# Numerical Weather Prediction Science and HPC at CAPS University of Oklahoma

Keith Brewster, PhD
Senior Research Scientist and CAPS Associate Director
kbrewster@ou.edu

Advanced Cyberinfrastructure Research &
Education Facilitators Virtual Residency 2017
31-July-2017

# CAPS Introduction

## Center for Analysis and Prediction of Storms

- **History**:
  - Established 1989 as NSF Science and Technology Center
  - A leader in convective-scale data assimilation and numerical weather prediction (NWP)
- **People:**
  - 2 Primary Faculty, ~10 Affiliated
  - 7 Research scientists
  - 15 Graduate students & Post-Docs
  - 2 Administrative & IT support staff
- **Research:**
  - ~50 publications per year (author or co-author)
  - ~ $3 million supported research per year

25th Anniversary 1989-2014

# CAPS Research Areas

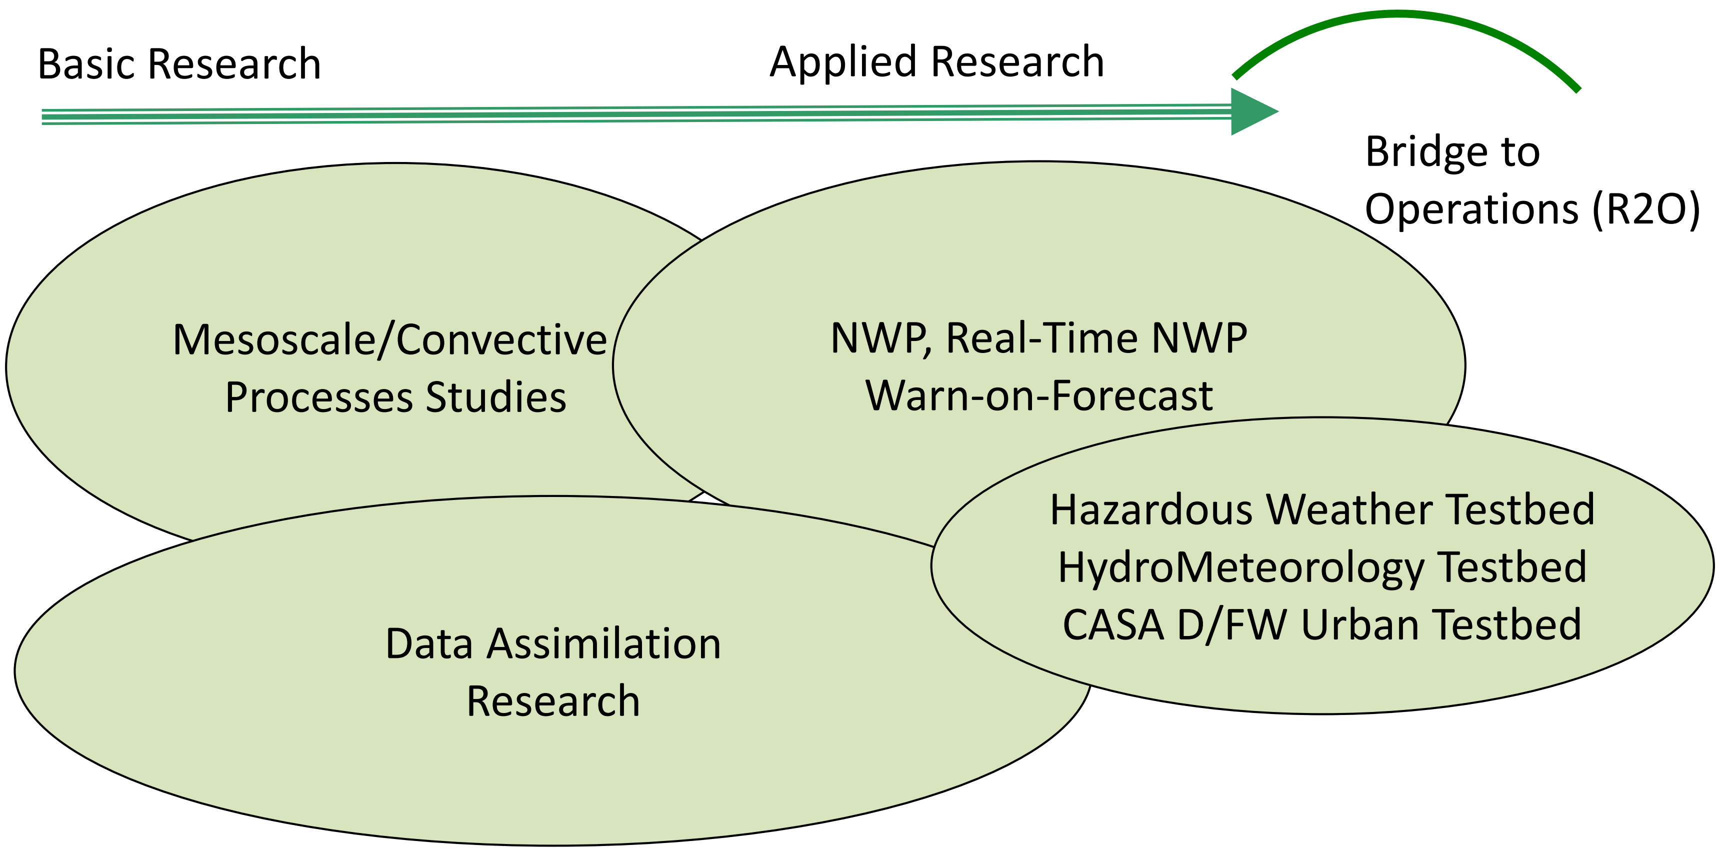


www.caps.ou.edu

# Scope of Research

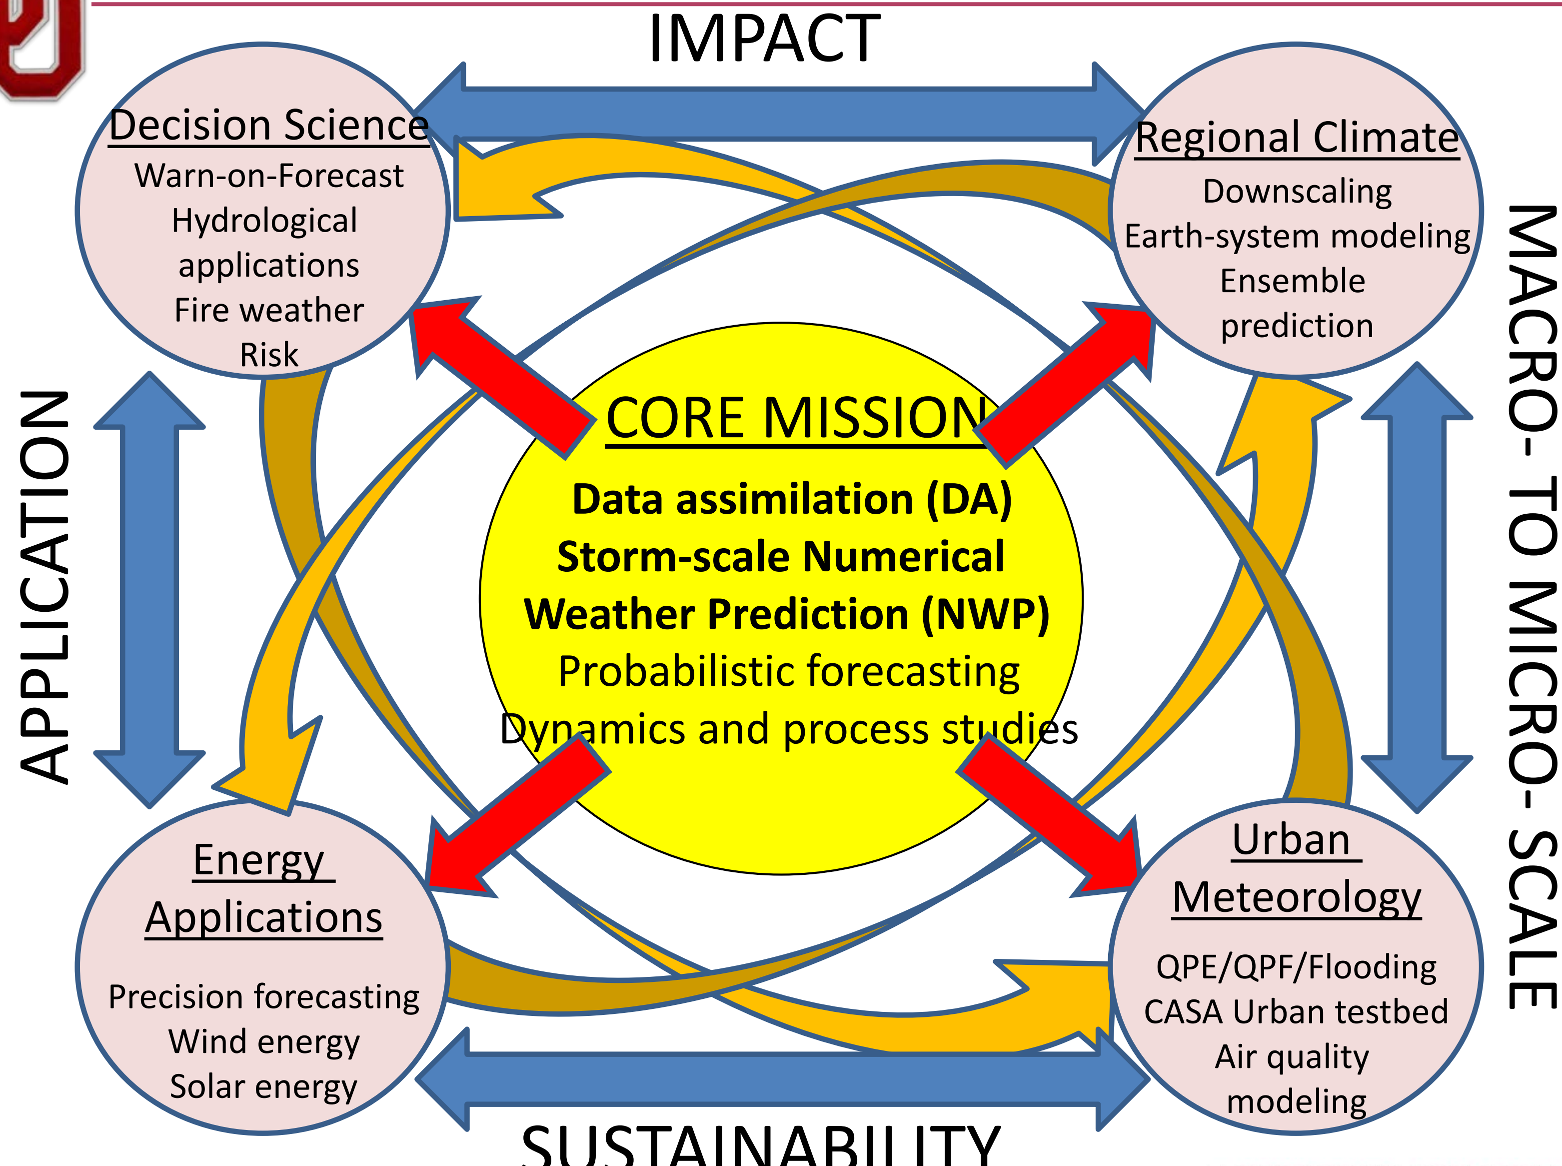

# Short-Term Weather Forecasting Continuum

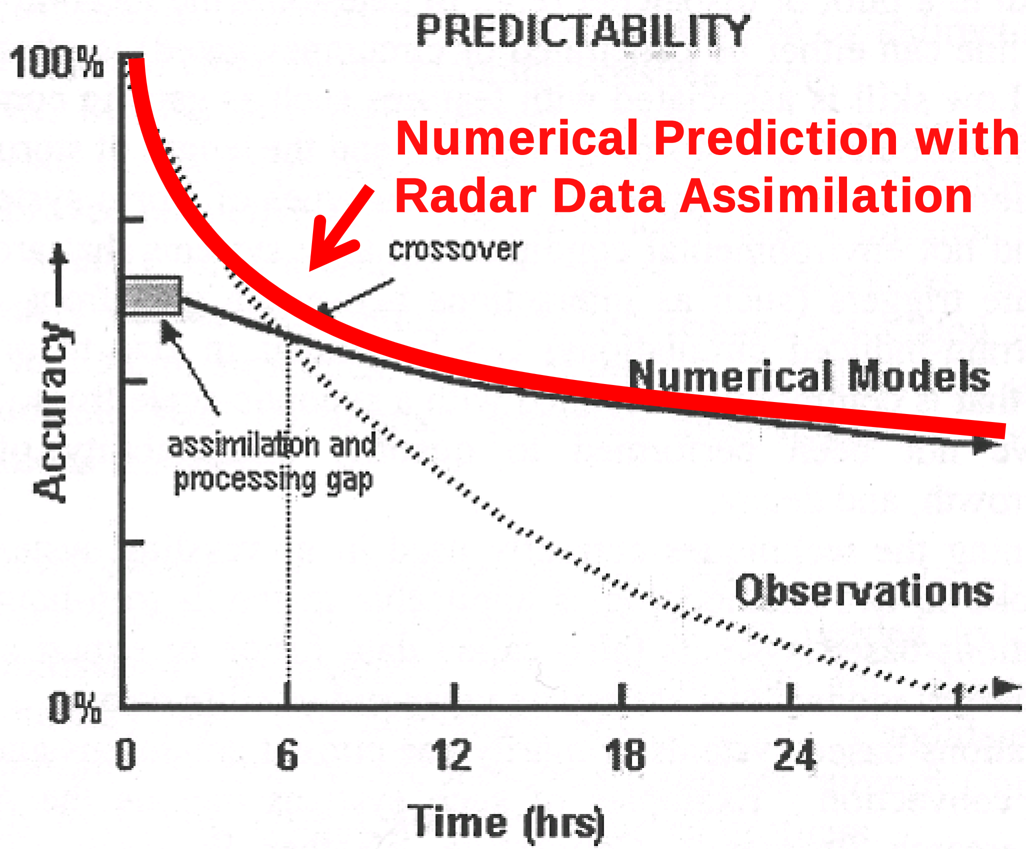


When deterministic predictability becomes low, ensembles are valuable.

# Current US Severe Weather Warning Paradigm

Warnings are a natural culmination of weather information generated and distributed over a period of several days...

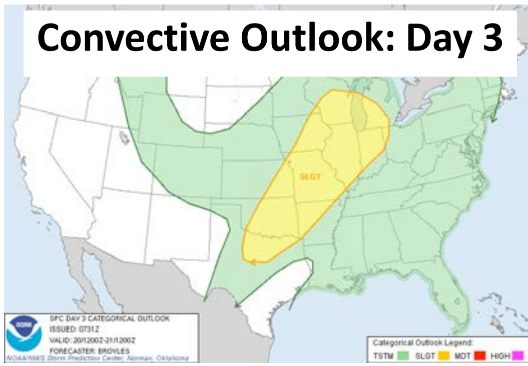


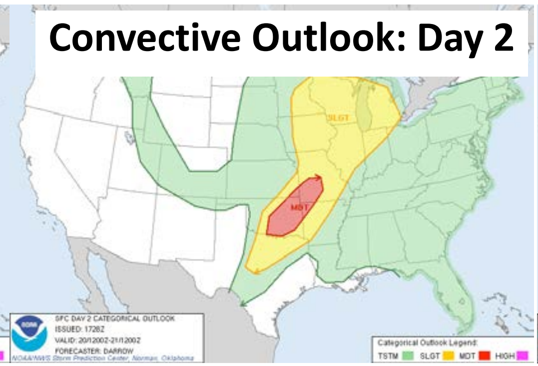


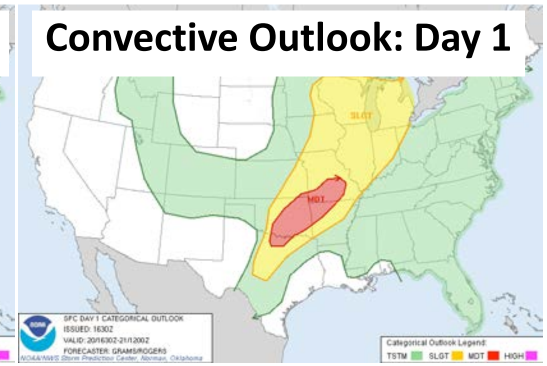


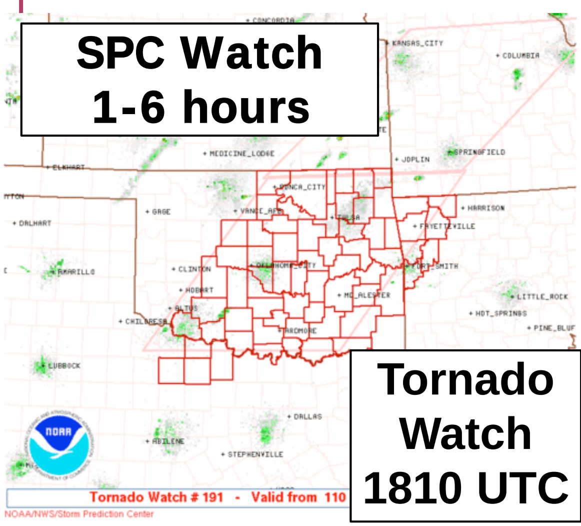


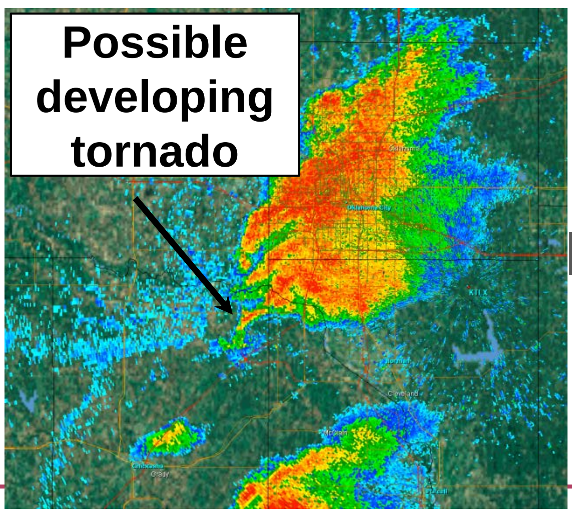


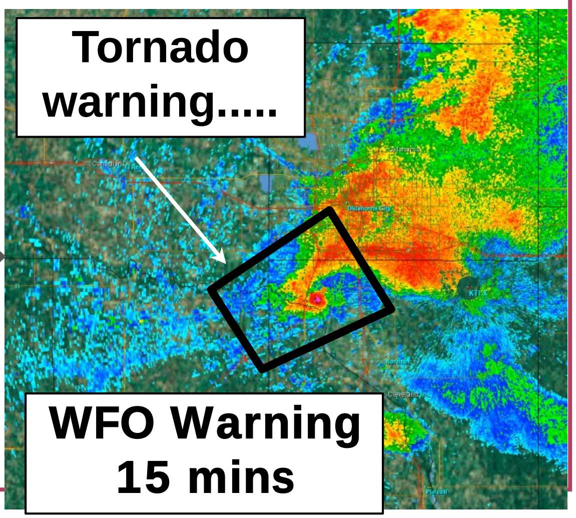

# CAPS Research-to-Operations (R20)

## Hazardous Weather Testbed (HWT)

- Collaboration with NOAA Storm Prediction Center, NSSL, Aviation Weather Center, others
- 2015 - 2017: 3-km CONUS Forecast Ensembles
- Probabilistic NWP for Severe Weather : 0-2.5 Day Forecasts

Probability of Storm Tops > 35 kft 24h Forecast

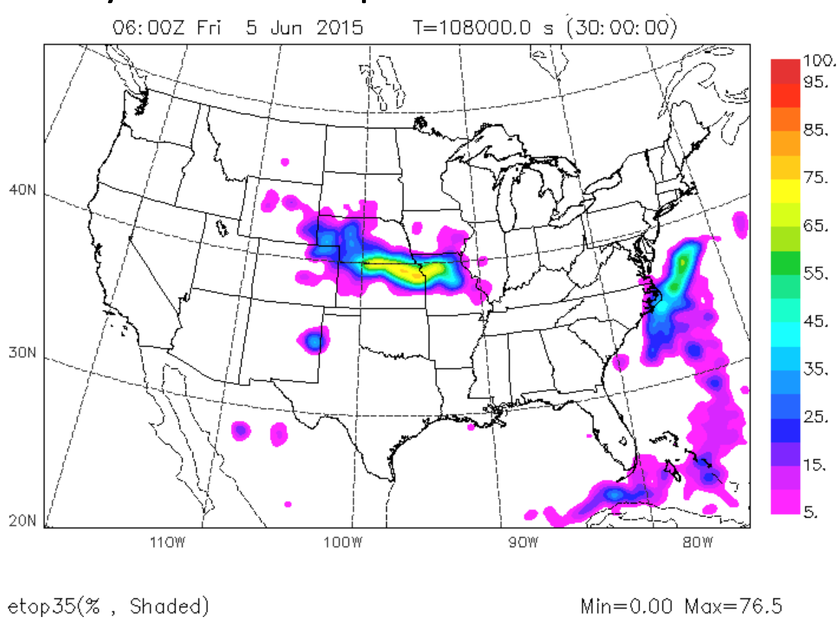


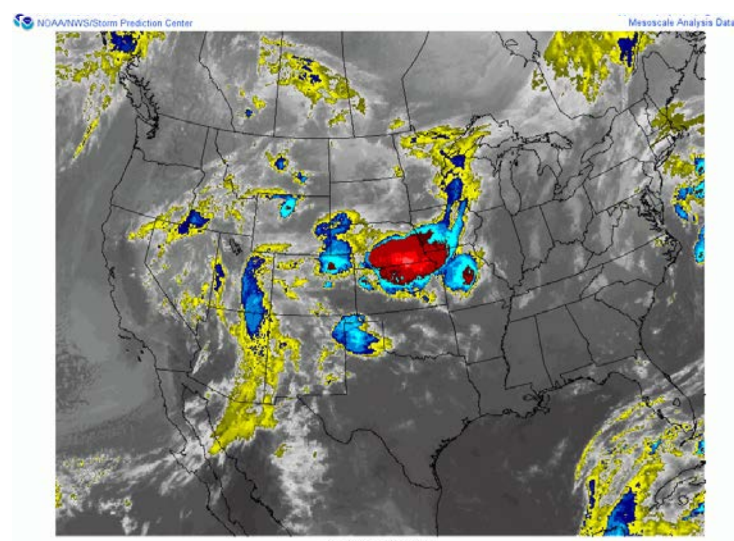

# Storm-Scale Ensemble Forecasts

- Weather Research and Forecasting Model (WRF)
  - Community model developed with goal of having researchers and operational forecasting based on same model framework
  - Finite Difference Model on a staggered grid
  - Fortran 90
  - MPI via Domain Decomposition
  - Different flavors of the core developed (ARW and NMM)
  - Multiple options for sub-grid parameterizations
- CAPS Perl-based job control
- CAPS Observation Analysis Programs (ARPS-based)
- Ensemble: Multiple runs with different initial conditions, core and/or sub-grid parameterizations

# Workflow (Readers Digest Version)

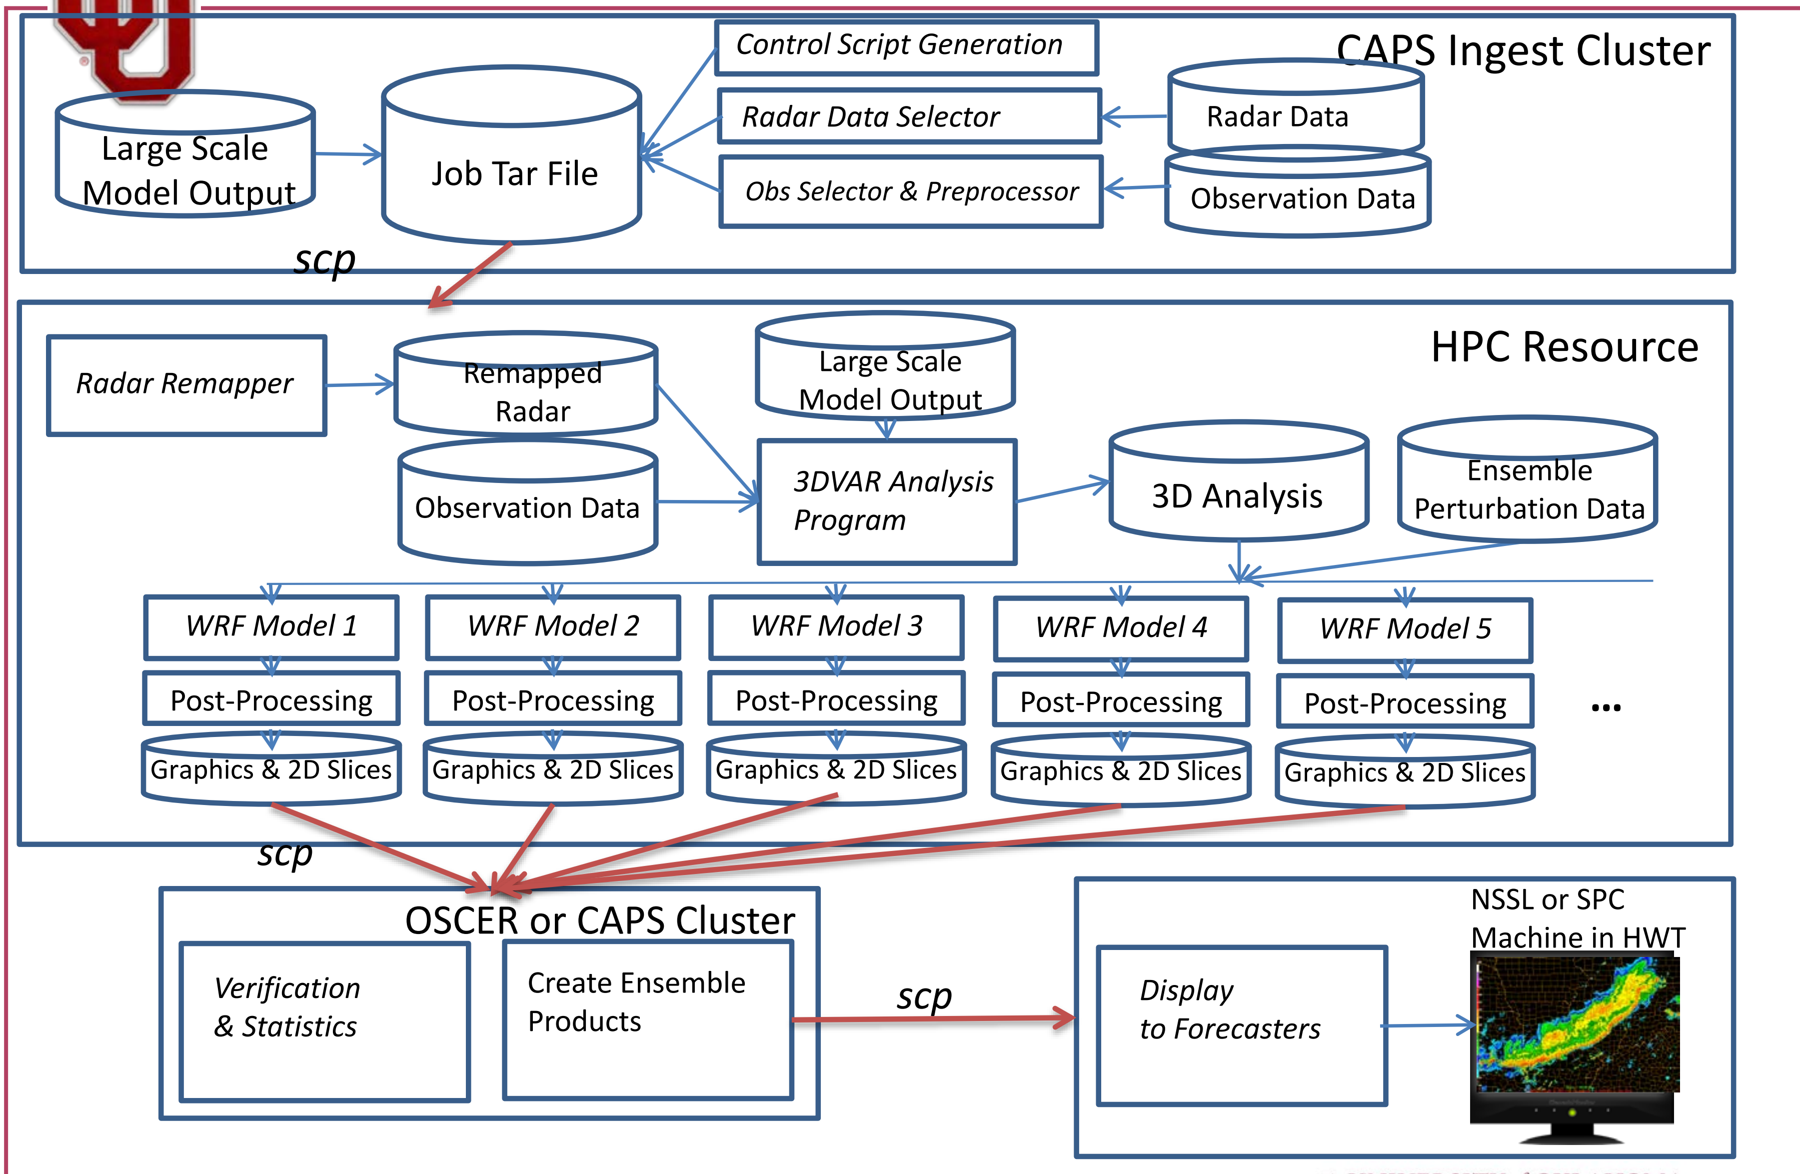

# 2016 CAPS SSEF for HWT Highlights

## Storm Scale Ensemble Forecasts (SSEF)

- 18 April 2016 through 3 June 2016
- 3-km horizontal grid spacing
  (ARW: 1680×1152; NMMB: 1568x1120)
- WRF version 3.7.1 (coupled with ARPS v5.4)
- 1) 3DVAR SSEF: 18 ARW members, 6 NMMB members, initiated with 3DVAR analysis & Cloud/Hydrometeor Analysis at 0000 UTC, with 36- to 60-h forecast, running on **Stampede at TACC**
- 2) GSI+EnKF SSEF: 40-member storm-scale ensemble background, a 5-hour GSI+EnKF hourly cycling and 1-hour EnKF radar data at 15 min interval, and a 9-member ensemble forecast starting at 0000 UTC. Running on **Darter at NICS**
- 3D Visualization demonstration (VAPOR)

Supported by NWS CSTAR & HWT grants and NSF XSEDE computing resources

XSEDE Extreme Science and Engineering Discovery Environment

The UNIVERSITY of OKLAHOMA

# 2016 CAPS SSEF for HWT Computing

- 1) 3DVAR SSEF: **Stampede at TACC**
  **686 nodes/10096 cores**
  Dell C8220 running CentOS
  Each Node: Intel E5 Sandy Bridge (Dual 8-core processor)
  MIC Co-processor

- 2) GSI+EnKF SSEF: **Darter at NICS**
  **360 nodes/5760 physical cores/11520 hyperthreaded**
  Cray XC30 running CLE 5.2
  Each Node: Intel E5 Sandy Bridge (dual 8-core processor)

- 3D Visualization demonstration (VAPOR)
  Data merge and segment on Stampede TACC on single node
  Transfer: Internet-2 and Oklahoma Friction-Free Research Network
  Visualization production on Intel Core-i7 Laptop

Supported by NWS CSTAR & HWT grants and NSF XSEDE computing resources

The UNIVERSITY of OKLAHOMA

Extreme Science and Engineering Discovery Environment

# 2017 CAPS SSEF for HWT Computing

- 1) 3DVAR SSEF: **Lonestar-5 at TACC**
  **483 Nodes**
  Cray XC40 Aries Dragonfly Interconnect
  Each Node: Intel E5 Haswell (dual 12-core processor)

- 3D Visualization demonstration (VAPOR)
  Data merge and segment on Lonestar-5 TACC on single node
  Transfer: Internet-2 and Oklahoma Friction-Free Research Network
  Visualization production on Intel Core-i7 Laptop

Supported by NWS CSTAR & HWT grants and TACC computing resources

# 2017 CAPS 3km SSEF Domain

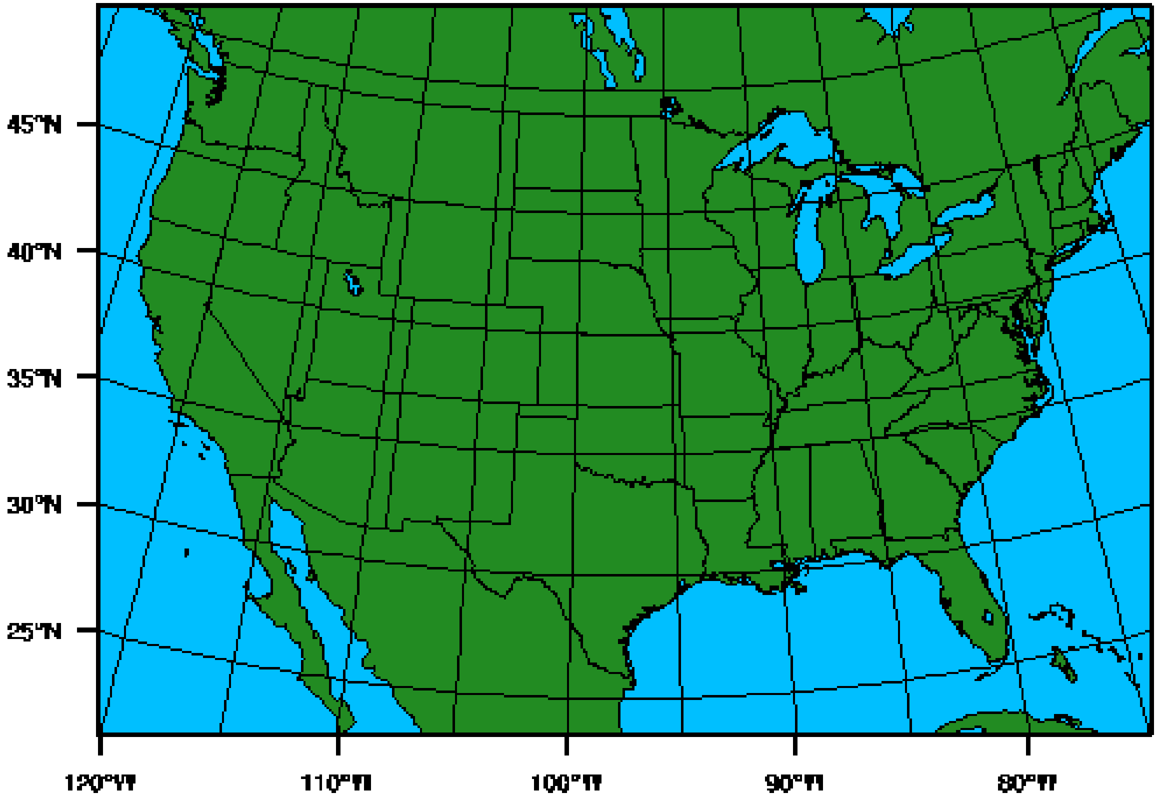


3 km WRF Grid  1620x1120

# 2016 CAPS SSEF Ensemble Members

| Member | IC | BC | Radar data | Microphy | LSM | PBL |
|---|---|---|---|---|---|---|
| arw_cn | 00Z ARPSa | 00Z NAMf | yes | Thompson | Noah | MYJ |
| arw_m3 | arw_cn + arw-p1_pert | 21Z SREF arw-p1 | yes | P3 | Noah | YSU |
| arw_m4 | arw_cn + arw-n1_pert | 21Z SREF arw-n1 | yes | MY | Noah | MYNN |
| arw_m5 | arw_cn + arw-p2_pert | 21Z SREF arw-p2 | yes | Morrison | Noah | MYJ |
| arw_m6 | arw_cn + arw-n2_pert | 21Z SREF arw-n2 | yes | P3 | Noah | YSU |
| arw_m7 | arw_cn + nmmb-p1_pert | 21Z SREF nmmb-p1 | yes | MY | Noah | MYNN |
| arw_m8 | arw_cn + nmmb-n1_pert | 21Z SREF nmmb-n1 | yes | Morrison | Noah | YSU |
| arw_m9 | arw_cn + nmmb-p2_pert | 21Z SREF nmmb-p2 | yes | P3 | Noah | MYJ |
| arw_m10 | arw_cn + nmmb-n2_pert | 21Z SREF nmmb-n2 | yes | Thompson | Noah | MYNN |
| arw_m11 | arw_cn + arw-p1_pert | 21Z SREF arw-p1 | yes | Thompson | Noah | MYJ |
| arw_m12 | arw_cn + arw-n1_pert | 21Z SREF arw-n1 | yes | Thompson | Noah | MYJ |
| arw_m13 | arw_cn + arw-p2_pert | 21Z SREF arw-p2 | yes | Thompson | Noah | MYJ |
| arw_m14 | arw_cn + arw-n2_pert | 21Z SREF arw-n2 | yes | Thompson | Noah | MYJ |
| arw_m15 | arw_cn + arw-p3_pert | 21Z SREF arw-p3 | yes | Thompson | Noah | MYJ |
| arw_m16 | arw_cn + nmmb-p1_pert | 21Z SREF nmmb-p1 | yes | Thompson | Noah | MYJ |
| arw_m17 | arw_cn + nmmb-n1_pert | 21Z SREF nmmb-n1 | yes | Thompson | Noah | MYJ |
| arw_m18 | arw_cn + nmmb-p2_pert | 21Z SREF nmmb-p2 | yes | Thompson | Noah | MYJ |
| arw_m19 | arw_cn + nmmb-n2_pert | 21Z SREF nmmb-n2 | yes | Thompson | Noah | MYJ |

ARW mixed (9)

ARW single-phys (9)

NMMB (6)

| member | IC | BC | Radar data | mp_phy | lw_phy | sw-phy | sf_phy |
|---|---|---|---|---|---|---|---|
| nmmb_cn | 00Z ARPSa | 00Z NAMf | yes | Ferrier-Aligo | GFDL | GFDL | Noah |
| nmmb_m1 | 00Z NAMa+ arw-p3_pert | 21Z SREF arw-p3 | no | Ferrier-Aligo | GFDL | GFDL | Noah |
| nmmb_m2 | 00Z NAMa+ nmmb-p1_pert | 21Z SREF nmmb-p1 | no | Ferrier-Aligo | GFDL | GFDL | Noah |
| nmmb_m3 | 00Z NAMa+ nmmb-n1_pert | 21Z SREF nmmb-n1 | no | Ferrier-Aligo | GFDL | GFDL | Noah |
| nmmb_m4 | 00Z NAMa+ nmmb-p2_pert | 21Z SREF nmmb-p2 | no | Ferrier-Aligo | GFDL | GFDL | Noah |
| nmmb_m5 | 00Z NAMa+ nmmb-n2_pert | 21Z SREF nmmb-n2 | no | Ferrier-Aligo | GFDL | GFDL | Noah |

The UNIVERSITY of OKLAHOMA

CAPS

# 20 May 2013 Moore Tornado - 20h forecast

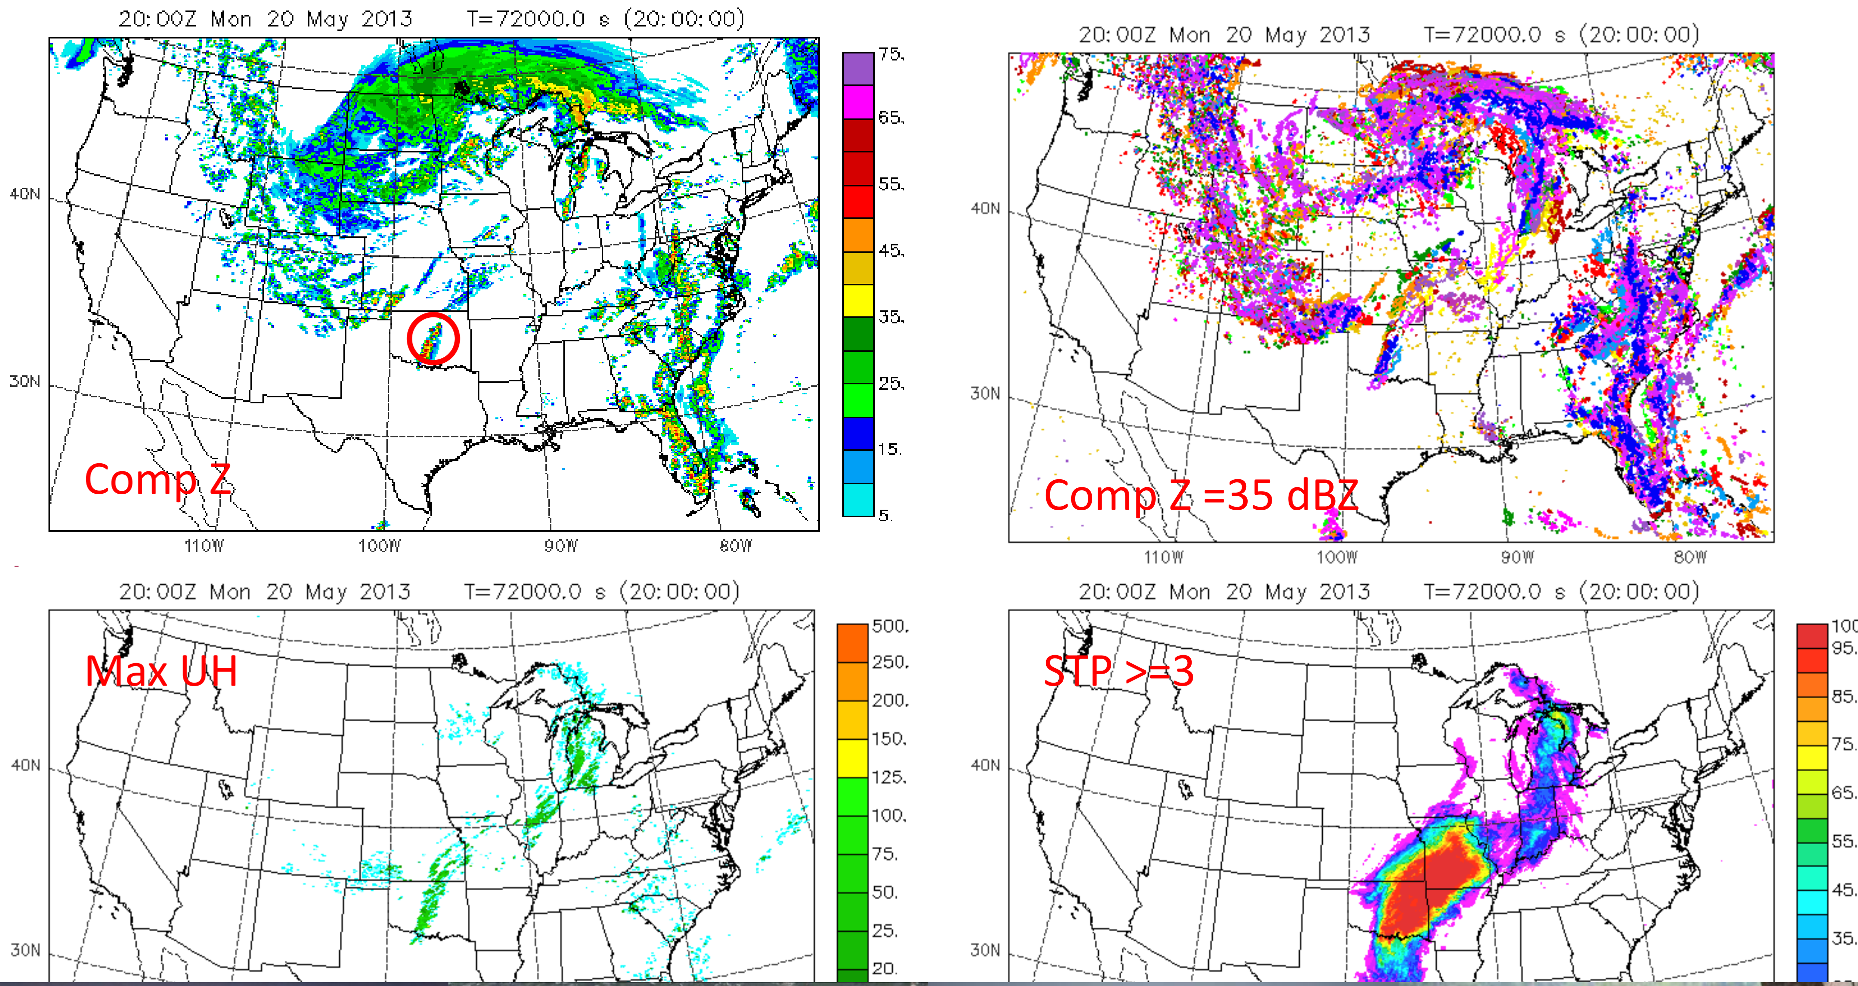


# Sample SSEF Ensemble Products

# 15-May-2017

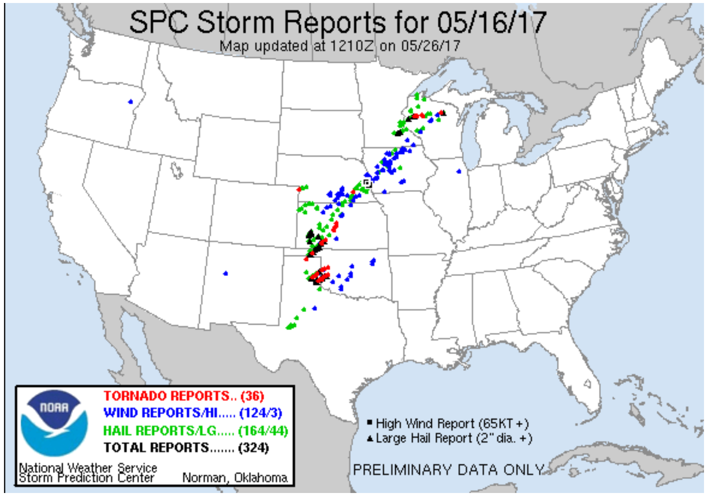


## Probability Updraft Helicity > 25

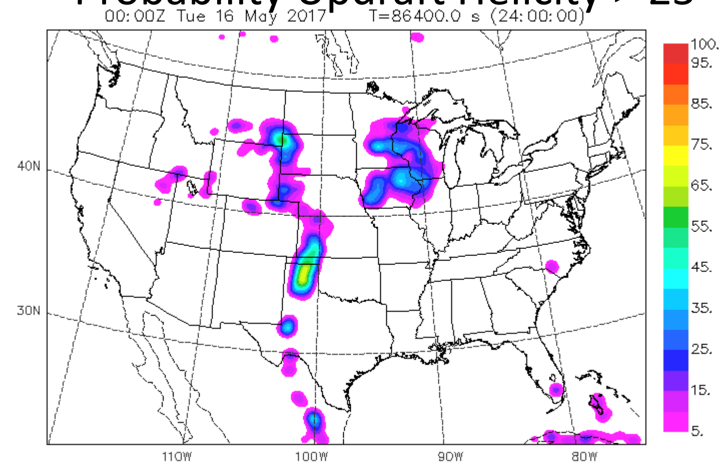


## Probability SuperCell Composite > 3

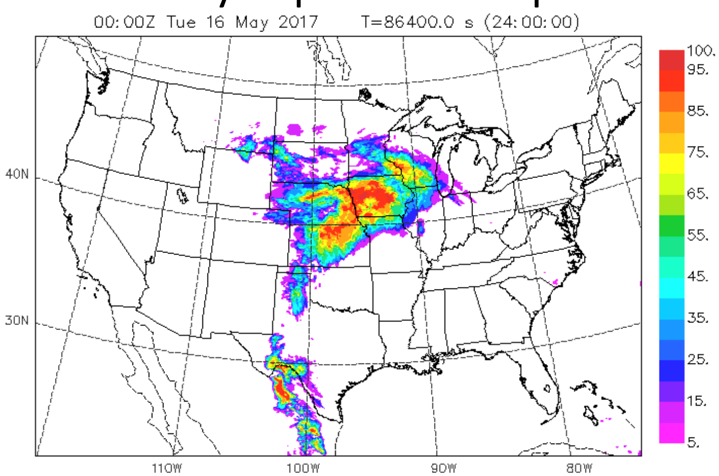


## Ensemble Max Sfc Wind Speed

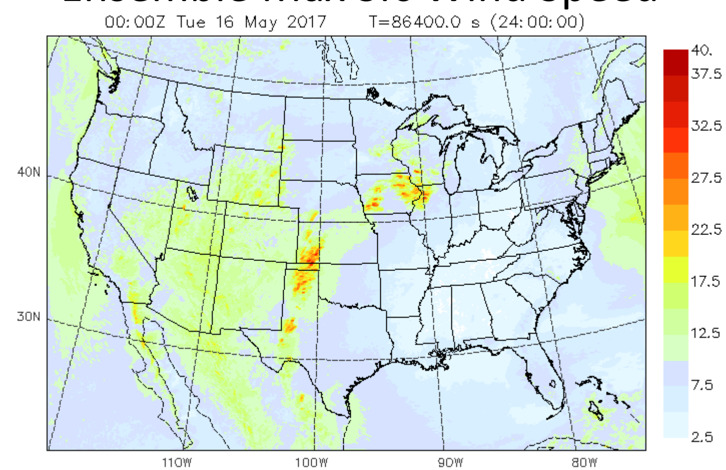

# ETS of 1-hourly QPF (mixed ARW)

>=0.01 inch

>=0.25 inch

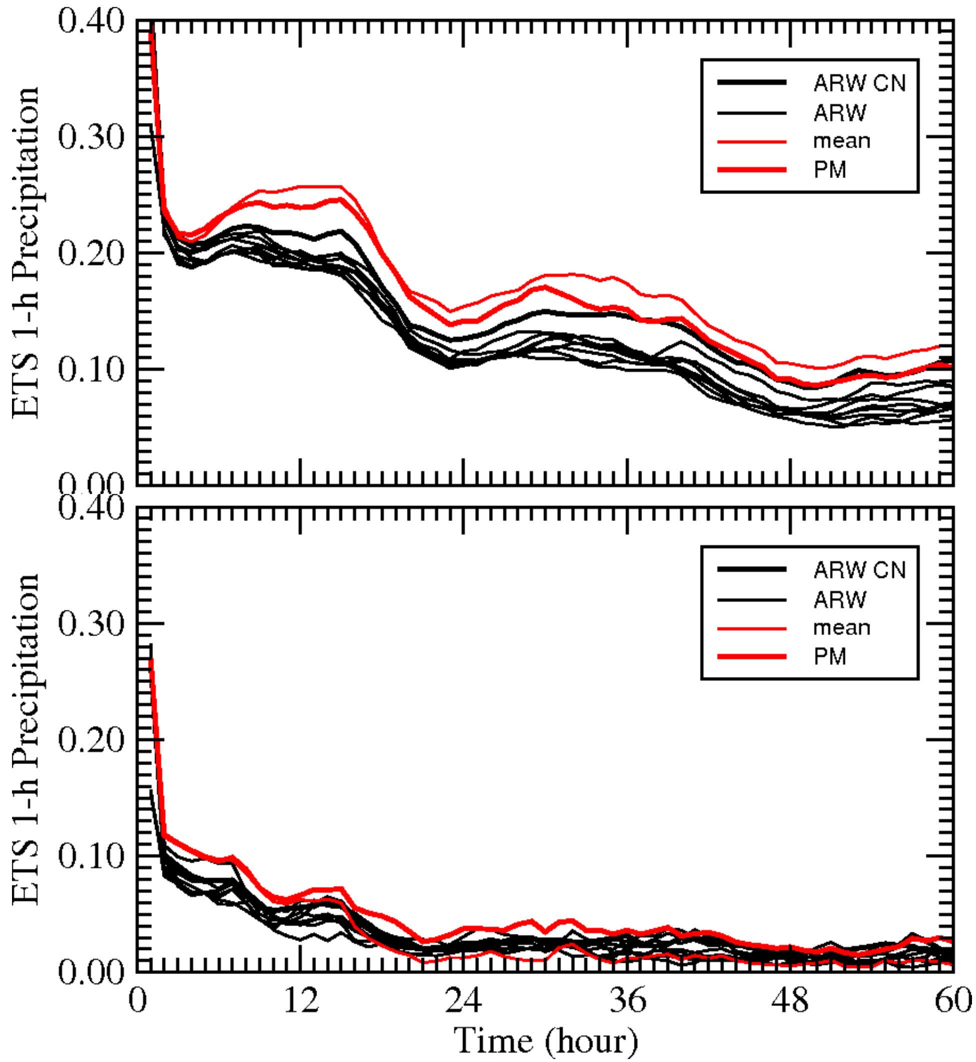


Averaged over all 2016 3DVAR HWT SSEF forecasts initialized at 0000 UTC

The UNIVERSITY of OKLAHOMA

# ETS of 3-hourly QPF (mixed ARW)

>=0.01 inch

>=0.50 inch

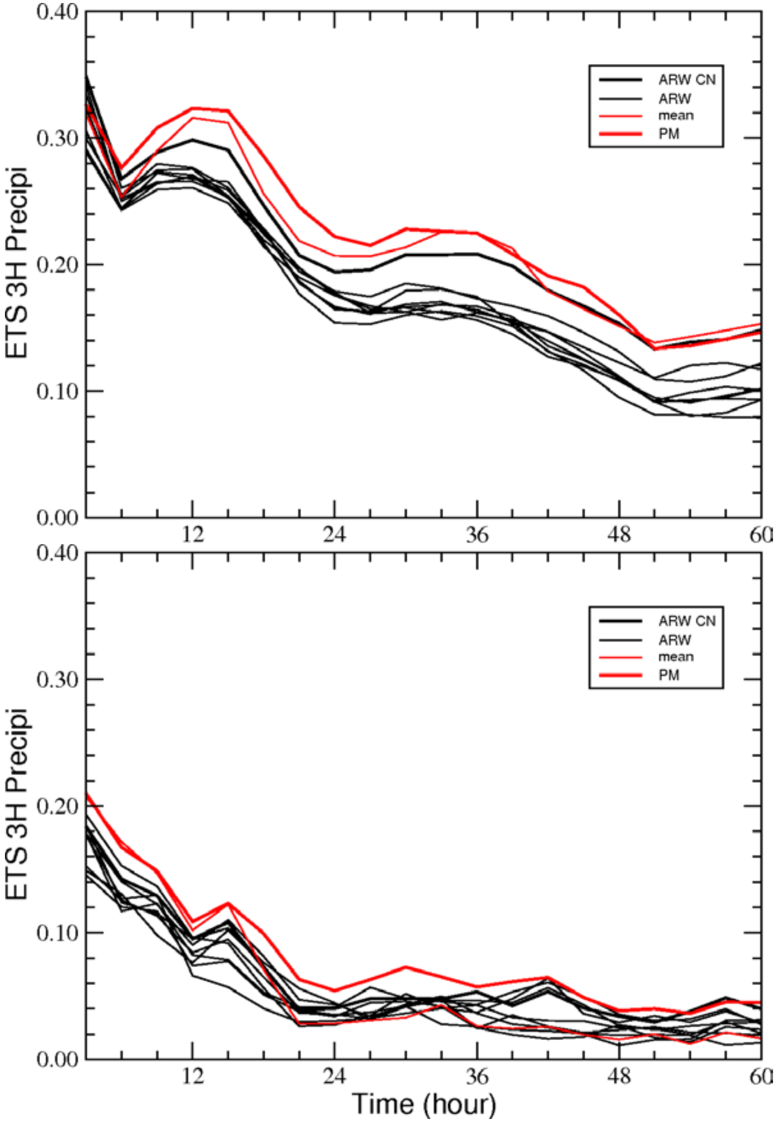


Averaged over all 2016 HWT 3DVAR SSEF forecasts initialized at 0000 UTC

The UNIVERSITY of OKLAHOMA

# ETS of 3-hourly QPF
# GSI+EnKF vs 3DVAR (12-days)

0.01”

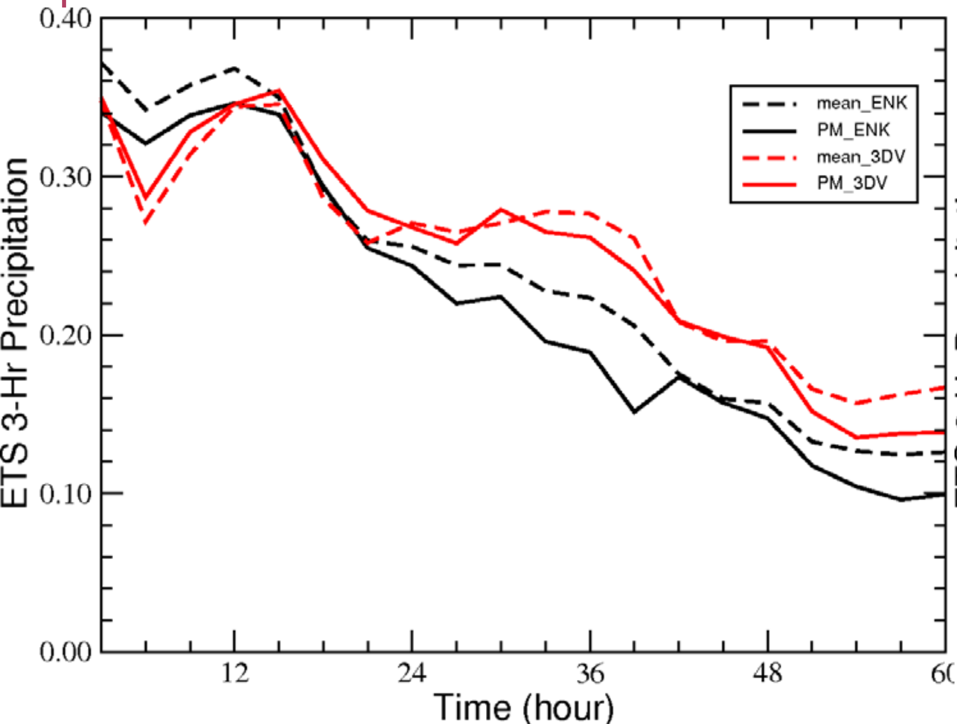


0.5”

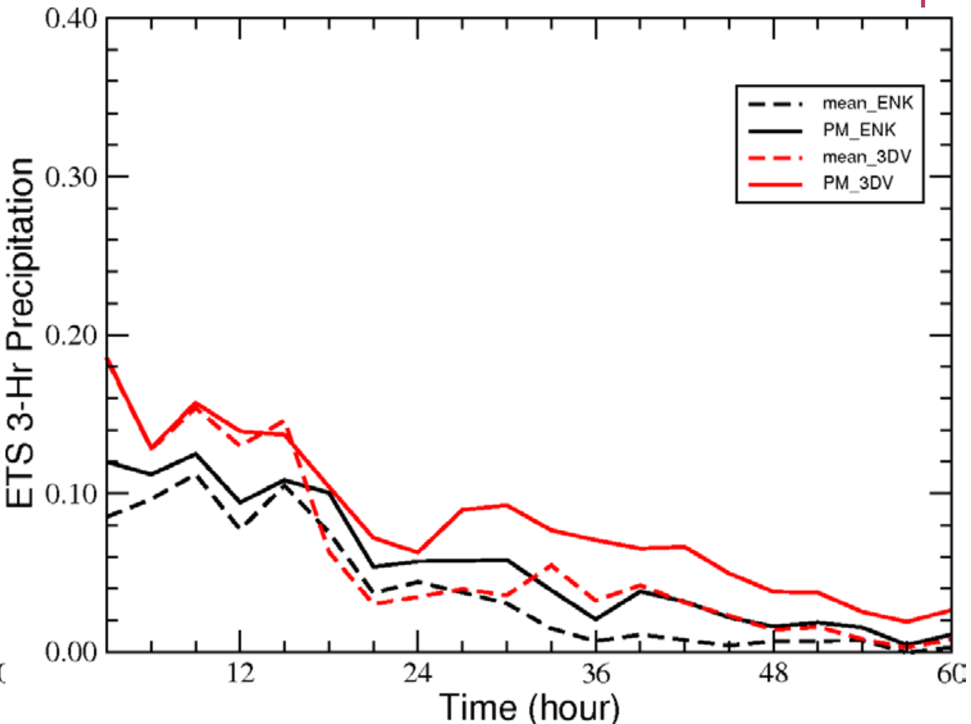

# FV3 Overview

- FV3 Finite-Volume Cubed Sphere
  https://www.gfdl.noaa.gov/fv3/
- Developed at NOAA Geophysical Fluid Dynamics Lab (Princeton, NJ) by former CAPS Post-Doc, S.J. Lin
- Chosen as the NWS "Next Generation Global Prediction System"
- Plan: Blend global-scale prediction and regional prediction
- Cubed-sphere grid (with stretching and nesting capabilities)
- No specialized data assimilation yet (initialized from T1534 GFS analysis)
- A stand-alone regional model is a work in progress (limited by human resources)

# FV3 CONUS grid configuration

- Global grid spacing: 13 km avg (9 km on CONUS face)
- Nest grid spacing: 3 km avg
- CAPS Added Thompson Microphysics Scheme

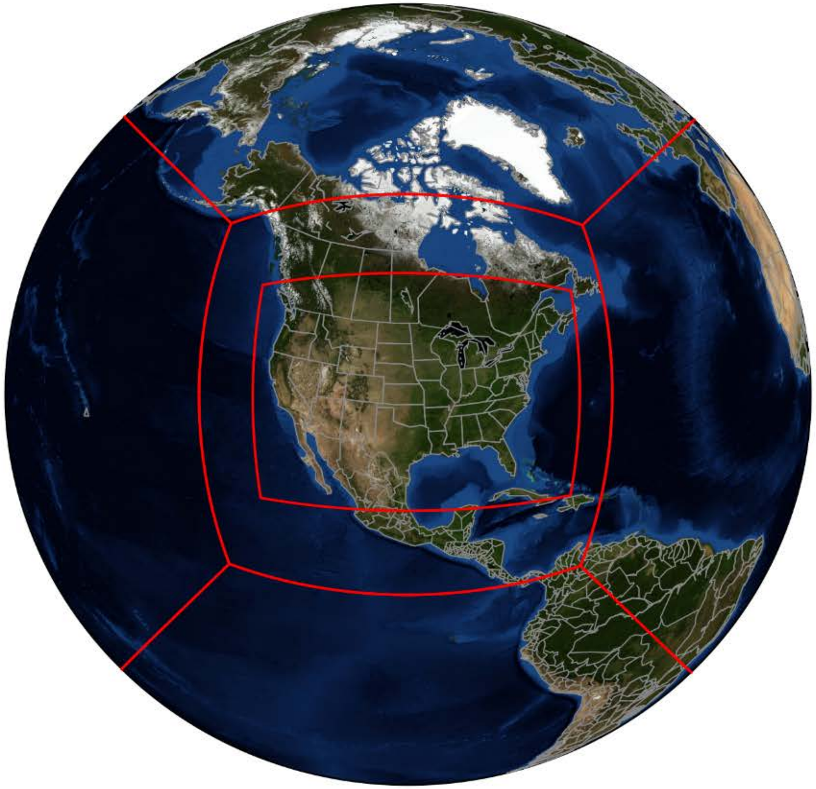

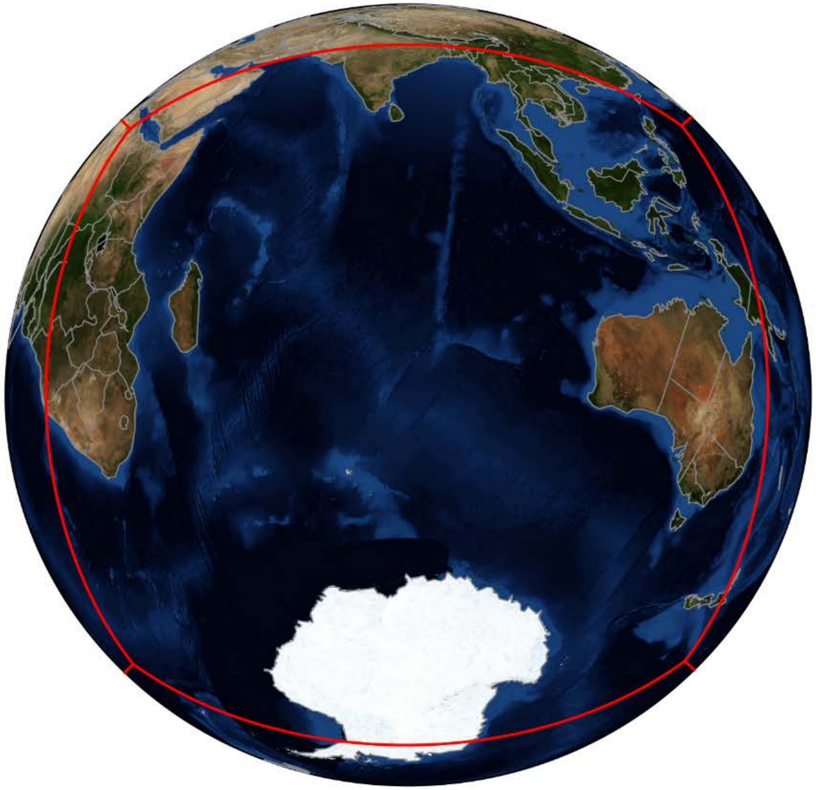

# CAPS FV3 Computing

- OSCER Schooner
  - Intel Xeon Haswell w/Infiniband Interconnect
  - Each node: Dual 10-Core processors
- 1920 Cores over 112-115 Nodes
- 9 Hours wallclock time for 5-day forecast

- First Operational Runs of FV3 Outside NOAA!
- Many thanks to OSCER Staff for making this possible!

# Preliminary Impressions

- GFDL Microphysics has a low-reflectivity bias
- Thompson Scheme seems better in that regard
- Model has trouble developing convection in weakly-forced regimes (possible consequence of the GFS PBL scheme?)
- Model sometimes evolves convection too quickly (as in the 5 April example)

# Visualization Experimentation

- VAPOR from NCAR
- 3D subsection of 4 selected members extracted
- Visualization domain is 600 km x 600 km
- 6-minute time resolution 18h-30h
- Python functions installed for variables such as Updraft Helicity
- Sample results for each day posted to web: http://www.caps.ou.edu/~kbrews/hwt_2017/

# Visualization Example
# 16-May-2016 Texas Panhandle

# CAPS Research-to-Operations

## Hydrometeorology Testbed

- Collaboration with NOAA Weather Prediction Center & others
- 2016-2017: 3-km CONUS Forecast Ensembles
- Probabilistic NWP for Flash Flood Guidance 1-60 hours

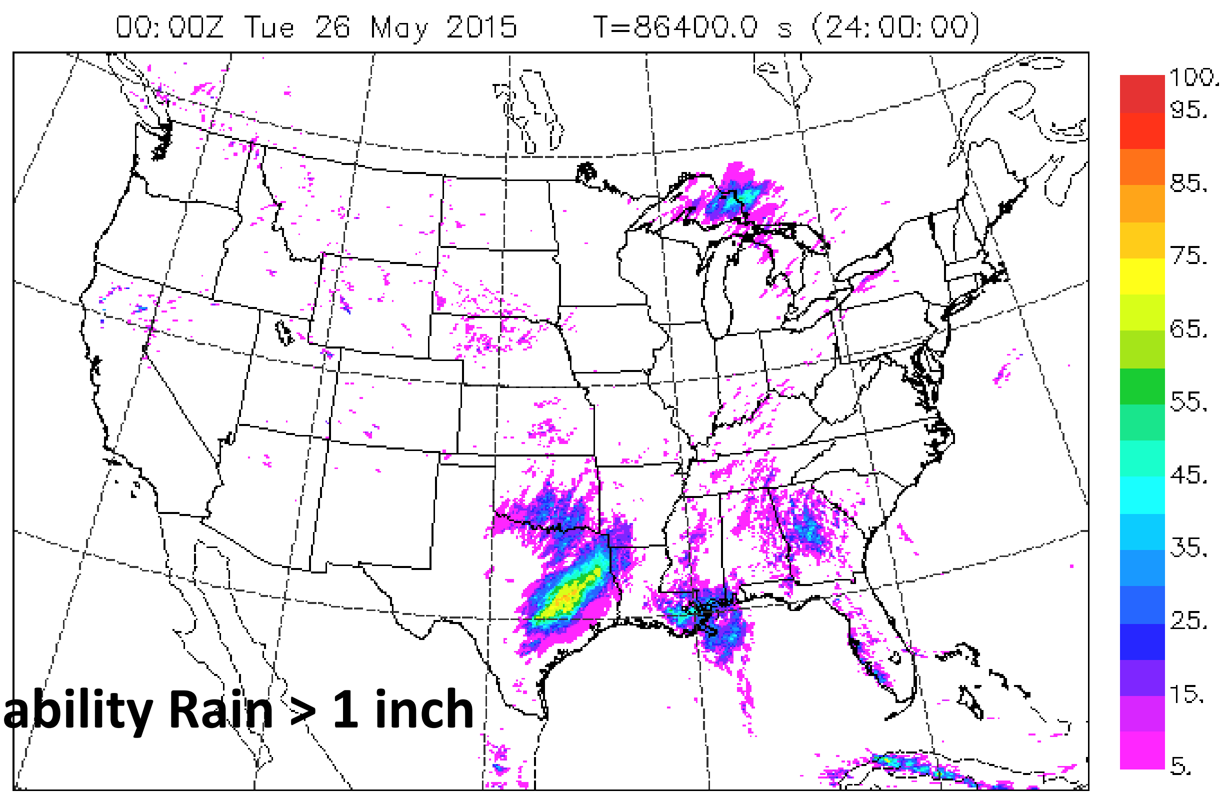


**Probability Rain > 1 inch**

# 2016 CAPS SSEF for HMT Highlights

## Storm Scale Ensemble Forecasts (SSEF)

- 20 June – 1 July ;  11 - 22 July (4 weeks)
- 3-km horizontal grid spacing (1680×1152)
- WRF version 3.7.1 (coupled with ARPS v5.3.6)
- 3DVAR SSEF, 15-member: 13 ARW members, 2 NMMB members initiated with 3DVAR analysis & Complex Cloud/Hydrometeor Analysis at 0000 UTC, with 60-h forecast

Supported by NWS CSTAR & HMT grants and NSF XSEDE computing resources

XSEDE
Extreme Science and Engineering
Discovery Environment

The UNIVERSITY of OKLAHOMA

# 2016 CAPS SSEF for HMT Computing

- First 2 weeks  on **Darter at NICS**
  **388 nodes/5488 physical cores/10976 hyperthreaded**

- Second 2 weeks on **Stampede at TACC**
  **528 nodes/7968 cores**

Supported by NWS CSTAR & HMT grants and NSF XSEDE computing resources

ARW

| Member | IC | BC | Radar data | Microphy | LSM | PBL |
|---|---|---|---|---|---|---|
| arw_cn | 00Z ARPSa | 00Z NAMf | yes | Thompson | Noah | MYJ |
| arw_m2 | arw_cn + arw-p1_pert | 21Z SREF arw-p1 | yes | Morrison | Noah | MYNN |
| arw_m3 | arw_cn + arw-n1_pert | 21Z SREF arw-n1 | yes | MY | Noah | MYNN |
| arw_m4 | arw_cn + arw-p2_pert | 21Z SREF arw-p2 | yes | Morrison | Noah | MYJ |
| arw_m5 | arw_cn + arw-n2_pert | 21Z SREF arw-n2 | yes | Thompson | Noah | YSU |
| arw_m6 | arw_cn + nmmb-p1_pert | 21Z SREF nmmb-p1 | yes | MY | Noah | MYNN |
| arw_m7 | arw_cn + nmmb-n1_pert | 21Z SREF nmmb-n1 | yes | Morrison | Noah | YSU |
| arw_m8 | arw_cn + nmmb-p2_pert | 21Z SREF nmmb-p2 | yes | Morrison | Noah | MYJ |
| arw_m9 | arw_cn + nmmb-n2_pert | 21Z SREF nmmb-n2 | yes | Thompson | Noah | MYNN |
| arw_m10 | 00Z ARPSa | 00Z NAMf | yes | P3 | Noah | MYJ |
| arw_m11 | 00Z ARPSa | 00Z NAMf | yes | Morrison | Noah | MYJ |
| arw_m12 | 00Z ARPSa | 00Z NAMf | yes | MY | Noah | MYJ |
| arw_m13 | arw_cn + arw-n2_pert | 21Z SREF arw-n2 | yes | Thompson | Noah | MYJ |

NMMB

| member | IC | BC | Radar data | mp_phy | lw_phy | sw-phy | sf_phy |
|---|---|---|---|---|---|---|---|
| nmmb_cn | 00Z ARPSa | 00Z NAMf | yes | Ferrier-Aligo | GFDL | GFDL | Noah |
| nmmb_m1 | 00Z NAMa+ arw-p3_pert | 21Z SREF arw-p3 | no | Ferrier-Aligo | GFDL | GFDL | Noah |

# West Virginia Flash Floods

- 23-24 June 2016
- Max Gauge 9.37 inches (238 mm) at Maxwelton, WV
- Elk River all time high 33.37 ft
- 23 Fatalities
  15 in Greenbrier Co.
- 44 of 55 WV counties
  placed in state of emergency


ABC News Image

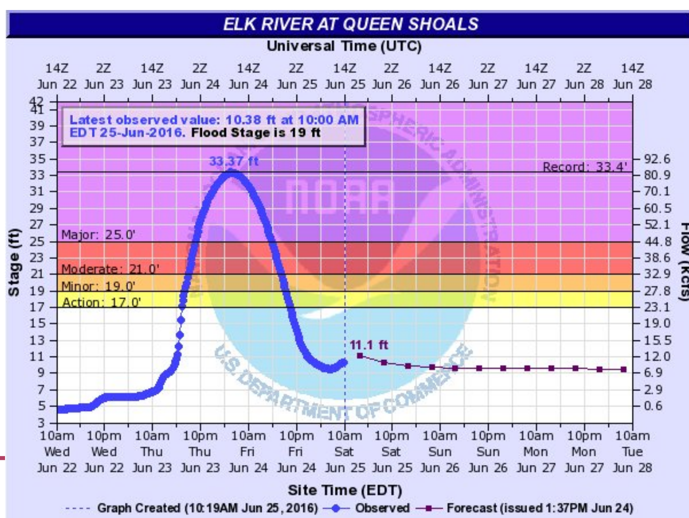


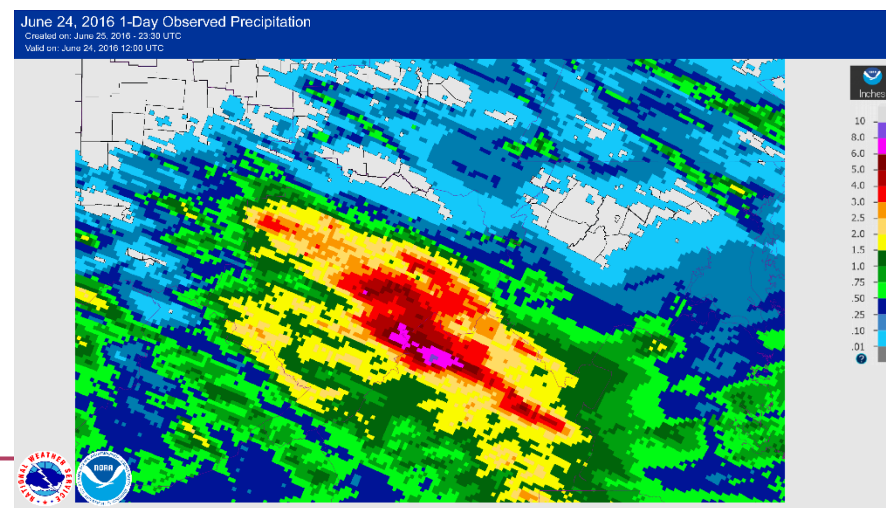

# 24-h Precip  00Z June 23 - 00Z June 24

MRMS
QPE

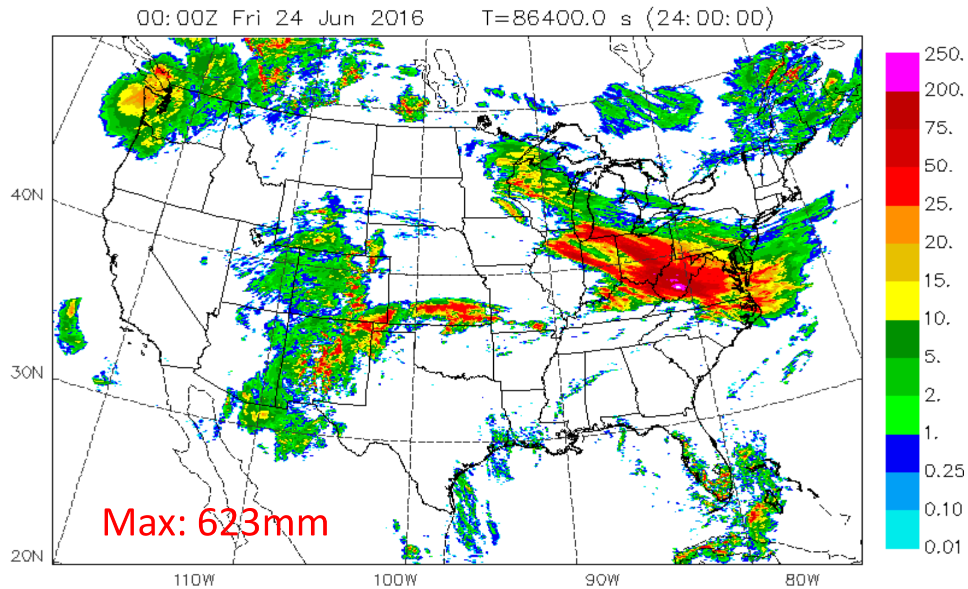


48 h fcst

24 h fcst

QPF-
PM

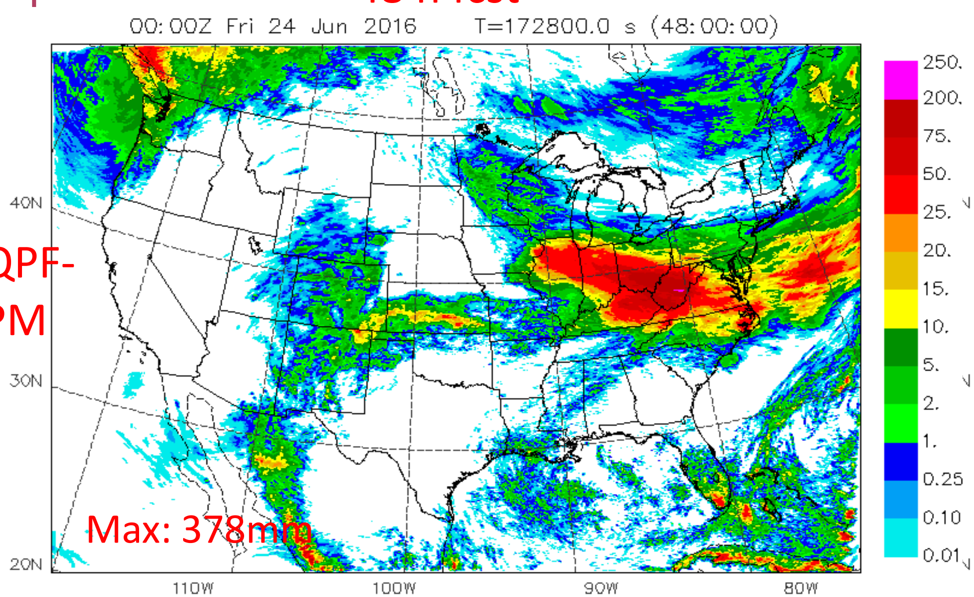


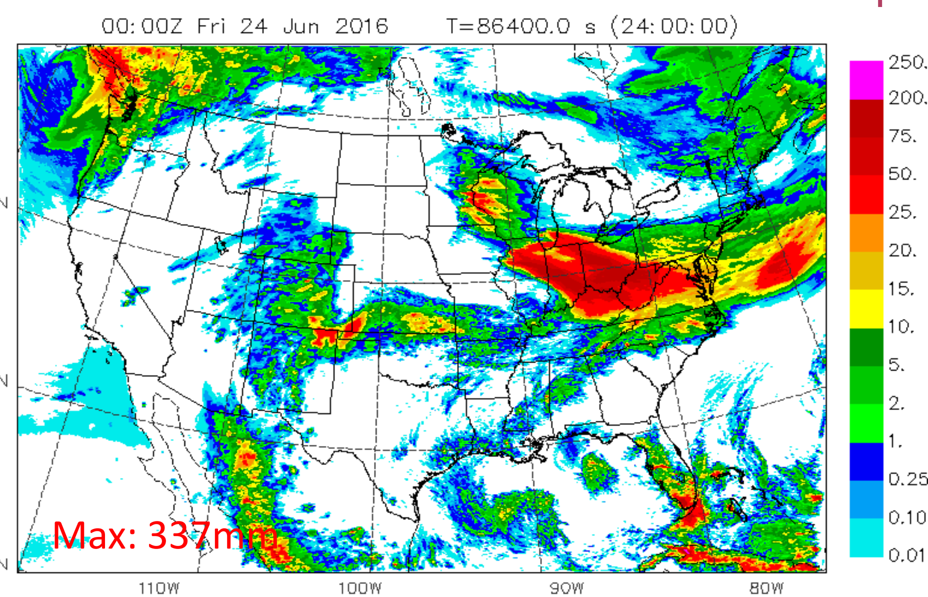

# 24-h Precip  00Z June 23 - 00Z June 24

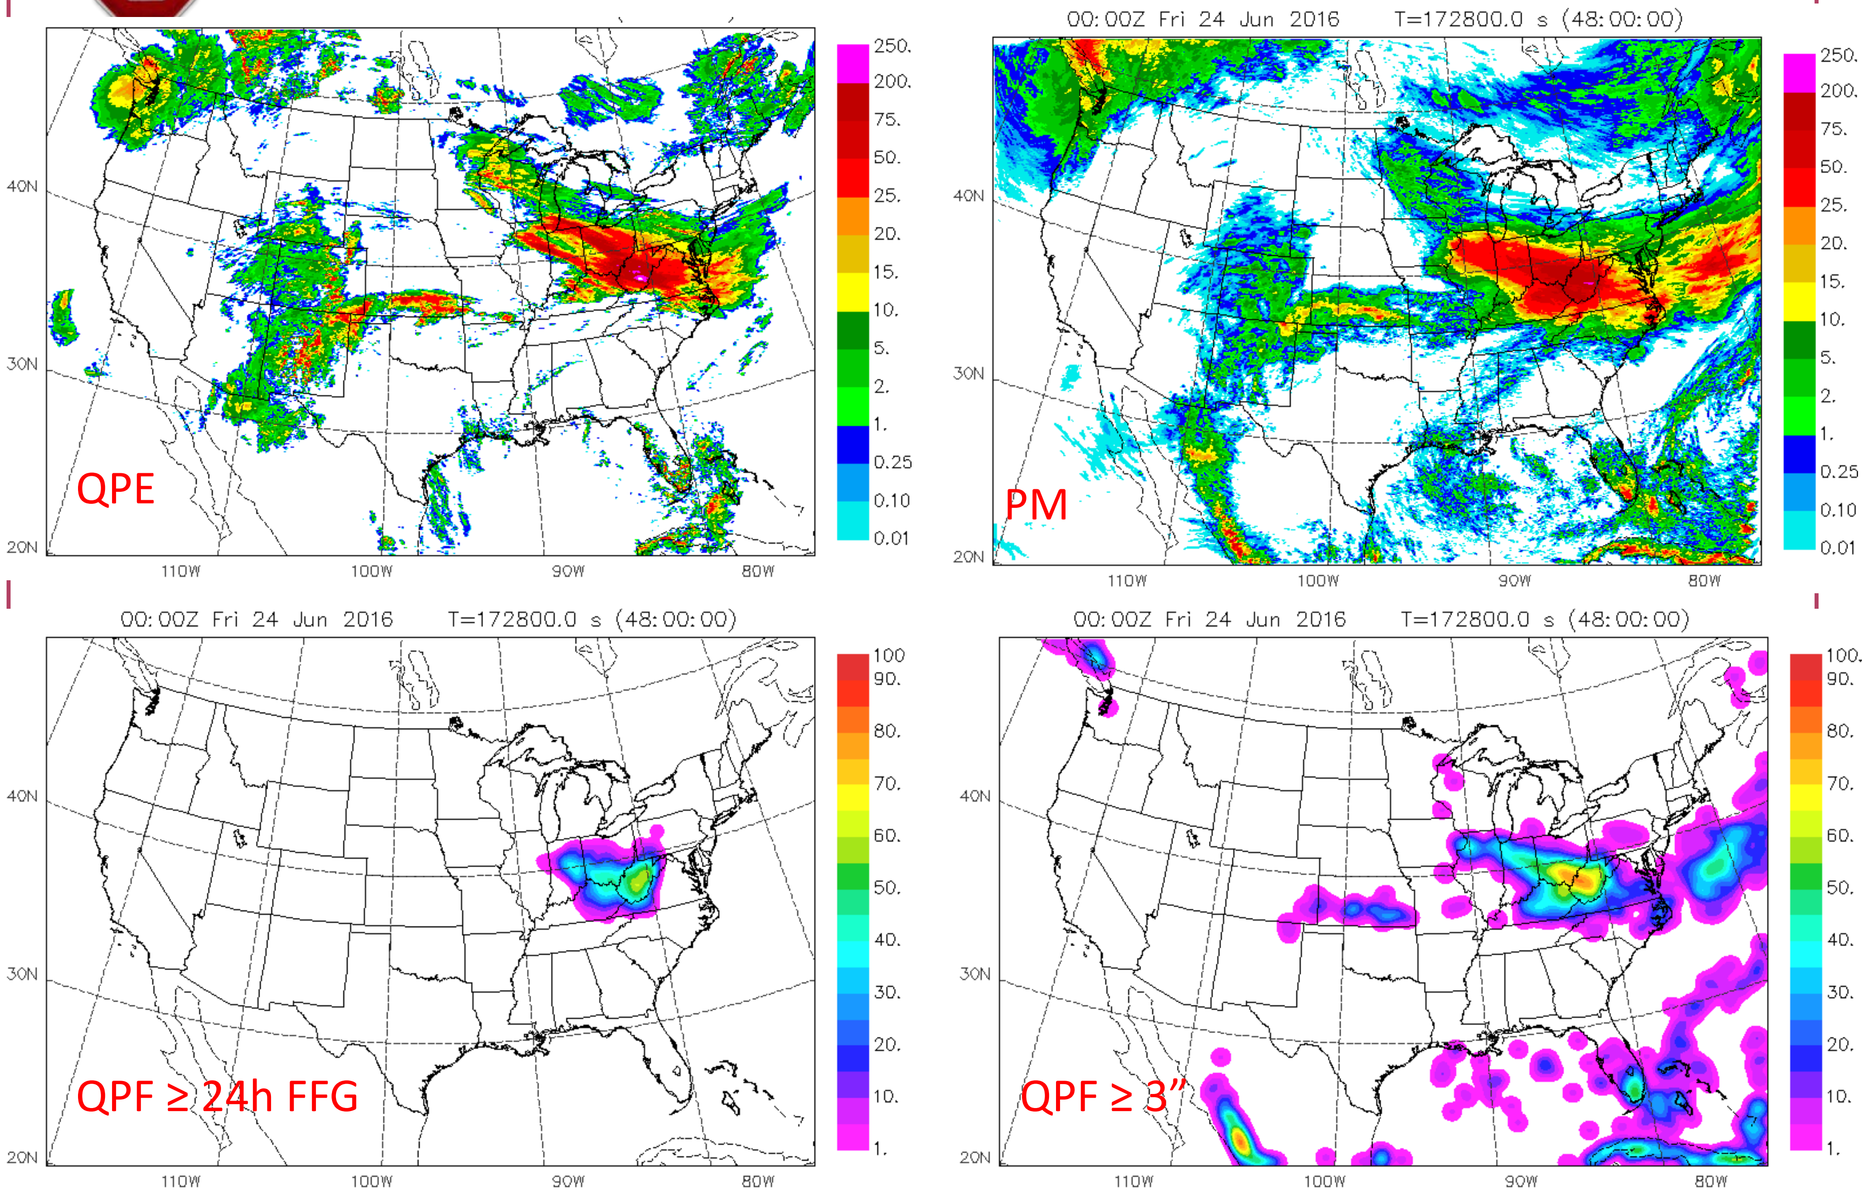

# 12-h Precip 12Z June 23 - 00Z June 24

MRMS QPE

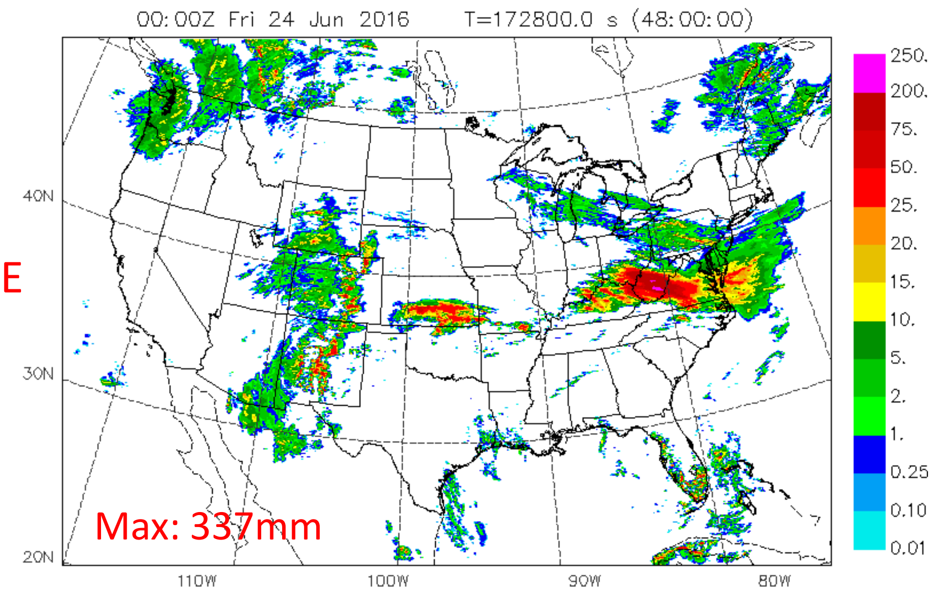


QPF-PM

48 h fcst

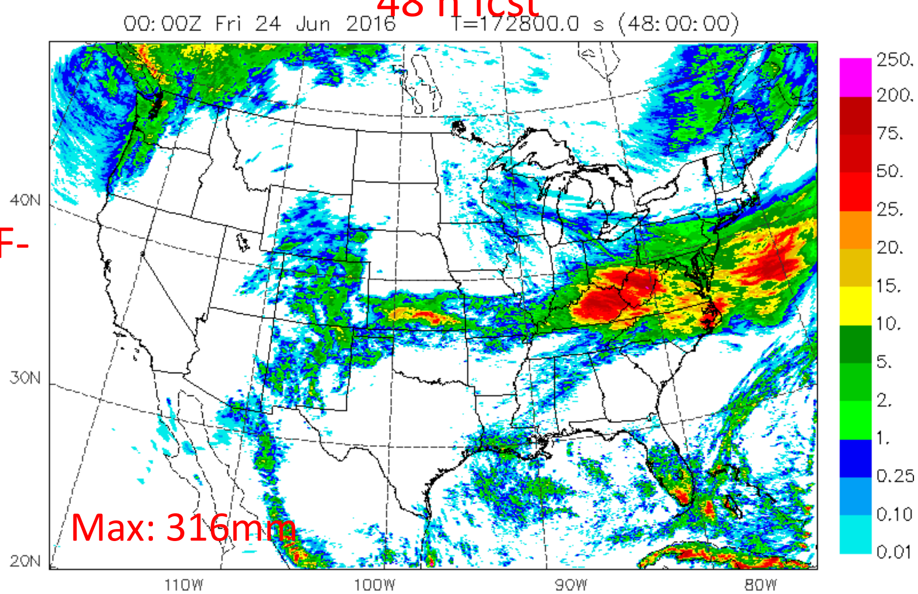


24 h fcst

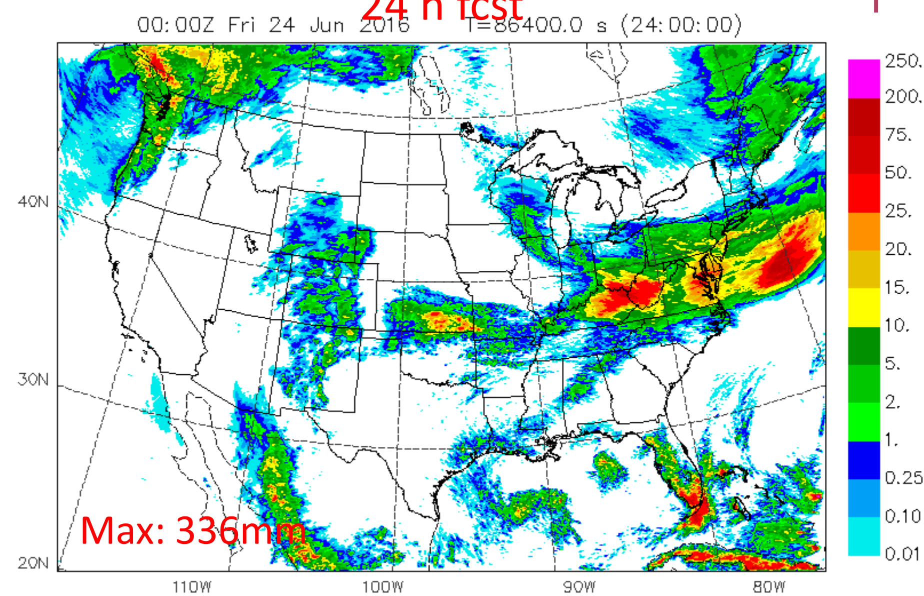

# 12-h Precip 12Z June 23 - 00Z June 24

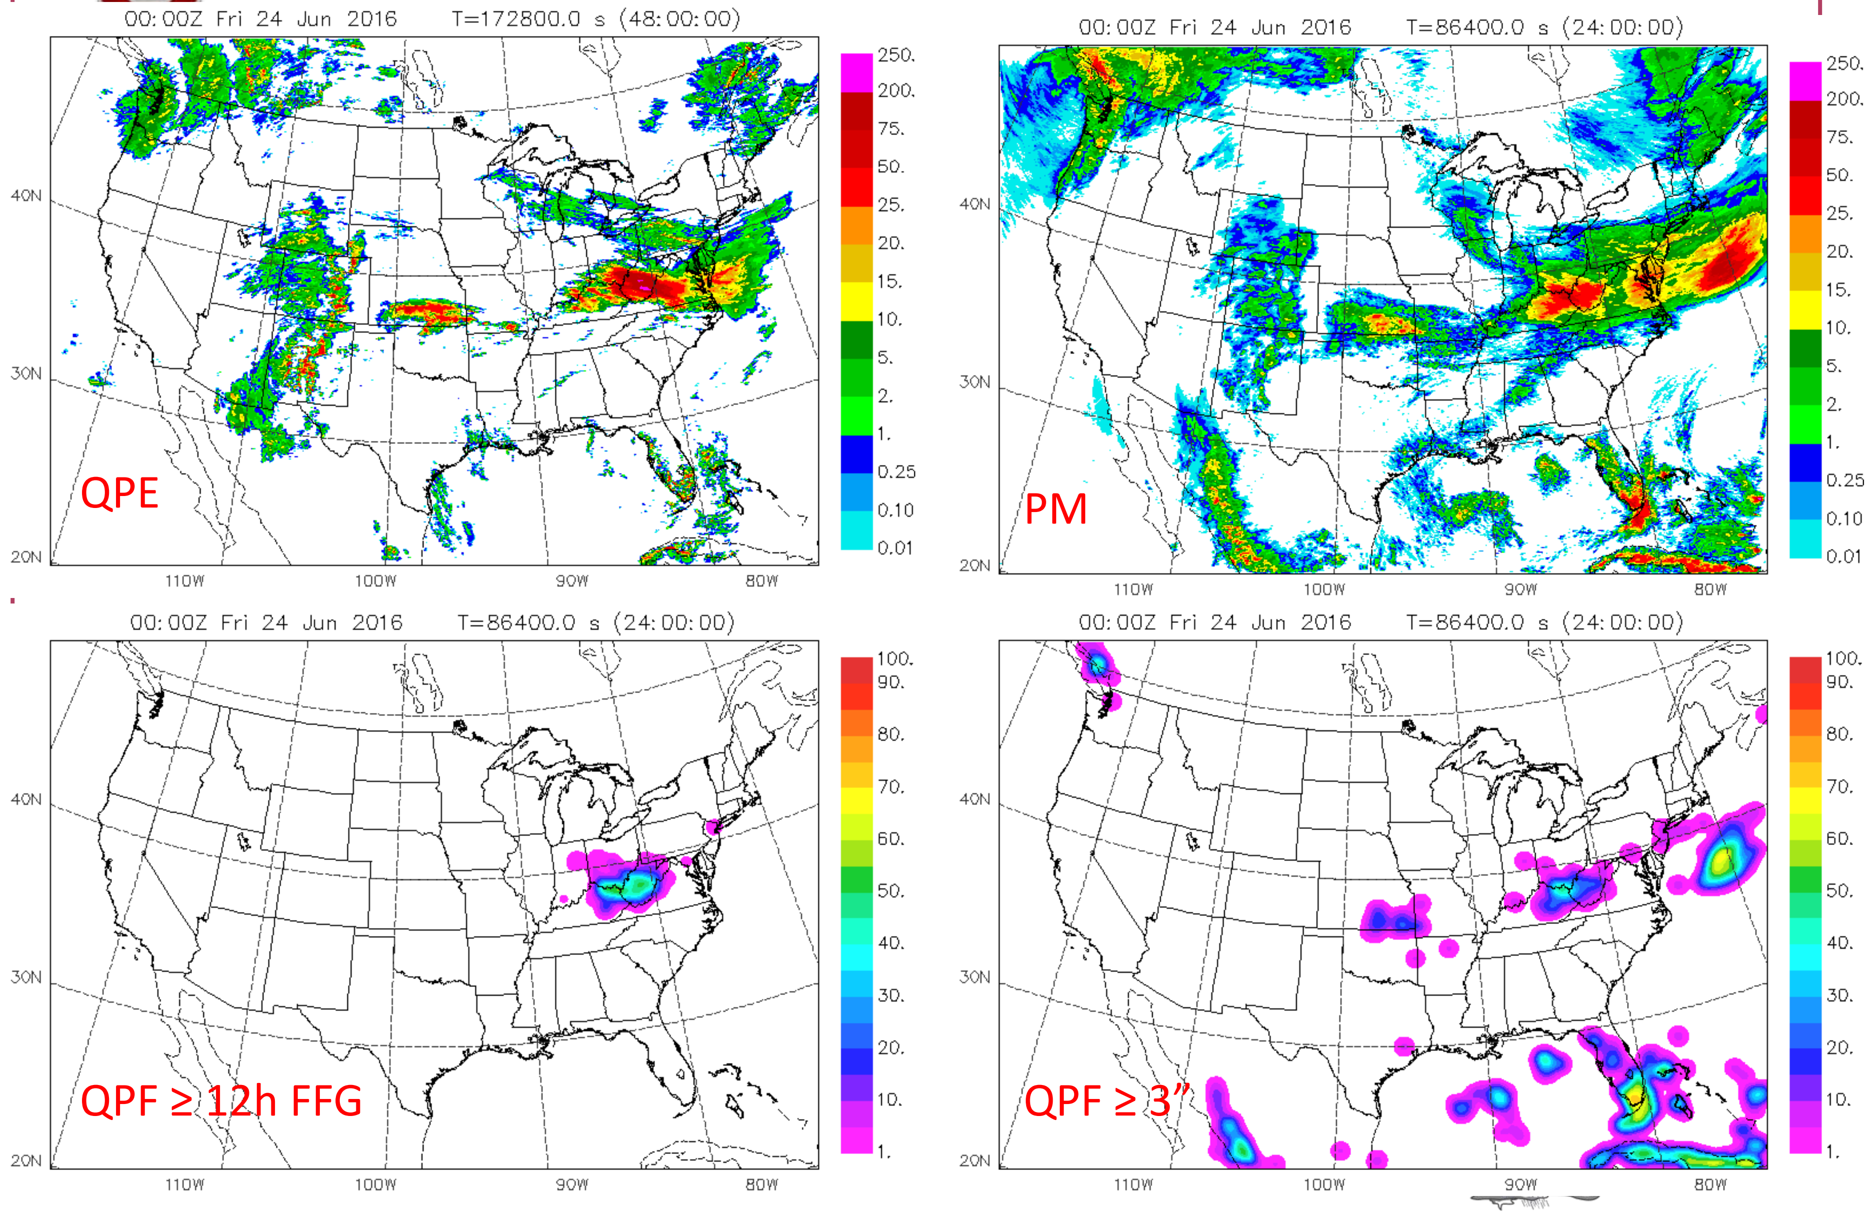

# Observing Systems Studies

## Collaborative-Adaptive Sensing of the Atmosphere

- Network of small X-band radars
- Low to the ground
- Adaptive scanning
- Low cost

Meant to ***complement*** operational radar networks.

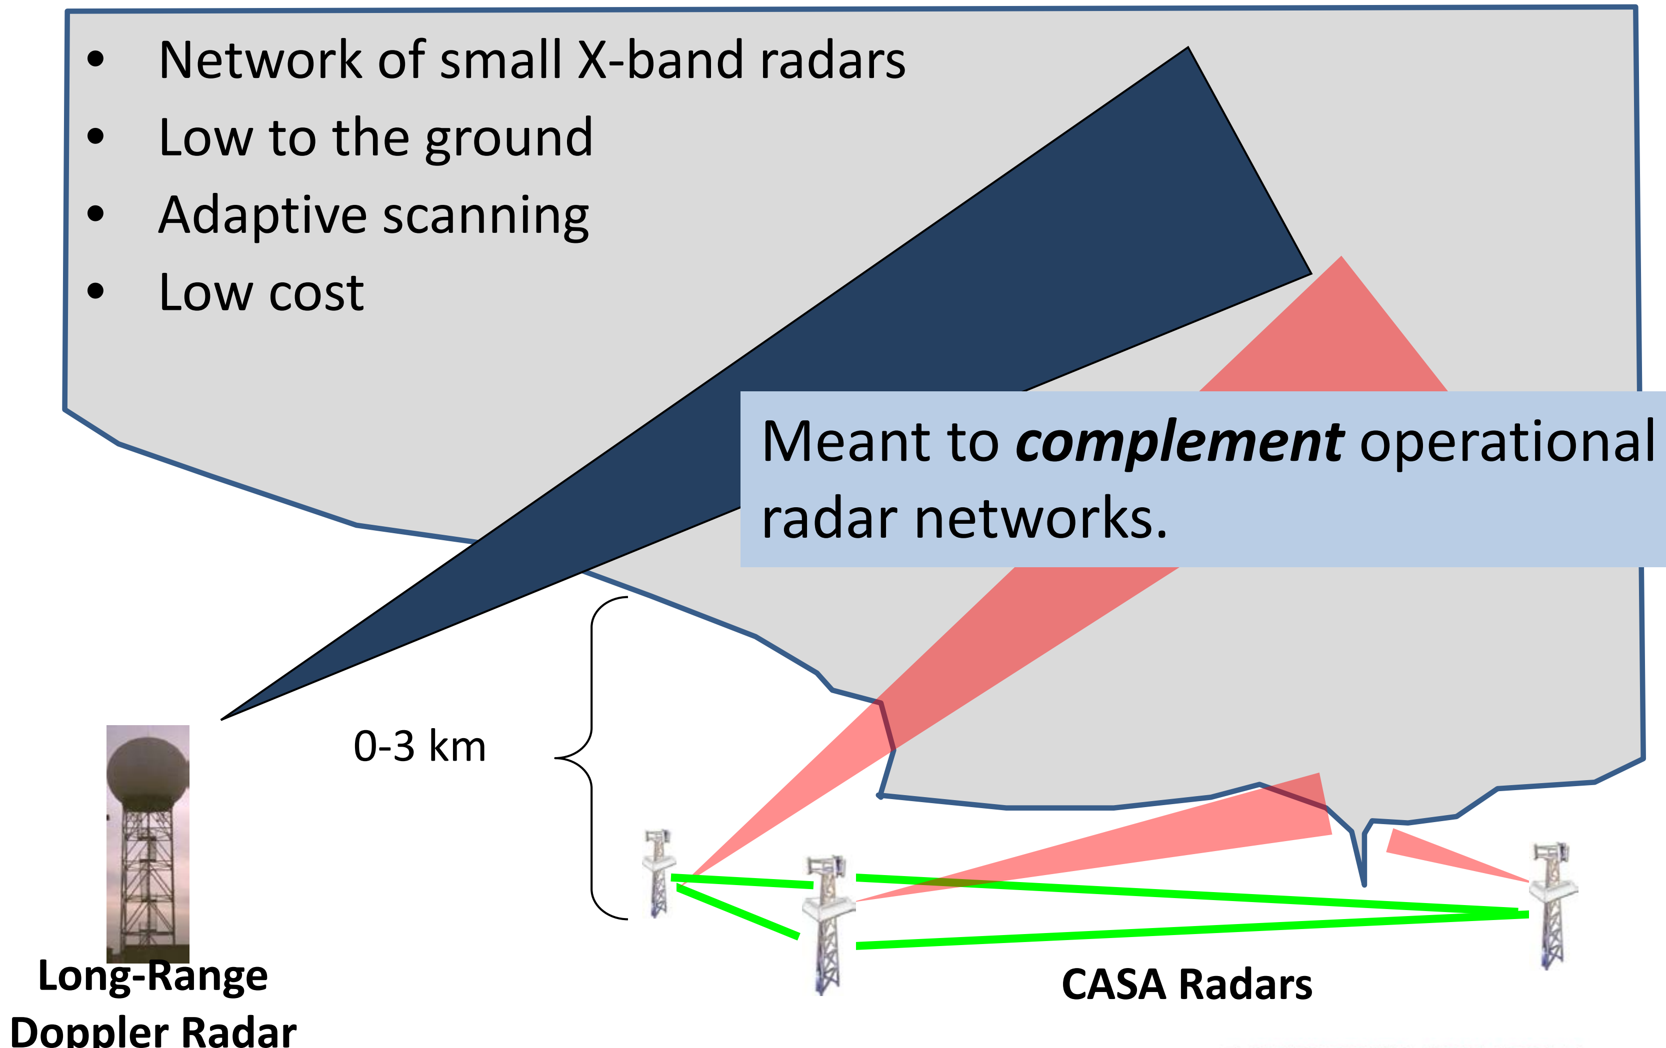

# CASA 2.0: Dallas/FW Urban Testbed

- Network of Networks approach
- Primary foci: Flooding, Severe Storms and Tornadoes

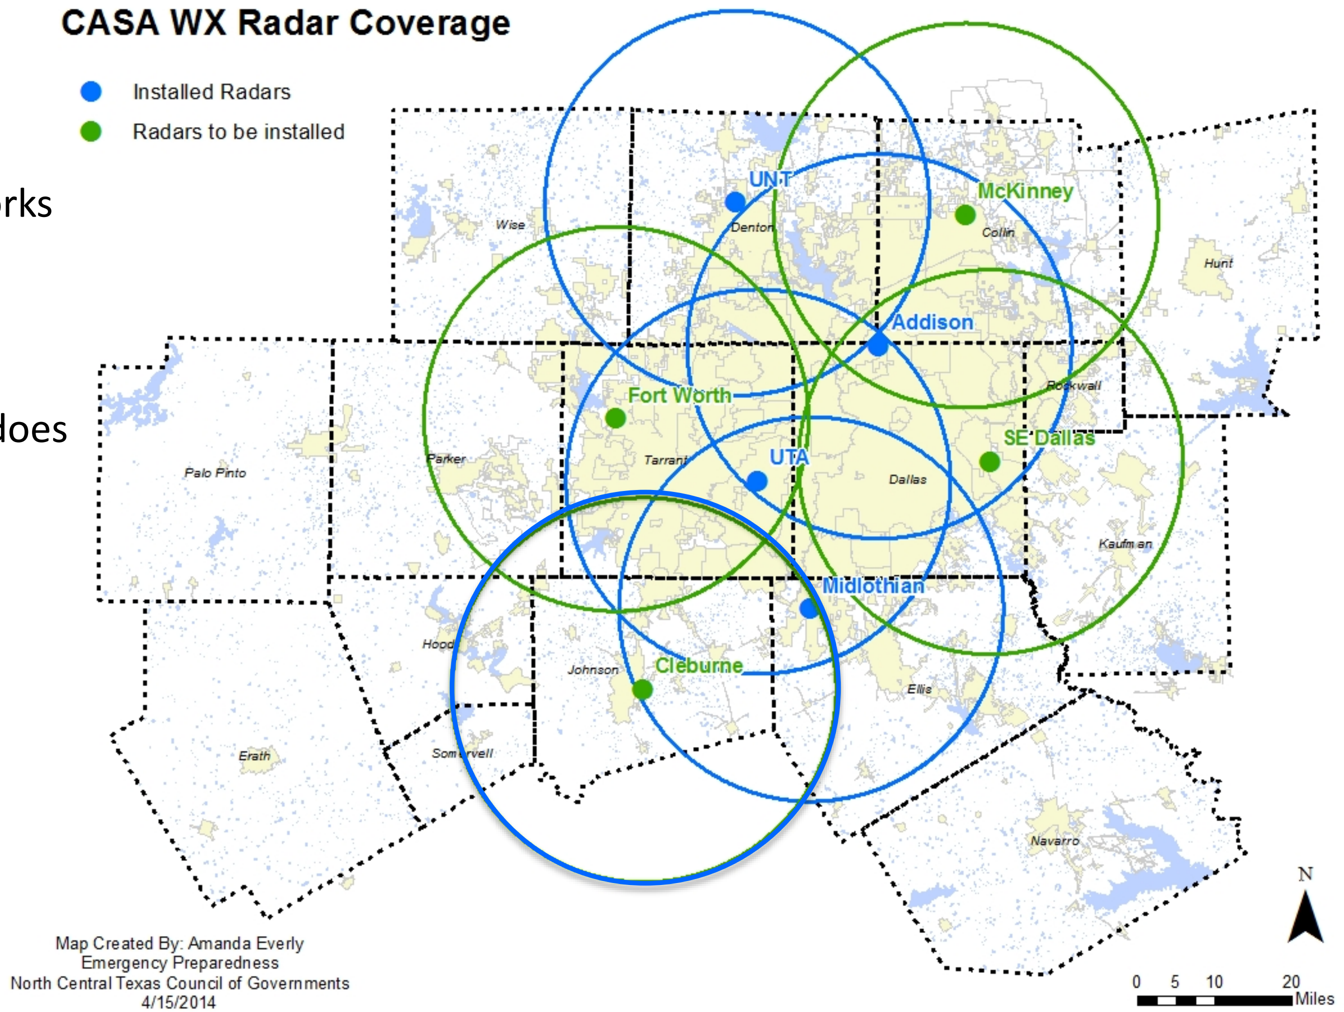

# CASA 2.0 Dallas/FW Urban Testbed

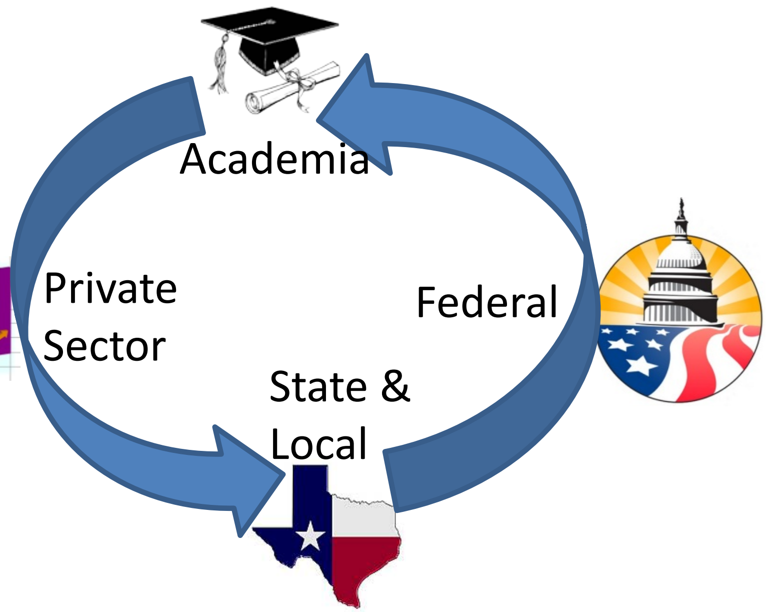

# Observation Summary

| Conventional Observations | Non-Conventional Observations |
| --- | --- |
| ASOS | EarthNetworks (WxBug) |
| AWOS | CWOP |
| | GST MoPED |
| | Oklahoma & W Texas Mesonets |
| S-band WSR-88D Radars | X-band Radars |
| | C-band TDWR Radars |
| Radiosondes | SODARs |
| | Radiometers |

# Surface Observations

Global Science & Technology, Inc.

al & Non-conventional Observations

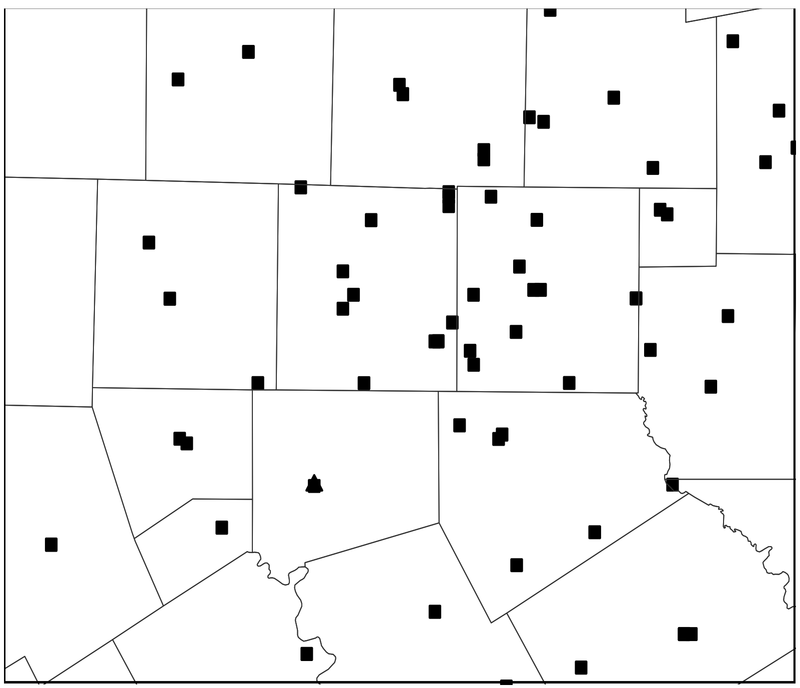

AWOS & ASOS

WxBUG & CWOP

# Short-Term Weather Forecasting Continuum

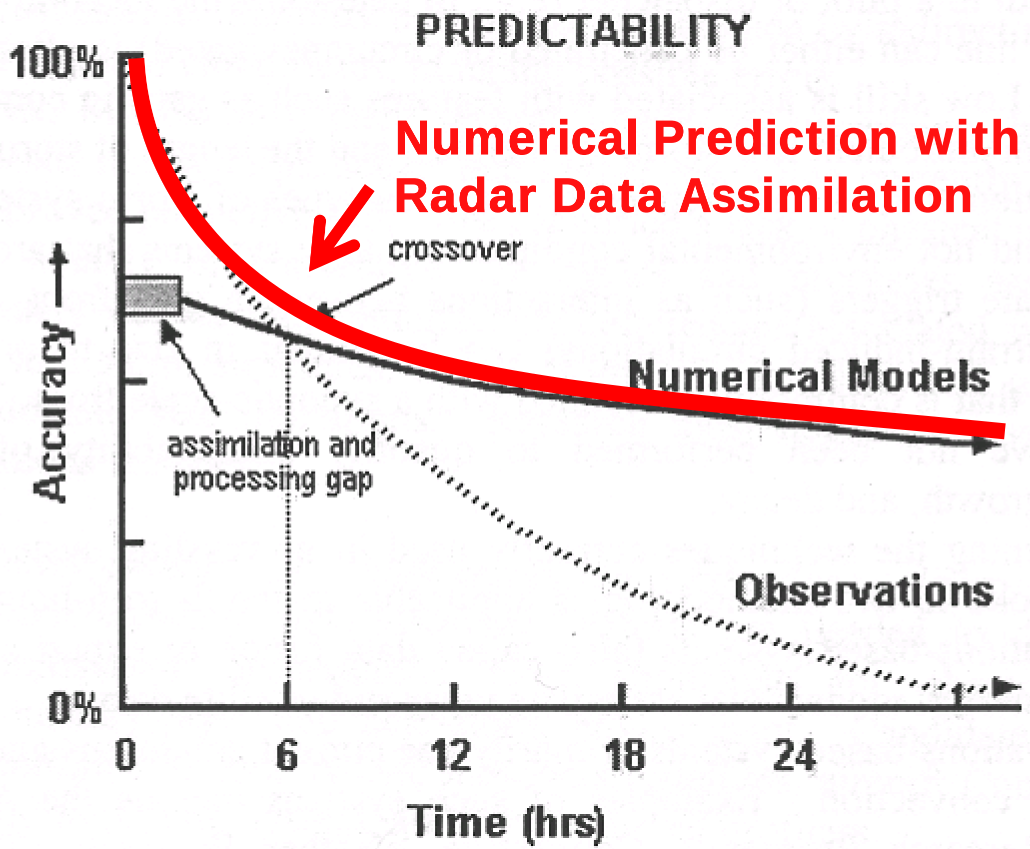


When deterministic predictability becomes low, ensembles are valuable.

# December, 2015
# Operational Configuration

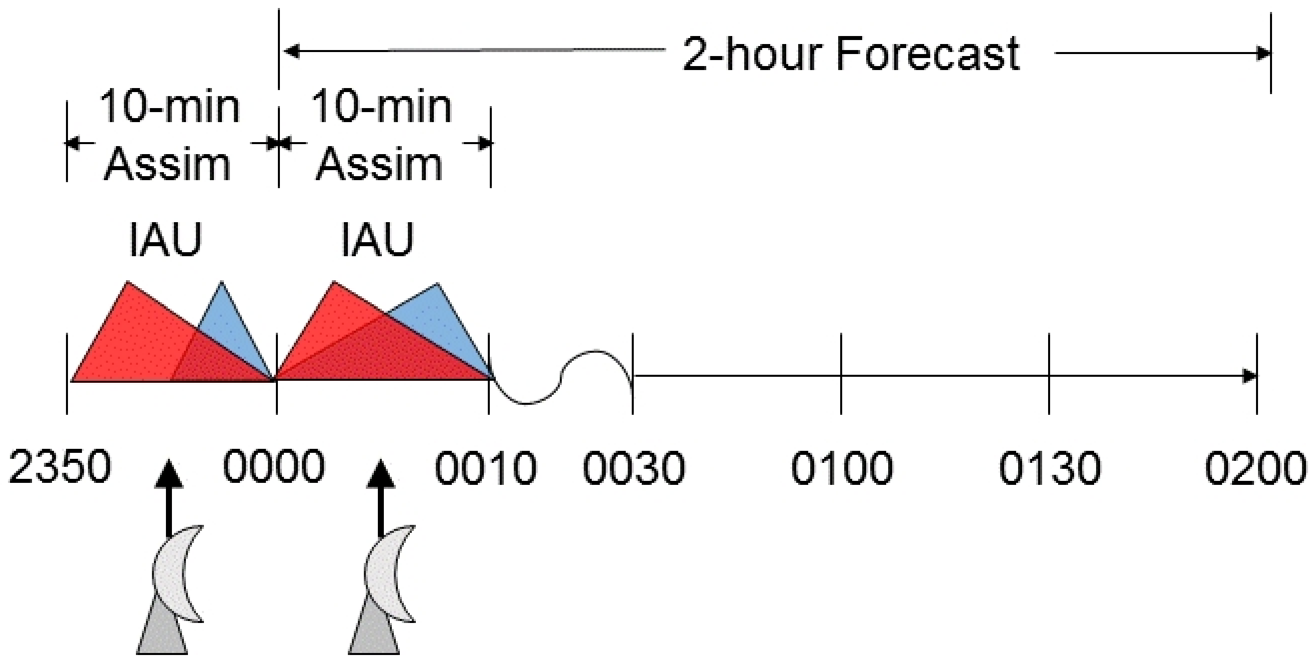

# Dec 2015 Operational Computing

OSCER Boomer

Xeon64 Oct-core SandyBridge 2.0 GHz

- Analyses at 400 m Resolution Dedicated Queue
  - 3DVAR and Cloud Analysis
  - Sfc, Profilers, VAD, Radar Wind and Reflectivity
  - 5-minute Interval
  - 400-m grid spacing Grid Size 448 x 456 x 28
  - Processors: 8 x 24 = 192
  - Obs Processing & Analysis Wallclock ~8 min
- Assimilation/Forecasts On-Demand
  - 3DVAR and ARPS with 10-min IAU
  - Sfc, Profilers, VAD, Radar Wind and Reflectivity Assimilation
  - 2-hour Forward Forecast
  - 15 minute interval
  - 1-km grid spacing Grid Size 363 x 323 x 53
  - Processors: 12 x 16 = 192
  - Obs Processing + Analysis + Forecast Wallclock ~20-25 min

# D/FW Metro Tornado Tracks, 26 Dec 2015

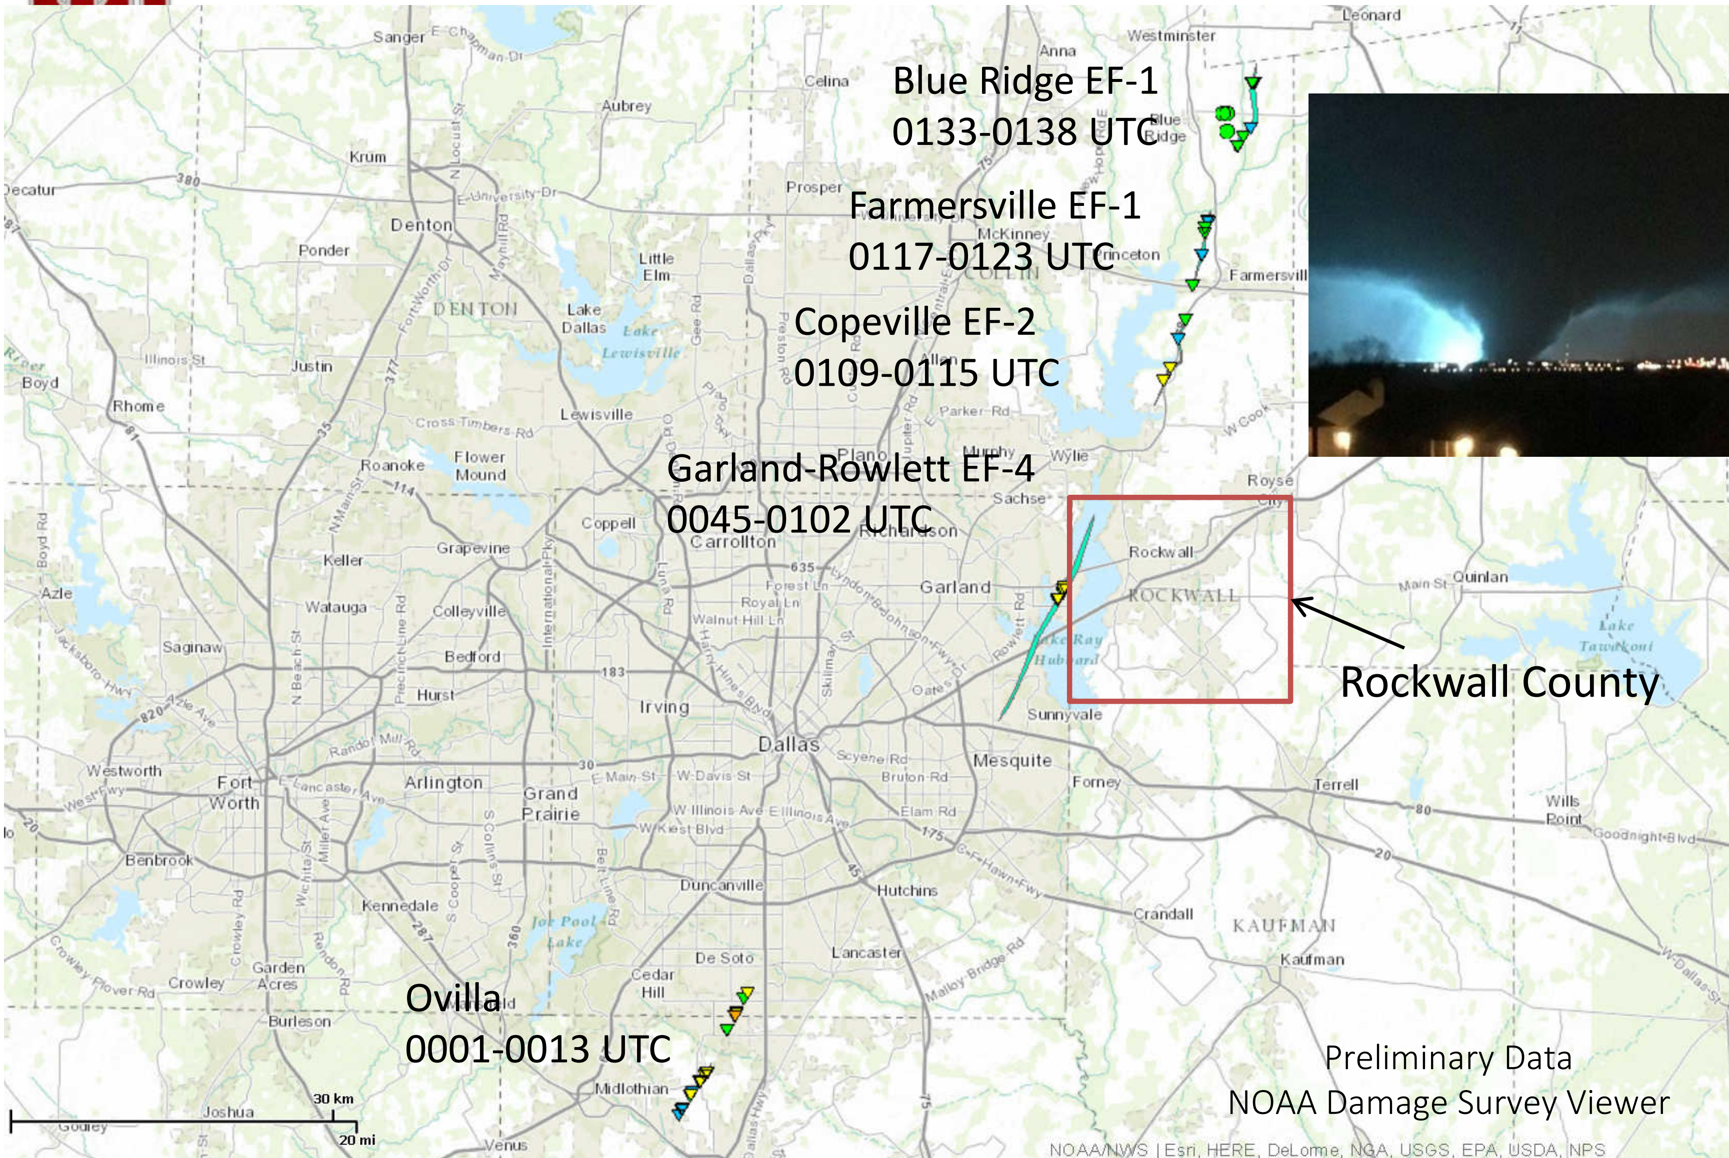

# 400-m Δx Analysis

400-m Grid Scale Analyses

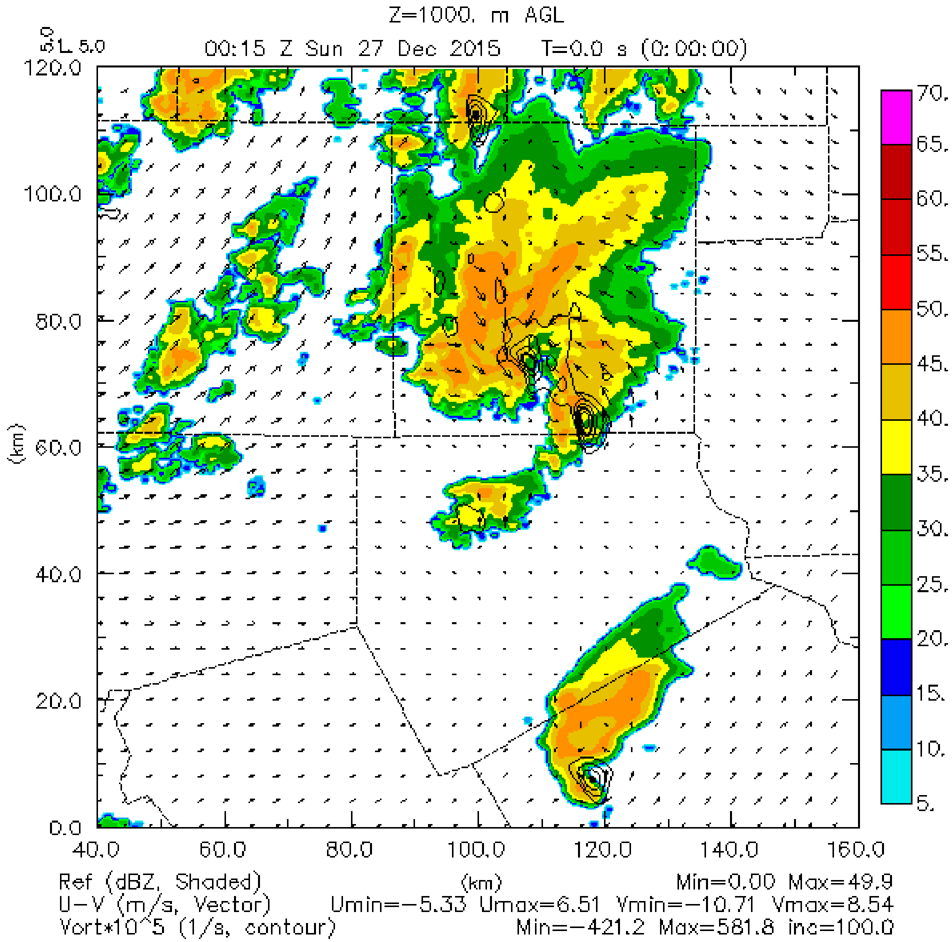


Research: Running OSE's to determine data impacts

The UNIVERSITY of OKLAHOMA

# 1-km Δx Forecast

1-km Grid Resolution 2-hour Forecast

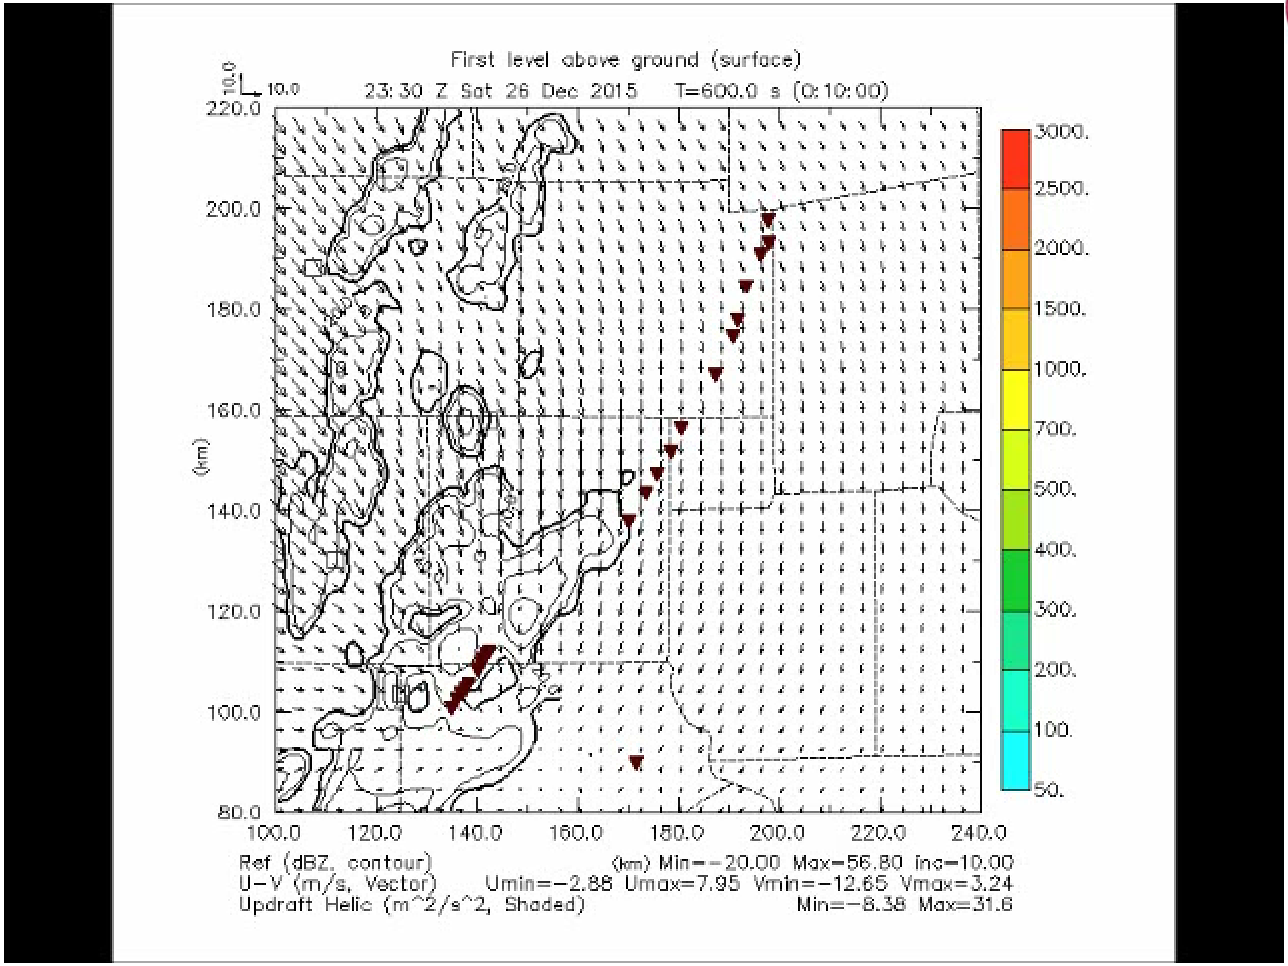

# CASA 2.0 Summary/Future Plans

- Realtime Data Analyses and Forecasts
  Operating with Low Latency
- Updated Hydrometeor Retrieval for
  Multiple Microphysics Schemes
- 2017: Implemented 2-Moment Microphysics
  on OSCER Schooner

- Thanks to CASA Colleagues from UMass, CSU, NCTCoG, and OSCER Staff
- Funded by NWS National Mesonet Program, NOAA OS&T

# Fine Scale Prediction Challenges

- **Data Latency**
  Many legacy observation and communications systems built for 12-h or 1-h data cycles.
- **Data Quantity vs Quality**
  Many new data sources but must handle non-optimal siting, non-independent errors, etc.
- **Poor Man's Ensembles**
  A well designed ensemble has dozens of members. Due to CPU limitations some operational systems today attempting ensembles with fewer than 10 and/or using time-lagged members.

# HPC Challenges

- **Multiple Hardware Platforms**
  Each Platform has unique features, compiler options.
  - Some help from community when using community software
  - Using cutting-edge computing sometimes means we are the beta-testers or even alpha testers
  - Research computing resources not ideally suited for real-time quasi-operational forecast experiments, but we have gotten good support
- **NWP Utilization of Co-processors Limited**
  - Domain decomposition is most often used for MPI.
  - Mixed results from combining MPI and OpenMP directives
  - Too much memory swapping and inter-processor communications required by mathematics of Navier-Stokes equations.

# HPC Challenges

- **All Steps Contribute to Latency**
  Model *system* must finish before the weather.
  - Includes observation data receipt and pre-processing
  - Post-processing steps: graphics and ensemble statistics
  - Transmission to end-users
    - Science DMZs help, but connection to NOAA a challenge
- **Education**
  - Fortran not taught to undergraduates
  - Recent grad students more experienced with MATLAB and Python
- **Our Focus is Science, but HPC Details Need to be Addressed**
  - MPI and optimization of code not normally supported by research grants.
  - MPI code often difficult to debug

# For more information…

CAPS Real-Time Forecasts and Ensemble Products

Online: http://forecast.ou.edu

Contact Info:
kbrewster@ou.edu
405-325-6115