

Docker WRF at NCAR / RAL

(publication in process)

John Exby (exby@ucar.edu)
Josh Hacker (hacker@ucar.edu)
Dave Gill (gill@ucar.edu)
Kate Fossell (fossell@ucar.edu)

July 25, 2016




NCAR



What are Docker Containers?

PACKAGE YOUR APPLICATION INTO A STANDARDIZED UNIT FOR SOFTWARE DEVELOPMENT

Docker containers wrap a piece of software in a complete filesystem that contains everything needed to run: code, runtime, system tools, system libraries – anything that can be installed on a server. This guarantees that the software will always run the same, regardless of its environment.



LIGHTWEIGHT


Containers running on a single machine share the same operating system kernel; they start instantly and use less RAM. Images are constructed from layered filesystems and share common files, making disk usage and image downloads much more efficient.

OPEN

Docker containers are based on open standards, enabling containers to run on all major Linux distributions and on Microsoft Windows -- and on top of any infrastructure.

SECURE BY DEFAULT


Containers isolate applications from one another and the underlying infrastructure, while providing an added layer of protection for the application.




Docker vs Virtual Machines

Containers:
smaller, faster, more secure, shares layers

Containers vs. VMs

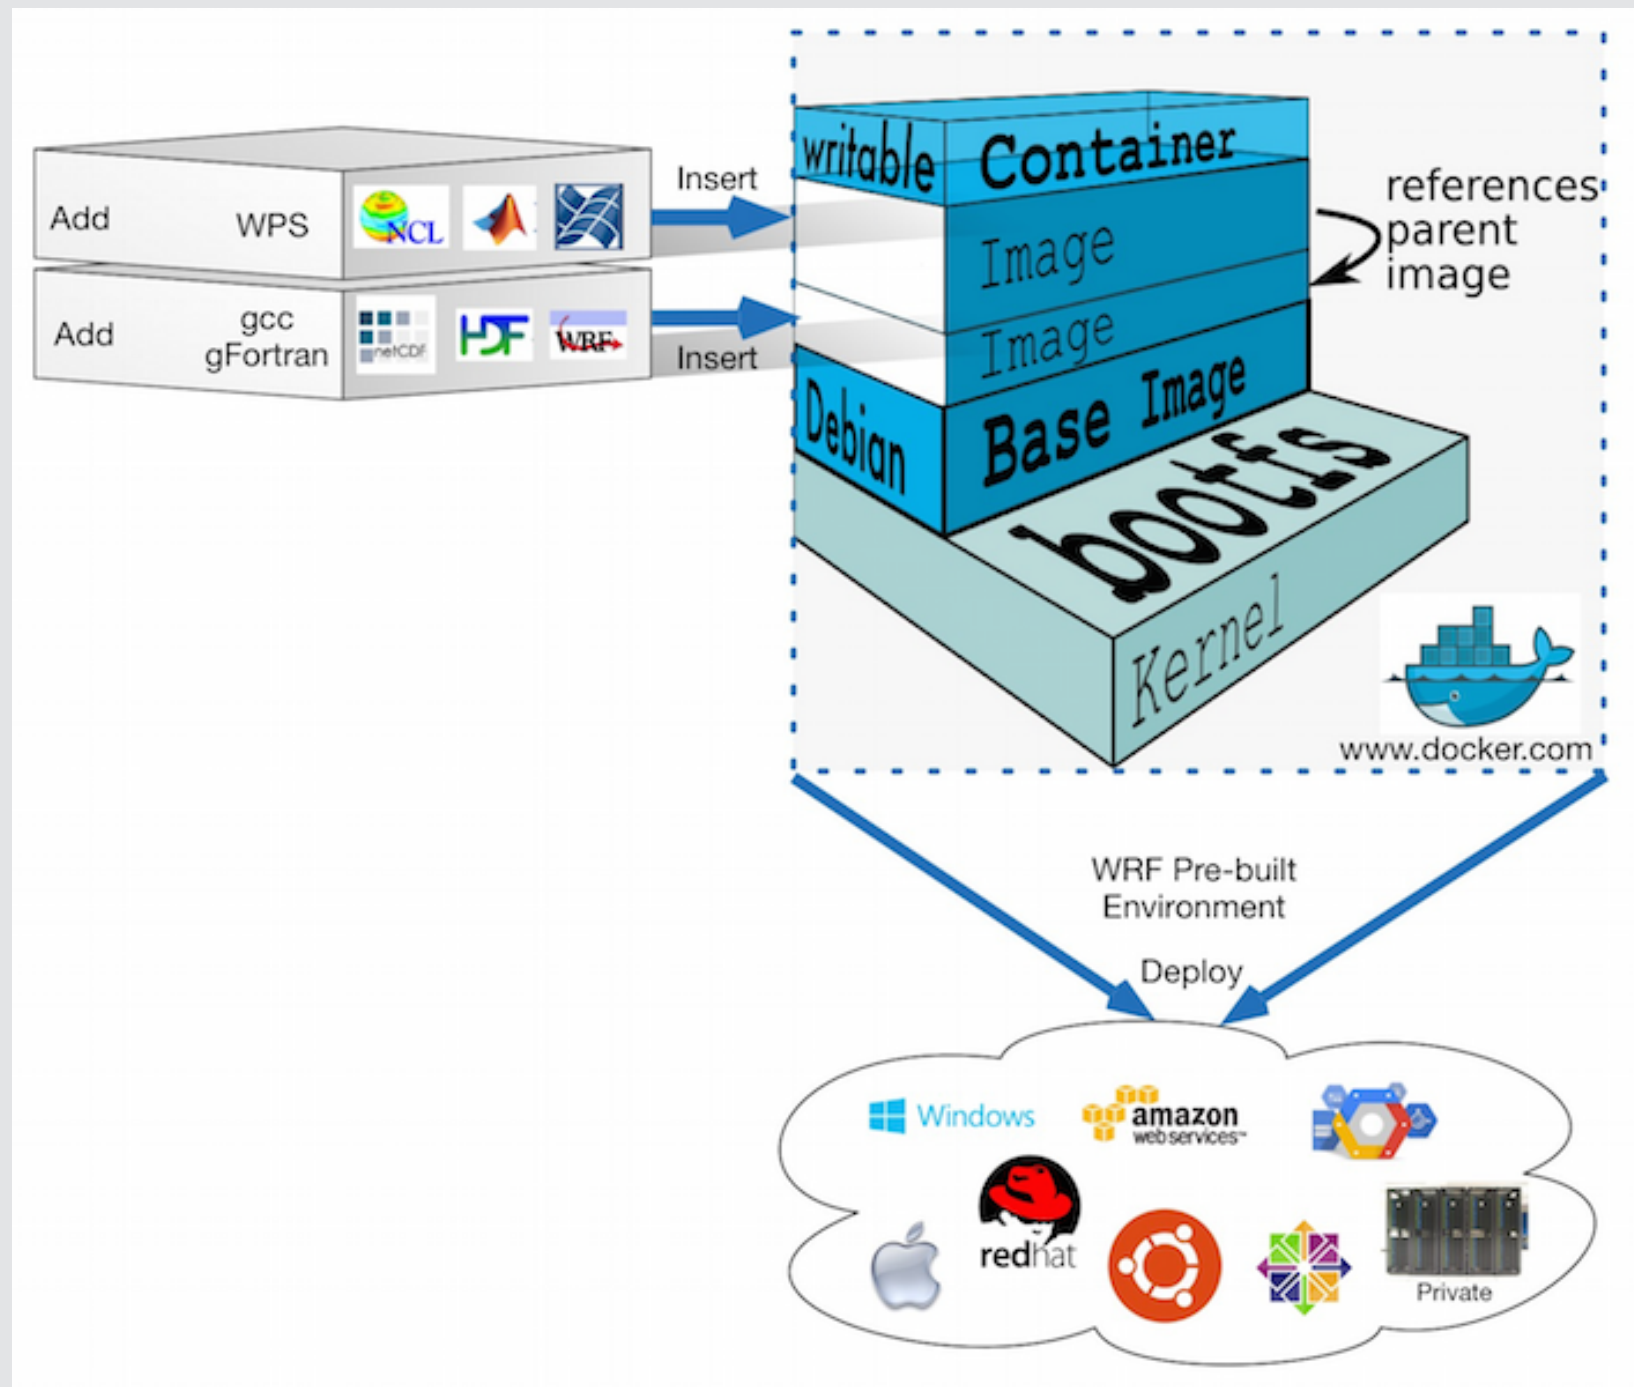


Containers are isolated, but share OS and, where appropriate, bins/libraries



WRF in Docker

Build (Dockerfile), Deploy (hub), Run (and compose)



NCAR Docker WRF Milestones

- WRF precompiled binaries run on any Docker engine/OS
- Allows scientific reproducibility and application portability from laptops to servers to clouds (Amazon, Google, Azure).
- Developed stable platform for research, case studies, tutorials and classroom curricula. Session 4.3 June 28.
- Constructed two example cases: 6 and 12 hour forecasts of Hurricane Sandy or Katrina (at 40km grid) initialized from NOAA global weather model.

WRF on laptop:

WRF in container using local data

- Install Docker for Mac/Windows/Linux
- <https://www.docker.com/products/docker>
- \$ docker run hello-world
- \$ docker ps -a
- \$ docker images
- view Dockerfile
- \$ docker build -t my-wrf . compiles WRF from source see Dockerfile
- \$ docker run -v mydir:containerdir my-wrf runs WRF binaries using data in local directories

See README-dockerbuild.txt for instructions

Complete WRF+WPS compiled in approximately 20 minutes.
WRF output and graphics computation (macbook pro, 2cpu):
5 minutes 30sec

WRF on single AWS using docker-compose

- Launch Amazon EC2 instance via docker-machine on laptop:

```
docker-machine -D create --driver amazec2 \  
--amazec2-access-key $AWS_ACCESS_KEY_ID \  
--amazec2-secret-key $AWS_SECRET_ACCESS_KEY \  
--amazec2-vpc-id $AWS_VPC_ID \  
--amazec2-region us-west-2 \  
--amazec2-instance-type c4.8xlarge \  
--amazec2-root-size 40 \  
--amazec2-zone b wrf-Large32
```

- `$ view docker-compose.yml` defines container images, Sandy data set, cores
- `$ docker-compose up` downloads container images from Hub to EC2 and runs WRF

Time for new EC2 to instantiate: 3 minutes

Full WRF output results completed (c4.8xlarge, 32cpu): 2min 5sec

EC2 cost for this demo: (\$1.65 per hour/60 x 305sec = 8.4 cents)

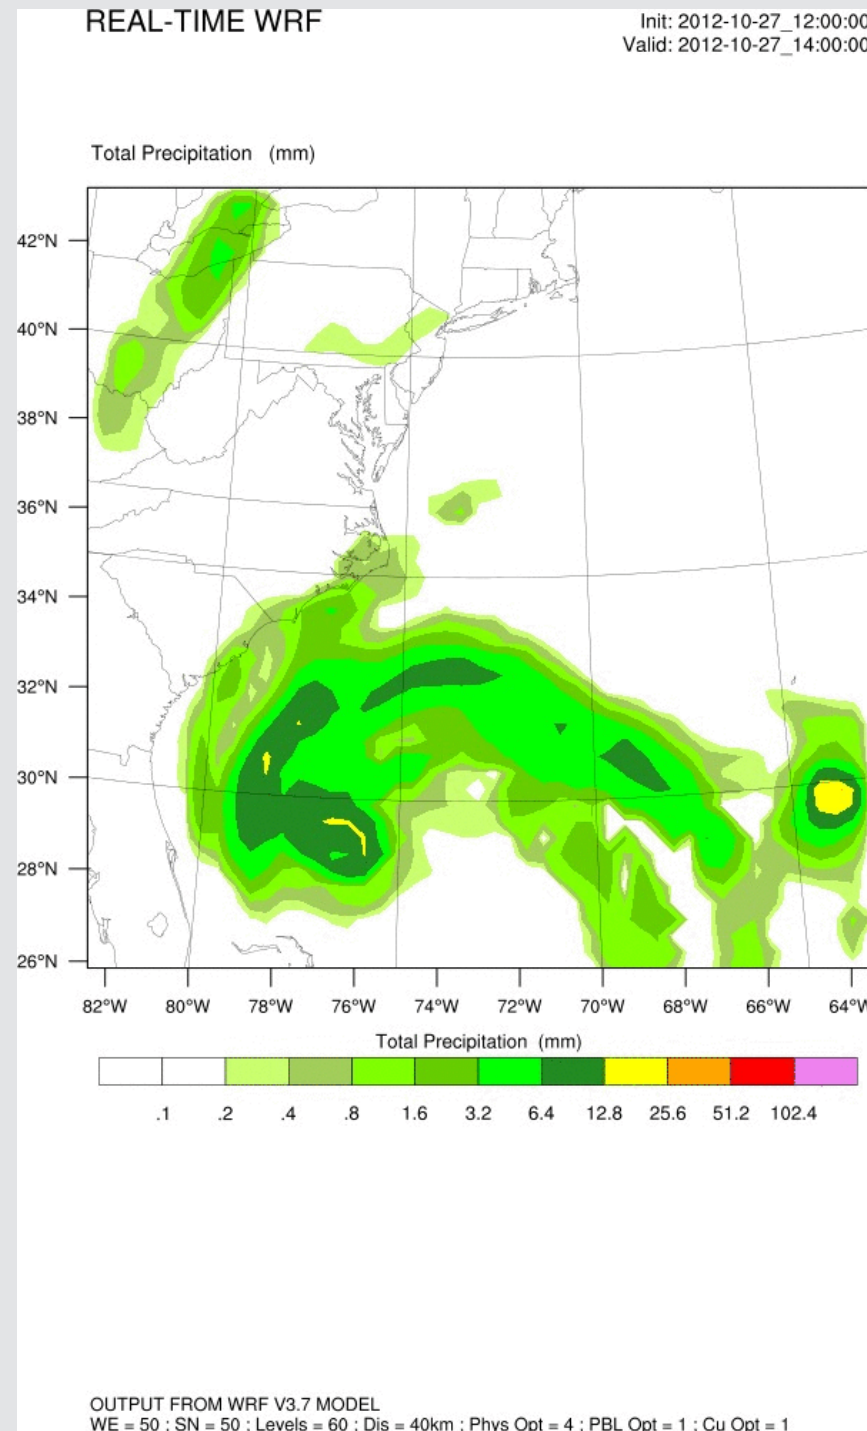
NCAR Docker WRF Demos

(we are brave let's try it!)

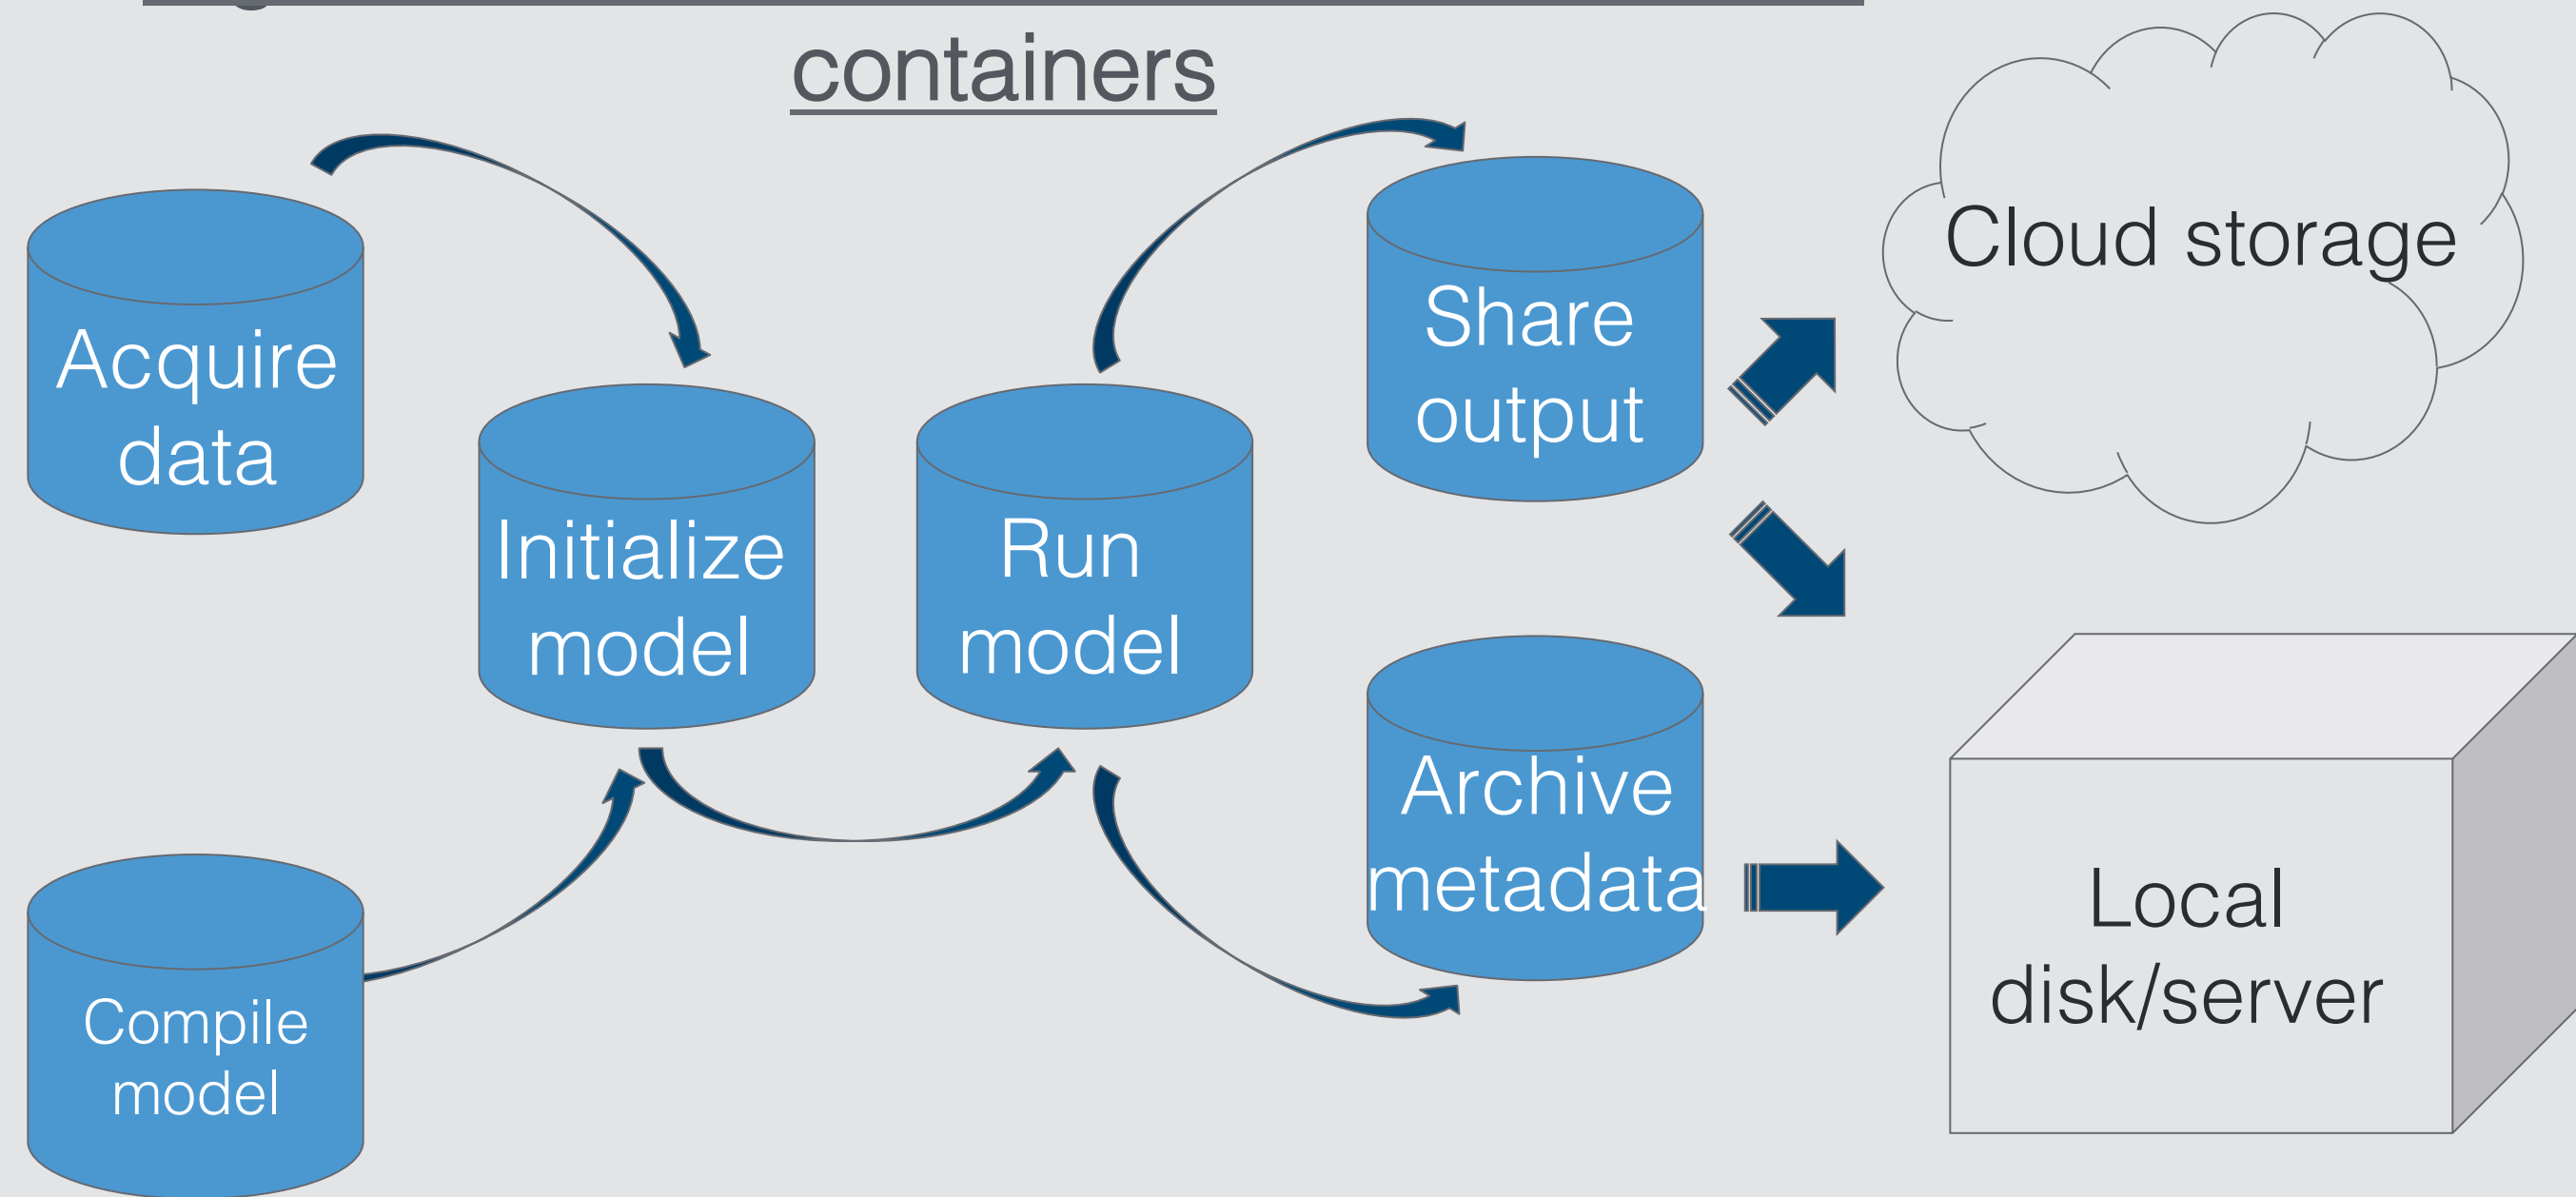
- Demo 1 – Using a laptop to run Docker WRF binaries with data files in local directories
- Demo 2 – Running Docker WRF in the cloud with data files in a container + automated workflow

Hurricane Results in 5 minutes on Amazon

150,000
grid points



BigWeather Web Vision: end to end linked containers




Run on your command line and link to your filesystem

Research Reproducibility

- Reproducibility is typically nearly impossible in NWP research
 - Simulations on different machines
 - Simulations with different compiler options
 - Simulations with different numbers of processors
 - Human error in setting up complex workflows and settings
- Now we can share containers: collaborators + classrooms
 - Identical model set up
 - Identical compilation
 - Identical results except when a modification is explicitly introduced
- Container repositories/revisions can be cited in a paper
- DOI a request feature in future, Docker tag values

Sandy container simulation different machines



12-h precipitation accumulation

Education Impact

- **Classroom**
 - Zero to WRF in 3 commands. Students can run an ensemble on their laptop during an NWP class
 - Good for both traditional and distance learning
 - Rollback to known working, tagged versions
 - Accelerate time to model output while learning to compile the model or compute environments
- **Research**
 - Minimizes mistakes when executing many simulations with small differences (i.e. sensitivity studies)

NCAR WRF Docker Next Steps

- Working with AWS Sci Engr on multi-node EC2, multi-container openmpi WRF runs. Scalability and Performance.
- Exploring further scientific reproducibility and application portability from laptops to servers to clouds.
(Amazon, Google, Azure)
- Develop Fall 2016 classroom curricula with Univ of North Dakota.
- Construct more complex case studies: dynamic user selectable datasets initialized from NOAA global model.
- Open beta access via public repositories hosted on dockerhub.com and github.com **June 28 2016!**

NCAR Docker Resources

- NCAR RAL Docker-WRF Project Web site:
- <https://www.ral.ucar.edu/projects/ncar-docker-wrf>
- NCAR Github repo: <https://github.com/NCAR/docker-wrf>
- NCAR Dockerhub repo: <https://hub.docker.com/r/bigwxwrf/>
 - bigwxwrf/ncar-wrf
 - bigwxwrf/ncar-wpsgeog
 - bigwxwrf/ncar-wrfinputkatrina
 - bigwxwrf/ncar-wrfbin
 - bigwxwrf/ncar-wrfinputsandy
 - bigwxwrf/ncar-ncl
- Slack channel for docker-wrf community discussion.
- email: exby@ucar.edu for invitations
- <https://ncar-dockerwrf.slack.com>

NCAR Docker WRF

- Thank you!
- All Questions accepted!