Collaborative WRF-based research and education enabled by software containers

J. Hacker, J. Exby, K. Fossell

National Center for Atmospheric Research

Contributions from Tim See (U. North Dakota)

Why the Weather Research and Forecasting (WRF) model in Docker?

- WRF is a state-of-the science numerical weather prediction (NWP) model for operations and research
- Compilation and execution can be an intensive effort, slowing time to results
  - Huge complex code
  - Numerous and non-trivial dependencies
  - Inexperienced users can take months to get WRF running for results
- Classroom opportunities for hands-on numerical weather prediction can be intensive to produce
- Research is almost never reproducible
- Collaboration is difficult and cumbersome

Goals of WRF-Docker

NCAR

- Lower the technical difficulty for new users
  - Graduate students can accomplish early results simultaneous to learning the Unix/Linux skills needed for more in-depth work
  - Provide a reference context
- Trivialize classroom and lab experimentation
- Provide a platform for reproducible numerical weather prediction research
- Facilitate efficient and easy collaboration

WRF in a container is not a black box.

Not a turn-key approach

NCAR

- Important use cases in an education
  - Change input data sets for land use
  - Change input data sets for initial and boundary conditions
  - Change physics, diffusion, time steps, etc
  - Change code and recompile in known environment

Our container development allows all of these, including deployments on cloud providers or local compute hardware.

Vision: End to end

NCAR

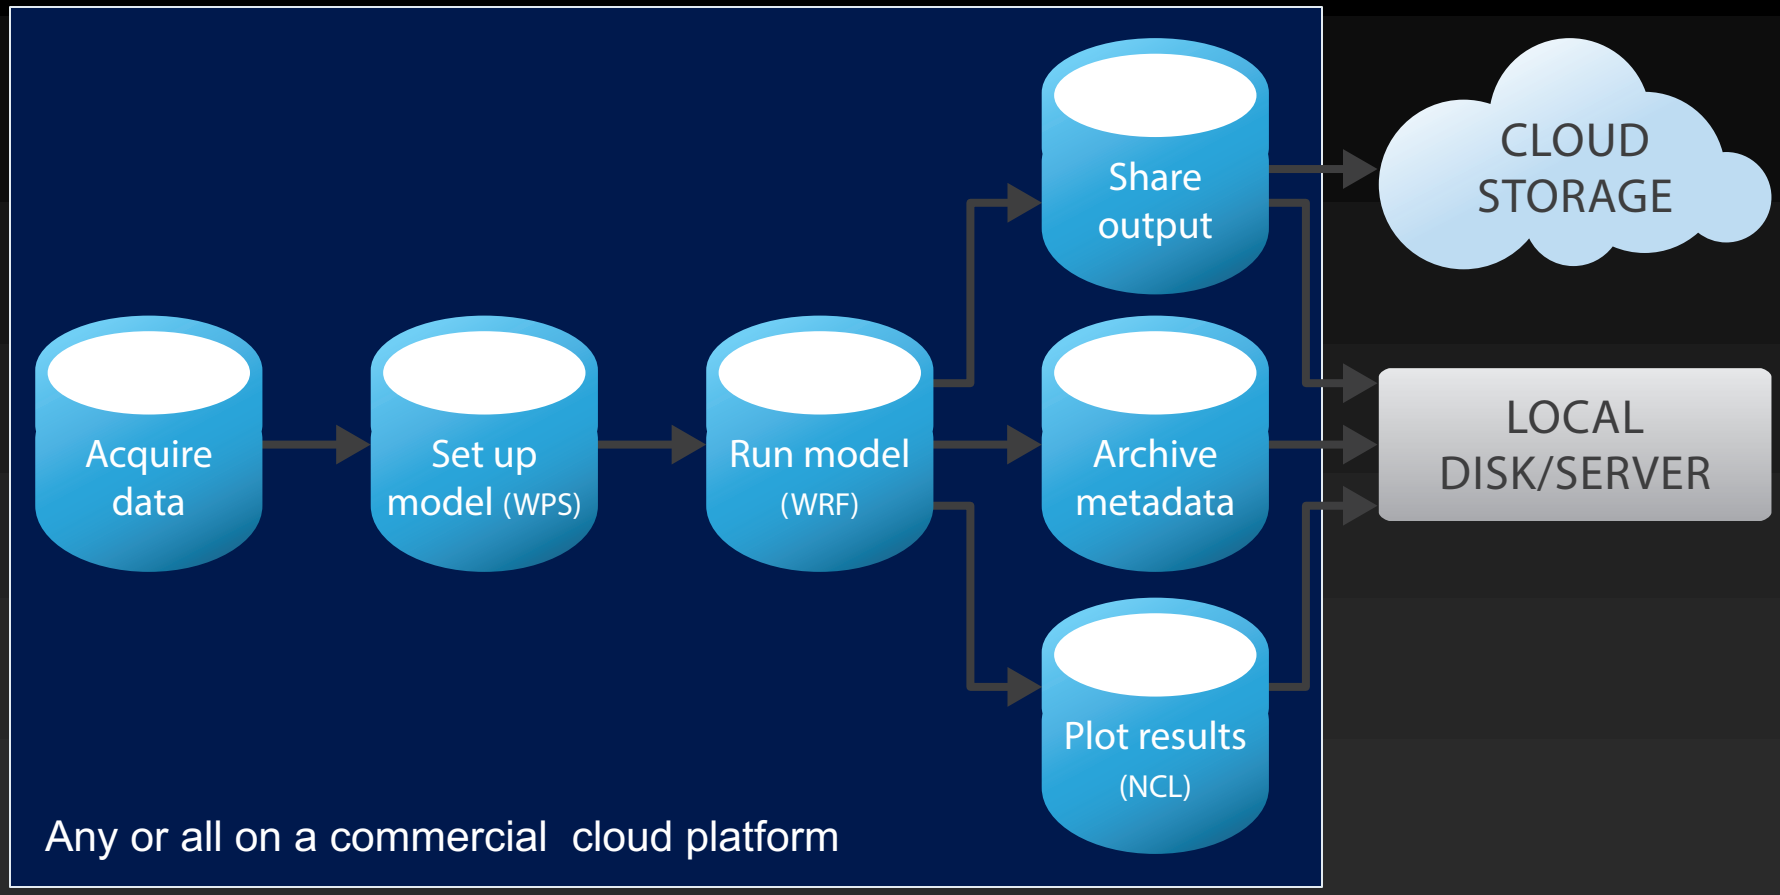

Run on your command line and link to your filesystem

Bit-wise reproducibility

NCAR

Incomplete list of test platforms

OS	Chip/CPU	Cores
OS X 10.10.5	Intel Xeon E5	6
OS X 10.9.5	Intel Core i7	2
Ubuntu 14.04	AMD Opteron 6320	16
Ubuntu 14.04	Intel Xeon E5	16
RHEL	Intel Xeon X5550	8
*Ubuntu 14.04	Intel Xeon E5-2666 v3	32
**Ubuntu 14.04	Intel Atom C2550	4

AWS EC2

Packet Tiny Atom

Reproducible on all platforms so far

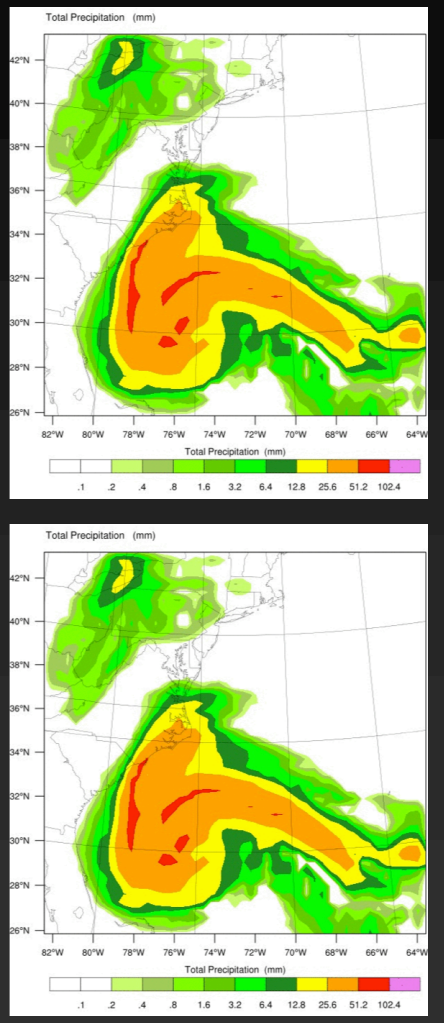

From Tim See, UND)

University of North Dakota Classroom Implementation

- Access Docker-WRF through Amazon Web Services.
- Students completed a classroom assignments to create an ensemble output of a tornadic supercell over North Dakota.
- Students personally changed the parameterization schemes within WRF.
- Classroom discussion generated through changing of parameterizations.
- Sample Plots below

CAPE

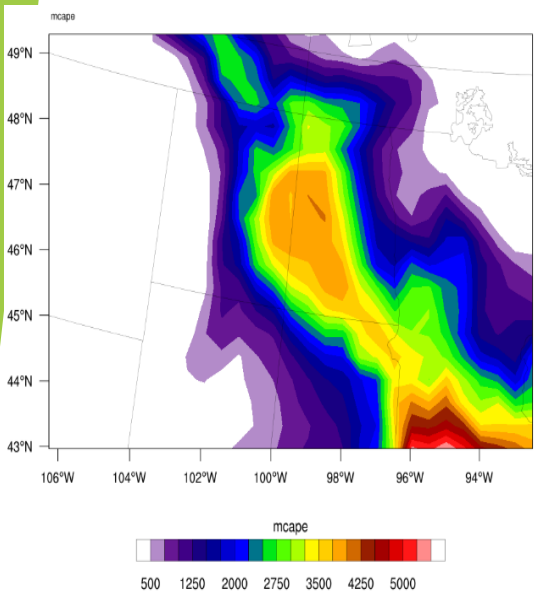

Accumulated Precipitation

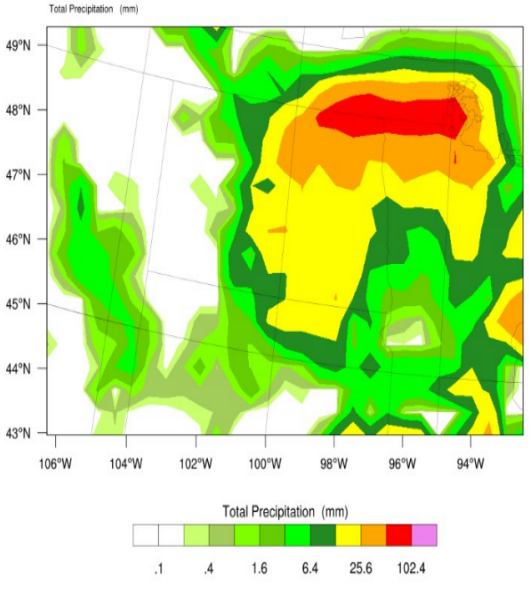

Surface Analysis

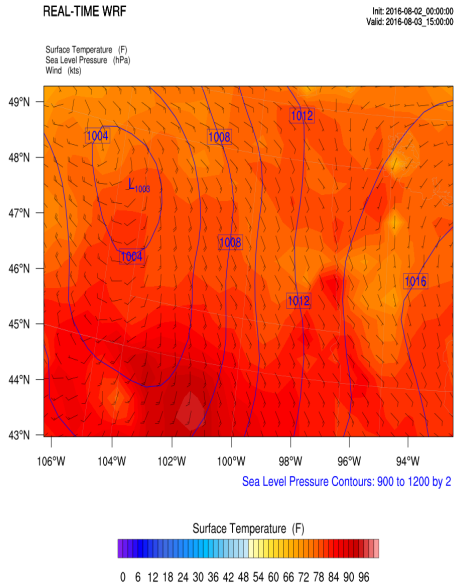

University of North Dakota Classroom Implementation

From Tim See, UND)

- Total Cost for Homework Assignment: $40.21 over 11 days
  - Inflated cost due to not shutting down instances properly first day.
- Reproducibility of Docker allows for plots to remain the same across all students.
- Students were able to complete their model runs from

Dynamic pull of terrestrial data sets (WPS_GEOG) – work in progress

Current practice

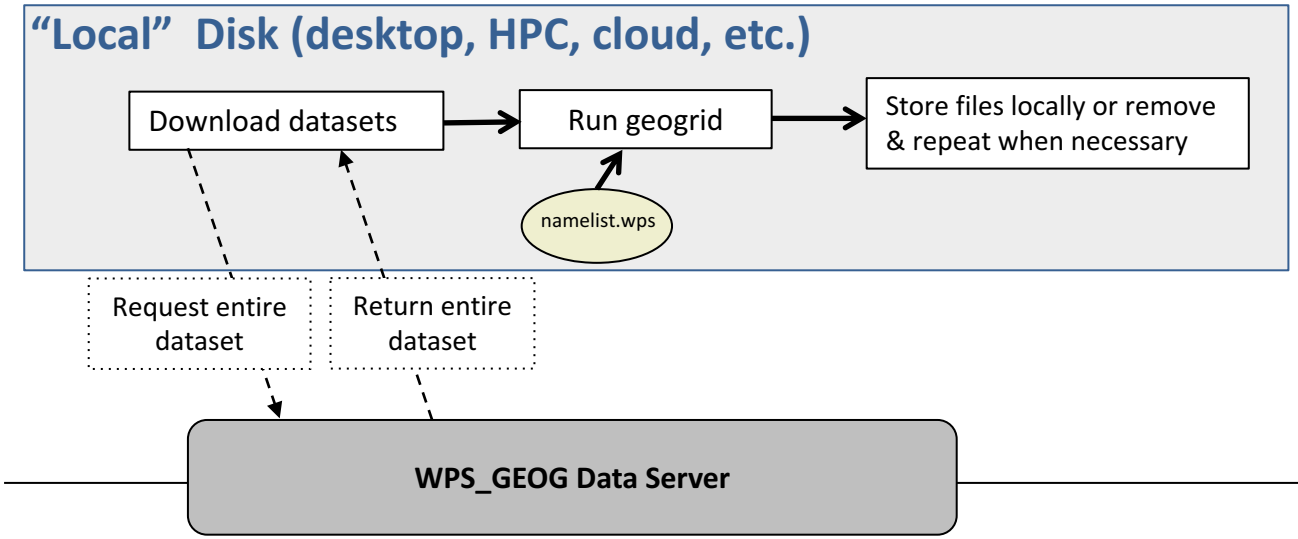

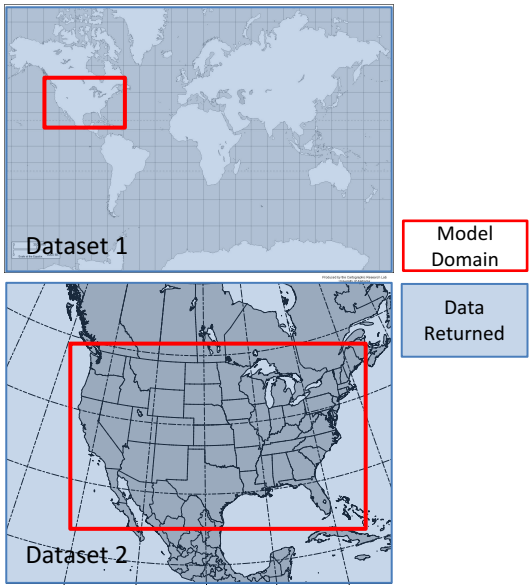

Downsides:

- Required to download and store entire datasets and unnecessary data
- Not conducive to cloud or container environments
- Large files to store and transfer
- Costs ($) associated with storing or downloading/transferring data in cloud
- Computational inefficiencies in containers due to size of files

Dynamic pull of terrestrial data sets (WPS_GEOG) – work in progress

Current practice

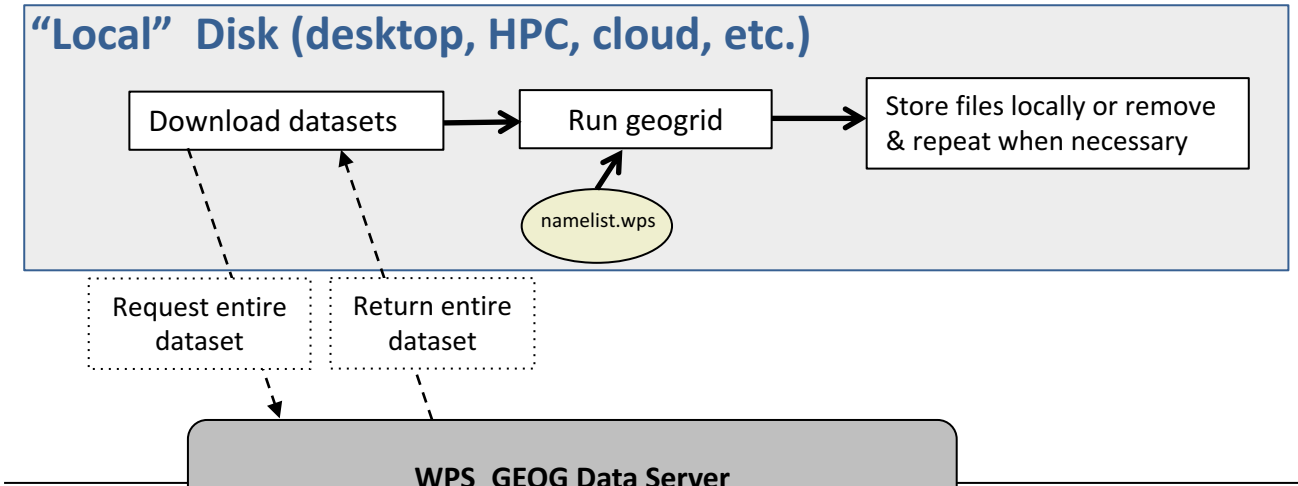

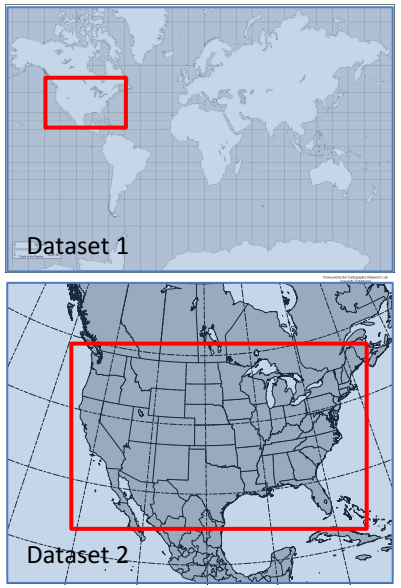

New utility

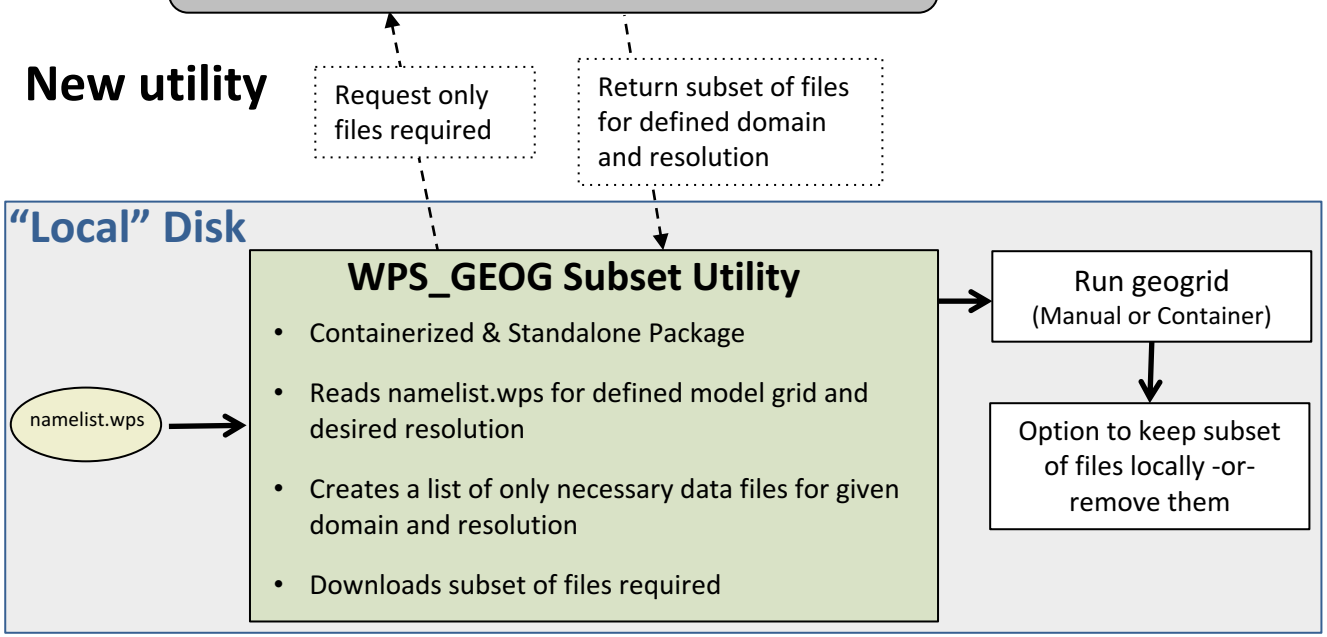

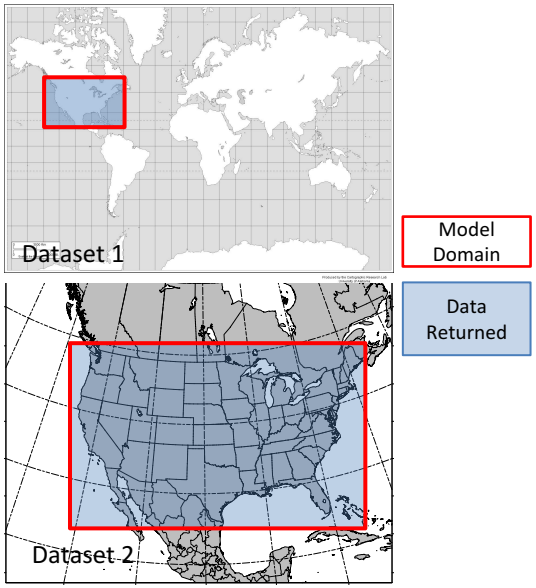

Zero to WRF!
in 3 Steps

NCAR

Using a relatively ”modern” laptop, workstation, server.
Linux, Mac, Windows:

- 1. Install docker for free ! https://docker.com/products/
- 2. git clone https://github.com/NCAR/container-wrf
  (Hurricane Sandy and Katrina)
- 3. cd 3.7.1/demos/local ; docker-compose up
  - (psssst. Windows users- please first edit docker-compose.yml for output DIR)

VOILA!

Demo

- So a University professor gives their student a dollar to “do some atmospheric science” in one hour, then write a paper to help toward graduation!

- Is this possible?

- 🏫👨‍🏫🤝$1💵👉👩‍🔬⏩AWS☘️◻️🆔💡🌐🌧️🌤️📈📉🎛️(10🔁⌛) 💾🖋️📄📄📄🏁🎓🍺💤

Demo Architecture

NCAR

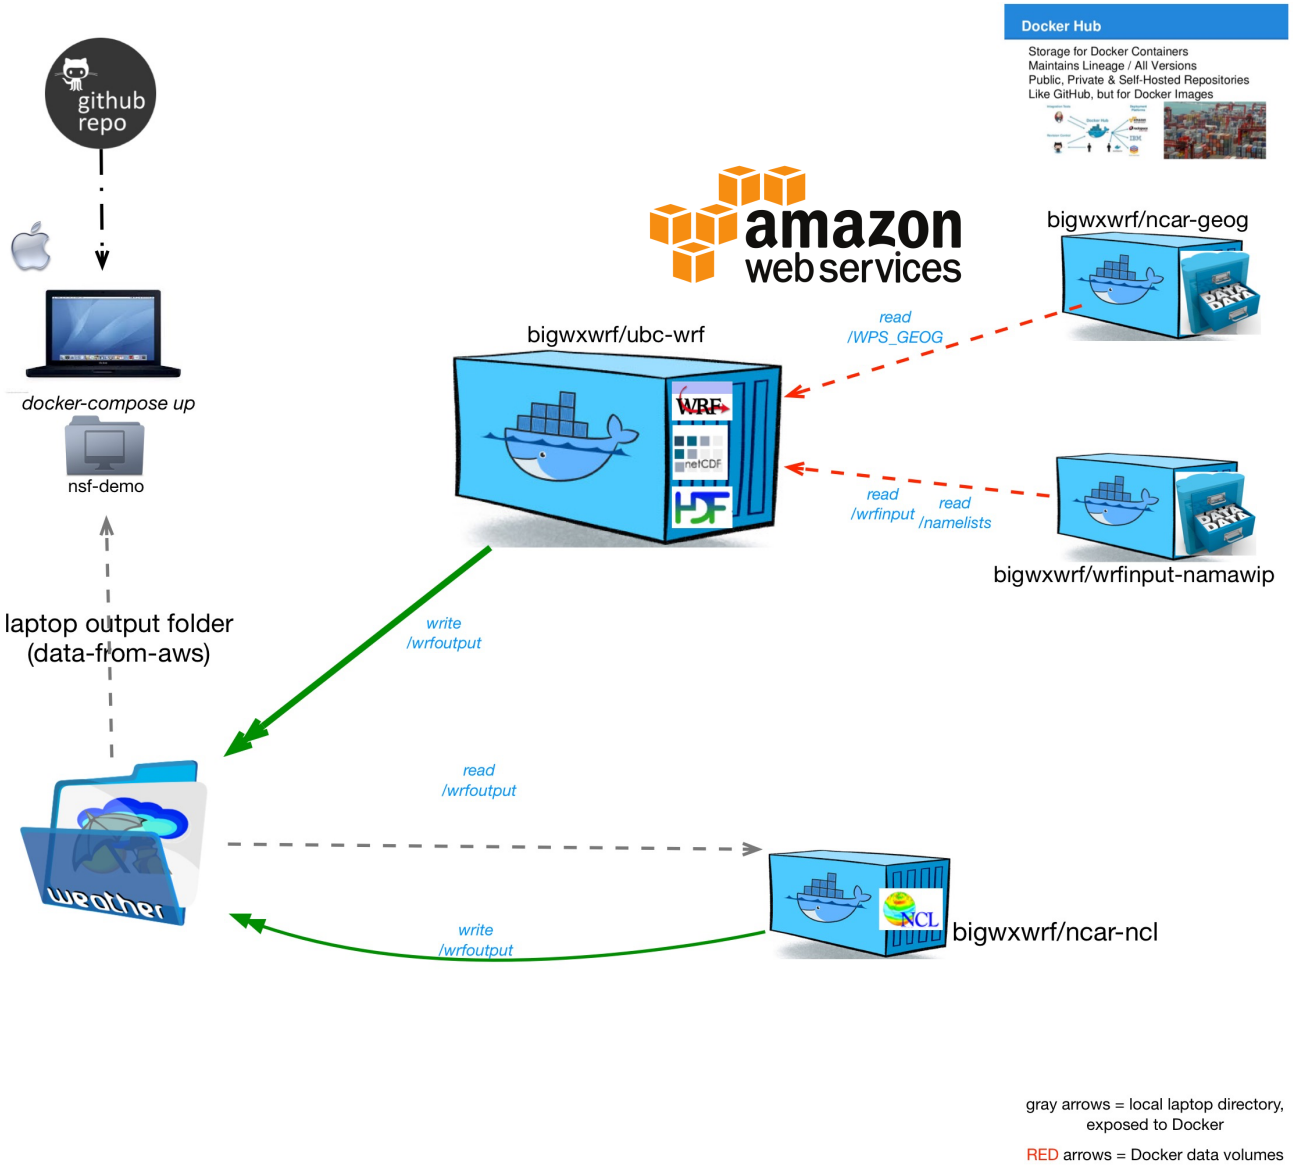

32 core WRF run on AWS

NCAR

- script to spin up a new AWS resource for our compute ac cess only

- script to launch the docker-compose.yml elements:
  - two containers with data
  - a container with wrf executables
  - a container with NCL scripts to post process
  - copy files from AWS back to macbook

- visualize results on macbook pro.

Live Demo Time

NCAR

- Fingers crossed wifi is fast!

Next Steps

NCAR

- Dynamic data query for smaller downloads
- AWS Batch
- Singularity and Swarms

NCAR Resources

- NCAR RAL Docker-WRF Project Web site:
- https://www.ral.ucar.edu/projects/ncar-docker-wrf

- NCAR Github repo: https://github.com/NCAR/container-wrf

- NCAR Dockerhub repo: https://hub.docker.com/r/bigwxwrf/
  - bigwxwrf/ncar-wrf
  - bigwxwrf/ncar-wpsgeog
  - bigwxwrf/ncar-wrfinputkatrina
  - bigwxwrf/ncar-ncl
  - bigwxwrf/ncar-wrfinputsandy
  - bigwxwrf/ubc-wrf

- Slack channel for docker-wrf community discussion.
- email: exby@ucar.edu for invitations
- https://ncar-dockerwrf.slack.com