy (for the dashboard)
- Seek advice and help from the developers

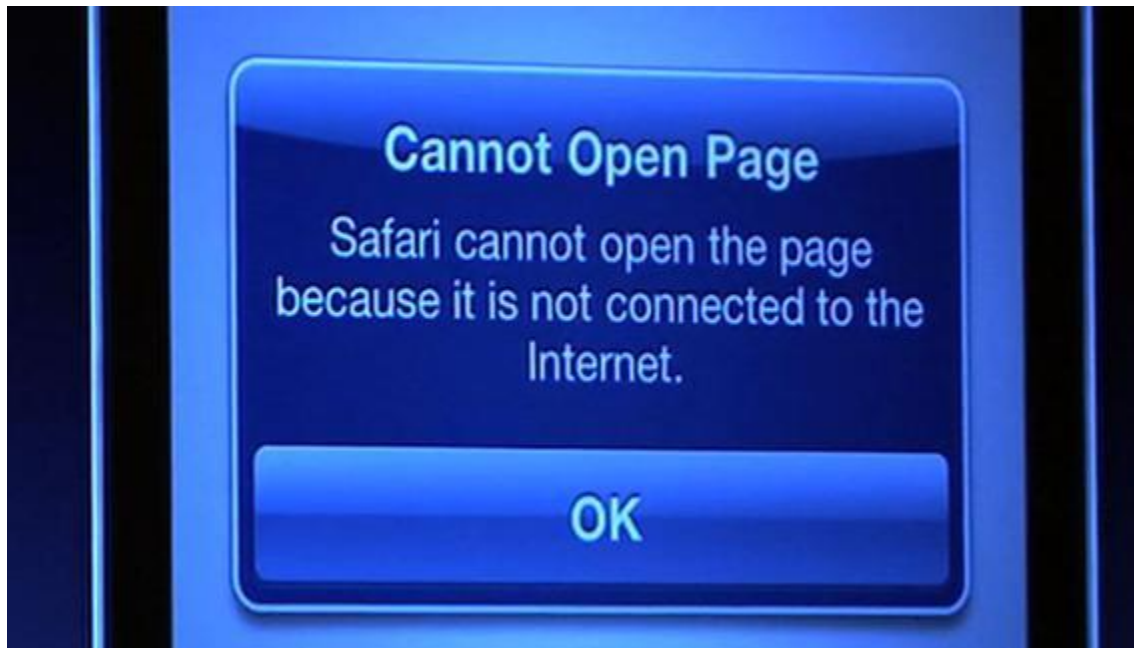
Demo consists of 3 notebooks:

- One way to start a Dask cluster on Cori compute in interactive QOS
 - How to connect to the cluster
 - How to start up and see the dashboard
- A simple parallel map-reduce example calculation
- A more complicated Dask dataframes based calculation on real data

These notebooks will be posted after the NUG videoconference today:

- They contain a lot of useful documentation and explanation
- You can adjust them and try to run them yourselves
- If you get stuck file a ticket and we'll help

These notebooks will be posted and linked from docs.nersc.gov



Why Dask and not say, mpi4py?



Dask's API tries to look familiar to Pandas, Scikit-Learn, NumPy users
... mpi4py tries to bring MPI to Python (that's great of course)

Dask scales out well (add workers, take away workers, it's still OK!)
... mpi4py processes are fixed, and if one goes down, they all go down

Dask scales down, you can test on a node or your laptop (and it's just pip install...)
... mpi4py can do that too but you need MPI built first

Dask delivers responsive feedback, it's meant for interactive exploration
... mpi4py is about parallelizing processes in jobs really

BUT!

While mpi4py can effectively use Aries via Cray MPICH and scale to all of Cori
... Dask can't do that, at NERSC it uses TCP

It has a single central scheduler with 1ms overhead per task
... this limits scalability of clusters, depending on work per task

What about Dask vs Spark



Dask is a lot more like Spark than a general library like MPI & mpi4py.

There are reasons why Dask or Spark would be just wrong instead of MPI.

There's a whole other page on why you might use Dask or Spark too:

<https://docs.dask.org/en/latest/spark.html>

The End



Today we have seen a way to use Dask on Jupyter at NERSC.

We will continue extending and enhancing the service:

- Better integration with the GUI to start clusters
- More documentation about configuration
- Profiles and contributions to Dask/Jupyter ecosystem
- It should get easier over time to use Dask at NERSC

Please contact us if you have questions:

- Please file a ticket (under Python or Jupyter) at <https://help.nersc.gov/>
- Read and contribute to our Dask documentation! <https://docs.nersc.gov/> (to contribute: <https://gitlab.com/NERSC/nersc.gitlab.io>)



Thank You



U.S. DEPARTMENT OF
ENERGY

Office of
Science

