NOAA Supercomputing Directions and Challenges

Frank Indiviglio
GFDL
frank.indiviglio@noaa.gov
MRC Workshop – June 1, 2017

NOAA Is Vital to American Economy

A quarter of the GDP ($4 trillion) is reliant on accurate weather and climate information.

Example of NOAA's role:

NOAA provides weather, water, and climate forecasts and warnings for the private and public sectors. Annually, NOAA provides 76 billion observations, 1.5 million forecasts, and 50,000 warnings.

-NOAA provides economic benefits of $240 million per year in mitigating flood losses

-NOAA's aviation forecasts reduce aviation delays and save the industry $580 million per year.

High Performance Computing Locations

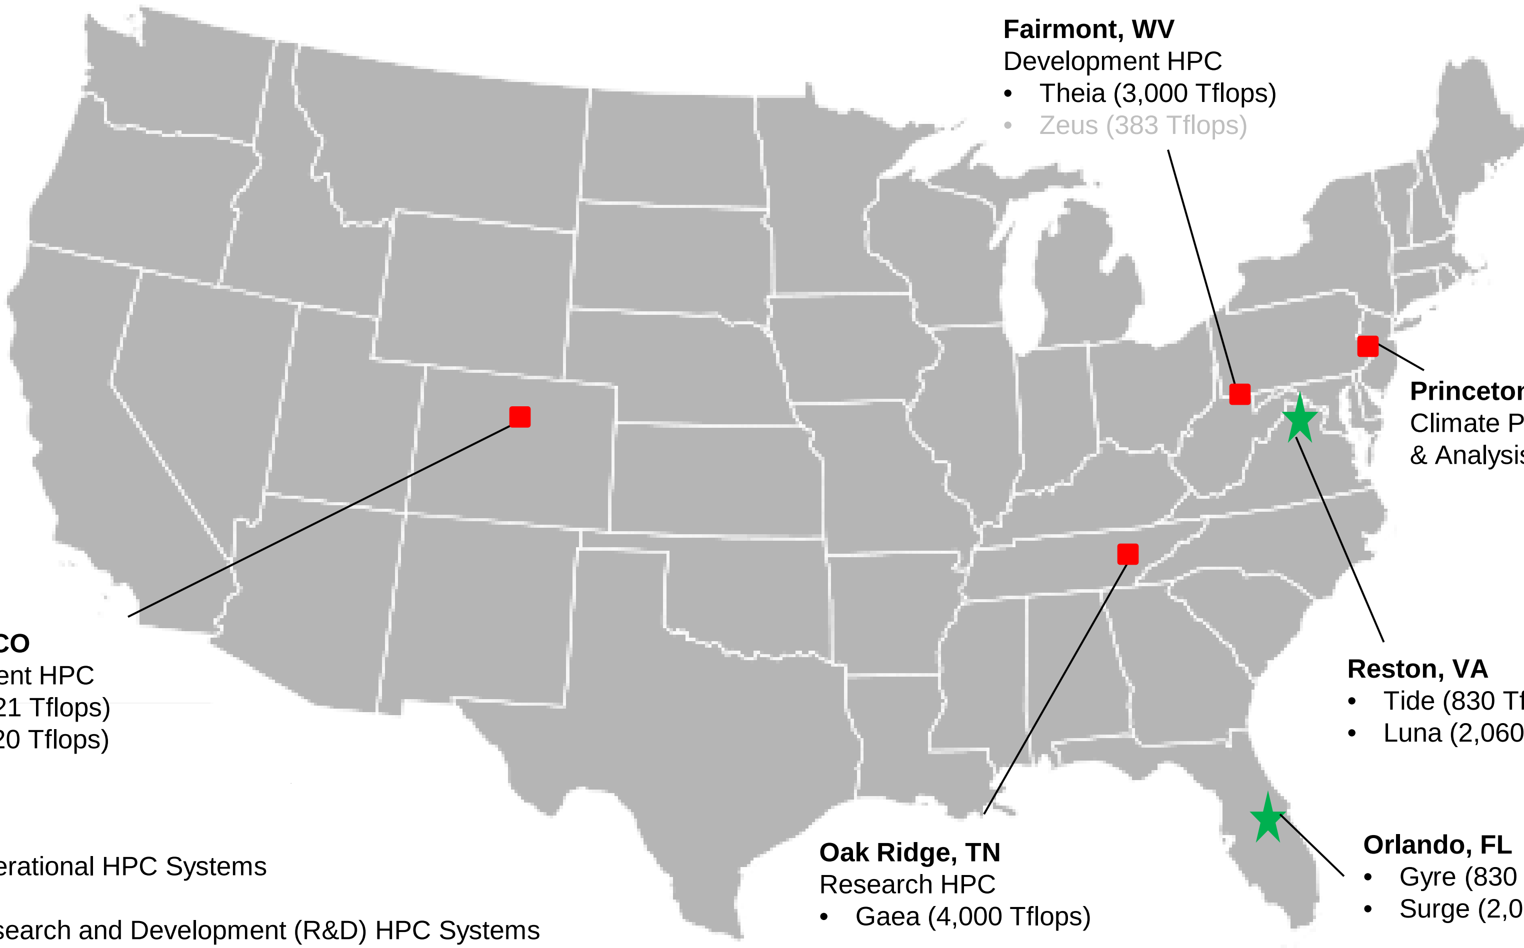

Fairmont, WV
Development HPC
- Theia (3,000 Tflops)
- Zeus (383 Tflops)

Princeton, NJ
Climate Post-Processing & Analysis

Boulder, CO
Development HPC
- nJet (421 Tflops)
- xJet (720 Tflops)

Reston, VA
- Tide (830 Tflops)
- Luna (2,060 Tflops)

Oak Ridge, TN
Research HPC
- Gaea (4,000 Tflops)

Orlando, FL
- Gyre (830 Tflops)
- Surge (2,060 Tflops)

★ Operational HPC Systems

■ Research and Development (R&D) HPC Systems

R&D Supercomputing

N-Wave Network

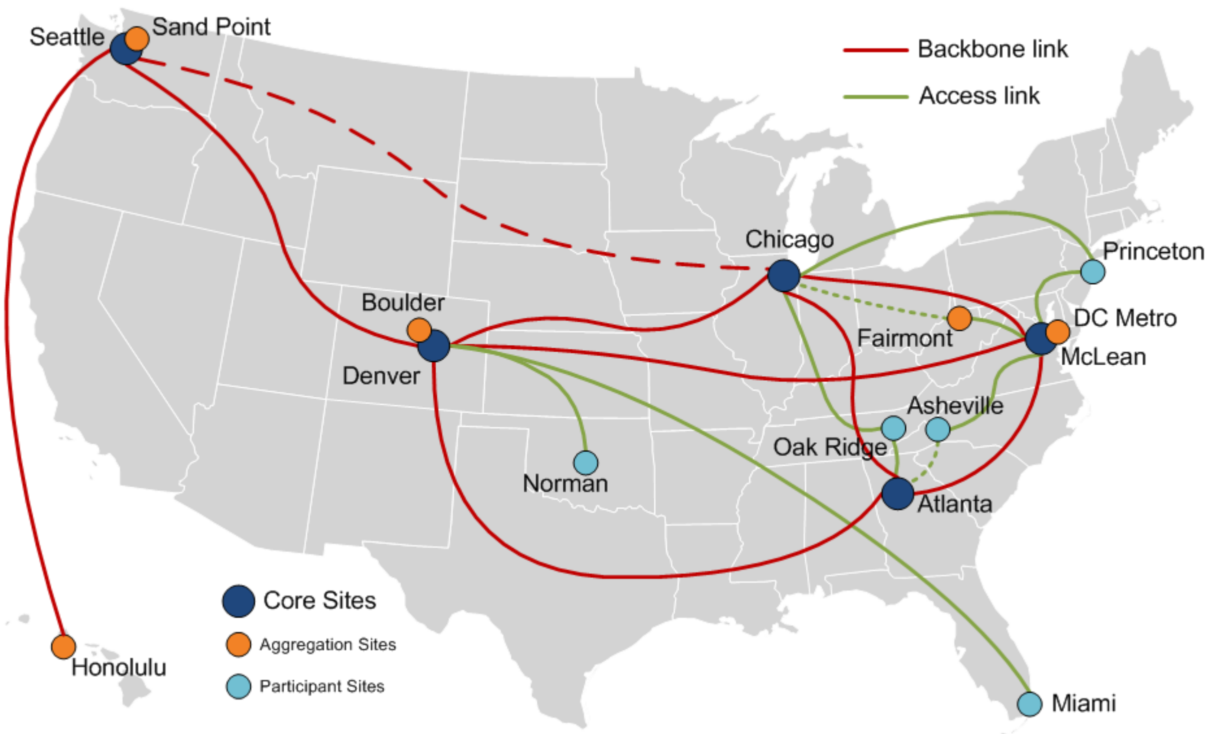

What is N-Wave?

- NOAA high bandwidth, low latency network connecting scientists to HPC
- Built using multiplexed fiber-optic links supplied by national Research and Education network community including Internet2 and university run regional network consortiums

SENA Project

Software Engineering for Novel Architectures

- Prepare codes for future production architecture
  - Monitor evolution
- Maintain codes in a way that subject matter experts can still work with the code
- Monitor evolving standards
- Codes should still be viable for current architectures
  - Performance is expected to increase on across new and old architectures
- Develop expertise within NOAA

Computational Profile

- Algorithms generally process weak scalability
- Physics components have predictable data dependencies associated with grids
- Intrinsic variability at all timescales from minutes to millennia
  - All space scales from microbes to megacontinents
- Adding processes and components improves scientific understanding
  - This complexity implies lots of diagnostic I/O
  - New physics and higher process fidelity at higher resolution
  - Ensemble methods to sample uncertainty
- Modeling requires long-term integrations of weak-scaling, I/O and memory-bound models of increasing intensity

Generating Parameterizations From High-Resolutions

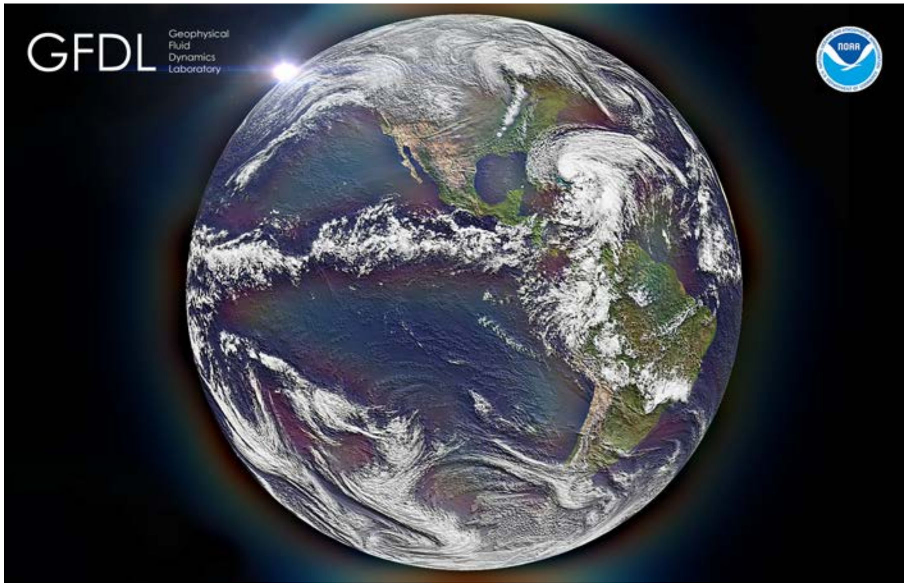

(Courtesy: S-J Lin, NOAA/GFDL).

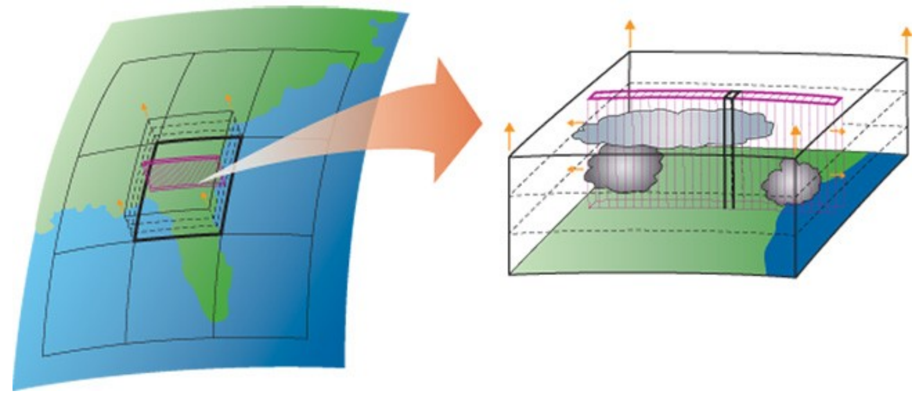

(Courtesy: D. Randall, CSU; CMMAP).

- Global-scale Cloud Resolving Models (e.g 7 km simulation on the left) and even super-parameterization using embedded cloud models (right) remain prohibitively expensive.
- Explore the use of machine learning (using GCM-resolution predictors) to emulate columns of a cloud field.

NOAA's Model Production Suite

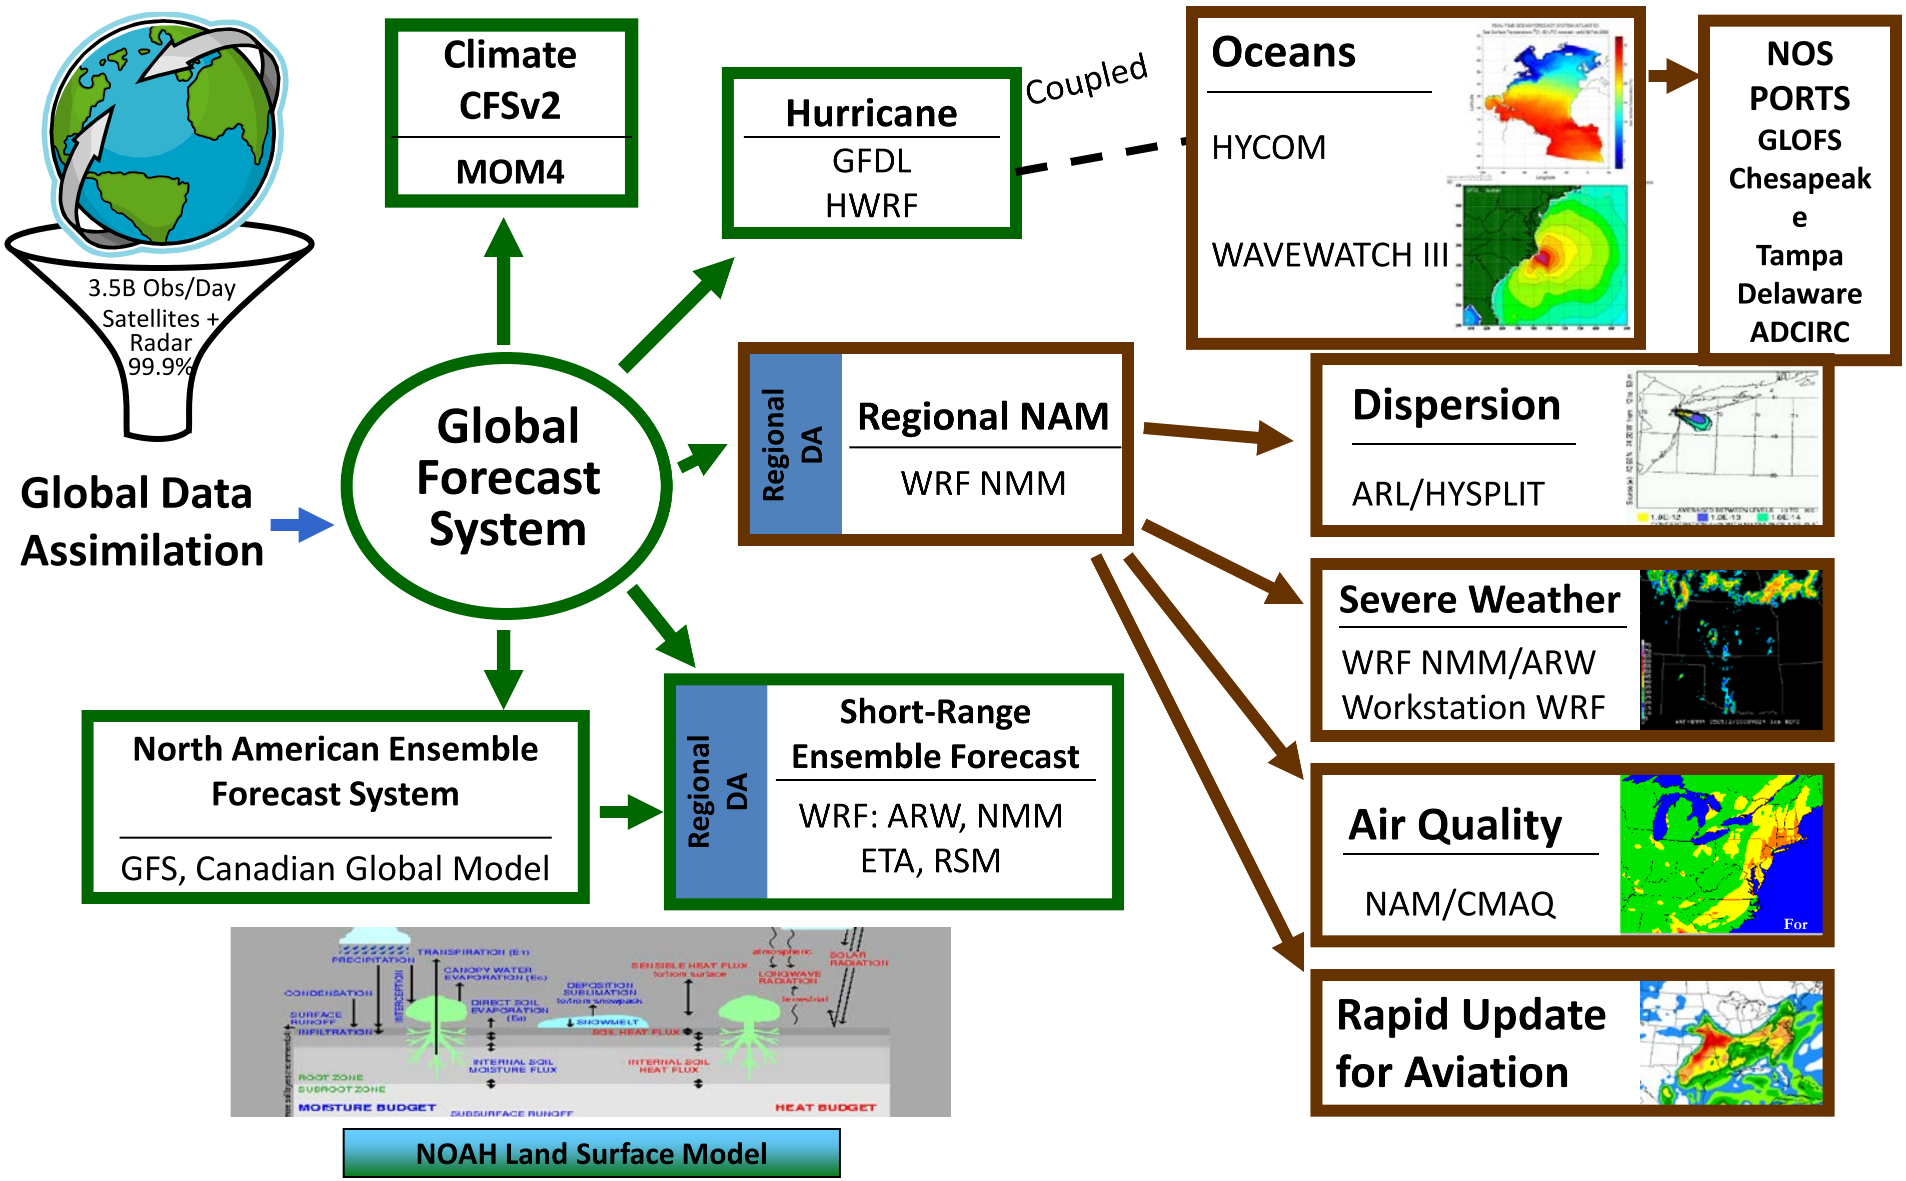

R&D Workflow

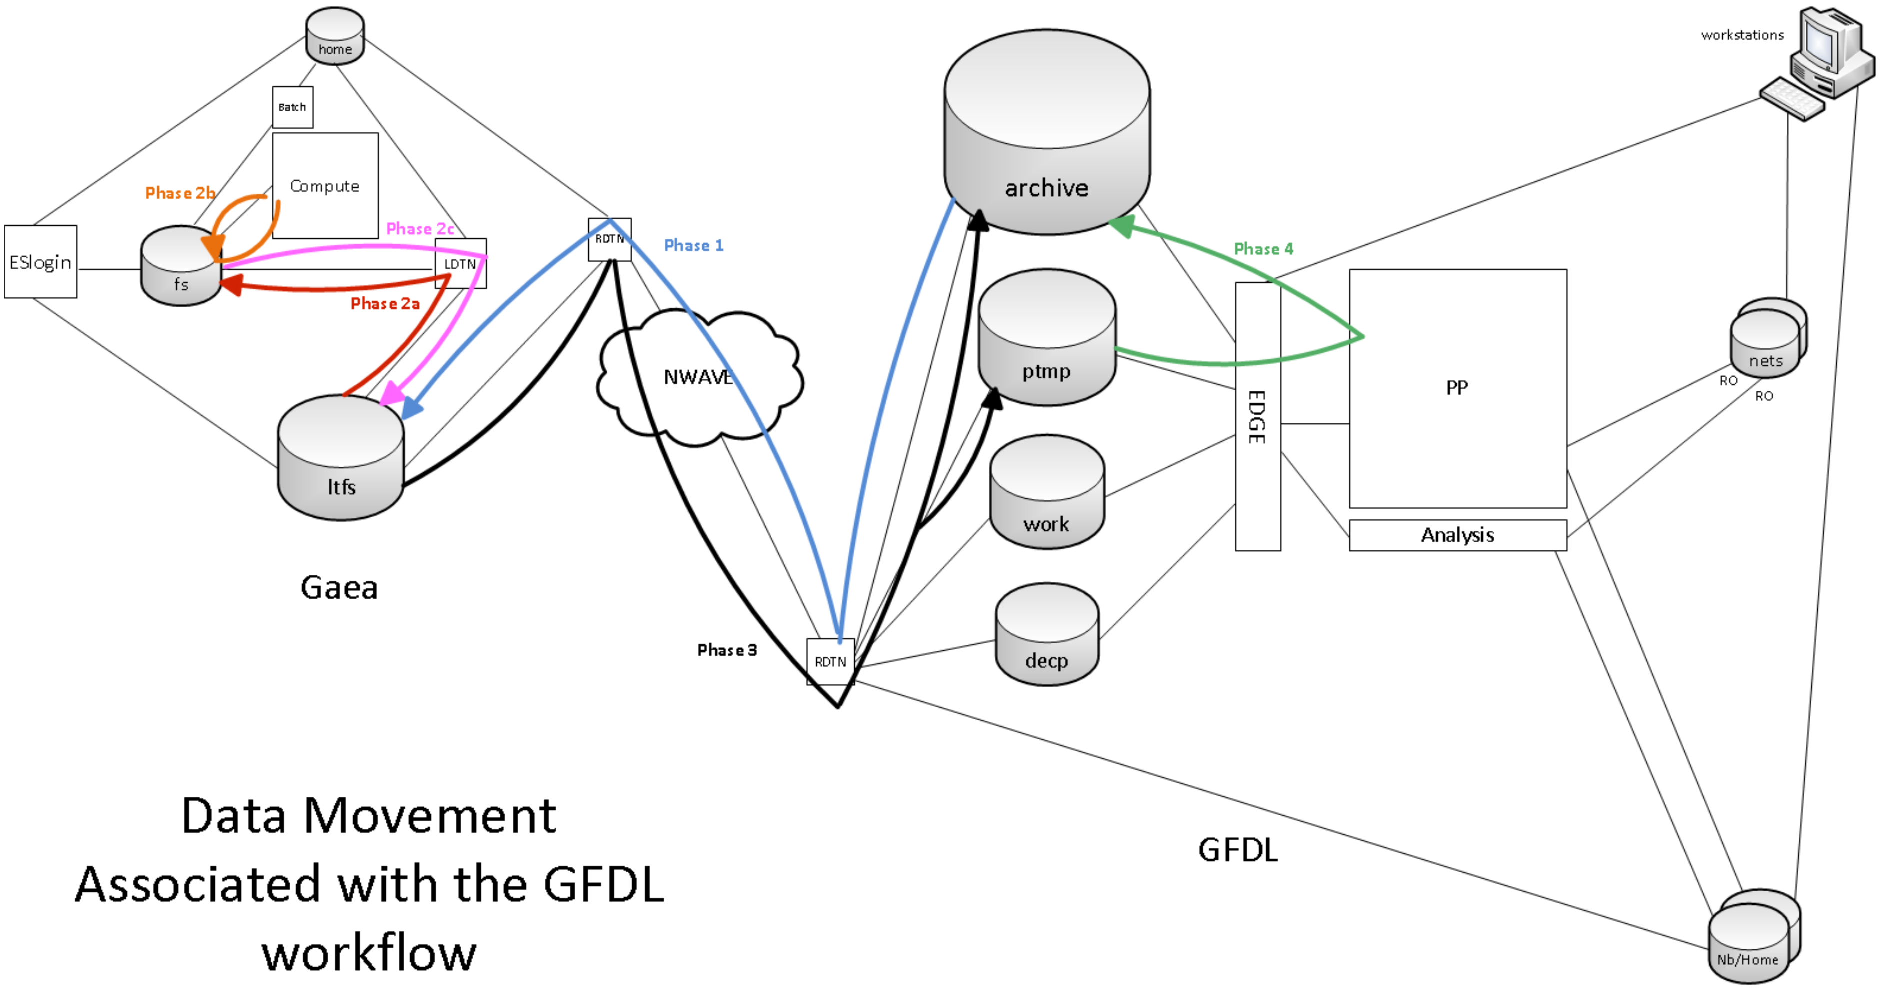

Data Movement
Associated with the GFDL
workflow

Challenges

HPC Reliability

- Requires data integrity
  - Complexity in components increases risk to data
  - Build in safeguards against silent data corruption
  - How do you know if you have data integrity? Check!
- HPC Systems are growing more complex and correspondingly have more failure points
- Users and support staff need to know what failed and where it failed to accurately continue the experiment
- Automated error detection and handling needs to be built into the systems, schedulers and user workflows

Challenges

Programming

- Technology directions are challenging even experienced users
  - HPC and analysis will have to rely on parallelism to keep up – which is not always a natural fit
  - Error handling cannot currently be left to the underlying infrastructure and is left to the programmer

- Fine Grain Computing (the next generation) offers no relief
  - Converting existing codes to support accelerators will take time and significant effort – for limited initial reward
  - The software environment for accelerators changes rapidly
  - Data movement to memory remains a bottleneck
    - Efficient use of MCDRAM is hard for most of our applications

- Waiting is not the answer! Start testing codes on new platforms now.

Challenges
Programming

- Lack of standards across divergent architectures
  - OpenMP and OpenACC are gaining momentum
- Large Amount of Legacy Code
  - Do we adapt or rewrite
- Access to developmental platforms
- Uncertainty of performance gains
  - The optimizations will still benefit our traditional architectures

Challenges

Data Movement

- An adaptable and capable network is required
- Moving and processing data entails a large number of tools and protocols
  - Not all tools provide adequate verification methods
  - We have found that a successful transfer is not always successful
- A need for a uniform method for transferring data
  - There are no one-size fits all tools, wrapping transfers with verification and retry ability has been required
- Required changes to NOAA's automated workflow
  - End-to-end data checksums
  - Error handling for many edge cases needed to be built in

Cloud Computing

- Cloud technologies are becoming the norm within HPC architectures
  - Image management and user environment management have been adopting these methods and will continue to grow into the future
  - Environment preservation is the most interesting of these use cases
  - The industry is driving towards hybrid HPC/Big Data platforms
    - Potential for greater user adoption as these software stacks converge

Thank You

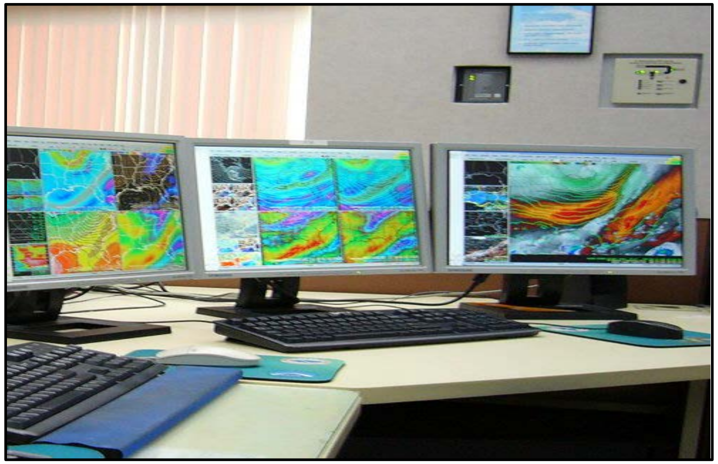

frank.indiviglio@noaa.gov