JUPITER

# Cloud-based WRF Downscaling

## Simulations at Scale using Community Reanalysis and Climate Datasets

Luke Madaus -- 26 June 2018
luke.madaus@jupiterintel.com

2018 Unidata Users Workshop

# Outline

- What is Jupiter?
- Challenges to running WRF at scale
- Jupiter's approach to this problem
  - Leveraging community datasets
  - Containerization
  - Distributed data-proximate analysis
- (Brief) demonstration
- Example use cases
- Looking ahead to broader access
- Questions for the group

# What is Jupiter?

# What is Jupiter?


≠

# What is Jupiter?


=

**Translating the Latest Climate Science to the Needs of Diverse Users**

# What is Jupiter?

Industry

Publish!

JUPITER

**Constraining Ensemble Forecasts of Discrete Convective Initiation with Surface Observations**

LUKE E. MADAUS[a] AND GREGORY J. HAKIM

*University of Washington, Seattle, Washington*

(Manuscript received 14 October 2016, in final form 17 February 2017)

ABSTRACT

Predicting when and where individual convective storms will develop remains an elusive challenge. Previous studies have suggested that surface observations can capture convective-scale features relevant to the convective initiation (CI) process, and new surface observing platforms such as crowdsourcing could significantly increase surface observation density in the near future. Here, a series of observing system simulation

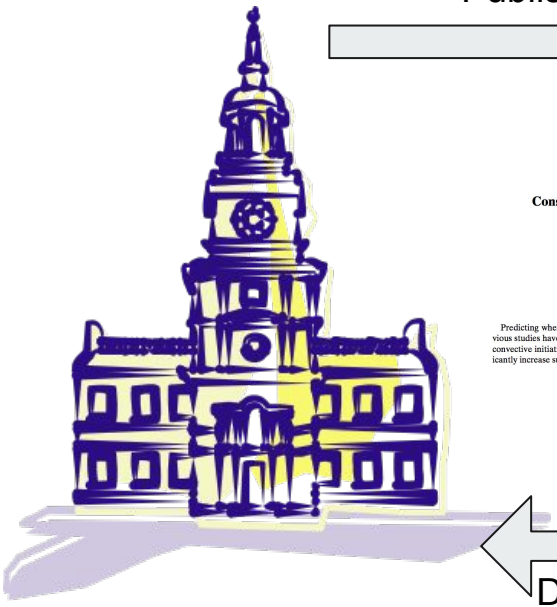

Don't perish!

Academia

Governments

JUPITER

# What is Jupiter?

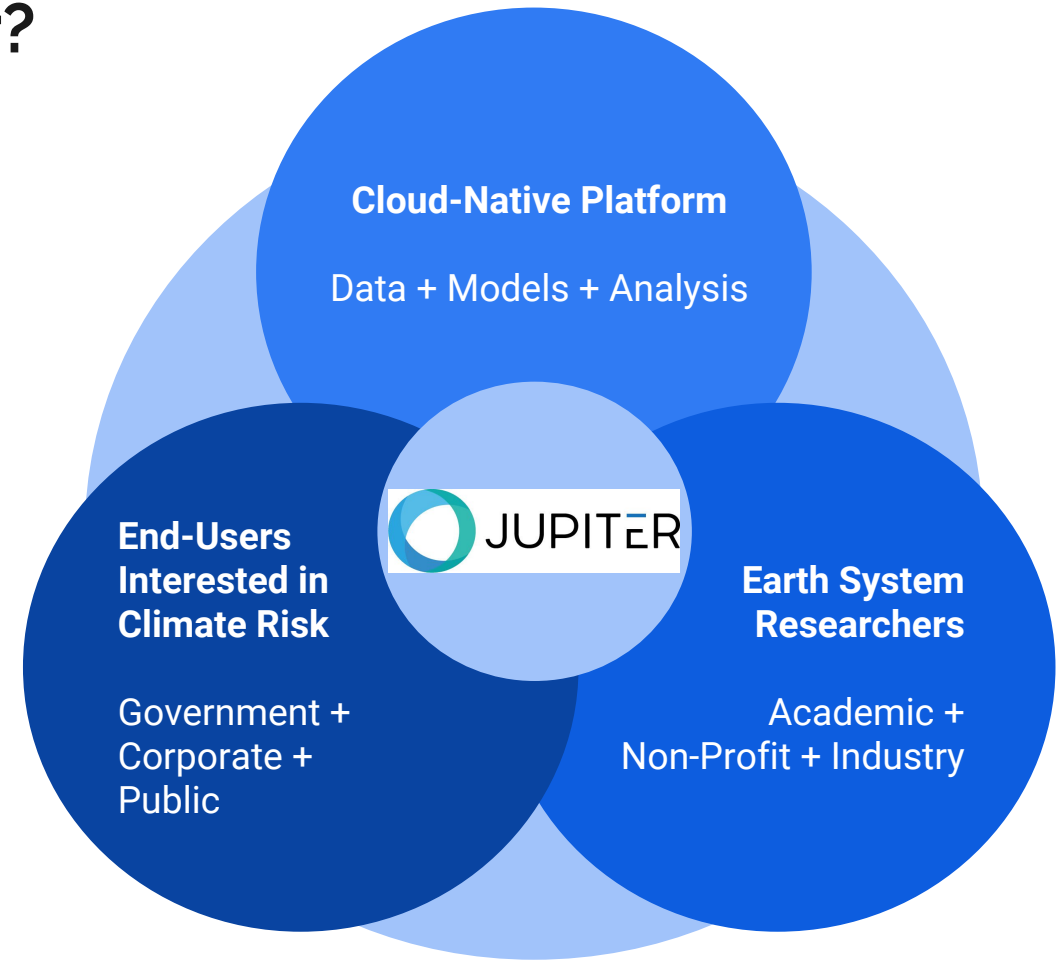

# What is Jupiter?

**FloodScore**
(Operations and Planning)

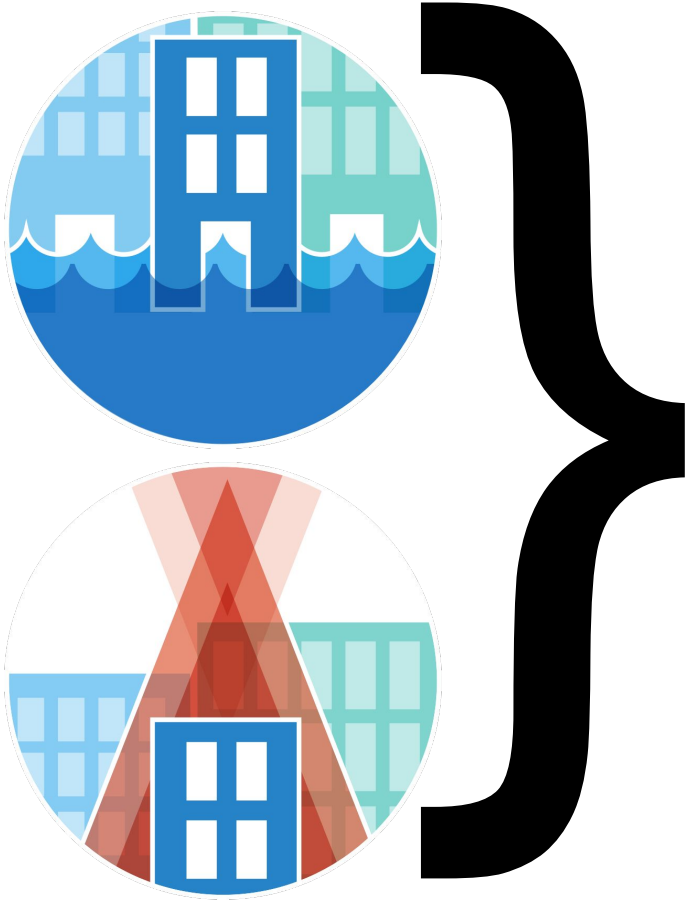

**HeatScore**

(Downscaling)

How are ***hyper-local-scale*** risks for flood and extreme heat hazards changing on ***daily to decadal timescales***?

JUPITER

# Regional WRF Simulations on the cloud -- Challenges

# Regional WRF Simulation

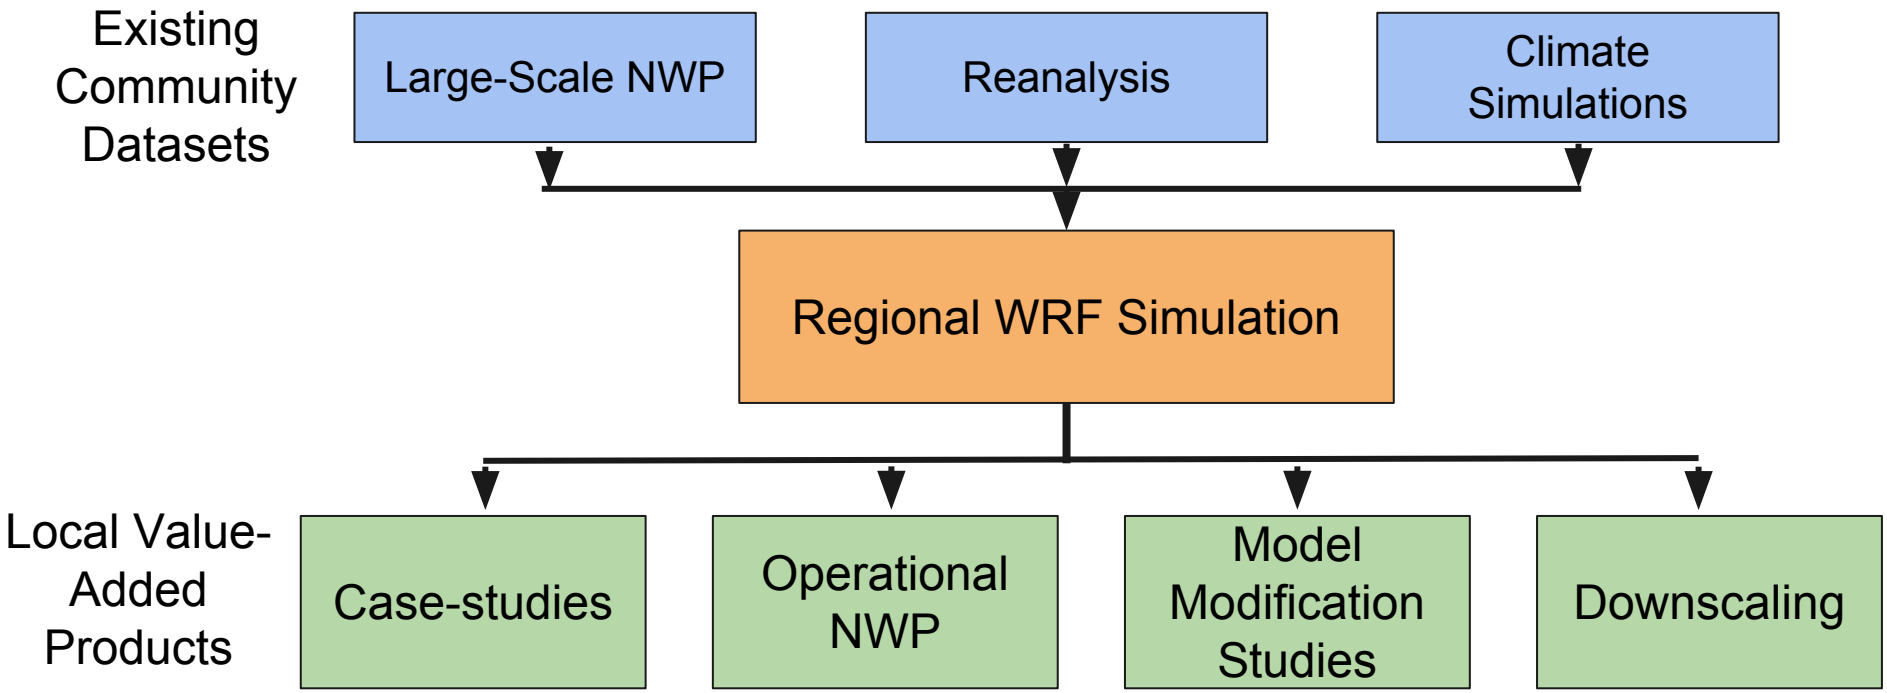


JUPITER

# Challenges

Existing Community Datasets

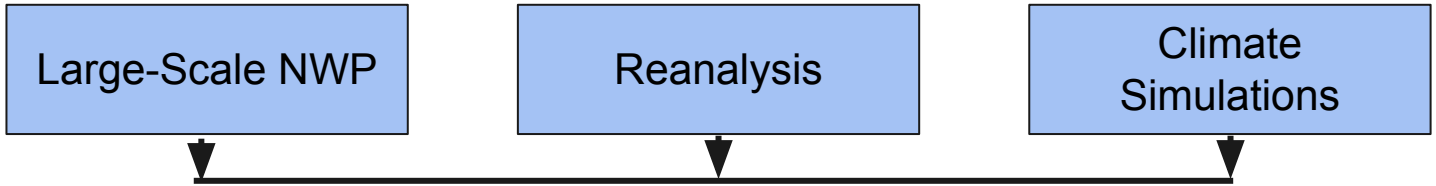


**Large files designed for bulk downloads!**

Current Solution: Have to download everything even if you just need a handful of dates/variables

**Distributed in netCDF, WRF designed for GRIB!**

Current Solution: Complicated reformatting and repackaging

**Variables not in the right format!**

Current Solution: Have to do some sort of calculations or manipulations on data to make it usable for WRF

JUPITER

# Challenges

**How do I make WRF work on my system?**

Current Solution: While WRF and supporting libraries have gotten easier to compile and install, still a challenge

**I want to simply swap out another version of my code?**

Current Solution: Changing dynamic links, keep track of pointers, making sure things aren't overwritten...

Regional WRF Simulation

**I need to move my work to a new cluster. How?**

Current Solution: Reinstalling libraries, recompiling, running baseline checks again...

**How do I share all the data I'm generating?**

Current Solution: Downloading to local machine, setting up own file server...

JUPITER

# Regional WRF Simulations on the cloud -- Jupiter's Approach

# Jupiter Platform

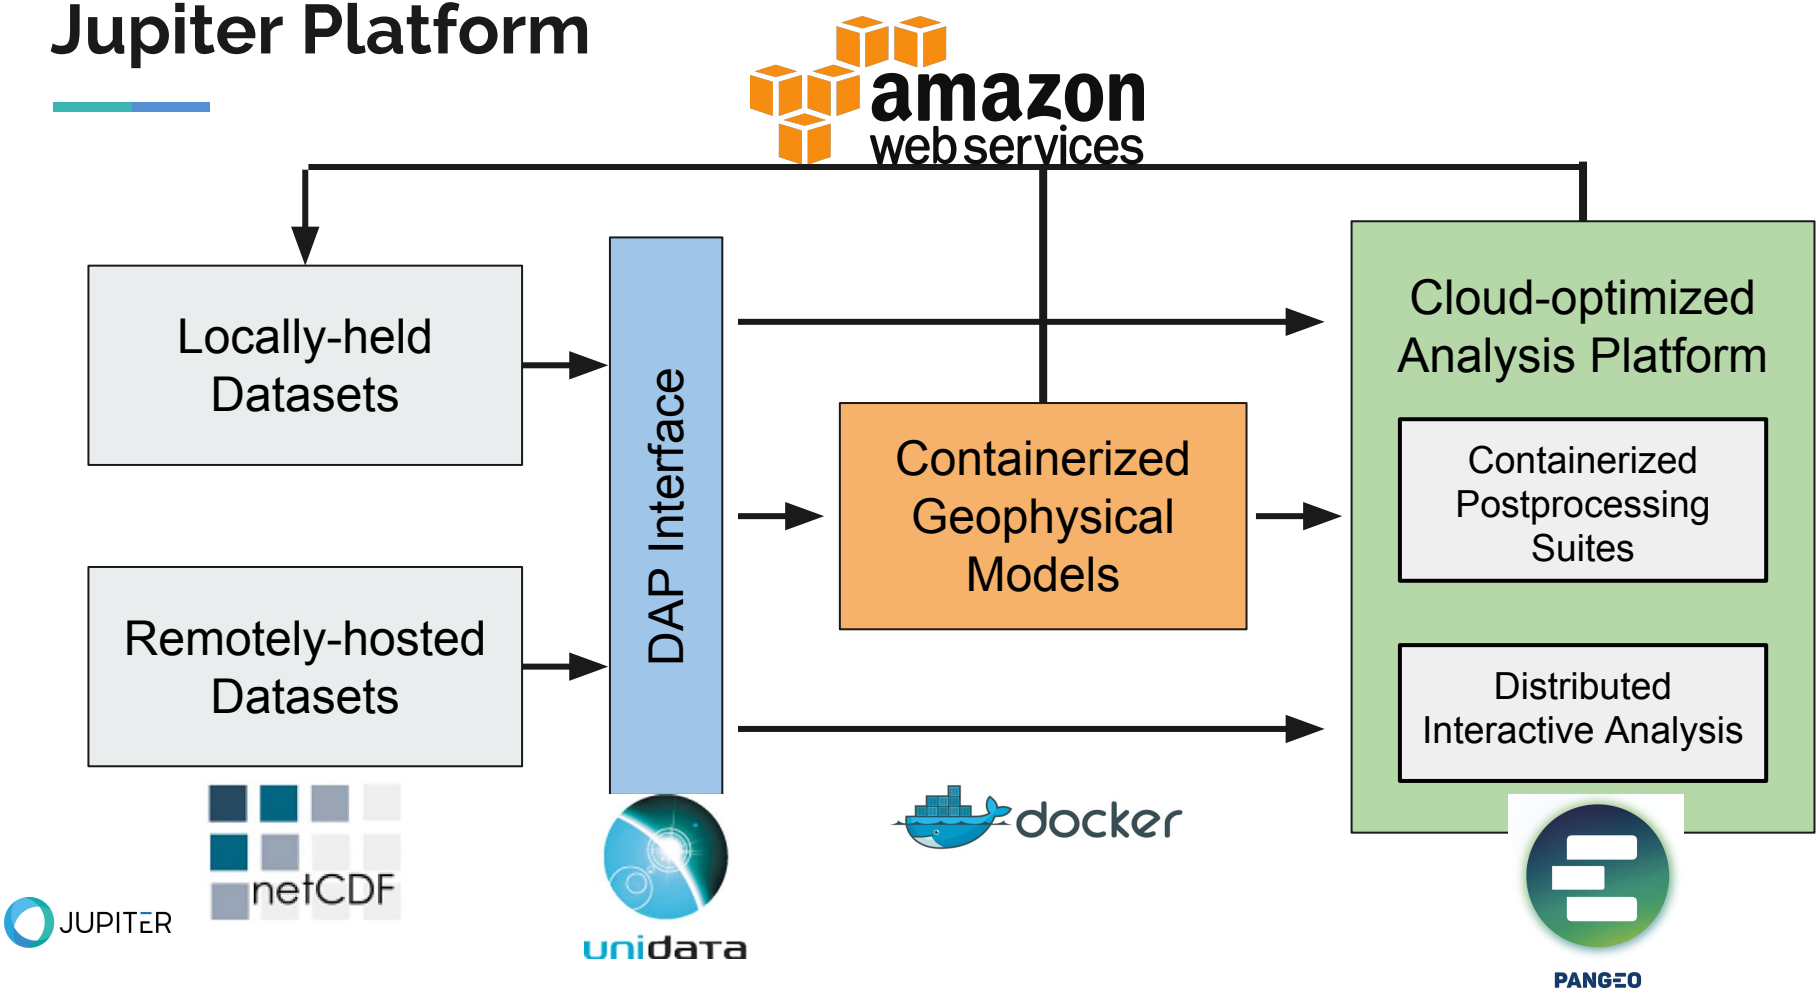

# Jupiter Platform

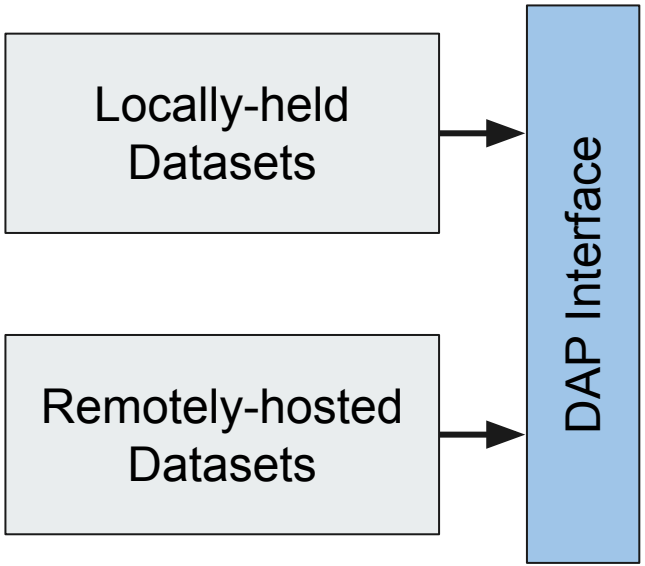


## The Data

- netCDF format as current standard
- Increasing number of weather and climate datasets hosted through streamable platforms (e.g., THREDDS)
- On-the-fly subsetting, minimize download size
- Jupiter exploring more cloud-optimized data formats and hosting methods (e.g. Zarr...more on this at the end)

# DAP-Fronted Community Datasets

NASA -- CREATE-IP
Reanalysis Datasets (Potter et al. 2018[1])

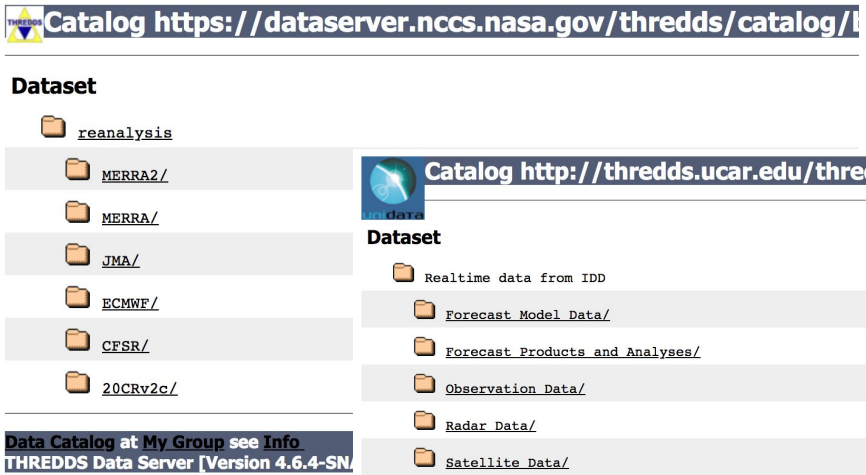


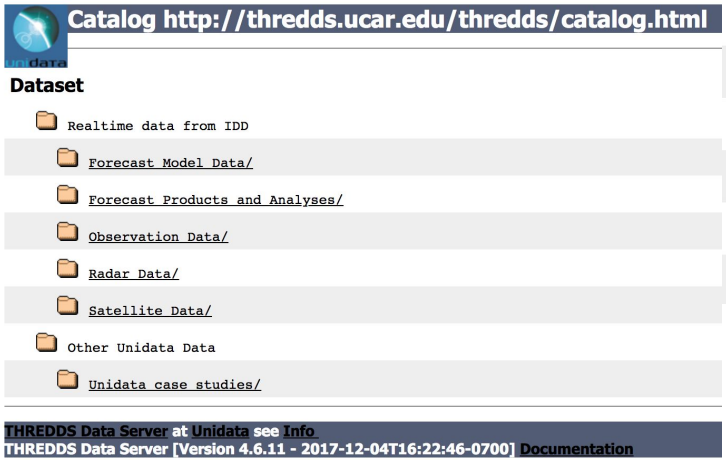


NOAA-ESRL

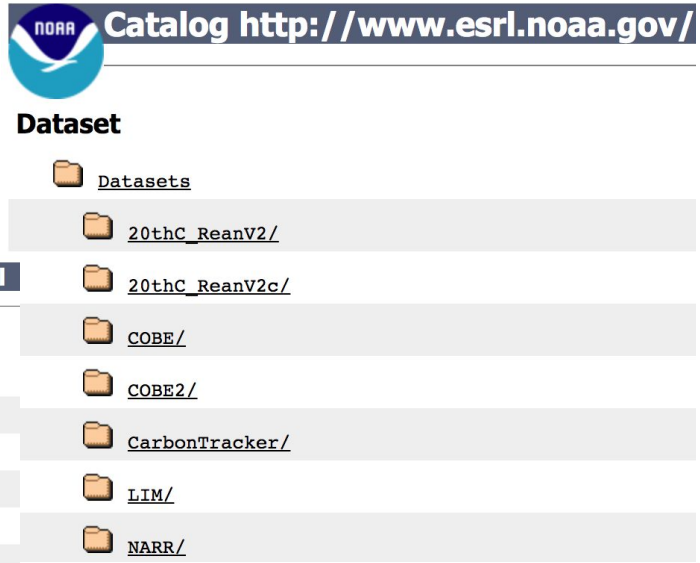


Unidata

Many more...

JUPITER

1- Potter, G.L., L. Carriere, J. Hertz, M. Bosilovich, D. Duffy, T. Lee, and D.N. Williams, 2018: Enabling Reanalysis Research Using the Collaborative Reanalysis Technical Environment (CREATE). *Bull. Amer. Meteor. Soc.*, **99**, 677–687, https://doi.org/10.1175/BAMS-D-17-0174.1

# Jupiter Platform

## The Models

- NSF BigWeatherWeb: Containerization allows confident deployment of models at scale (Hacker et al. 2017[1])
- Preprocessing automated -- ***direct from netCDF to WRF***
- Batch and on-demand simulation for research and real-time forecasts
- Pushing computational limits of current cloud compute architectures

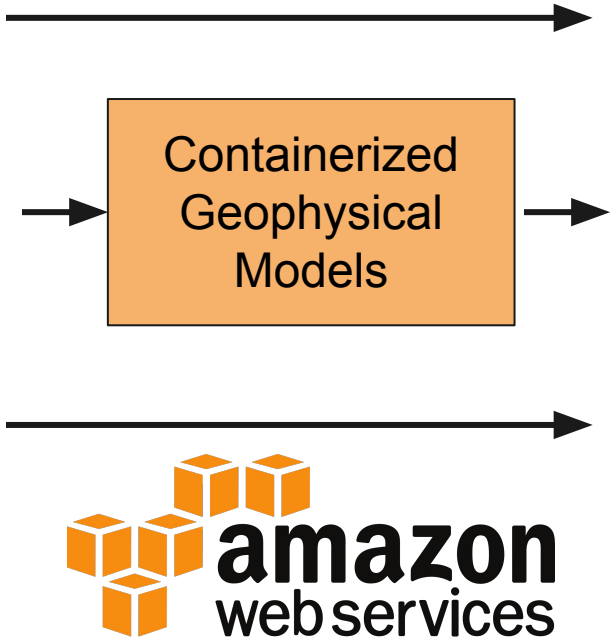


JUPITER

1 - Hacker, J.P., J. Exby, D. Gill, I. Jimenez, C. Maltzahn, T. See, G. Mullendore, and K. Fossell, 2017: A Containerized Mesoscale Model and Analysis Toolkit to Accelerate Classroom Learning, Collaborative Research, and Uncertainty Quantification. *Bull. Amer. Meteor. Soc.,* **98**, 1129–1138, https://doi.org/10.1175/BAMS-D-15-00255.1

# What is containerization?

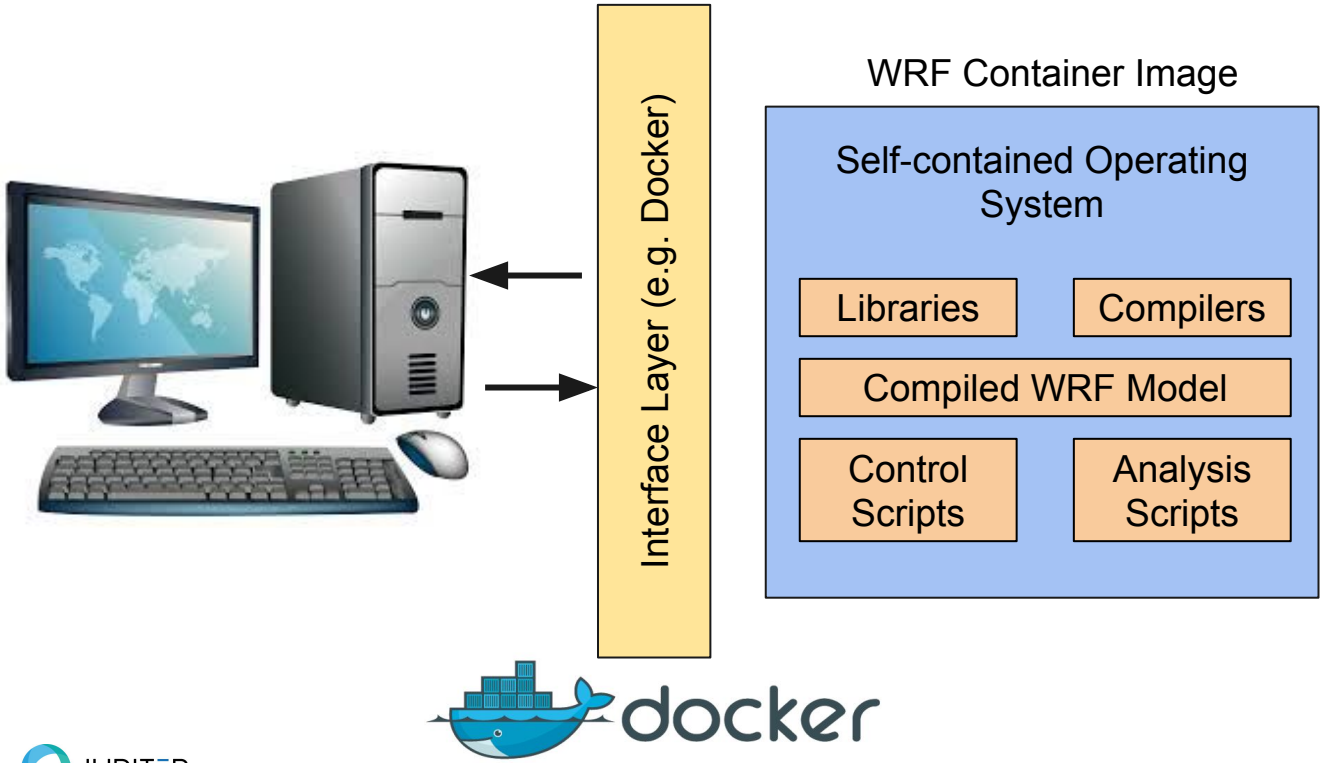


A portable, self-contained, only what is needed package that can perform some function on ***any machine*** running an ***interface layer***.

JUPITER

# What is containerization?

Pre-built container images can be “pulled” from online repositories

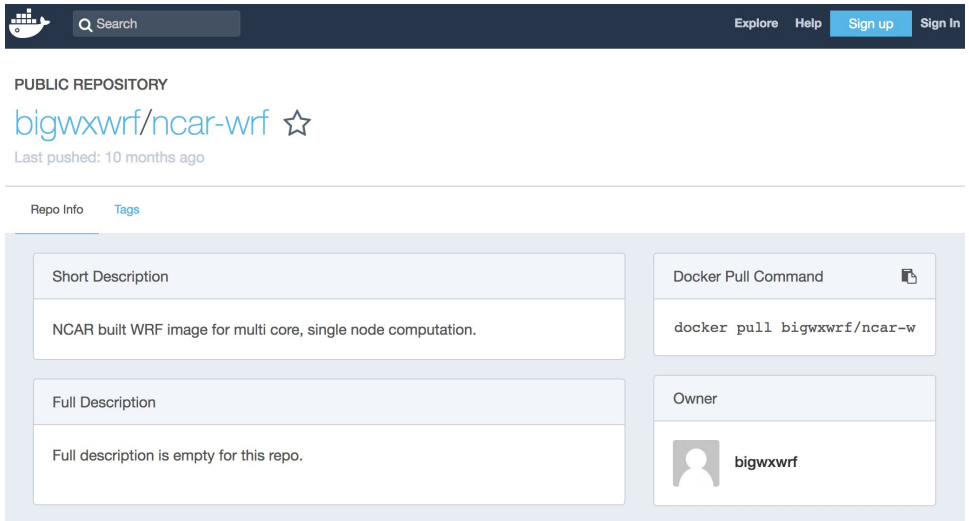


Or, a single text file that contains a “recipe” for building the container can be downloaded for local builds

```
FROM fedora:latest
LABEL maintainer="Luke Madaus <luke.madaus@jupiterintel.com>"
#
# This Dockerfile compiles WRF from source during "docker build" step
#ENV WRF_VERSION 3.7.1
ENV WRF_VERSION=4.0 \
    CONDA_VERSION=latest
RUN yum -y update && yum -y install file gcc gcc-gfortran gcc-c++ glibc.i686 libgcc.i686 libpng-devel \
    jasper jasper-devel hostname m4 make perl tcsh time wget which zlib-devel \
    openssh-clients openssh-server net-tools openssl openssl-devel hdf5 hdf5-devel \
    mpich mpich-devel netcdf-mpich-devel netcdf-fortran-mpich-devel geos-devel \
    redhat-rpm-config bzip2      You, 7 months ago • run-wrf.py now checks for time deltas in metgrid and WPS-format

# Install python3 libraries for reading netcdf and grib files
#RUN pip3 install numpy NETCDF4 pyproj boto3
#RUN pip3 install https://github.com/matplotlib/basemap/archive/v1.0.7rel.tar.gz
#RUN pip3 install pygrib
# Switch to using conda
RUN wget http://repo.continuum.io/miniconda/Miniconda3-${CONDA_VERSION}-Linux-x86_64.sh -O ~/miniconda.sh && \
   bash ~/miniconda.sh -b -p /miniconda && \
   export PATH="/miniconda/bin:$PATH"
ENV PATH=/miniconda/bin:$PATH
# Use conda to install the required libraries
RUN conda install -y -c conda-forge numpy netCDF4=1.3.1 pyproj xarray dask esmf xesmf wrf-python requests \
 && pip install Pydap

# Get the license from NCAR
RUN curl -SL https://ral.ucar.edu/sites/default/files/public/projects/ncar-docker-wrf/ucar-bsd-3-clause-license.pdf
#
# Set some environment variables to the bashrc file for runtime
RUN echo export CC=gcc >> /etc/bashrc \
    && echo export FC=gfortran >> /etc/bashrc \
    && echo export CXX=g++ >> /etc/bashrc \
    && echo export FCFLAGS=-m64 >> /etc/bashrc \
```

JUPITER

# Scalable WRF Simulations in <10 lines of Python

Python-based interface to platform

Cloud-based storage keeping simulation data organized and sharable

Simulation “job” options self-contained in Python object

Explore costs of running simulation with various cloud options

Command-line submission of configured jobs to cloud compute

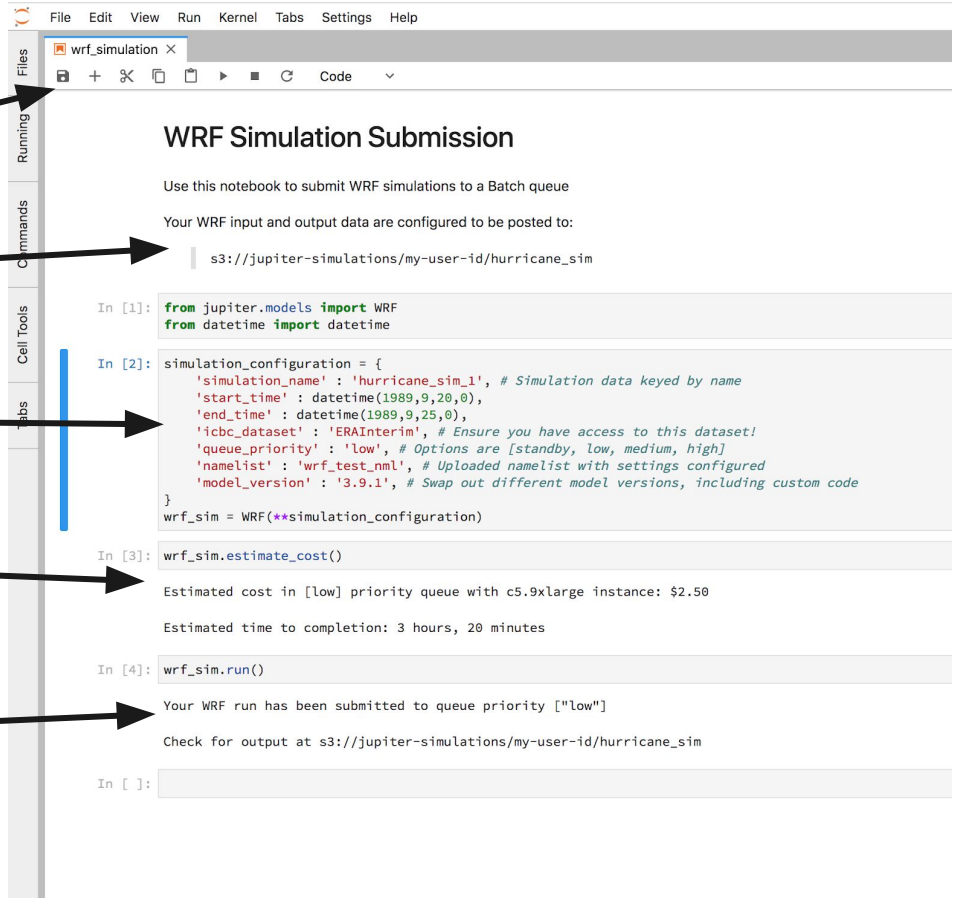


JUPITER

# Demonstration

# Jupiter Platform

## The Analysis

- Containerization allows standardized post-processing or analysis chains to be deployed
- Distributed analysis of large datasets through interactive tools (Pangeo stack)
- Allow end-users to dynamically query data for questions of interest

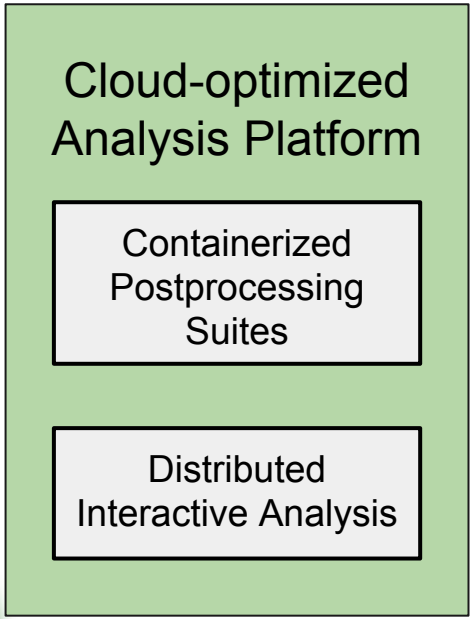


PANGEO

# Example Use Cases

# Example Uses

Extensive WRF benchmarking tests across a variety of cloud resource configurations

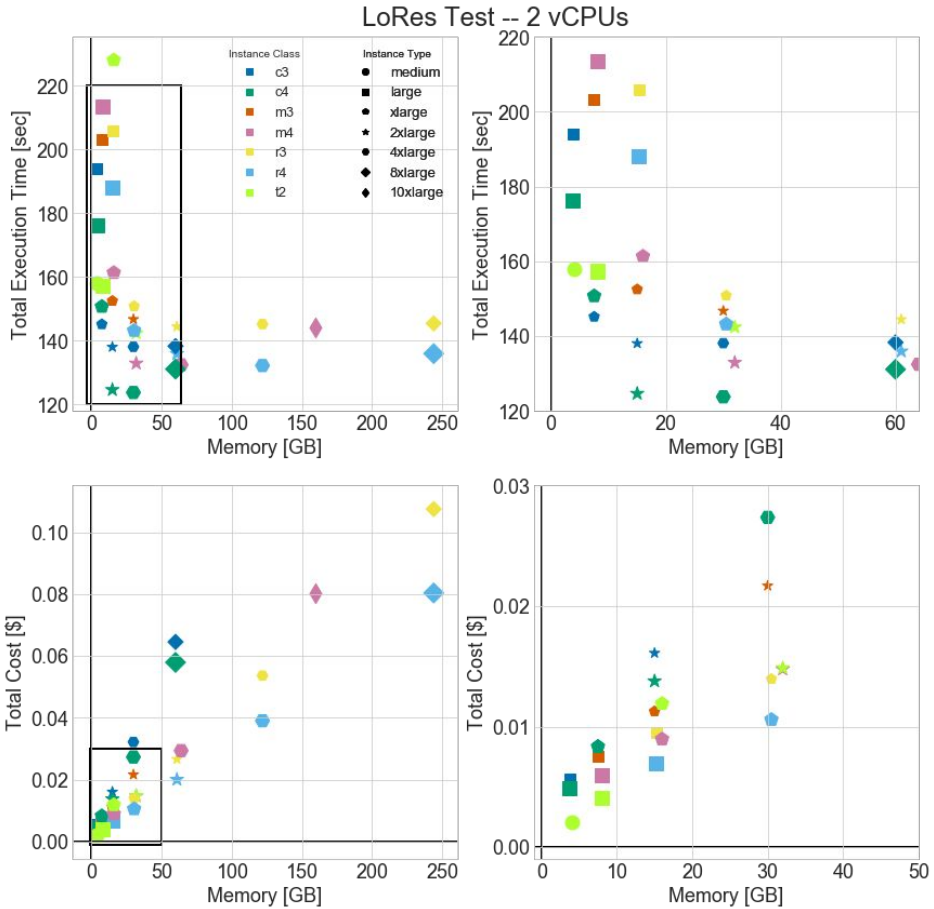


JUPITER

# Example Uses

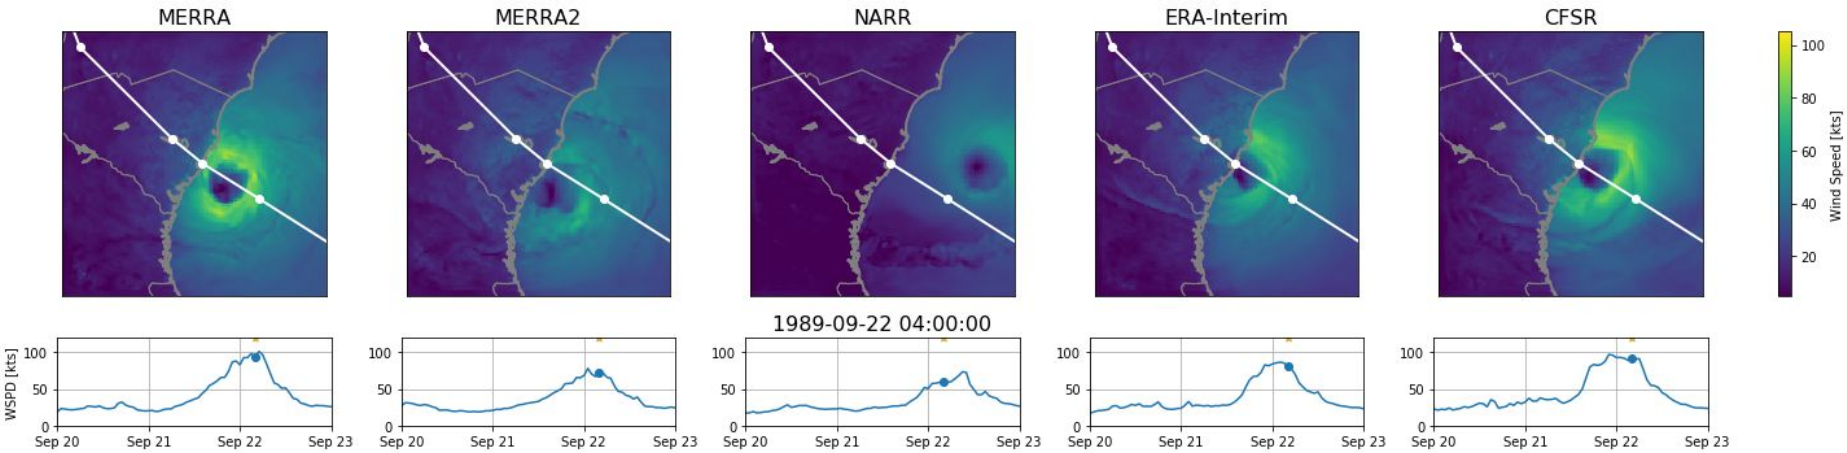


**Hurricane Hugo (1989) from the perspective of multiple reanalyses**

- None of this reanalysis data held locally
- Total pre-processing time: 1 minute per simulation (in parallel)
- WRF execution: 5 day simulation, nesting down to 3 km, 1hr 30min (in parallel)
- Conception to execution ->  about 10 lines of Python on Jupiter platform

JUPITER

# Example Uses

- Dynamical downscaling simulations to project extreme precipitation in Charleston, SC using 10 member of the CESM-LENS
- Thousands of simulations of extreme precipitation events; completed in one weekend

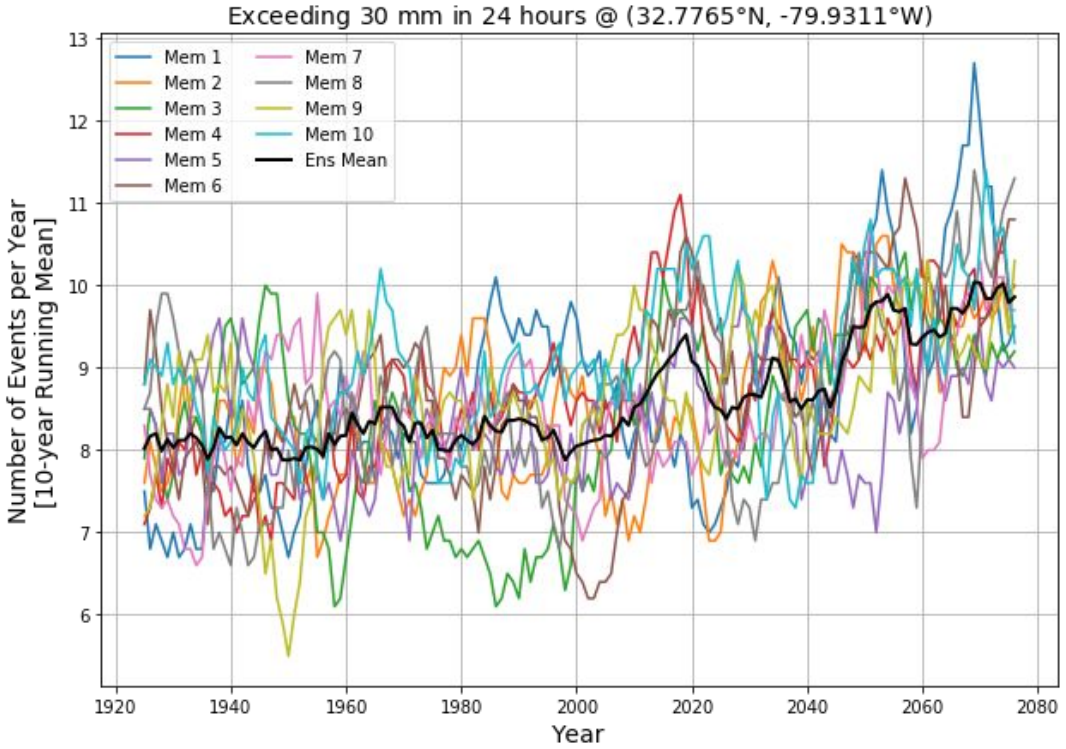


JUPITER

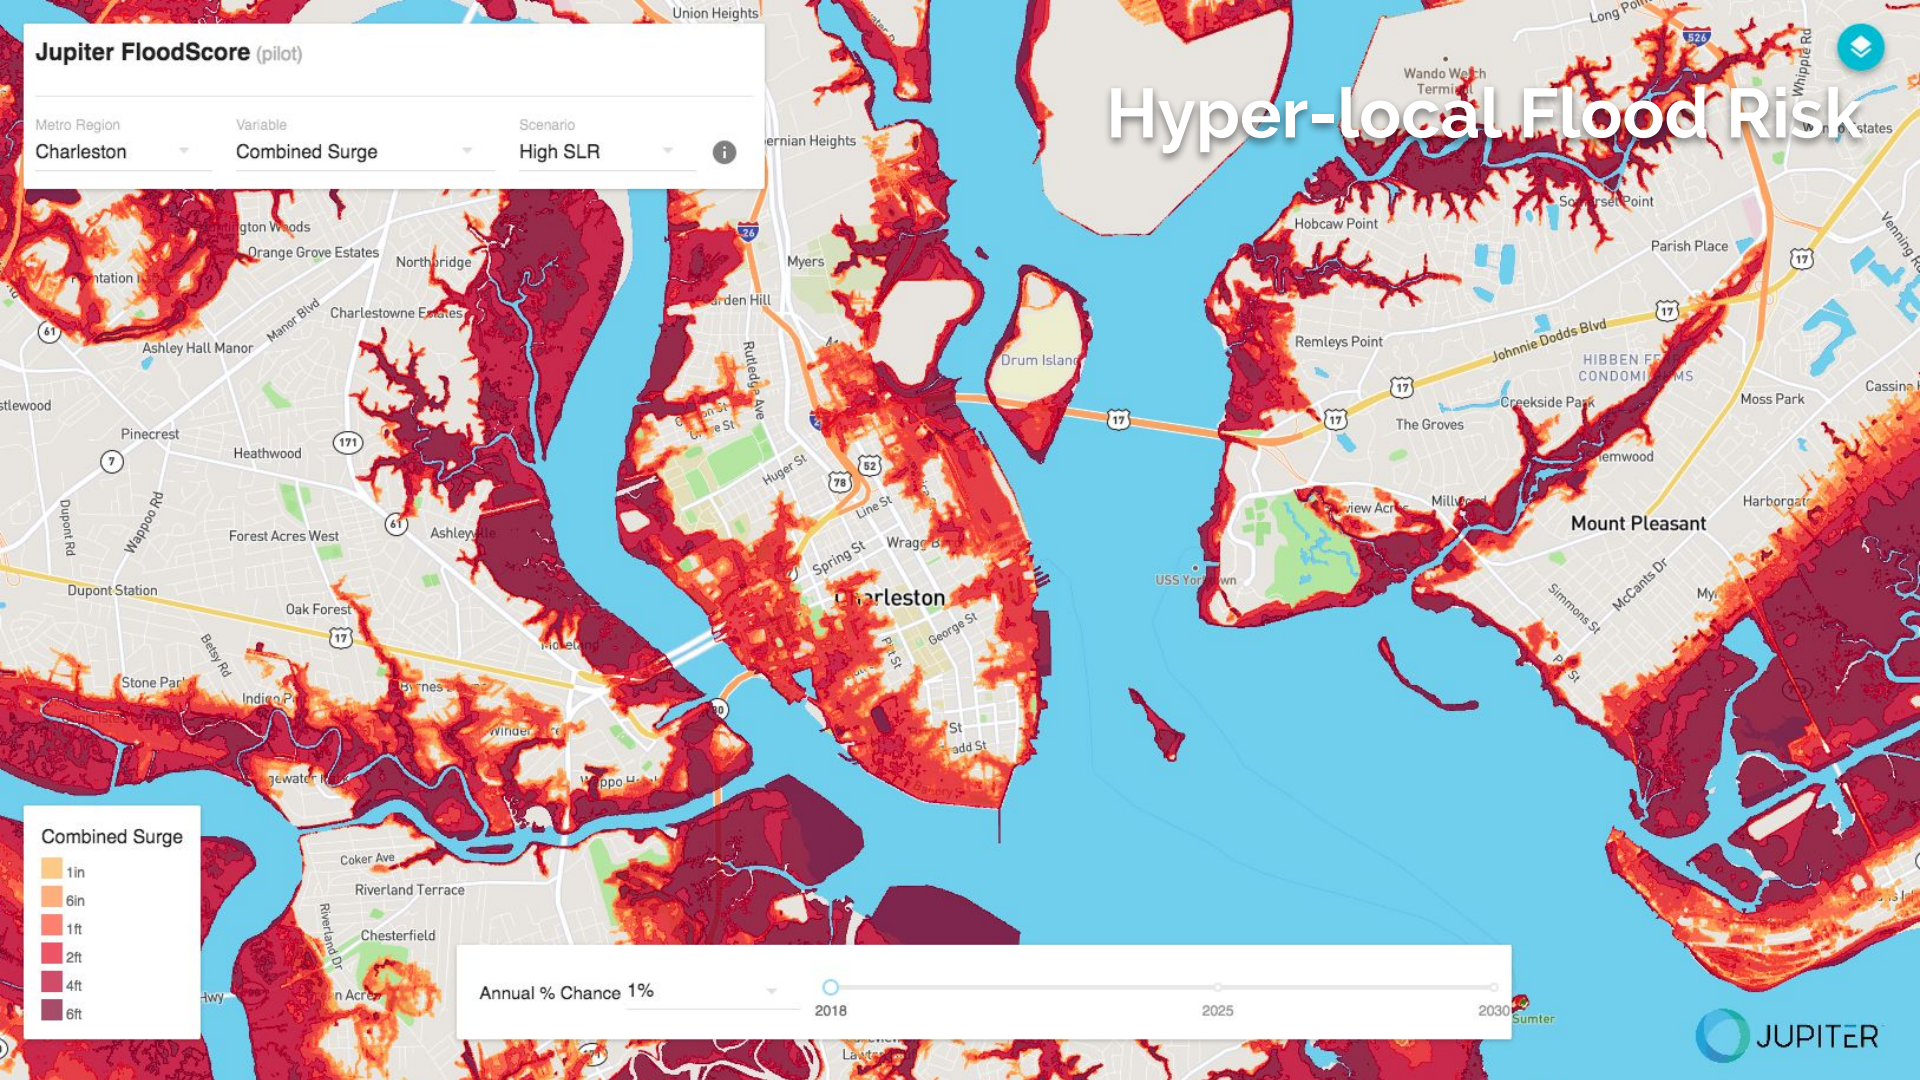
Jupiter FloodScore (pilot)
Metro Region
Charleston
Variable
Combined Surge
Scenario
High SLR
Hyper-local Flood Risk
Combined Surge
1in
6in
1ft
2ft
4ft
6ft
Annual % Chance 1%
2018
2025
2030
Charleston
Mount Pleasant
Union Heights
Orange Grove Estates
Northbridge
Charlestowne Estates
Ashley Hall Manor
Pinecrest
Heathwood
Forest Acres West
Dupont Station
Oak Forest
Stone Park
Riverland Terrace
Chesterfield
Coker Ave
Myers
Drum Island
Hobcaw Point
Parish Place
Remleys Point
Johnnie Dodds Blvd
Creekside Park
The Groves
Moss Park
Harborgate
Wando Welch Terminal
JUPITER

# Looking Ahead to Broader Access

# BigWeatherWeb


NSF Award #1450488
Overview at:
http://bigweatherweb.org/Big_Weather_Web/Home/Home.html

BAMS Article: **A Containerized Mesoscale Model and Analysis Toolkit to Accelerate Classroom Learning, Collaborative Research, and Uncertainty Quantification**

Containerized WRF available!
https://github.com/NCAR/container-wrf
https://hub.docker.com/r/bigwxwrf/ncar-wrf/

Coming soon: Code for generating WRF initial conditions from arbitrary netCDF files

JUPITER

# Working on the Jupiter Platform

- Jupiter Containerized WRF targeted for Autumn 2018 release (GitHub)

- Jupiter platform access on a project-by-project basis
  - Let me know if you think you have a good use case!

# Questions for this group

1. Cloud-optimized netCDF
    a. An ever-growing need -- case for it is only strengthening
    b. Is this the way to go?
    c. How easy would it be?
2. Containerized WRF simulations
    a. Could you envision using this given your current compute platforms/resources?
    b. What features would you find useful in a standalone, containerized WRF?
3. Cloud-hosted datasets and data-proximate workflows
    a. What would it take to motivate you to move some (or all) of your work to the cloud?

luke.madaus@jupiterintel.com | @lmadaus | http://www.jupiterintel.com

JUPITER