NOAA Global Systems Laboratory

Advanced Technologies

Mark Govett
Advanced Technologies Division Chief

Global Systems Laboratory

Introduction

- Advanced technologies support the end-to-end capabilities needed to improve weather prediction
  - Earth System Prediction - from observations to prediction
  - Decision Support – tools to translate data into information, knowledge, actions

- Advanced technologies drive improvements in the end-to-end prediction system
  - Computing to enable faster, more accurate models
  - Data technologies to improve handling, analysis, delivery, visualization

Excellence and Innovation at GSL

- Data Systems
  - 2016 - NOAA Admin Award for transition of MADIS to NWS - Greg Pratt
- High Performance Computing
  - 2015: Demonstrated simplicity, performance and portability across CPU, GPU and Intel-MIC processors for weather model s - Mark Govett
- Machine Learning
  - 2016 – First ML development, use of multi-GPUs at OAR - Jebb Stewart
  - 2020 – Co-chair of NOAA AI Executive Committee - Jebb Stewart
- Cloud Computing
  - 2018 – First to run FV3 model on AWS - Jebb Stewart
- Visualization
  - 2000 - SOS demonstrated innovative display of planetary data – SOS team

Technology R&D - from research to operations

- Explore and track potential technologies
- Investigate promising technologies to understand potential
- Develop prototypes to understand, measure value, capability
- Sustain investment to improve capabilities
  - performance, ease-of-use, data handling, flexibility, understanding
- Transition capability to an operational entity

2015 Review Panel Recommendations

- Transition
  - E4.1: *Continue to mature the highly successful MADIS and SOS developments*
  - E4.2: *Work with NWS Operations to improve the R2O process as part of the continued evolution of MADIS*
- High Performance Computing
  - E4.3: *Align with NOAA-wide efforts to develop the next-generation global and regional modeling capabilities*
  - E4.4: *Work with NCEP /EMC to support improvement of the NGGPS (UFS) modeling system*
- Visualization
  - E4.5: *Incorporate visualization technologies into existing GSL tools to avoid creating niche applications*

Schedule

• Data Systems - MADIS, AQPI	Greg Pratt
• High Performance Computing	Isidora Jankov
• Cloud, AI, Visualization	Jebb Stewart
• Wrap Up	Mark Govett

NOAA Global Systems Laboratory

Observations, Data Collection, and Delivery Systems

Greg Pratt
Innovative Weather Delivery Systems Branch Chief

Global Systems Laboratory

The Need

Observations



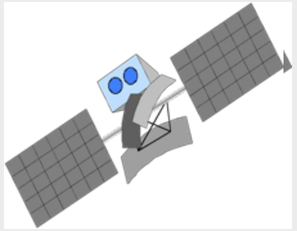


Data Collection

Data Delivery

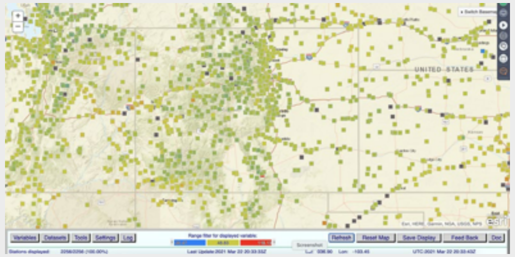
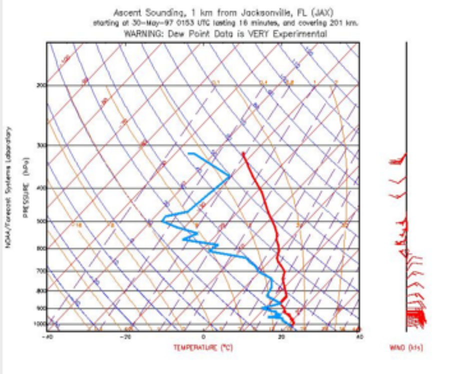

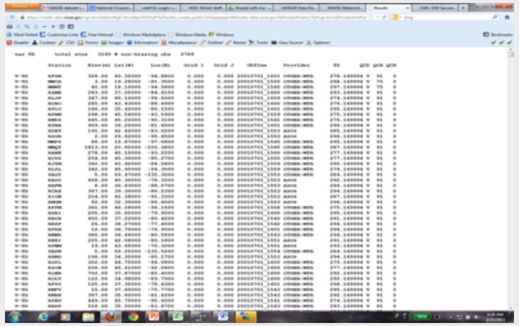

NOAA Global Systems Laboratory

Meteorological Assimilation Data Ingest System (MADIS)

Need for MADIS

NOAA surface stations with MADIS

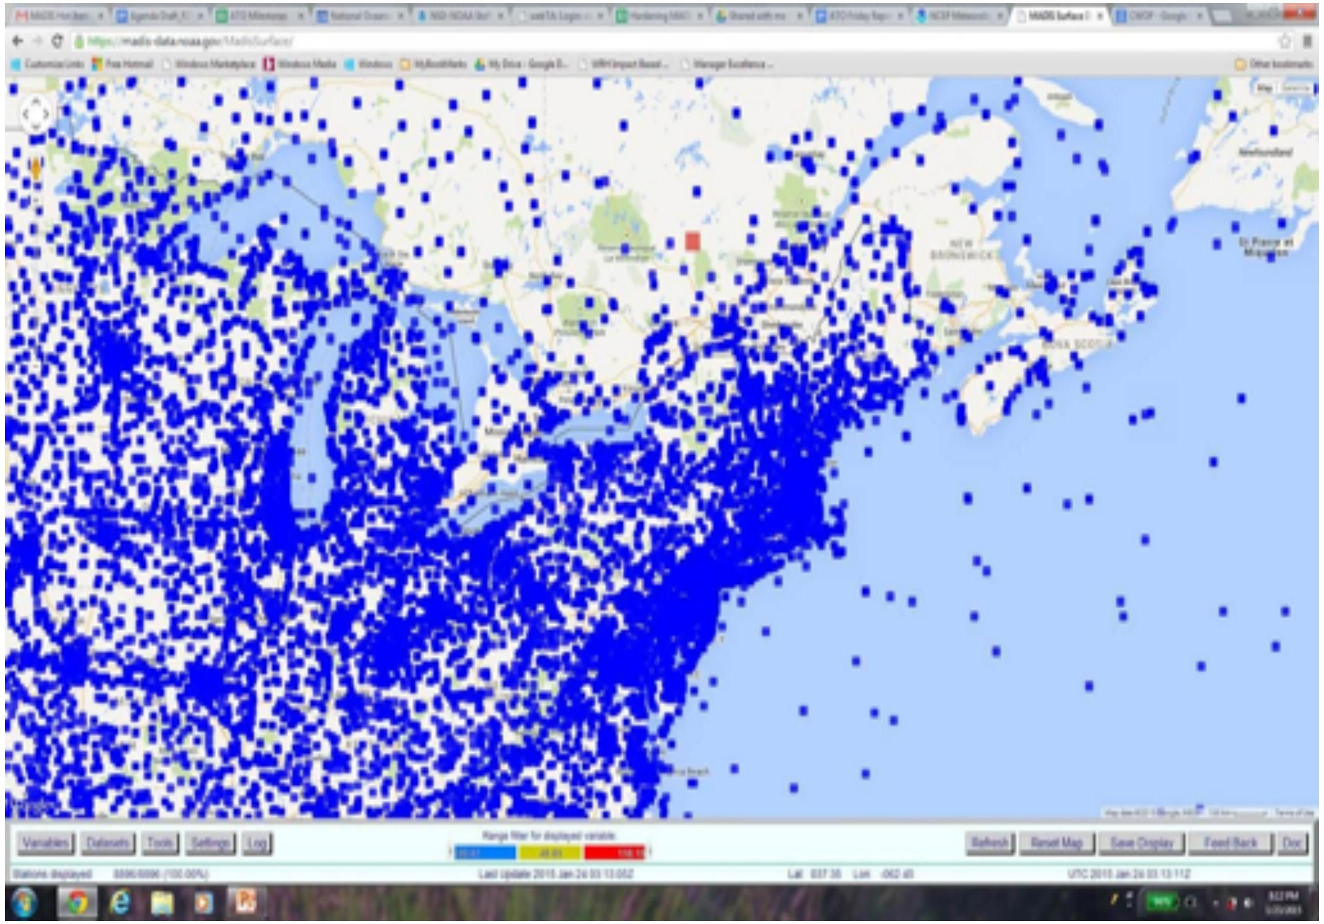

MADIS Operations and R&D budget

- $3.5M/Y

NOAA saves

- Equipment Installation costs
  - $464.5M
- Communication costs
  - $92.6M/Y
- Maintenance costs
  - $46.4M/Y

10

MADIS Collaborators

Operational Support Provided by NWS

- Observations
- Dissemination

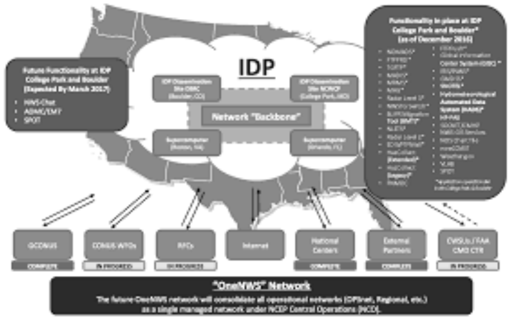

Operationally maintained:

NCEP Central Operations

Simplified MADIS Operations
NCO/IDP (College Park and Boulder)

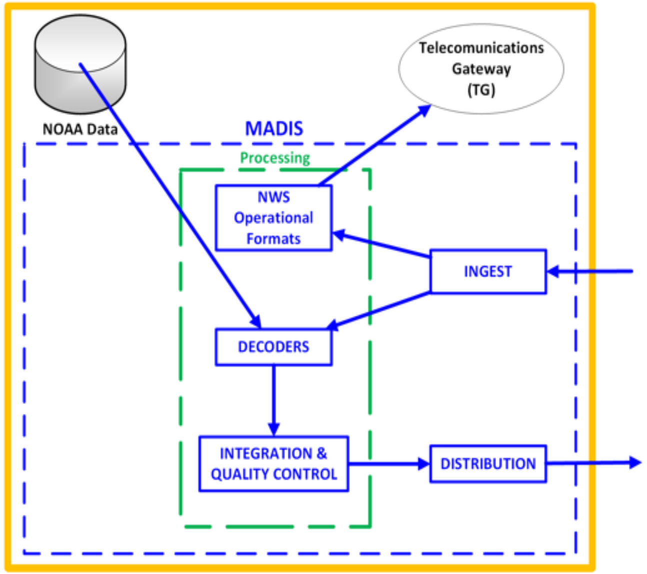

MADIS archive:

National Centers for Environmental Information

National Environmental Satellite Data and Information Service

DEPARTMENT OF COMMERCE

MADIS Research & Development:

Global Systems Laboratory

What is MADIS

Framework for collecting, integrating, assessing the quality, and delivering observations to the meteorological community

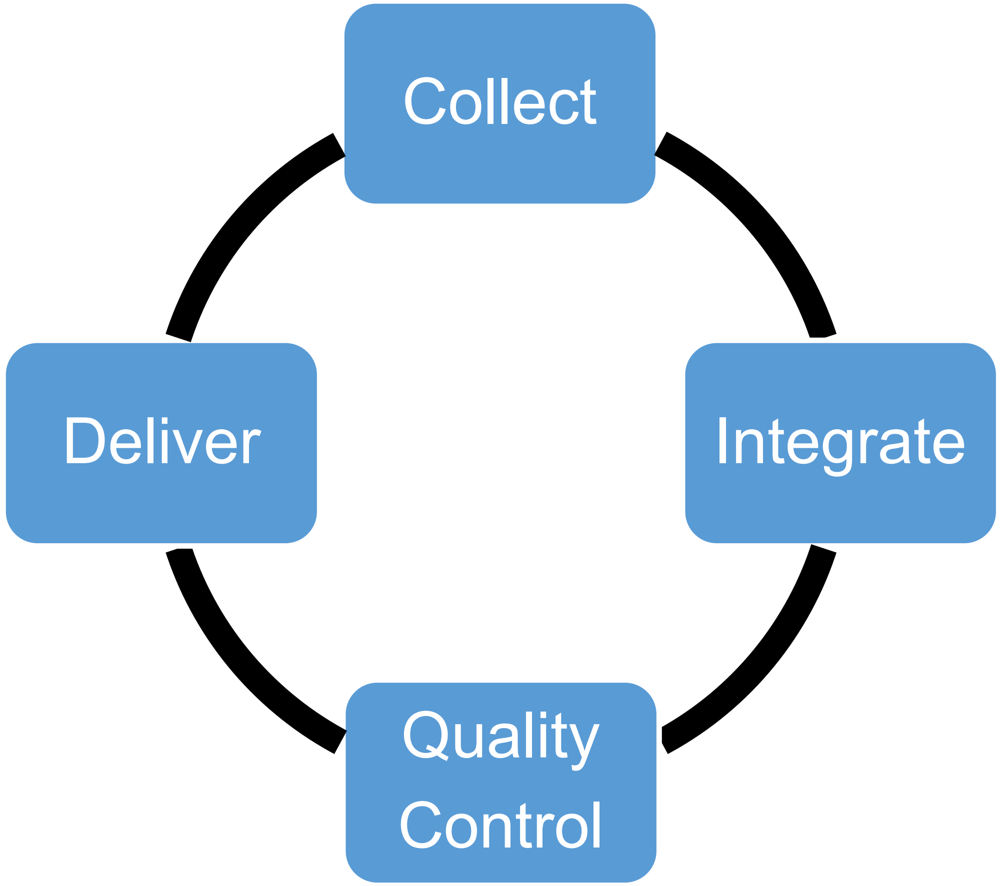

MADIS Accomplishments 2016-2020

4 Major releases between 2016-2020

Added new data systems

- Hydrometeorological Automated Data System
- Automated Flood Warning System
- Clarus (DOT RWIS)

Automated Flood Warning Systems

HADS Hydrometeorological Automated Data System

MADIS Accomplishments 2016-2020

Collection Improvements
- Adding new data systems
- Adding new data sets
- Setting a MADIS data and metadata standard

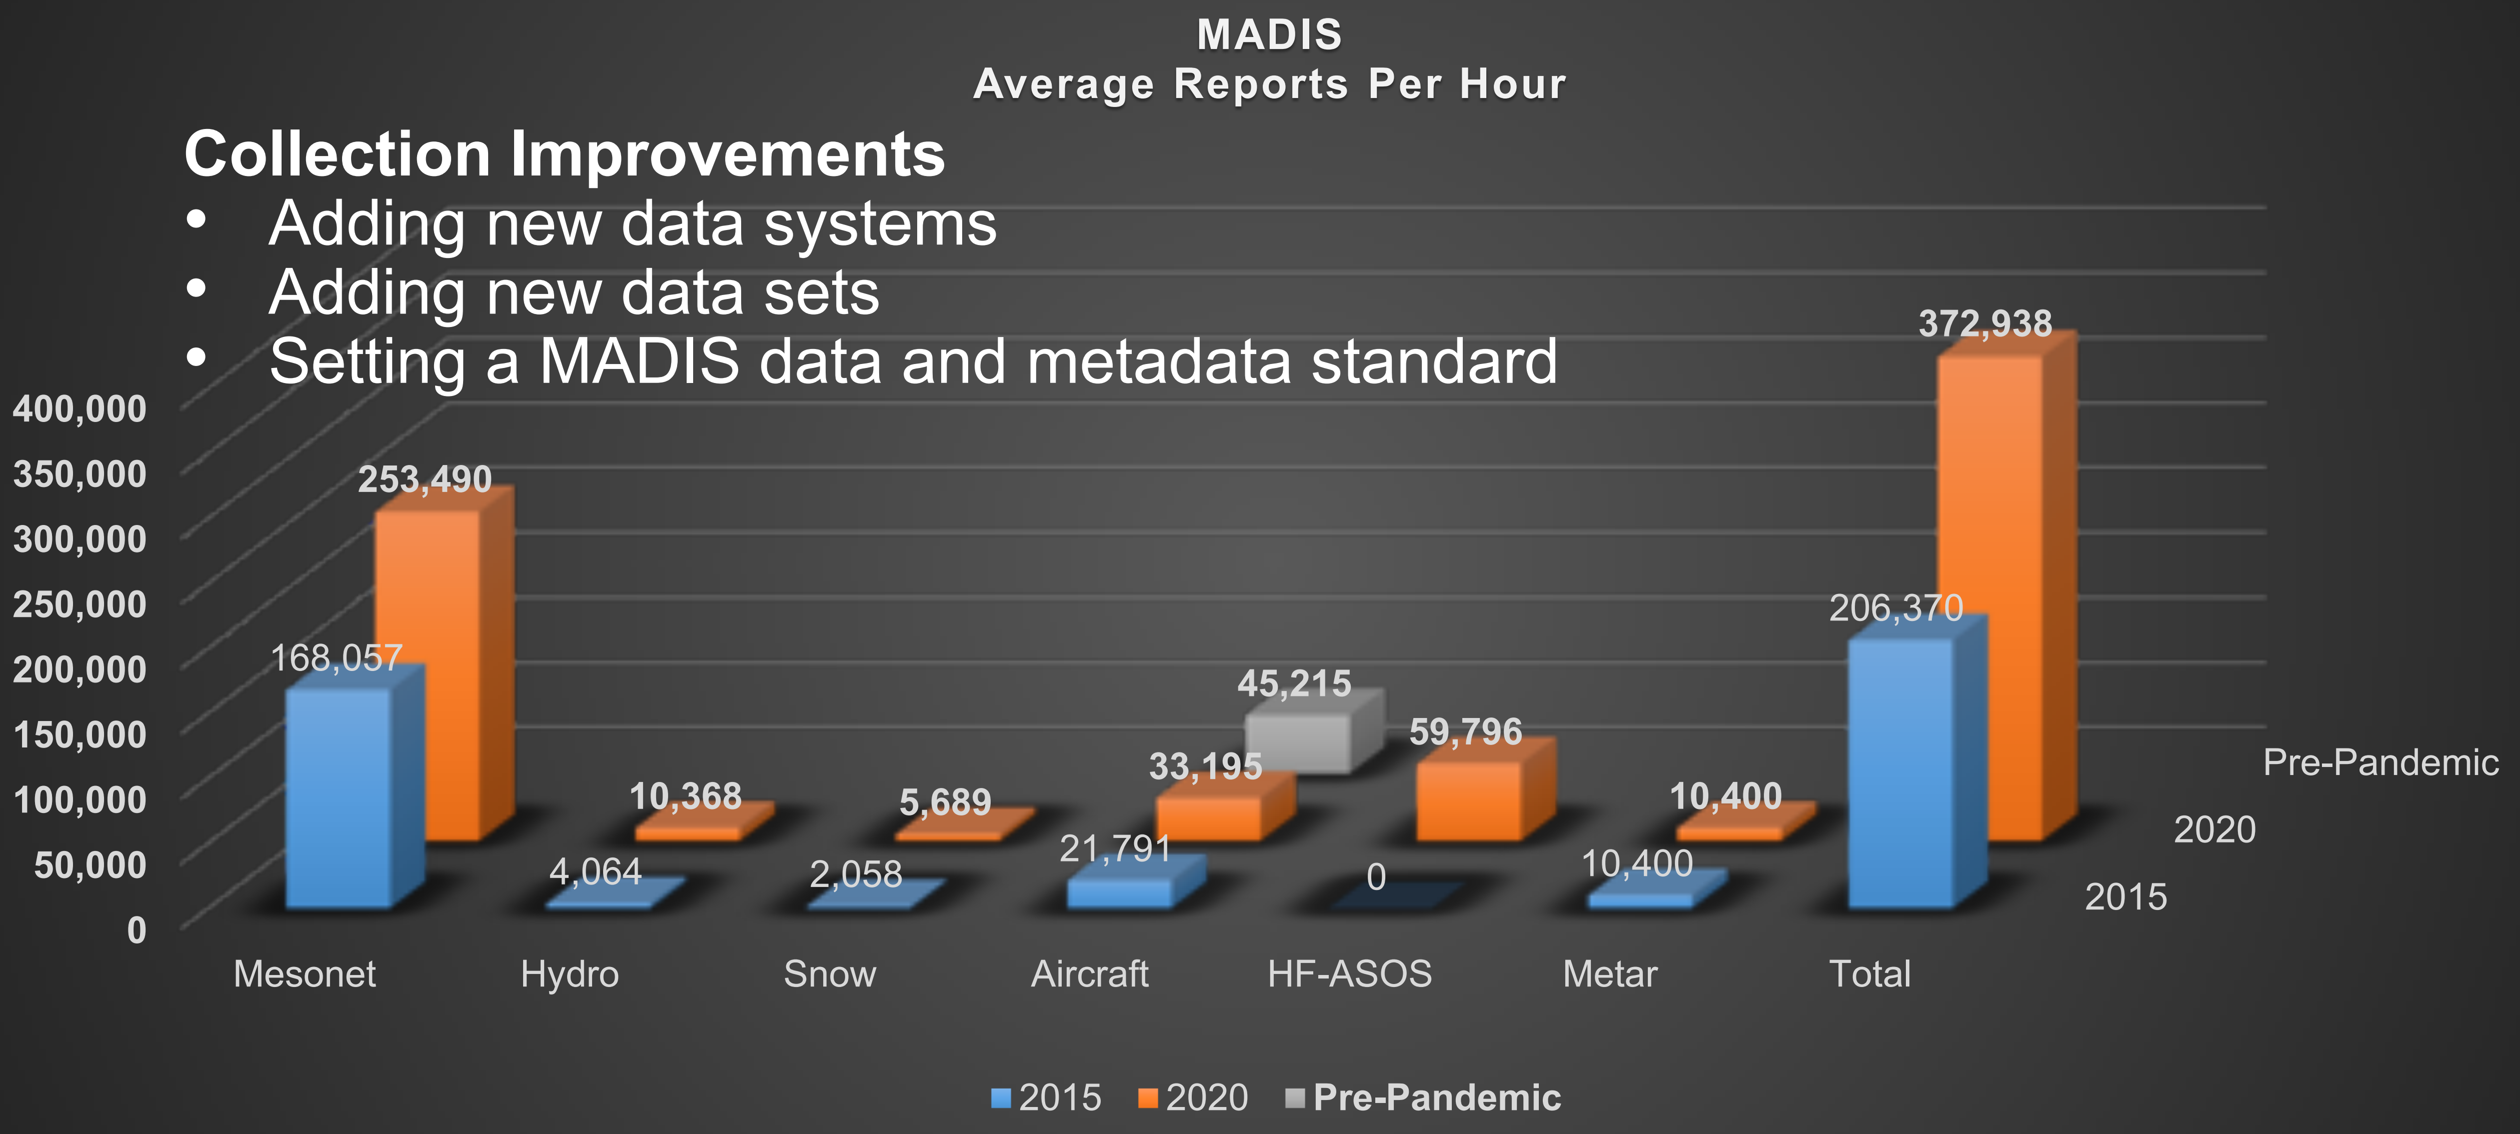

MADIS Accomplishments 2016-2020

Global Data Center (GDC) for WMOs to acquire Aircraft Based Observations

MADIS accomplishments are aligned with the 2015 laboratory review recommendations

- OAR continued to support the growth of MADIS (E4.1)
- Extended MADIS by adding new data sets and improving data latency (E4.2)
- Integration into existing operational systems (E4.5)

NOAA Global Systems Laboratory

Advanced Quantitative Precipitation Information (AQPI) System

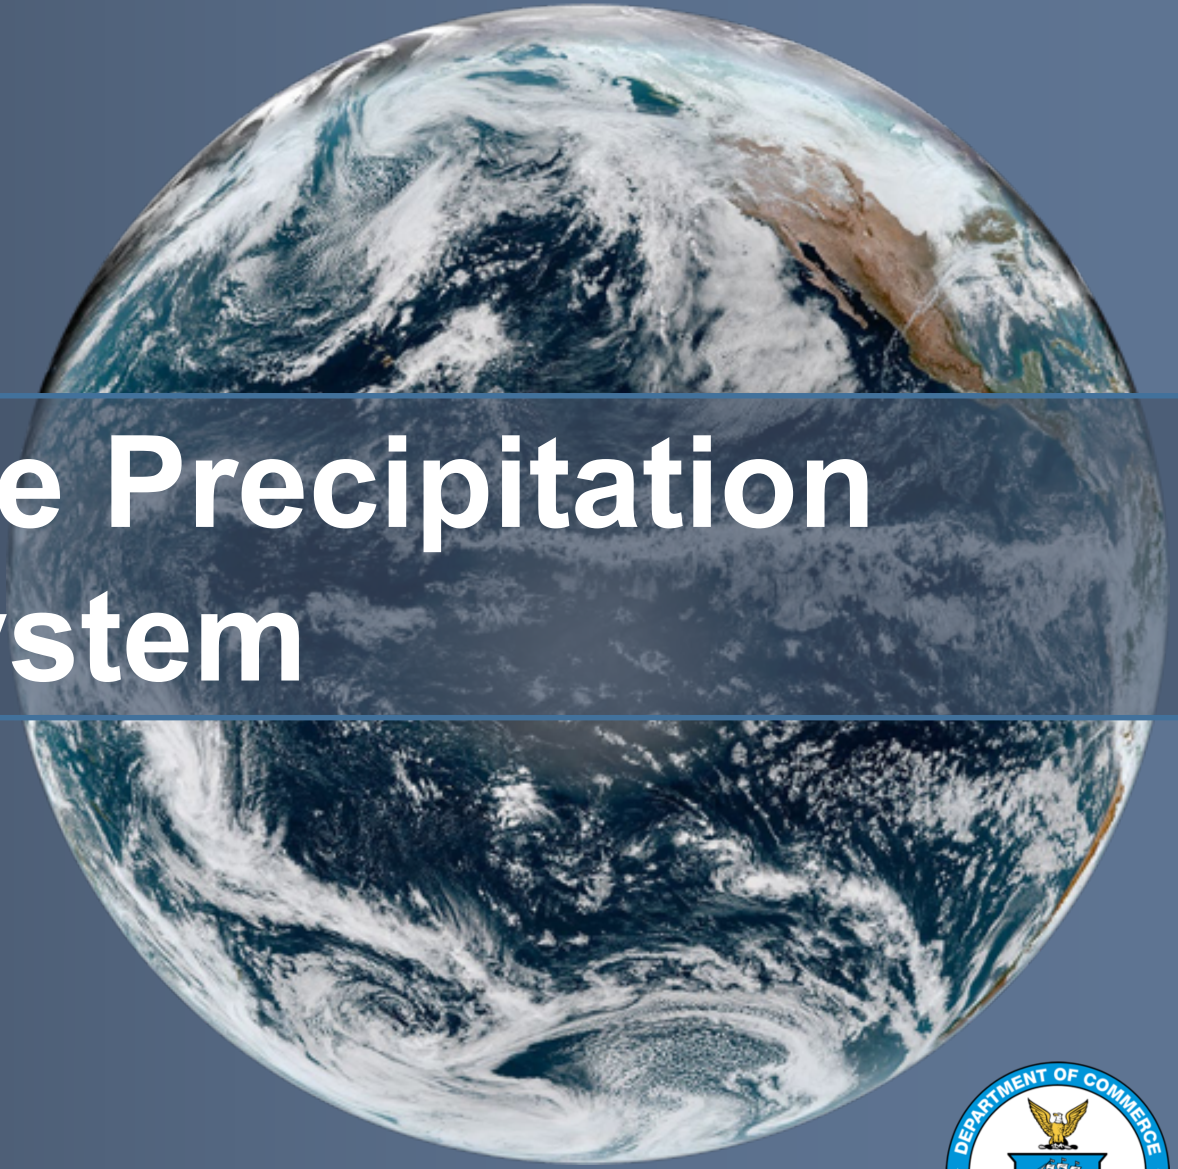

Global Systems Laboratory

AQPI The Challenge of Extreme Precipitation

30-50% of CA's annual precipitation comes from a handful of Atmospheric Rivers (AR) events that are vital for replenishing water supplies but…

- ARs often cause flooding
  - Responsible for>80% of flood damages in the western US, including CA
  - On average, >$1B in annual damage costs
- SF Bay Area is particularly vulnerable
  - 7^{+} million people (5th largest in US)
  - >350,000 people in 100 year flood plain ($46B in exposed structures)
- Sea level rise and urbanization will exacerbate the problem

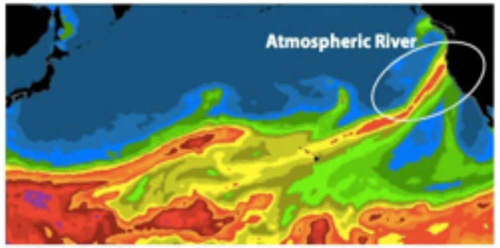

Satellite image of an atmospheric river in February 2017.

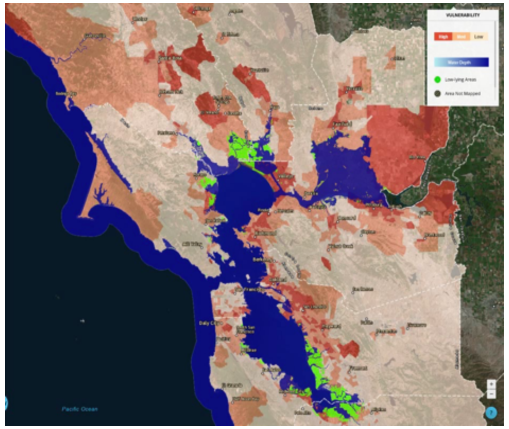

AQPI Collaborators

- California Department of Water Resources (DWR) Funding
- NOAA Research
  - Physical Science Lab
  - Global Systems Lab
- US Geological Survey
- CIRA
- Scripps Center for Western Weather and Water Extremes
- 9 Counties in the San Francisco Bay Area

The AQPI System

What is it?

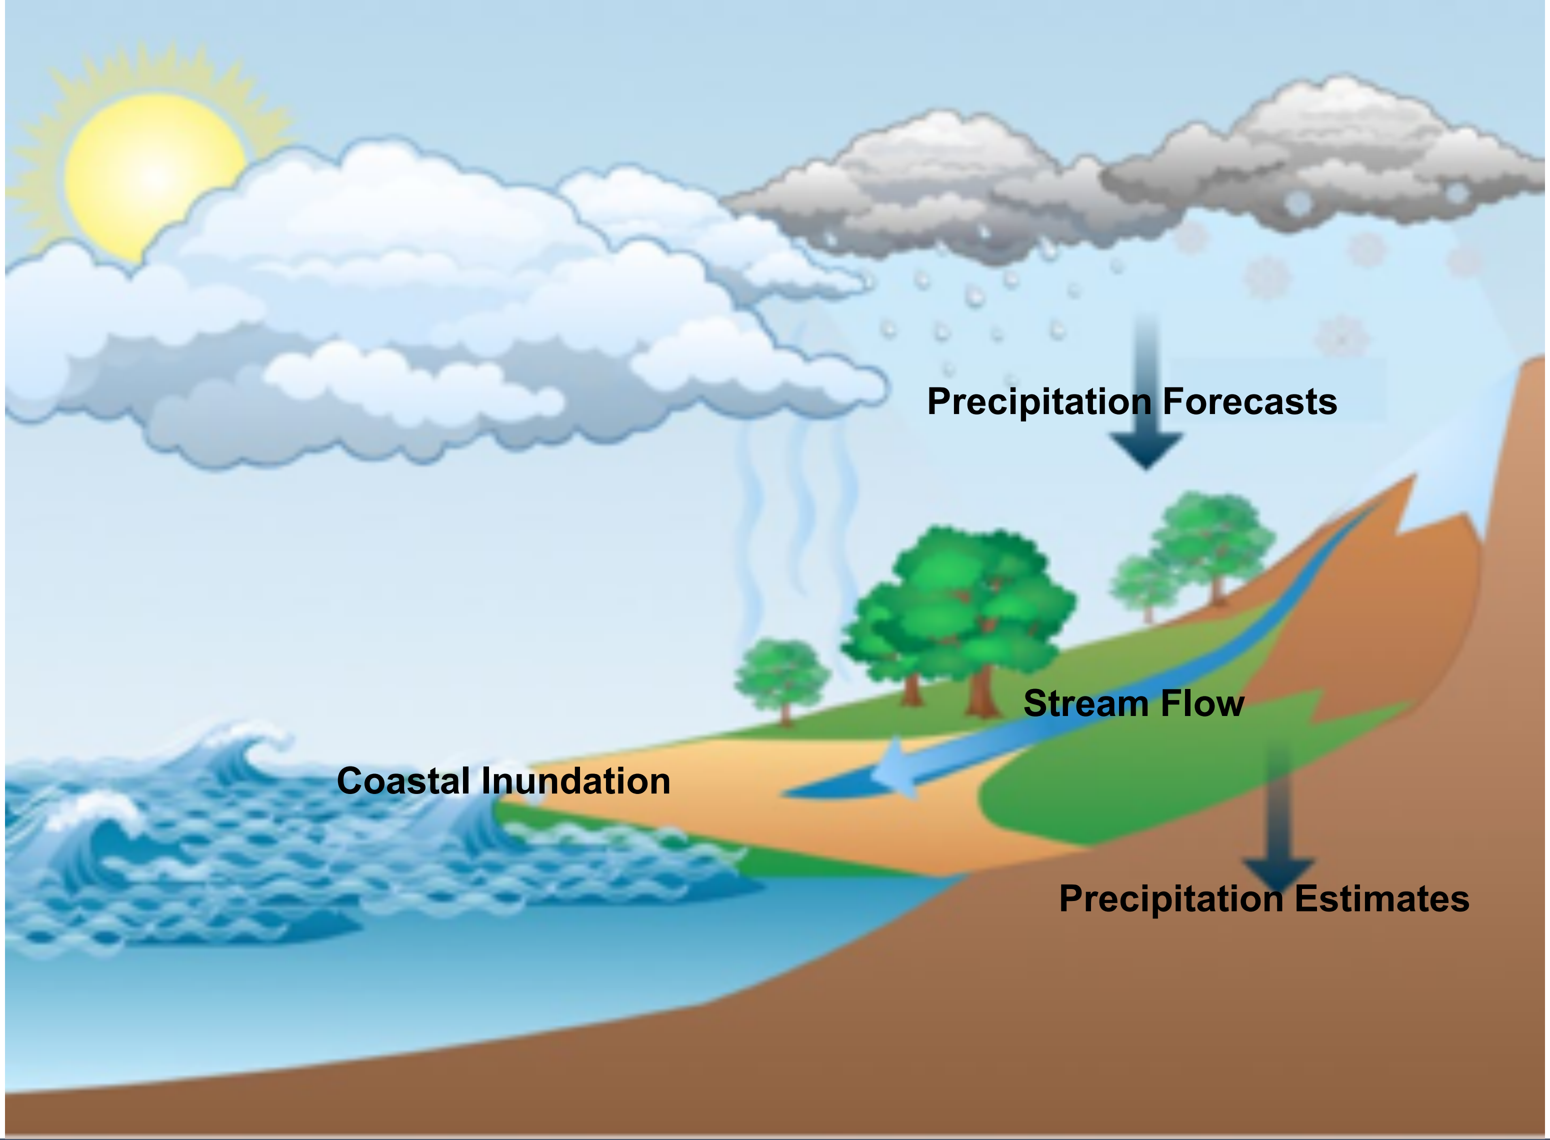

The AQPI System

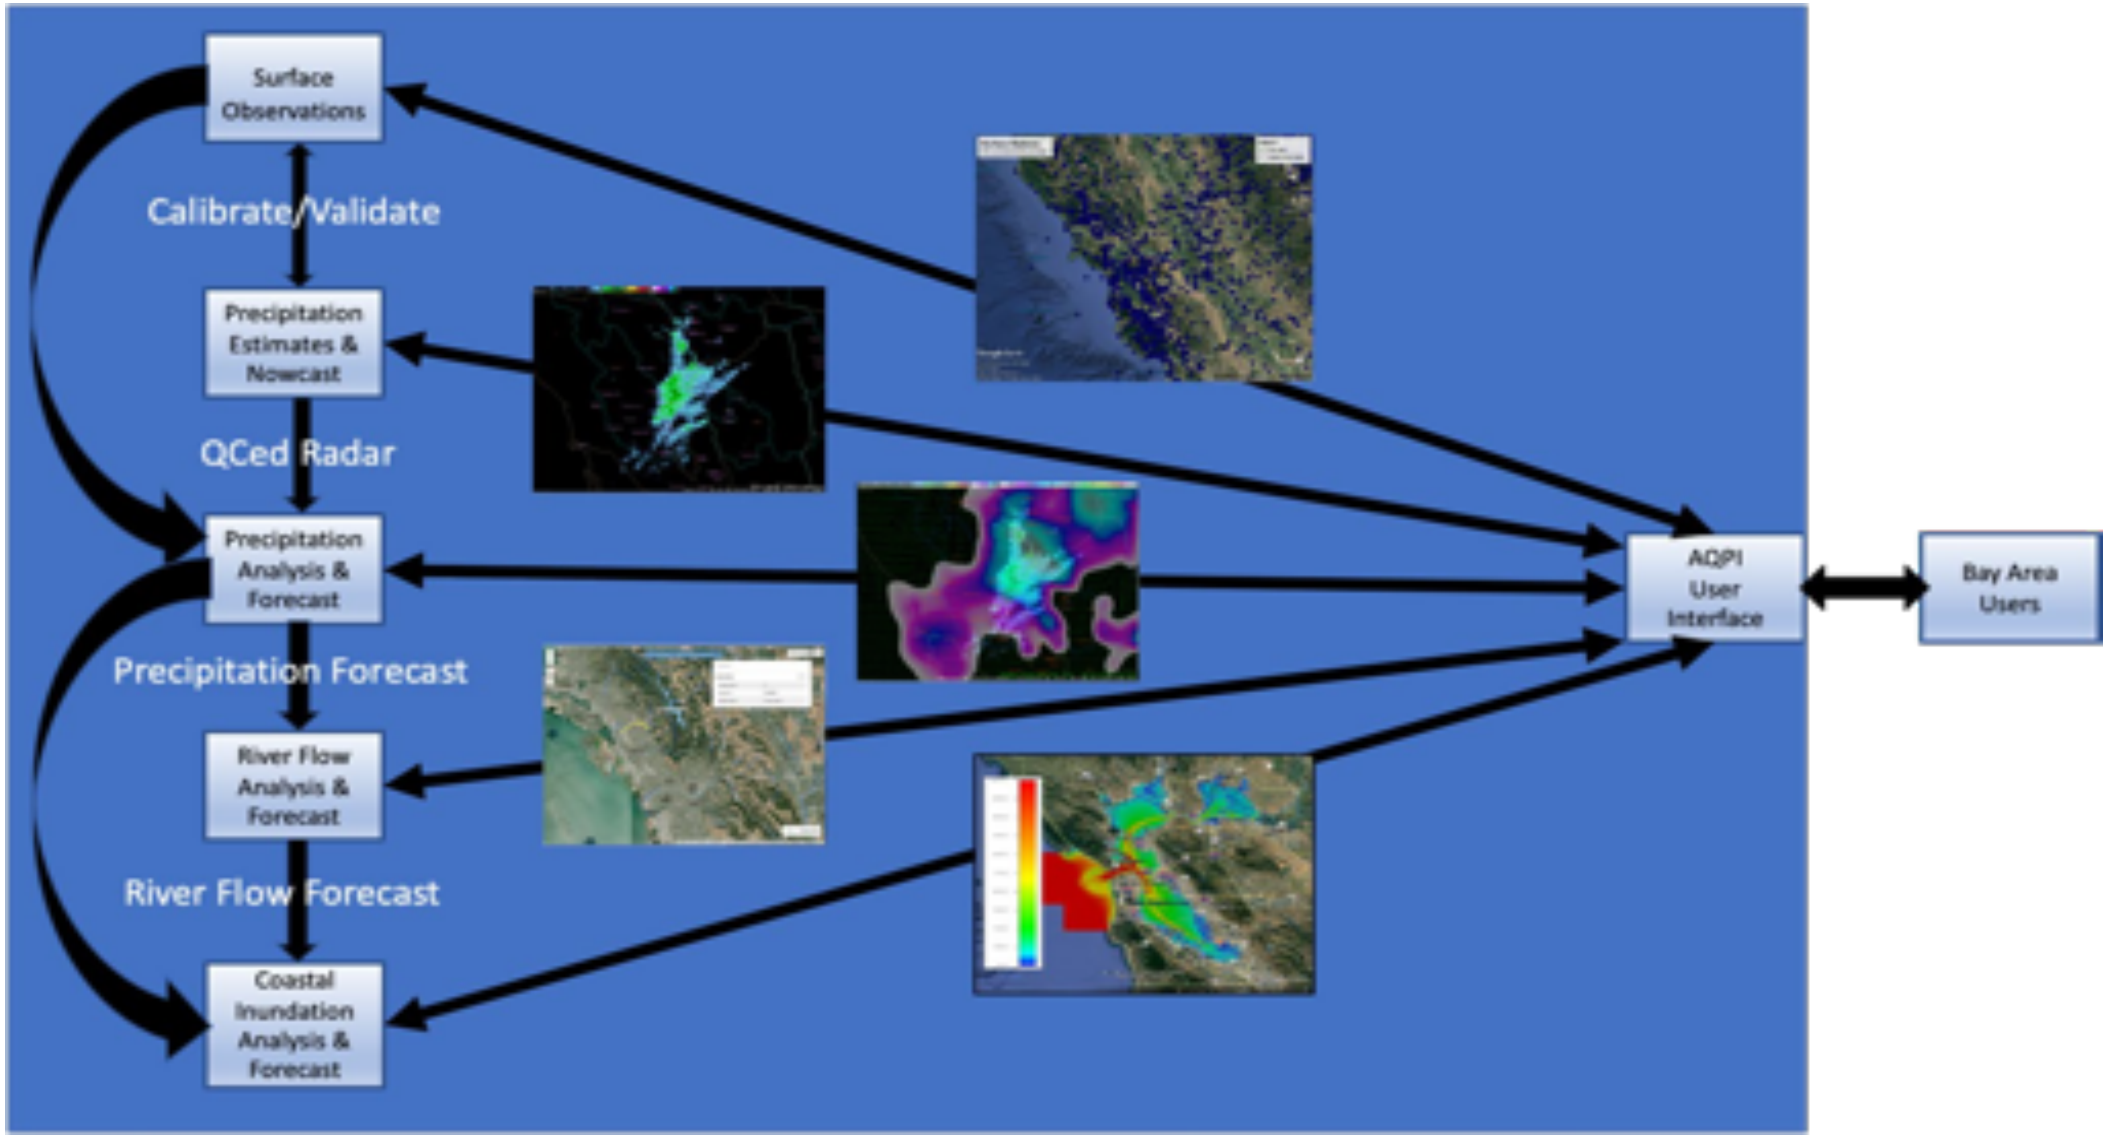

Accomplishments

- Tied to operational systems
- Collecting real-time surface observations from 9 counties feeding operational MADIS system
- Two radars of the 5 promised installed and sending products to experimental MRMS and working to deliver products to AWIPS-II
- MRMS and MADIS data used to initialize Experimental HRRR
- MADIS and HRRR data initializing NWM

Accomplishments

- USGS Coastal Storm Modeling System (CoSMoS)
  - Running in real-time hourly creating inundation products
  - Initialized with HRRR, NWM, and other NOAA data

- User interface for data selection

- User engagement through entire process

What is Next

Research and develop an agile observation ingest, collection, and delivery framework that provides a continuous update to the current state of the atmosphere

- Larger volumes of data both temporally and spatially continually updating
- Dynamic reactive data additions
- Dynamic reactive data collection
- Dynamic reactive adaptable Quality Control
- User informed delivery
- Continuous framework improvements

Cloud computing, machine learning, and work with end-users provide the pathway for this research and development

NOAA Global Systems Laboratory

Advanced Technology High Performance Computing (HPC)

Isidora Jankov
High Performance Computing Branch Chief

GSL's Leadership in HPC R&D

GSL has been leading exploration of HPC technologies for NOAA since 1990's

Areas of GSL's leadership in HPC research and developments within NOAA and among International Partners:

- Computing with Graphics Processors (GPUs)
- Unified Forecasting System - computational performance improvements
- Exascale computing

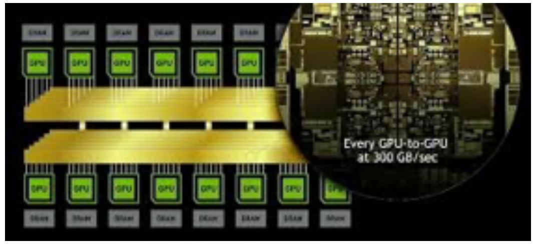

NVIDIA DGX-2
IBM CPU, 16 NVIDIA GPUs
80,000 processing cores

UFS

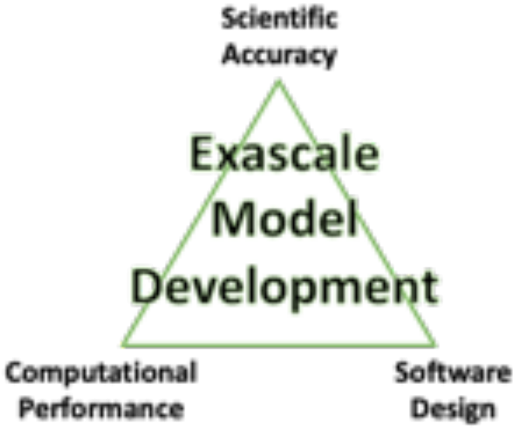

NOAA Global Systems Laboratory

GPU Computing at GSL

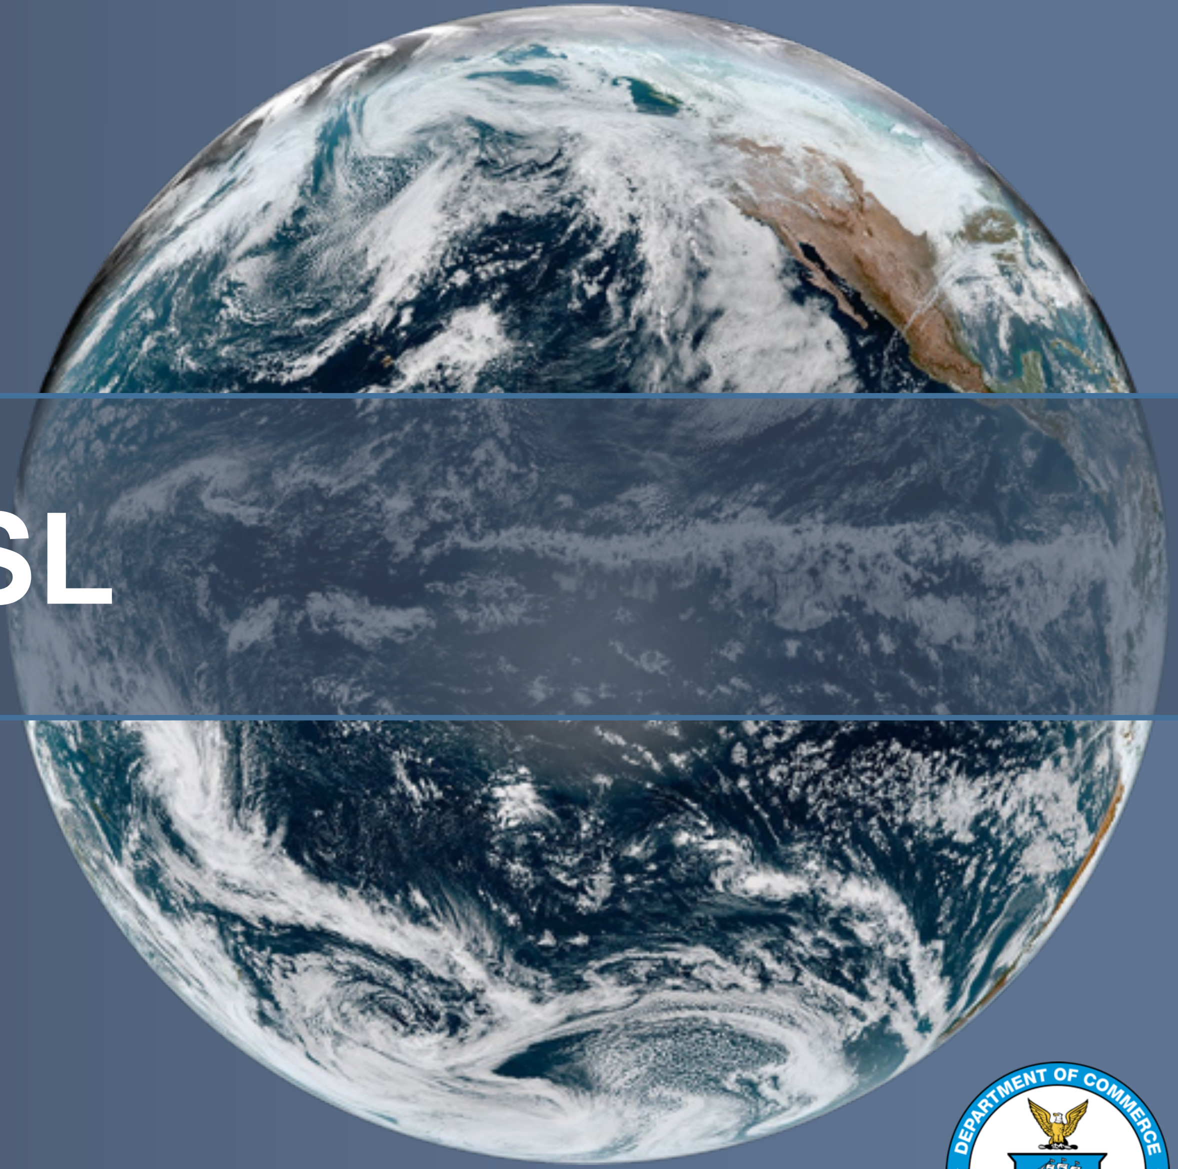

Global Systems Laboratory

GPU Computing

- 2008: GSL started investigating GPUs for computing applications
  - Focus on both performance and portability
- 2010: GSL built directive-based GPU compiler
  - Worked with the industry to develop, improve openACC standard
- 2015: Larger community started looking into GPUs
- 2017: BAMS paper demonstrated performance and portability between CPUs and GPUs with Nonhydrostatic Icosahedral Model (NIM)
  - Approach adopted by the Model for Predictions Across the Scales (MPAS) team

NOAA Global Systems Laboratory

UFS Computational Improvements

UFS Computational Improvements

Software Environment for New Architectures (SENA)

- An effort to ensure that NOAA’s operational code inventory is ready for expected evolution of HPC
- Collaboration between GSL, EMC and GFDL
- Improvements in Unified Forecasting System (UFS):
  - Extracting kernels from UFS and Joint Effort for Data Assimilation Integration (JEDI) to evaluate their performance on emerging CPU and GPU processors
  - Exploring different HPC computing languages (e.g. Julia) and software processes to improve efficiencies, readability, usability of NOAA codes

UFS Computational Improvement

- Legacy code, designed for CPU
- Selected several routines for evaluations

Level of code change	Level of Speedup
No change	~10x slower on GPU
Modest change	2x slower on GPU
Substantial change	1.5x faster on GPU
Major change (NVIDIA)	4x faster on GPU

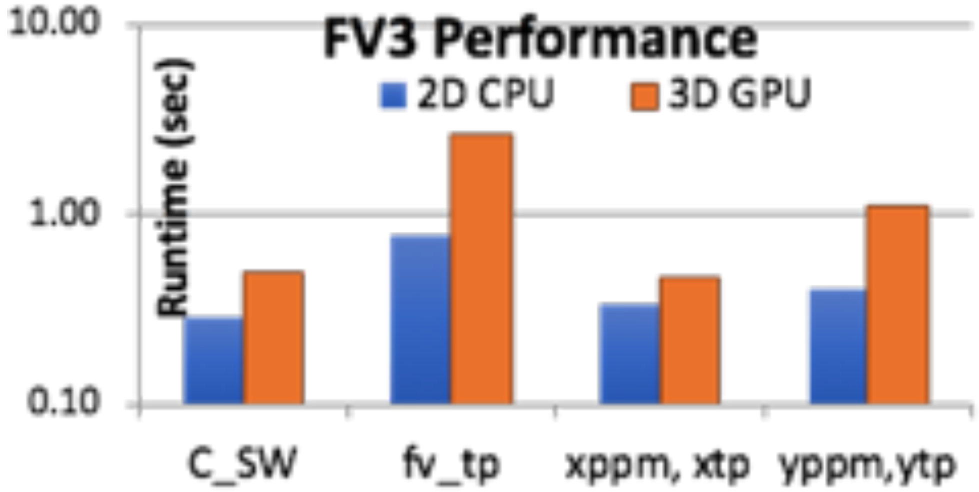

Outcome: rewrite of the code to address both performance and portability issues

JEDI Computational Improvements

- Shallow water model
  - First simple model with MPI functionality
  - Useful as a guide for new JEDI users
- Community Radiative Transfer Model (CRTM)
  - Too expensive to be executed frequently
  - Parallelization over both profiles and channels
  - Achieved speed up: ~7x faster
- Next steps:
  - GSL is committed to help optimize all of JEDI

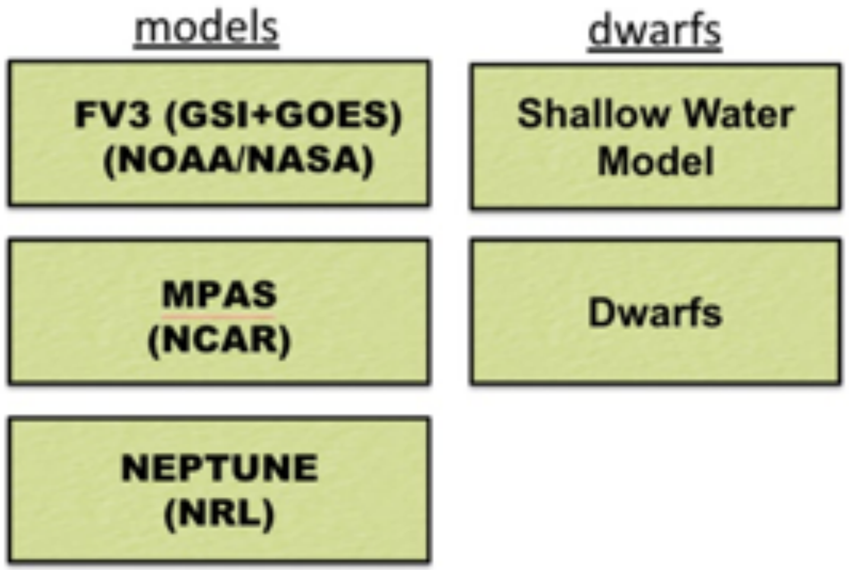

NOAA Global Systems Laboratory

Exascale Computing at GSL

Global Systems Laboratory

Exascale Computing

- Exascale systems
  - Millions of processors
  - GPUs are dominant technology
- Numerical models are getting more complex
  - Increasing resolution - towards cloud resolving scales on global domain
  - Changing physics schemes to explicitly resolve physical processes
  - Addressing model uncertainty - more ensemble members needed
- How to use these systems effectively to meet NOAA goals?
- 2017: GSL launched the exascale computing project
  - To evaluate algorithms/approaches in terms of accuracy and computational efficiency

GeoFLuid Object Workbench (GeoFLOW)

- Focus on evaluation of algorithms in terms of scientific accuracy and computational performance
- Grids supported
  - 2D & 3D Cartesian, 2D & 3D icosahedral grids
  - Cube-sphere grid
- Spectral elements and finite volume approaches
- Preliminary testing has begun
  - traditional dy-core tests
  - performance at fine scales

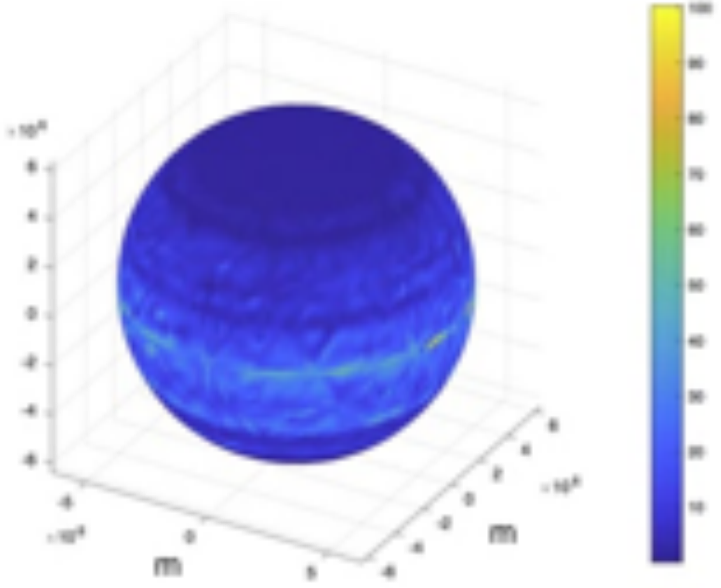

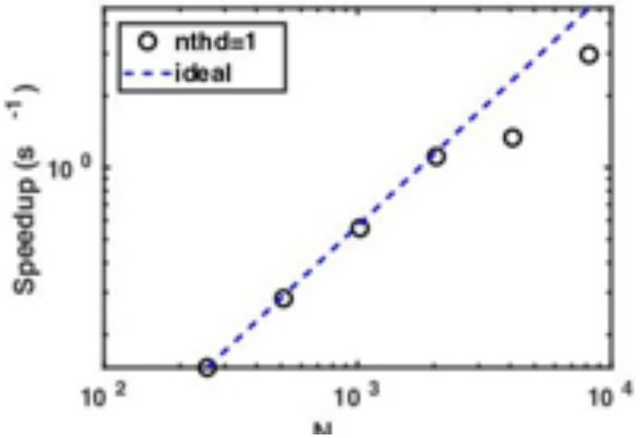

Kinetic energy density (J/m^3) for examination of planetary wave propagation on idealized model earth. Computations made with GeoFLOW convection solver using dry dynamics highlight initial global modeling capability (upper), "speedup" on a parallel computer up to 10,000 processors is almost ideal (lower).

Algorithm Evaluation for Exascale

- Comparison of numerical accuracy for shallow water icosahedral A vs. C grid
- Improved algorithm resulted in better stability and higher accuracy
- Focus on computational efficiency

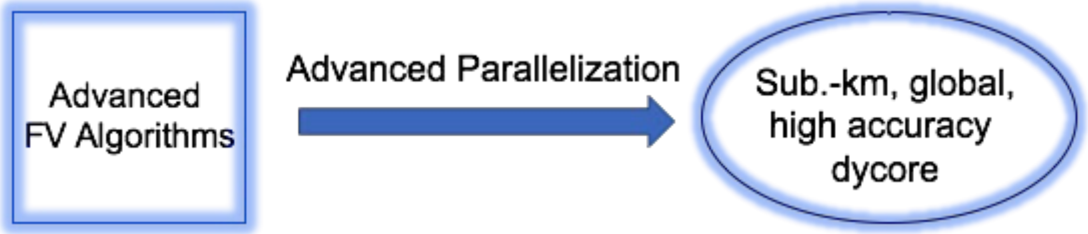

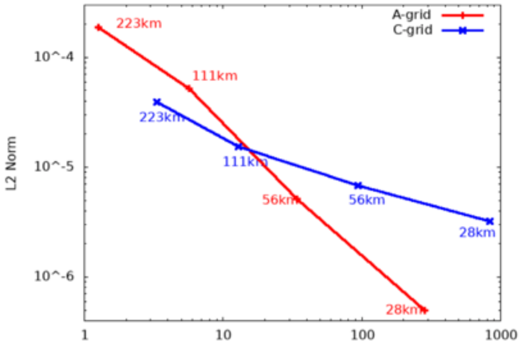

Solution error vs forecast time for two different grids for the SWM model, showing higher order convergence (larger slope) for the A-grid formulation.

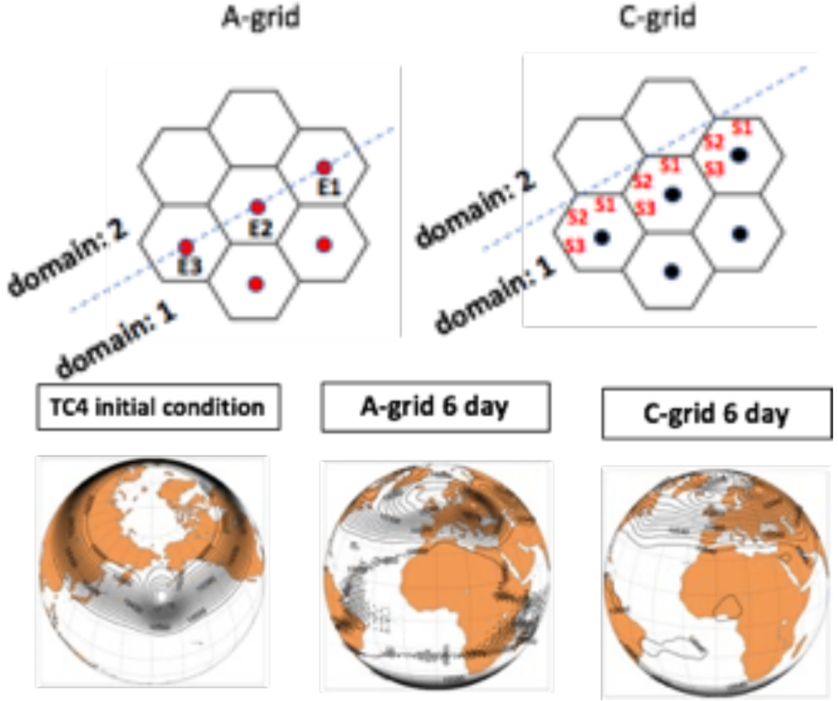

Yu, Y. G., Wang, N., Middlecoff, J., Peixoto, P. S., & Govett, M. W. (2020). Comparing Numerical Accuracy of Icosahedral A-Grid and C-Grid Schemes in Solving the Shallow-Water Model. Monthly Weather Review, 148(10), 4009-4033.

Going Forward

- Continue exploring advanced HPC technologies to provide the most benefit to NOAA
  - Performance, cost and capability
- Continue advancing computational performance of UFS and JEDI
- Continue leading exascale computing effort
  - Evaluate exascale ready model capabilities based on scientific accuracy and computational efficiency
  - Apply new improvements to UFS and JEDI

NOAA Global Systems Laboratory

Advanced Technology Information Delivery

Jebb Q. Stewart
Informatics and Visualization Branch Chief

Global Systems Laboratory

Information Systems and Visualization

- Key to the understanding of science, scientific data, and complex concepts
- Drives communication of what we do and why it’s important
- Allows us to deliver information on risks and impacts

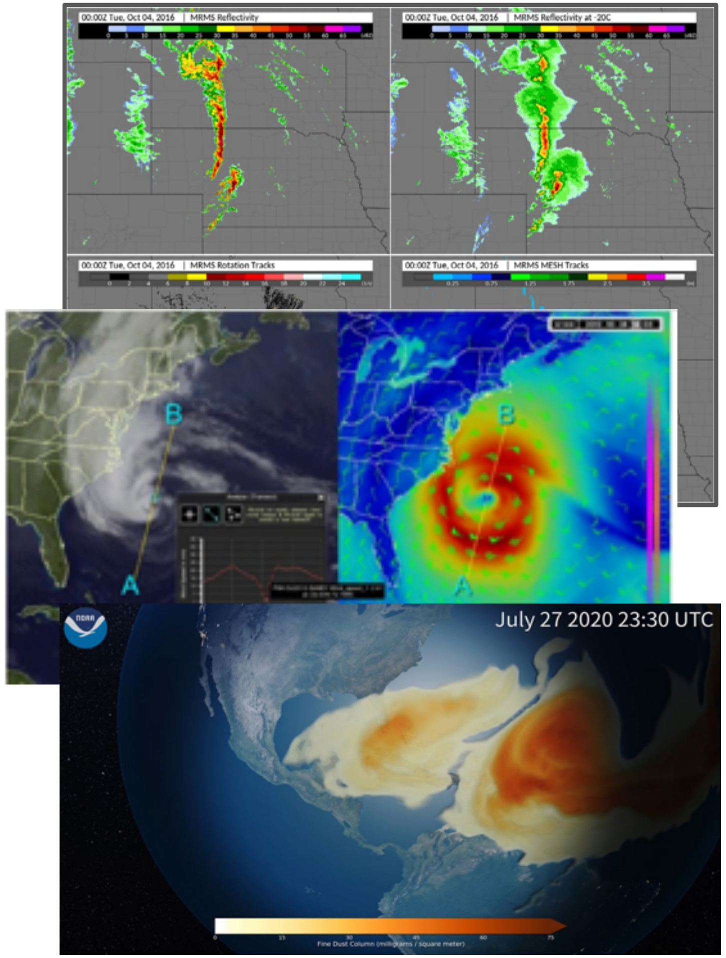

Needs
- As data volumes increase, we need the tools to extract and visualize relevant information.
- Communicating uncertainty and confidence
- Blend information from multiple sources to enhance understanding and decision making
- Situational Awareness and data mining

Accomplishments

Science on a Sphere
NOAA's Premier Education Tool

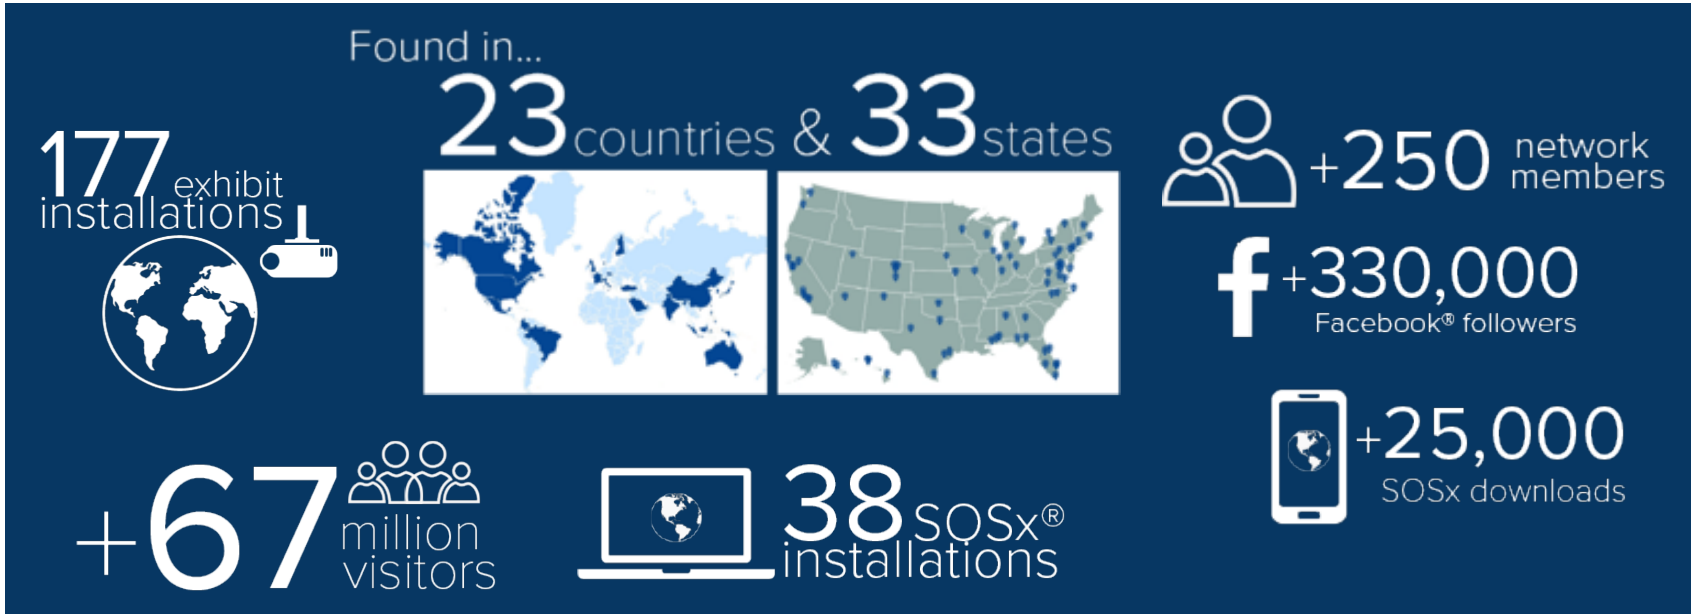

(Pre-COVID #)

SOSx Mobile 25K+ Downloads

Accomplishments

Engaging Visualization, Driving Readers to Stories

Arctic sea ice minimum

African dust across the Atlantic

Accomplishments

Widely used visualization capabilities

Interactive smoke visualization
(Peak ~ 1 million request per hour)

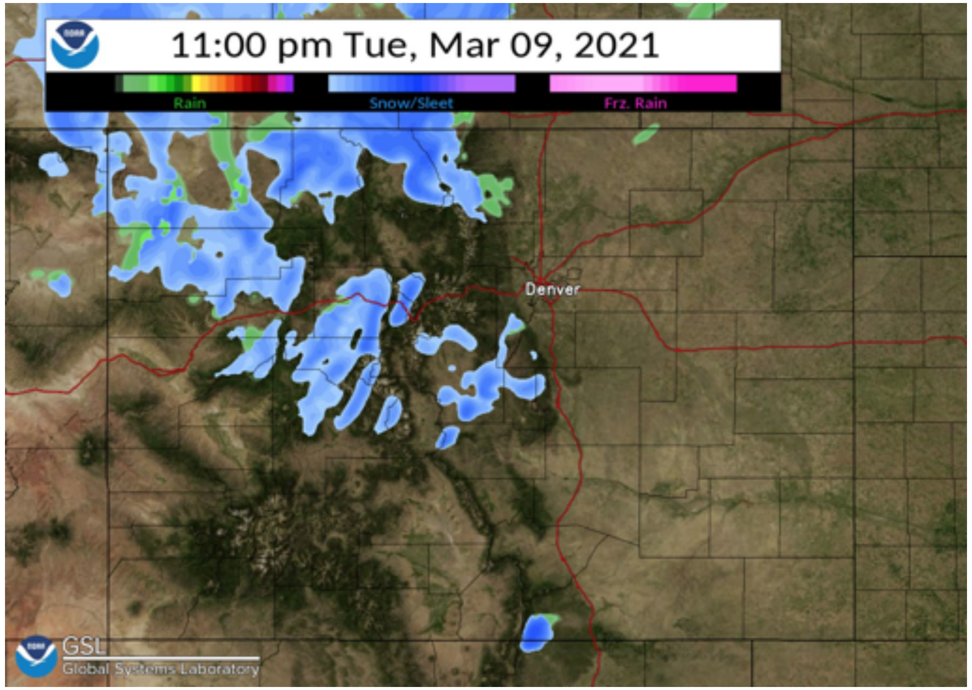

Weather Archive and Visualization
Environment (WAVE)

Future of Information Delivery

New ways of interacting with information

- Augmented Reality Apps and LED Spheres

Enhance Communication and Understanding

- Impact-based Decision Support Services (IDSS) Engine
- Social Scientists and User Experience Development

Big Data and Exascale Information Systems and Visualization

- Data mining for impacts and extracting actionable information
- Process diverse observations

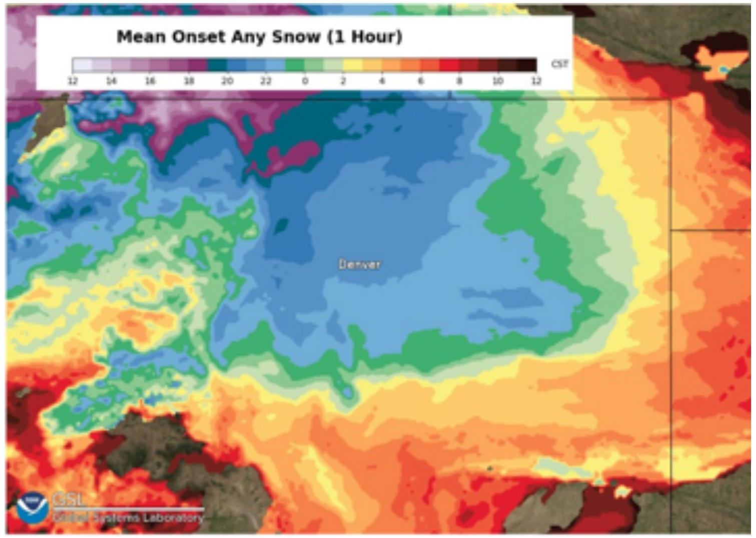

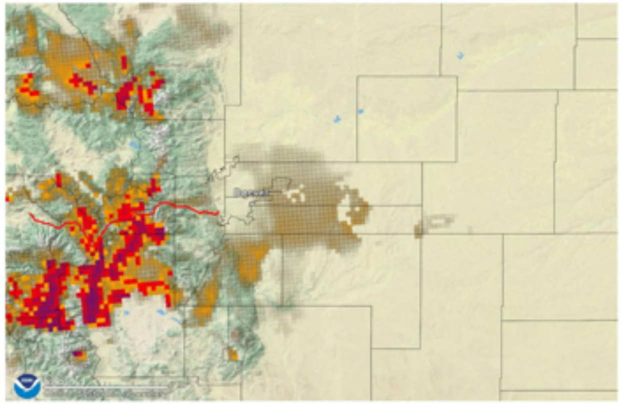

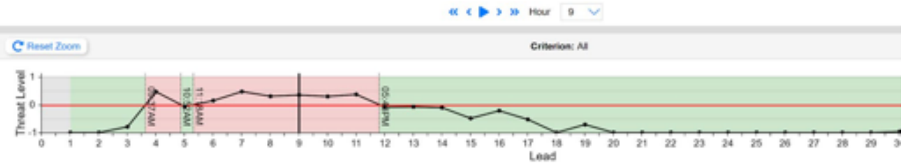

NOAA Global Systems Laboratory

Cloud Computing at GSL

Global Systems Laboratory

Cloud Technologies

Cloud Services provide new opportunities

- Increase operational efficiency and agility
  - Dynamic scaling of applications and surge computing needs
  - Ad-hoc and surge capacity
  - retrospective model evaluation
- Collaboration on shared data and code

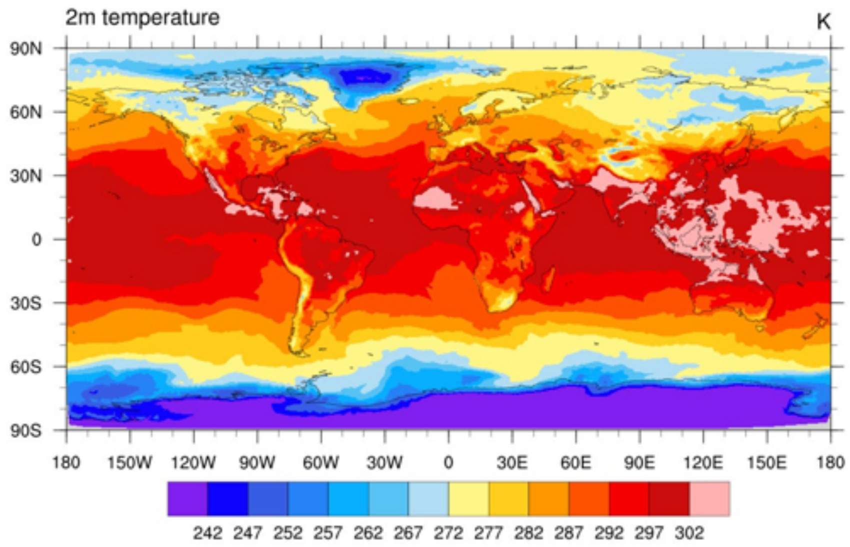

Needs

- Understand and build systems to use HPC cloud capabilities
- Evaluate cost effective uses for applications
- Create new capabilities and collaboration tools leveraging data in the cloud

Cloud Development

Modeling

- Development of UFS and Global Workflow
  - Ability to run end-to-end workflow using cloud services
- Rapid Refresh Forecast System (RRFS) in the cloud (Partnership with EMC)

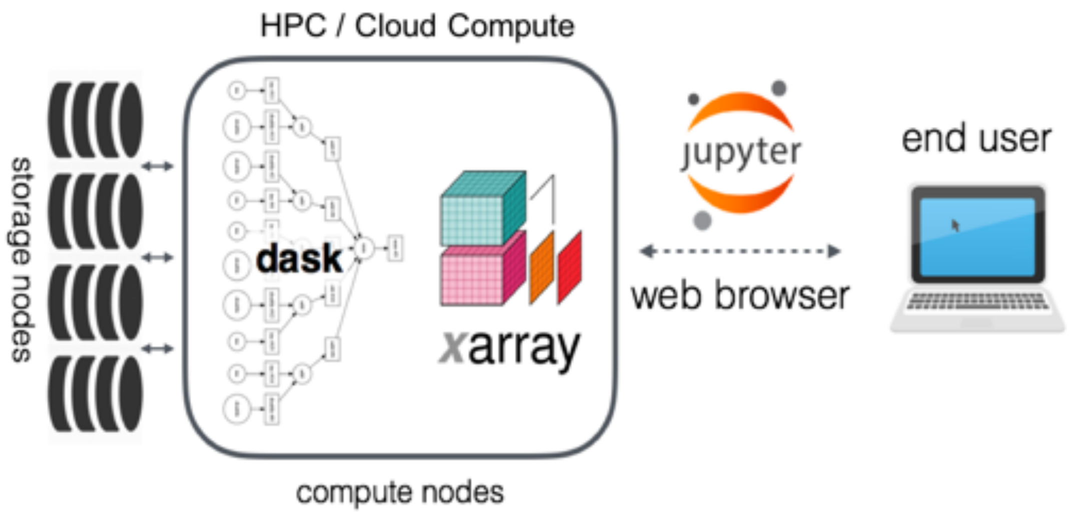

Visualization and Analysis

- Advanced multithreaded data processing for exascale data analysis and visualization

Leadership

- GSL co-lead of OAR Cloud Tiger Team

Future in the Cloud

Modeling

- Continue model development allowing retrospective, ad hoc, and surge computing capability
- Performance and cost trade-offs

Visualization and Analysis

- Continued research into dynamic scalable applications to leverage cloud based data and computing

Leverage new cloud capabilities

- Serverless computing, only using compute when needed.

NOAA Global Systems Laboratory

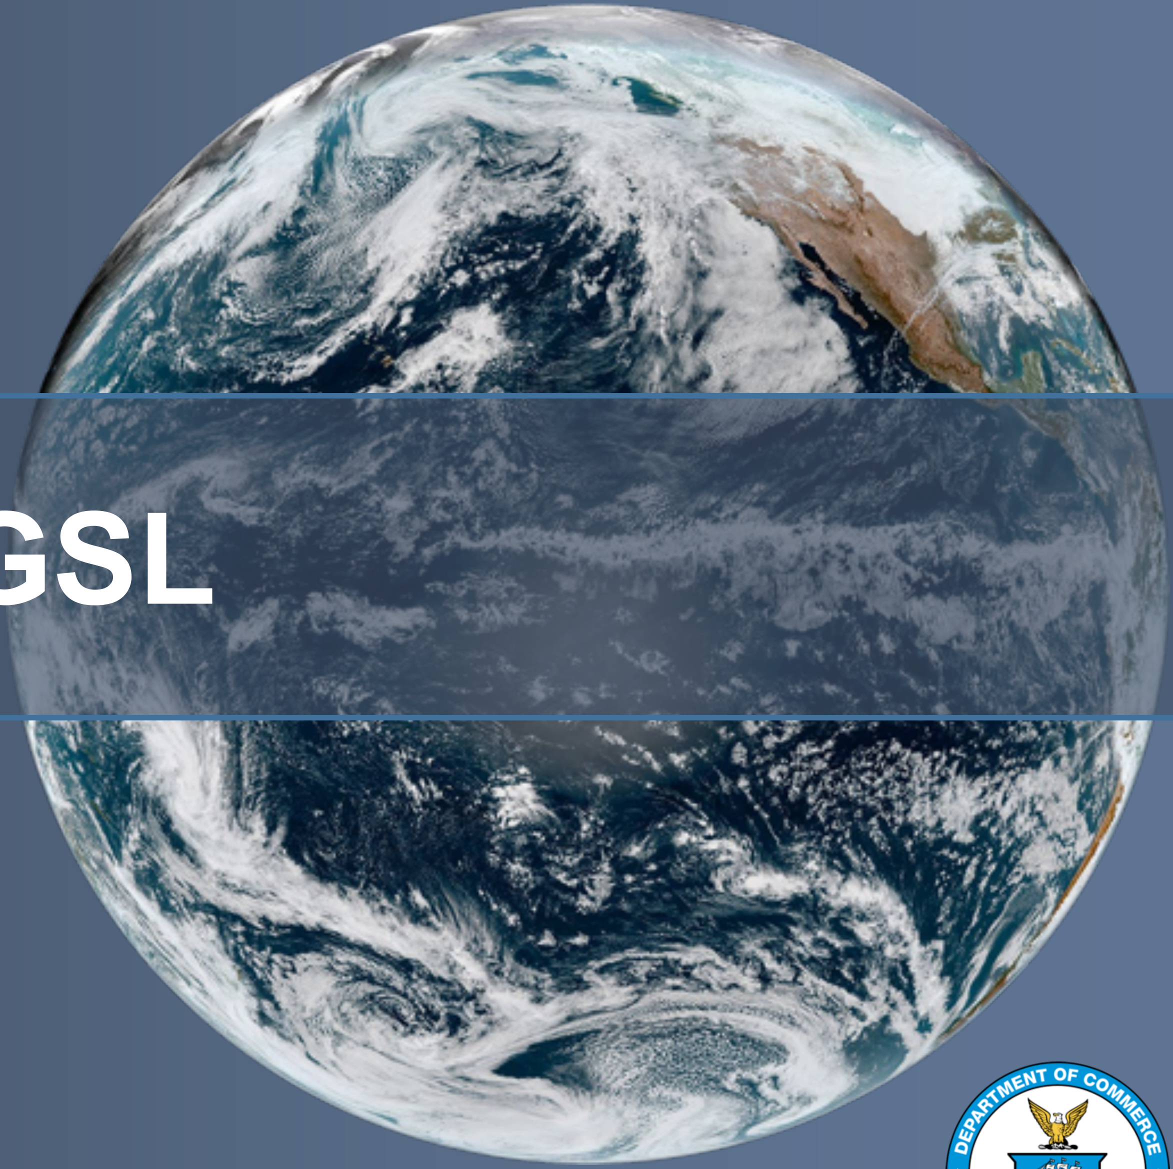

Machine Learning at GSL

Global Systems Laboratory

Artificial Intelligence and Machine Learning

Proven capabilities

- Significantly reduce computation time and costs
  - Needed as models move to higher resolution (spatial and temporal)
  - Efficiently and intelligently process large volumes of data

- Increase skill of forecasts and forecast tools
  - Advanced high dimensional statistics

Performance Driven

1000x Speed up over traditional methods

Physics Guided Approach

Neural Network where during training process, network is penalized when physical constraints are violated.

Example: Radiative Transfer

1. Flux at surface and Top of Atmosphere cannot be negative
2. Radiative Heating Rate = function of upwelling and downwelling flux for the layer

Explainable and Trustworthy AI

Example Tropospheric HR Profiles

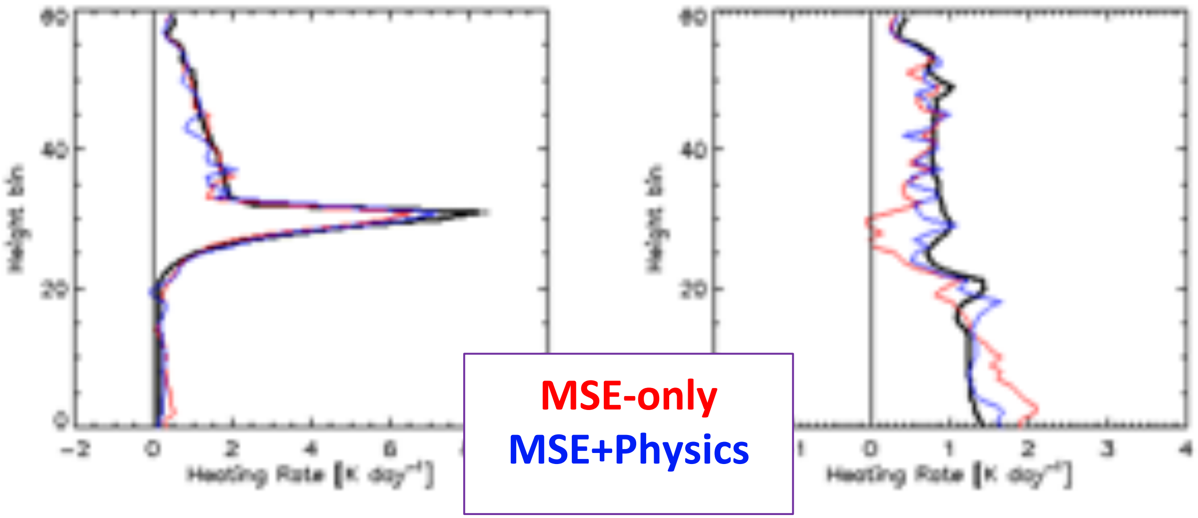

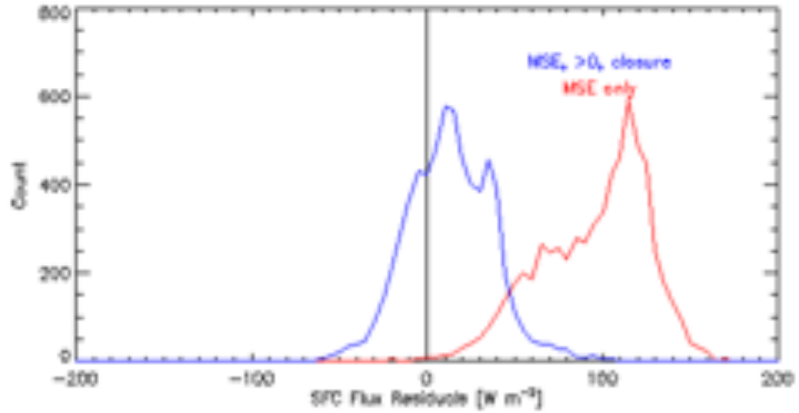

Feature Detection, Tracking, and Prediction

Improve performance and speed of traditional methods

Convection Initiation Detection and Monitoring

- Using AI to create probabilistic forecast of areas of likely convection initiation with various lead times (+30 to +180 minutes)
- Collaborating with Taiwan Central Weather Bureau

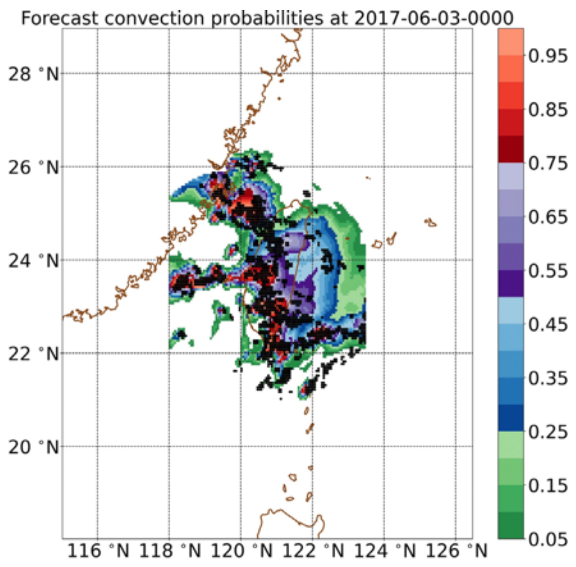

Cyclone Detection

- Using AI to detect and target areas for further data extraction
- Working with ECMWF to test in pre operational capability

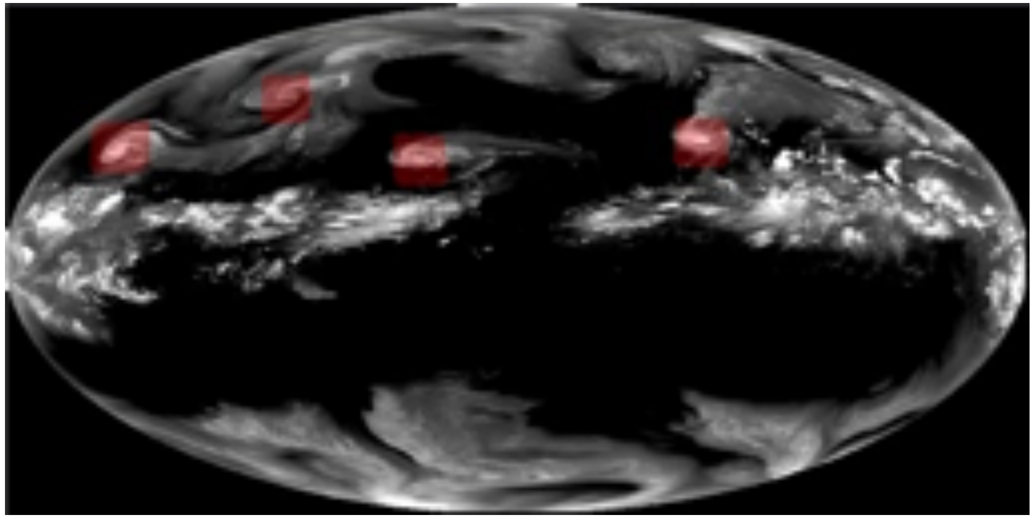

Artificial Intelligence - Future

Performance Optimization

- Physical parameterization optimization

Efficient Data Processing

- Intelligent data targeting and thinning
- Improved Data Assimilation

Increase Forecasting Skill

- Cyclone Detection and Intensification
- Applications in verification

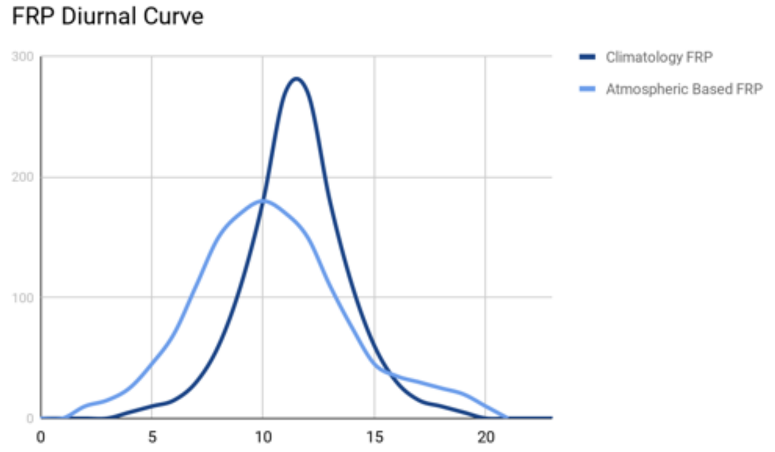

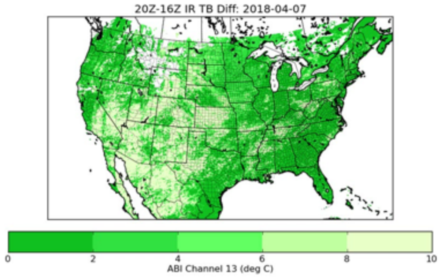

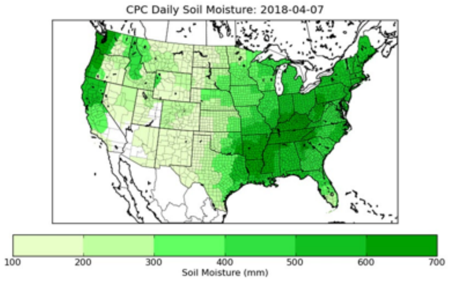

NOAA Global Systems Laboratory

Wrap Up

Mark Govett
Advanced Technologies Division Chief

Global Systems Laboratory

Relevance and Leadership

- GSL is building on its long history of innovative work using technology to advance weather prediction capabilities
- GSL is a leader in using technologies for NOAA
  - GPU computing - 2008
  - Machine learning - 2016
  - Cloud computing - 2017
  - Visualization - long standing
  - Exascale computing - 2017
- GSL works with national and international partners using technologies
- MADIS and SOS are highly valued and impactful developments that benefit NOAA and the community

Quality and Performance

- MADIS, AQPI systems
  - MADIS handles diverse types of data and operation use
  - AQPI is meeting growing demands for data by diverse users
- SOS, SOSx
  - Delivers enormous impact to diverse audiences worldwide
  - Millions of people worldwide see SOS as the window into NOAA research
- Cloud Computing and AI
  - Machine learning algorithm to detect hurricanes adopted by ECMWF
  - Develop ML algorithm to replace radiation that is 1000 times faster than HPC
- Visualizations
  - Science is effectively communicated to large and diverse audiences

Quality and Performance

- Using HPC more effectively
  - Improve UFS and JEDI performance
  - Extract routines from operational codes and improve them
- Leading GPU research
  - NIM demonstrated performance portability in single code
    - Modeling team and computational team succeeded by working together
  - Collaborate on UFS, CCPP physics
- Preparing for exascale computing
  - Investigate new algorithms and approaches
    - GeoFLOW to evaluate scientific accuracy and computational efficiency
  - Explore new languages and software approaches

Future Plans - Big Data

- How can NOAA systems work with 10,000 times more data?
  - Cloud computing
    - Continued improvements in data handling, storage, visualization
  - Machine Learning
    - Further advancements in data thinning to improve use of satellite, radar, model data
  - Data and Visualization
    - SOS, AQPI to demonstrate advanced delivery and display capabilities
- Collaborations
  - Industry, academia, agencies, laboratories

Future Plans - Exascale Computing

- Develop models to run efficiently on systems with millions of processors
    - UFS may need to be rewritten or replaced
        - Science, software, computational performance will become a barrier
    - Use GeoFLOW to evaluate new algorithms and approaches
    - Explore alternative development strategies and languages
- Collaborate on operational, kilometer-scale weather and climate prediction models
    - WMO research teams on exascale
    - International teams: DoE, ECMWF, UKmet, GFDL, NASA, others

Thank you!