Empowering Earth System Science Research: Federated Data and Compute for Earth System Predictability

Thomas Hauser

Director, Computational and Information Systems Lab, NSF NCAR

December XX, 20XX

This material is based upon work supported by the NSF National Center for Atmospheric Research, which is a major facility sponsored by the U.S. National Science Foundation under Cooperative Agreement No. 1852977.

NSF NCAR Provides Research, Facilities & Services

weather • water • space weather
climate • air quality
Earth system

- Collaborative Research
- World-class Computer Models
- Supercomputing

- Data Archive
- Observing Systems & Sensors
- Field Campaigns
- Education & Outreach

NSF NCAR
DIRECTOR
Everette Joseph

NSF NCAR Directorate

- ATMOSPHERIC CHEMISTRY OBSERVATIONS & MODELING LAB — **Pieternel Levelt**, Director
- CLIMATE & GLOBAL DYNAMICS LAB — **Jon Petch**, Director
- COMPUTATIONAL & INFORMATION SYSTEMS LAB — **Thomas Hauser**, Director
- EARTH OBSERVING LAB — **Allison McComiskey**, Director
- HIGH ALTITUDE OBSERVATORY LAB — **Holly Gilbert**, Director
- MESOSCALE & MICROSCALE METEOROLOGY LAB — **Gretchen Mullendore**, Director
- RESEARCH APPLICATIONS LAB — **Matthias Steiner**, Acting Director
- EDUCATION, ENGAGEMENT & EARLY-CAREER DEVELOPMENT — **Rebecca Haacker**, Director

LABS

CISL, a key CI provider for the Earth System Science Community

Transforming Science through Cyberinfrastructure

- Dual role for CISL
  - Research CI provider
  - CI research

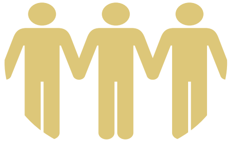

People, organizations, and communities

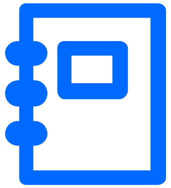

Data Infrastructure

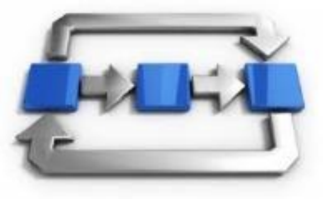

Software and Workflow Systems

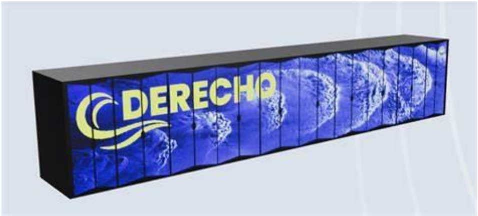

Computing Resources

Cloud Resources & Services

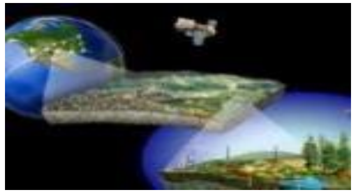

CI-Enabled Instrumentation

Gateways, Hubs, and Services

R&E Networks, Security Layers

Pilots, Testbeds

ACCESS Advancing Innovation

Coordination & User support

NSF NCAR-Wyoming Supercomputing Center (NWSC)

- Maximal energy efficiency of facility infrastructure
- Implementation of sustainable design practices for office and other personnel spaces
  - Harvesting of natural daylight and ventilation
  - Use of engineered, reclaimed, or recycled materials
- Conservation of potable water
  - Specialized cooling tower equipment
  - Low-flow plumbing fixtures
- **LEED certification Gold**
- Utilize the region's cool, dry climate to minimize energy use
- Utilize the waste heat generated from computing to provide:
  - Building heating
  - Generator block heating
  - Reheating of coils

- *Expandability*
- *Flexibility*
- *Upgradeability*
- *Sustainability*

- Focus on the biggest losses
  - Compressor based cooling
  - UPS losses
  - Transformer losses
- Minimizing energy use makes good environmental and economic sense
- PUE
  - Current Operations 1.07 - 1.10

NCAR's HPC Systems

2017 — *2023*

SGI/HPE
4032 Nodes, ***145,152 Cores****, 313 TB total memory,* ***4.79 PFlop/s***
#21 Supercomputer in the world at debut, #109 presently

New Exhibit at NWSC

Delivery & Installation: 01/23
Acceptance Testing Complete: 5/23
ASD Project Concludes: 7/23
***Production*: 7/23**

Cray/HPE
2570 Nodes, ***323,712 CPU Cores****, 680 TB total memory,* ***3.5X performance vs Cheyenne***
328 NVidia A100 GPUs *providing 20% of overall performance,* ***19.87 PFlop/s*** *(projected)*

mid-2023

Earth System Science Leadership: NSF NCAR

World leader in **Earth system science <u>with</u> and <u>for</u> the community**

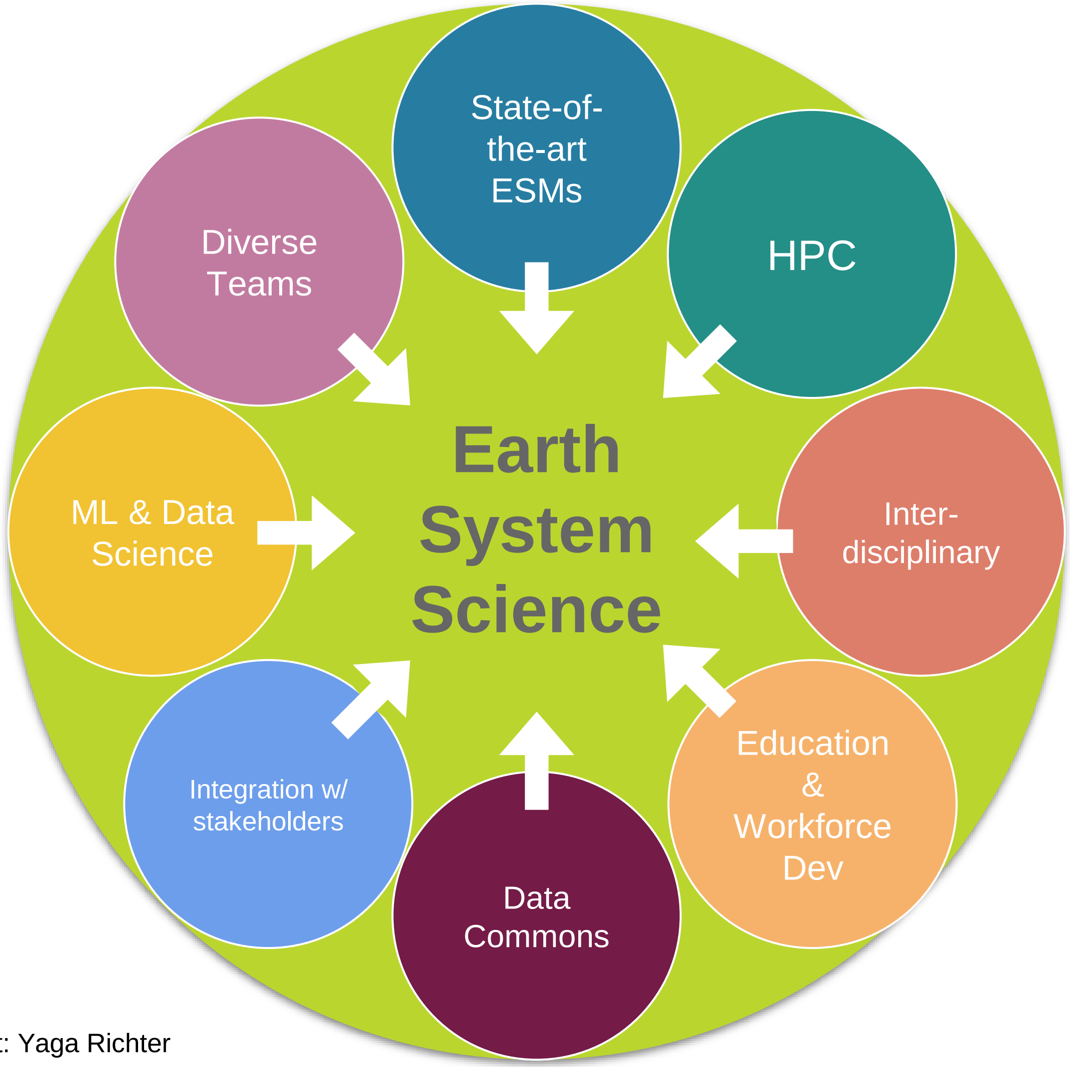

CISL as an enabler and leader through Cyberinfrastructure services, research, and development

Credit: Yaga Richter

What is Earth System Predictability Across Timescales?

Focused NSF NCAR Priority: in response to urgent calls to advance Earth system predictability research across scales

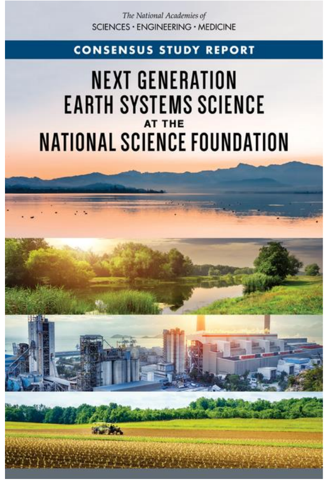

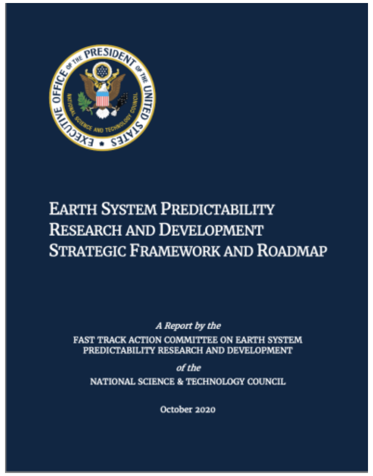

National Needs from Community Reports

- Foundational understanding & interdisciplinary/convergent research
- Advancement of ESP modeling & tools
- Computational advancements
- Reduction of observational gaps
- Partnerships across sectors & agencies
- Workforce development

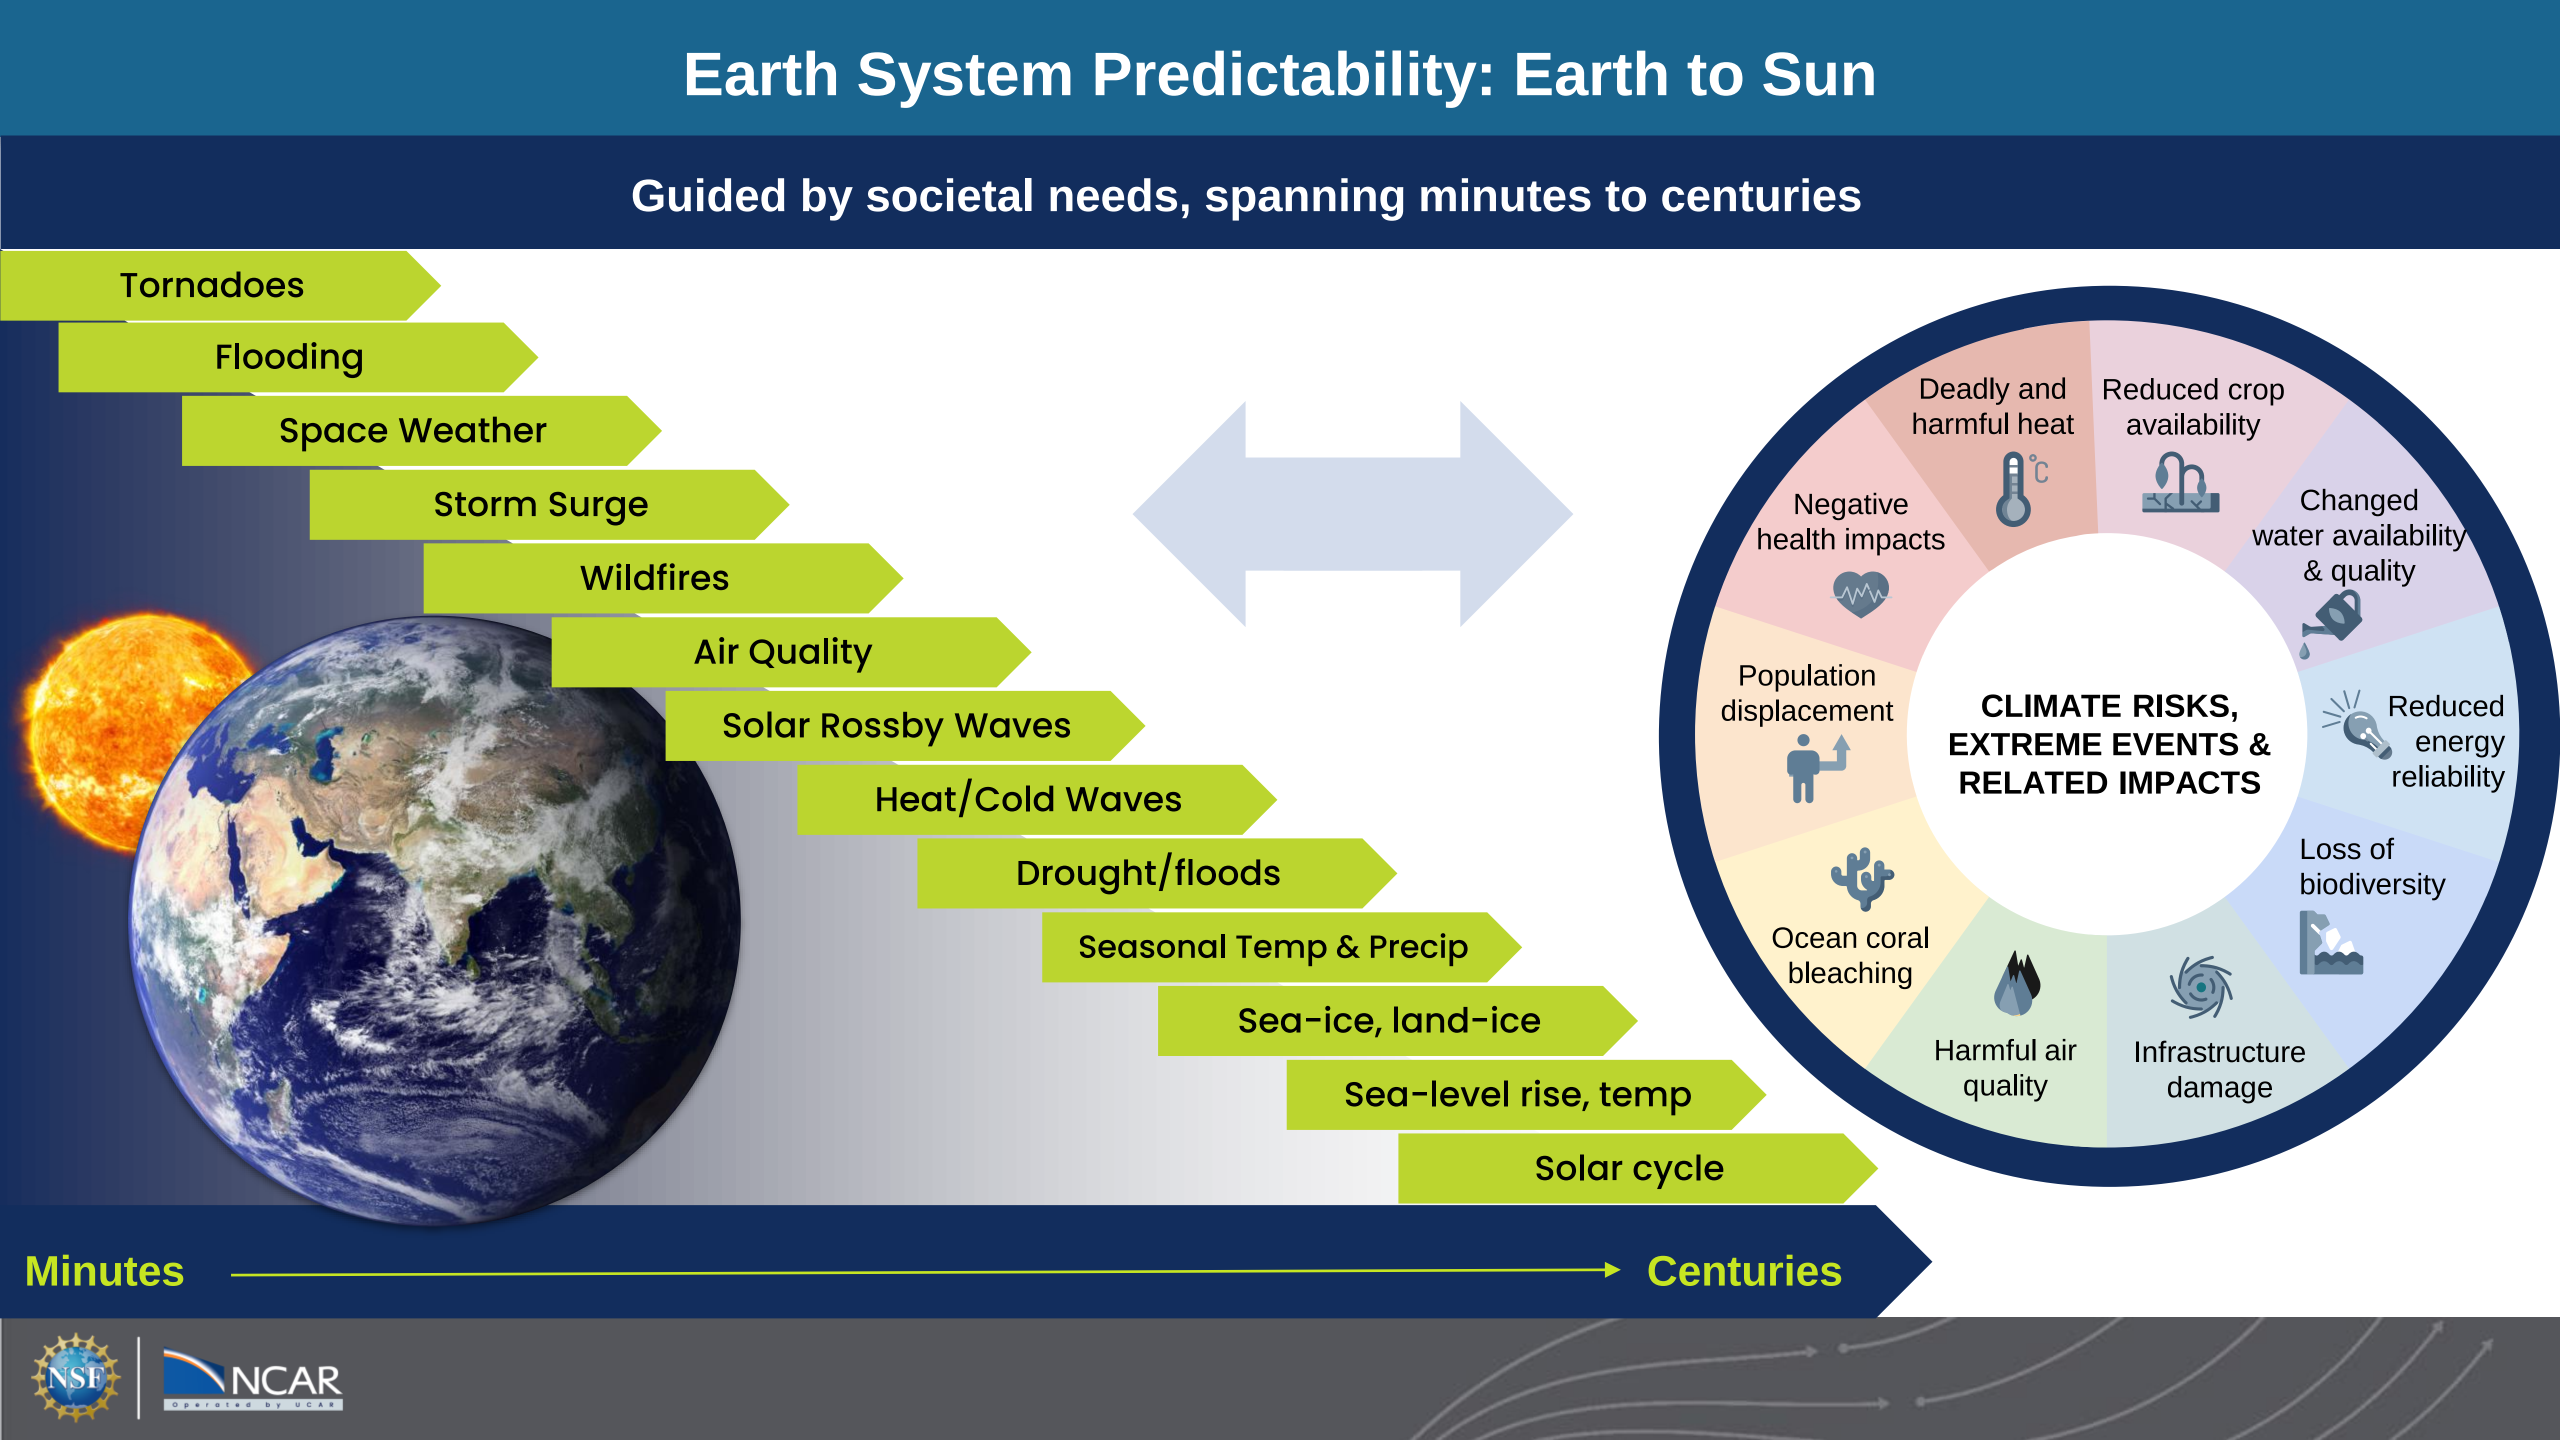

Earth System Predictability: Earth to Sun
Guided by societal needs, spanning minutes to centuries
Tornadoes
Flooding
Space Weather
Storm Surge
Wildfires
Air Quality
Solar Rossby Waves
Heat/Cold Waves
Drought/floods
Seasonal Temp & Precip
Sea-ice, land-ice
Sea-level rise, temp
Solar cycle
Minutes
Centuries
CLIMATE RISKS, EXTREME EVENTS & RELATED IMPACTS
Deadly and harmful heat
Reduced crop availability
Changed water availability & quality
Reduced energy reliability
Loss of biodiversity
Infrastructure damage
Harmful air quality
Ocean coral bleaching
Population displacement
Negative health impacts
NSF
NCAR

ESPP: Key Science Questions

How do **coupled interactions give rise to predictability** across the Earth system (atmosphere, ocean, land, sea ice, terrestrial, geospace, and marine ecosystems)?

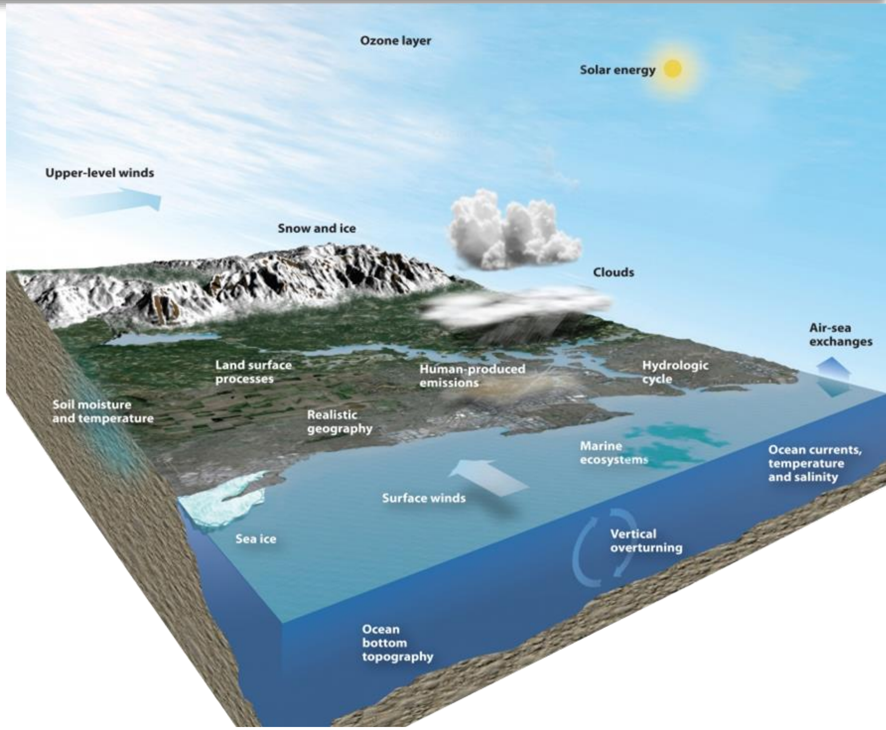

What is the **needed observational and model complexity** to increase prediction and projection skill?


NSF/NCAR C-130 Right Side Instrumentation

How do the **interactions across spatial scales impact** predictability and predictions of Earth and Space weather and climate?

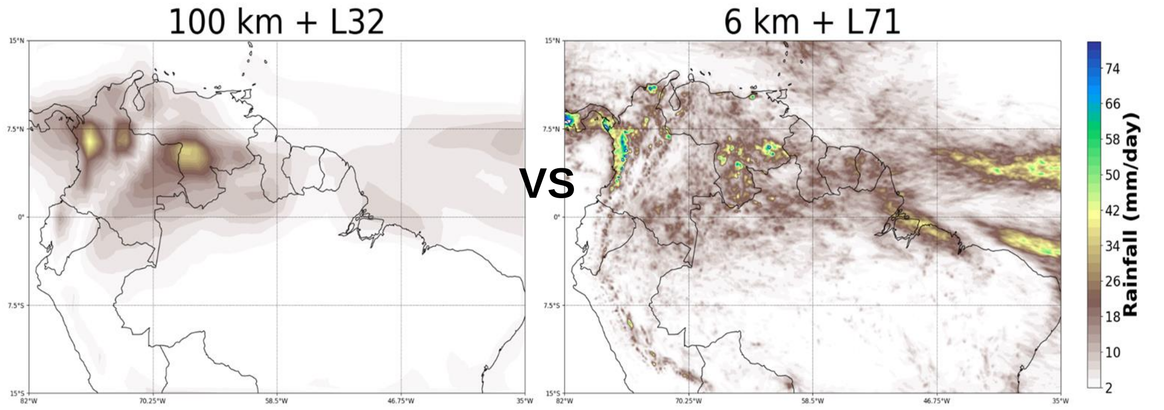

NSF NCAR'S Modeling & Observational Ecosystem

Enables interdisciplinary Earth system predictability research across

- global to local scales
- timescales from minutes to decades
- array of complex Earth system interactions

Fosters Community

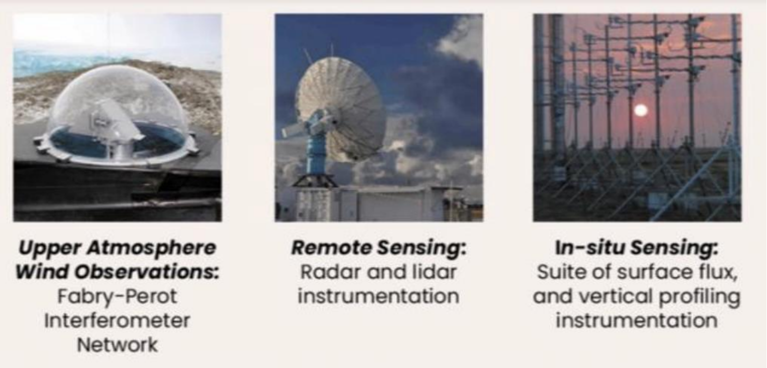

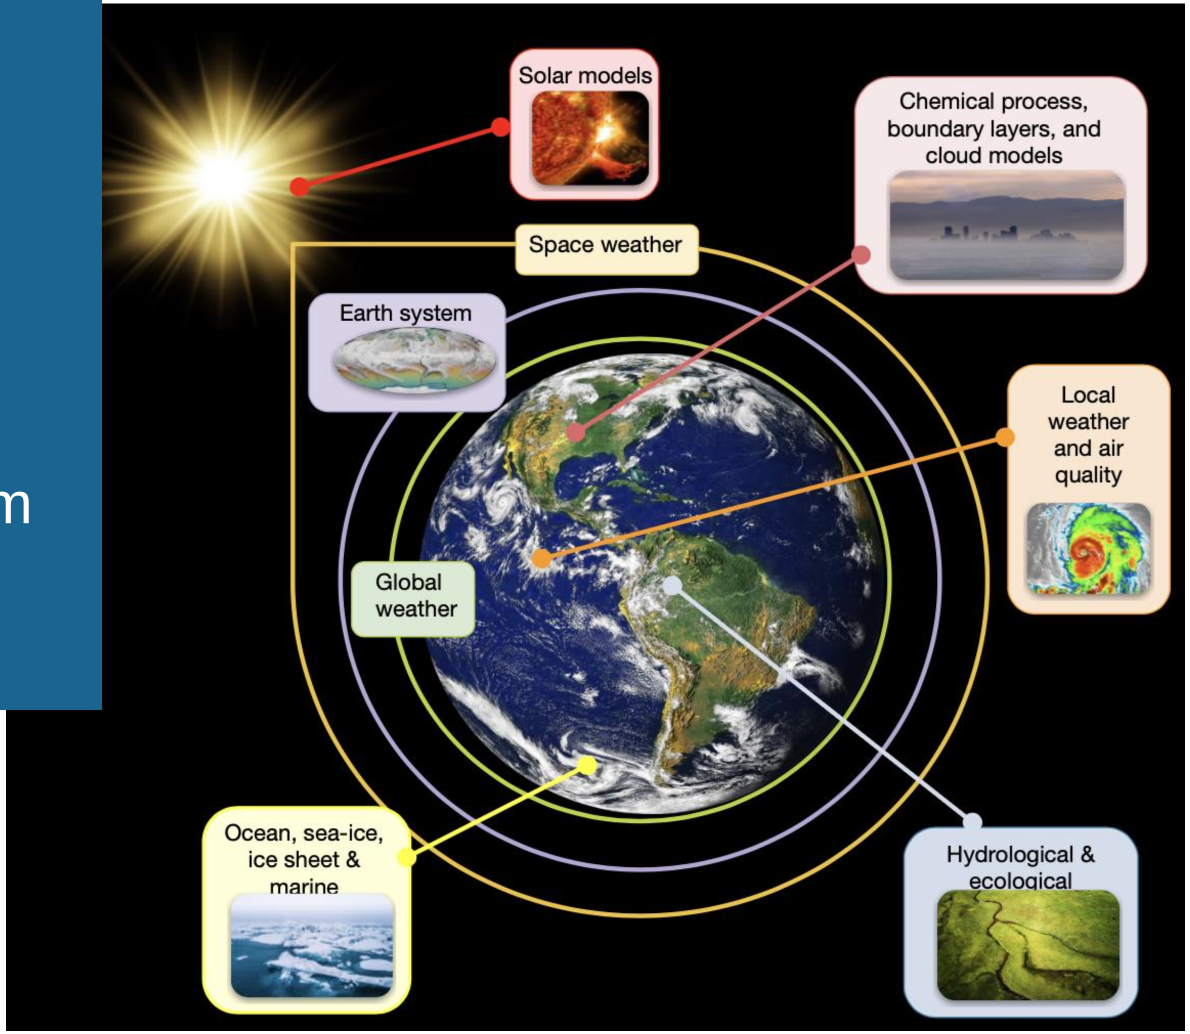

NSF NCAR Computational Science Strategy

- Enabling and supporting the community through a digital Earth System Science lab
- Democratizing Research Capabilities in support of Earth System Science
- Integration and Federation with the National and International Cyberinfrastructure (NAIRR, National Discovery Cloud for Climate, …)

Image created by DALL-E, OpenAI, 2023

NSF NCAR Community Software Facility

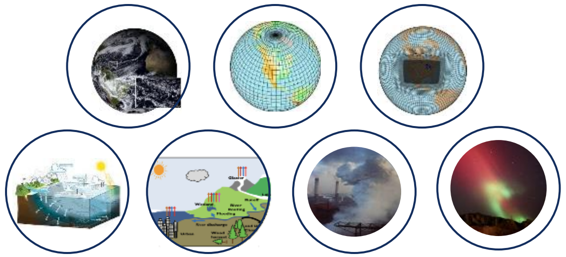

NSF NCAR Compute Infrastructure

Translational Computational Science

NSF NCAR Research Data Commons

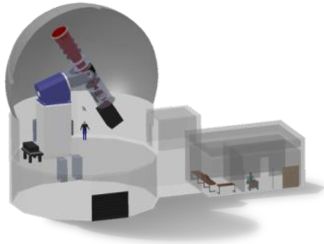

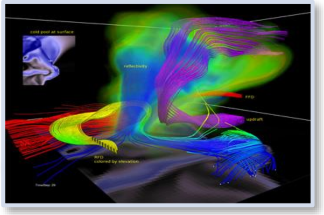

The changing Cyberinfrastructure Landscape – Disruption Ahead

Broad, growing, and diverse user community and needs

Disruption along multiple dimensions

Evolving Science Requirements

- Extreme scales & complex workflows
- Data-driven & AI/ML revolution

Changing Values

- Ease of use and access
- Time to science
- Transparency, robustness, security, trust
- Societal benefits

Evolving Technology Landscape

- End of Moore/Dennard Scaling impacts all aspects of computing
- Importance of software and its sustainability
- Novel paradigms
  - Cloud computing
  - Edge computing

Opportunity to reinvent!

Credit: Manish Parashar

NSF NCAR's computational science strategy - Models and Diagnostics

Community Software Facility

- **Models as scientific instruments**
  - Opportunity to rethink how we develop and support our models
    - Unified support across NSF NCAR
    - Co-Design with computational facility
    - Technology evaluation - Domain Specific Abstractions (DSA)
    - Verification approaches (non bitwise)
    - Standardization and APIs (Destination Earth, EVE, Earth 2)
  - Democratize access to our models - Ease of Use
    - Science gateways
    - Cloud technologies
- AI/ML disruption

How GPUs are accelerating solar physics…

Solar simulations need faster throughput *and* higher resolution. MURaM OpenACC will help meet these requirements.

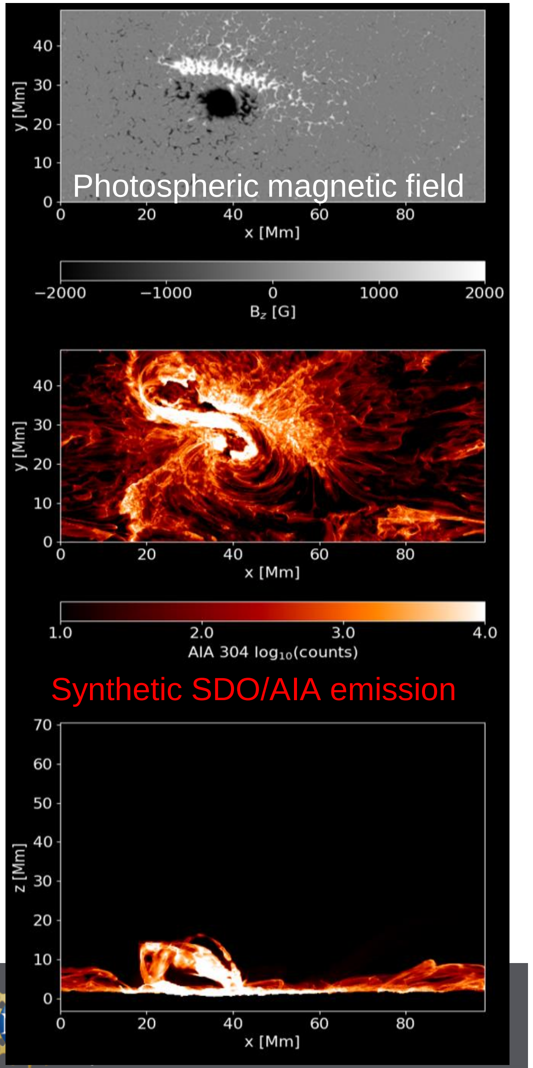

Simulation of a solar flare resulting from sunspot collision

- **Max Planck University of Chicago Radiative MHD (MURaM) models the solar atmosphere from upper convection zone to lower solar corona**
- **Goals & Actionable Science for MURaM-OpenACC:**
  - **Short-term:** Solar models capable of running models at the resolution of DKIST telescope observations.
  - **Long-term: Better prediction of space weather events** using data-driven models of solar eruptions
- **Refactoring of MURaM for GPU using OpenACC**

DKIST Telescope

MURaM OpenACC project is an HAO/CISL collaboration with the University of Delaware and the Max Planck Institute for Solar System Research & Lockheed

Diverse Team Success Story - MURaM

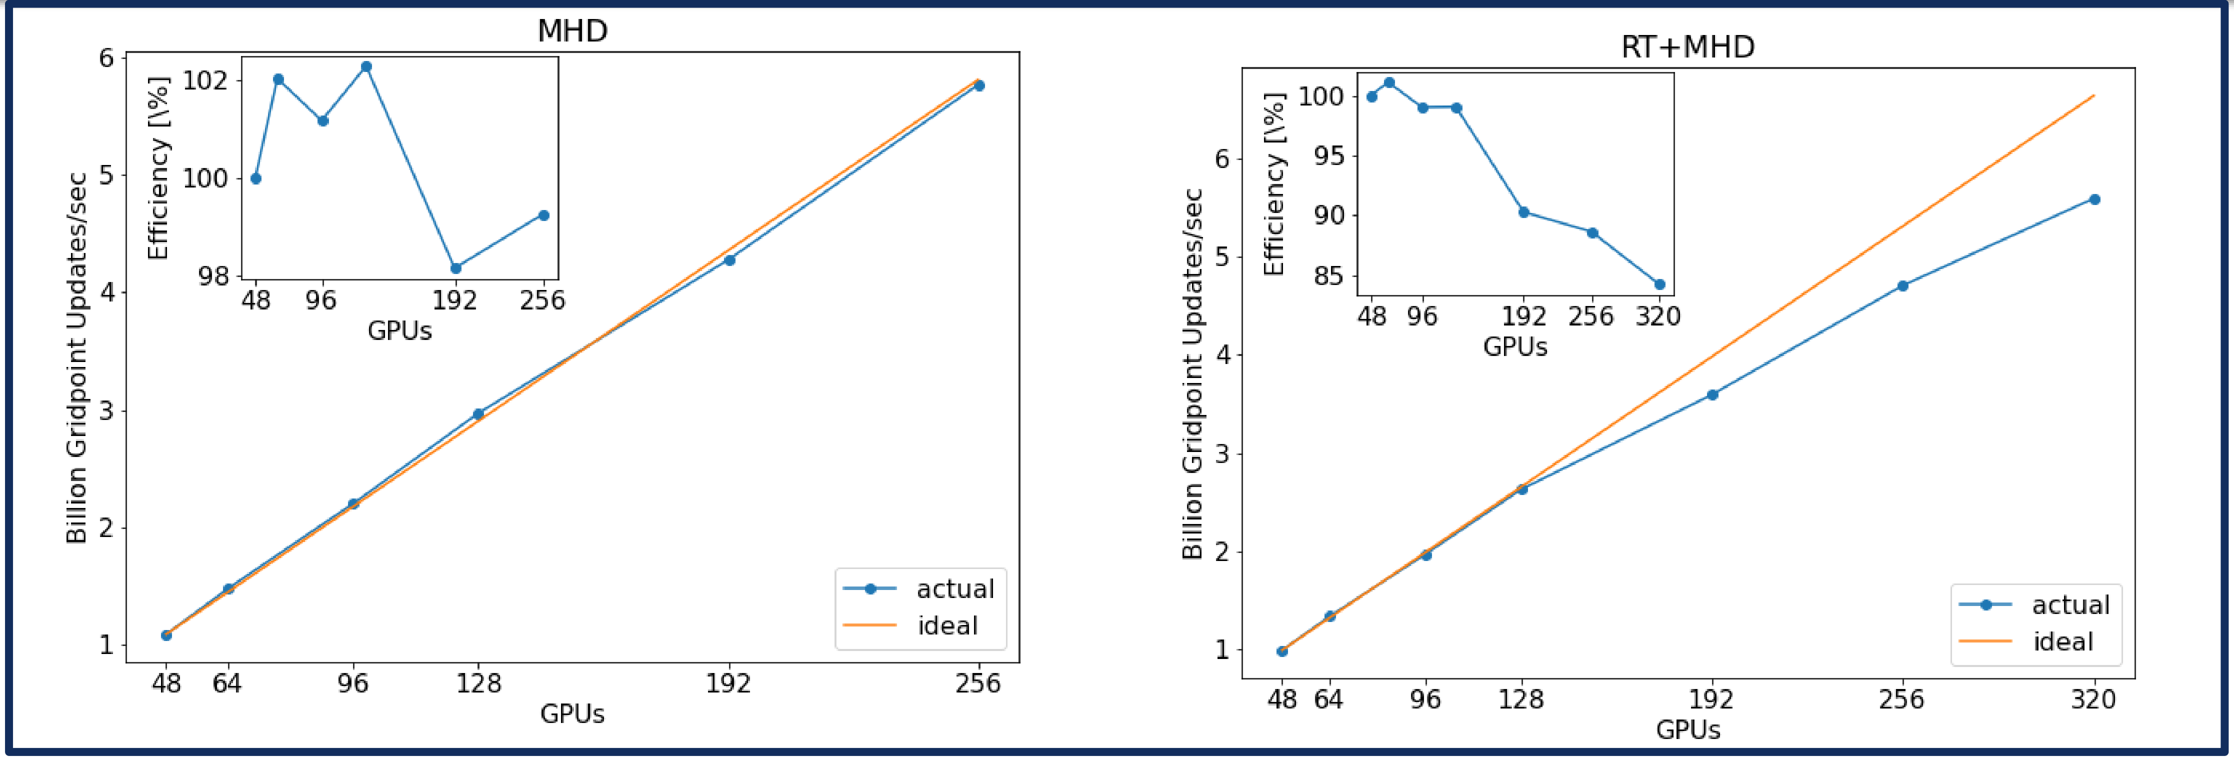

Strong scaling for MHD and MHD+RT physics within MURaM

CPU scaling for same case falls off above 16K cores, cannot catch up!!

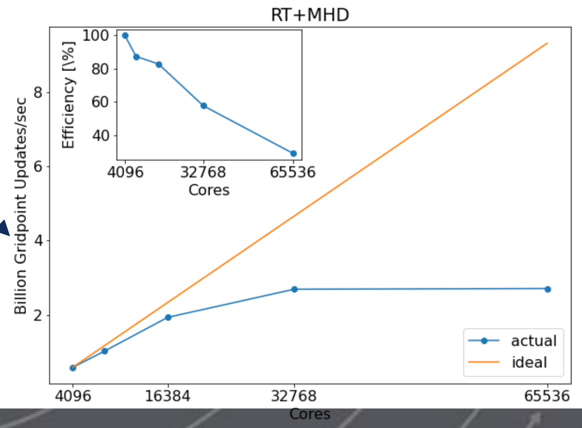

***MURaM can run the same problem efficiently on 192+ of Derecho's GPUs than currently possible on any number of CPU cores**, courtesy Matthias Rempel et al.*

NSF NCAR's computational science strategy - Applied Research

Translational Computational Science

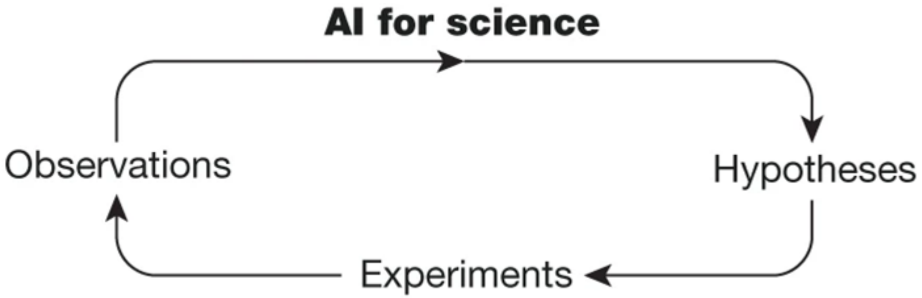

AI will reshape discovery by accelerating research at any stage of this process

Wang, H., Fu, T., Du, Y. et al. Scientific discovery in the age of artificial intelligence. Nature 620, 47–60 (2023). https://doi.org/10.1038/s41586-023-06221-2

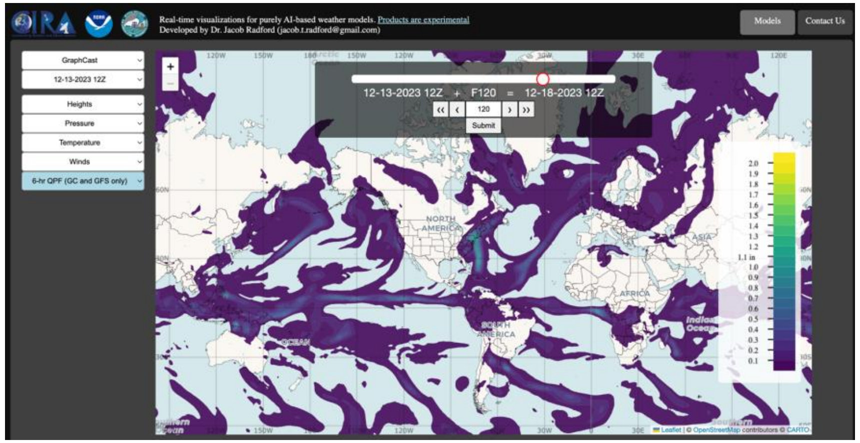

AI2ES

Funded by the National Science Foundation

- Generative AI will transform how we do science
  - Google, NVIDIA, Microsoft, ECMWF, and NOAA are pivoting to supporting AI for Earth Science
  - AI based atmosphere and ocean models are advancing rapidly
  - Copilots for software engineering and writing
  - Data analysis and management
- We need an AI-savvy workforce
- Innovate across all areas of the computational science strategy
- Extending our capability through connections and partnerships
- Student and PostDoc programs (collaboration with EDEC)

LEAP

Learning the Earth with Artificial Intelligence and Physics (LEAP)

SIParCS – Summer Internships in Parallel Computational Science – BOULDER • CO

NSF NCAR's computational science strategy - AI

ARCO Community Datasets

High-Resolution Models
- CONUS 404
- FastEddy

Climate Ensembles
- CESM Large Ensemble
- CESM PPEs

Field Observations
- HOLODEC cloud imager
- In-situ and remotely sensed soundings
- Radar and satellite

Accelerating and Emulating Earth System Models

AI Digital Twins
- Global (RDA ERA 5)
- Regional (High-Res)

Parameterization Emulators
- Bin Microphysics
- Surface-Air Flux

Parameter Estimation
- Community Land Model
- Bayesian Microphysics

Adding Value for Community Users

Scalable Analysis
- Feature Tracking
- Explainable AI
- Uncertainty Quantification

Interactive Visualization
- Real-time cloud workflows
- Low-latency web graphics

Convergence Research
- Impact Estimation
- Social Science Integration

Community standards, APIs, and federated resources

NSF NCAR's computational science strategy - Data

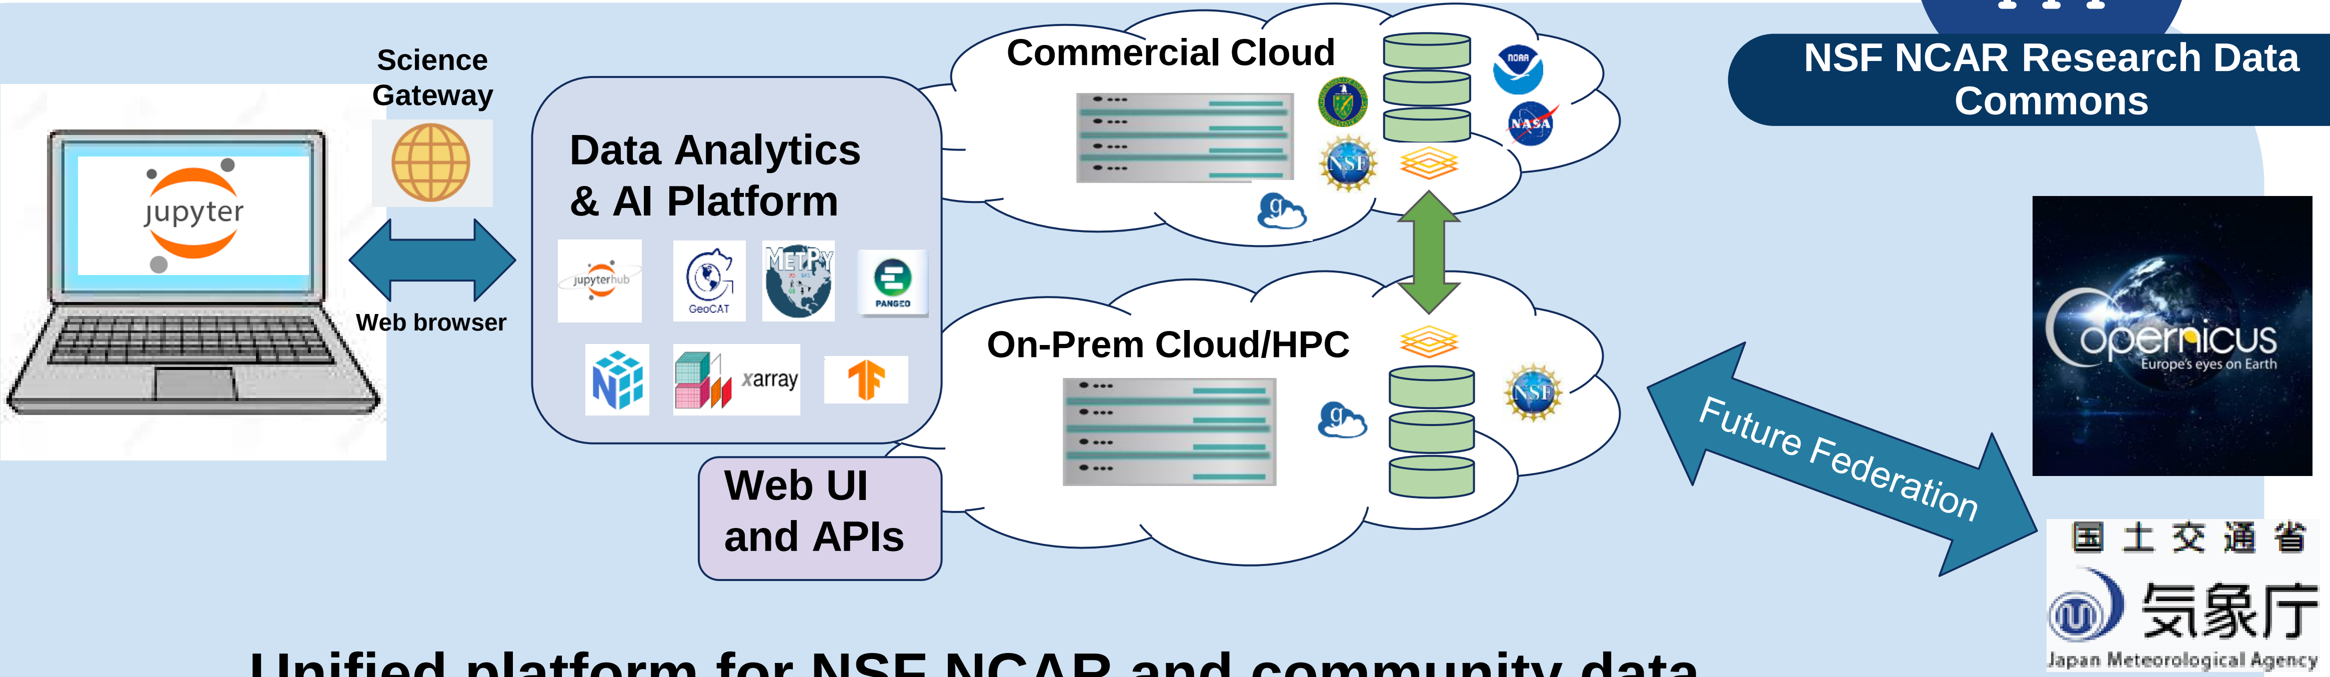

Unified platform for NSF NCAR and community data

- Sustainable business model
- Support the Geo community with data management services
- Cloud platform services for data analysis, AI/ML, and visualization
- AI disruption (How Large Language Models Will Disrupt Data Management https://doi.org/10.14778/3611479.3611527)
- Data Federation with other international centers (DKRZ, ECMWF, …)

Future NSF NCAR Research Data Commons

Existing - Loosely connected data systems

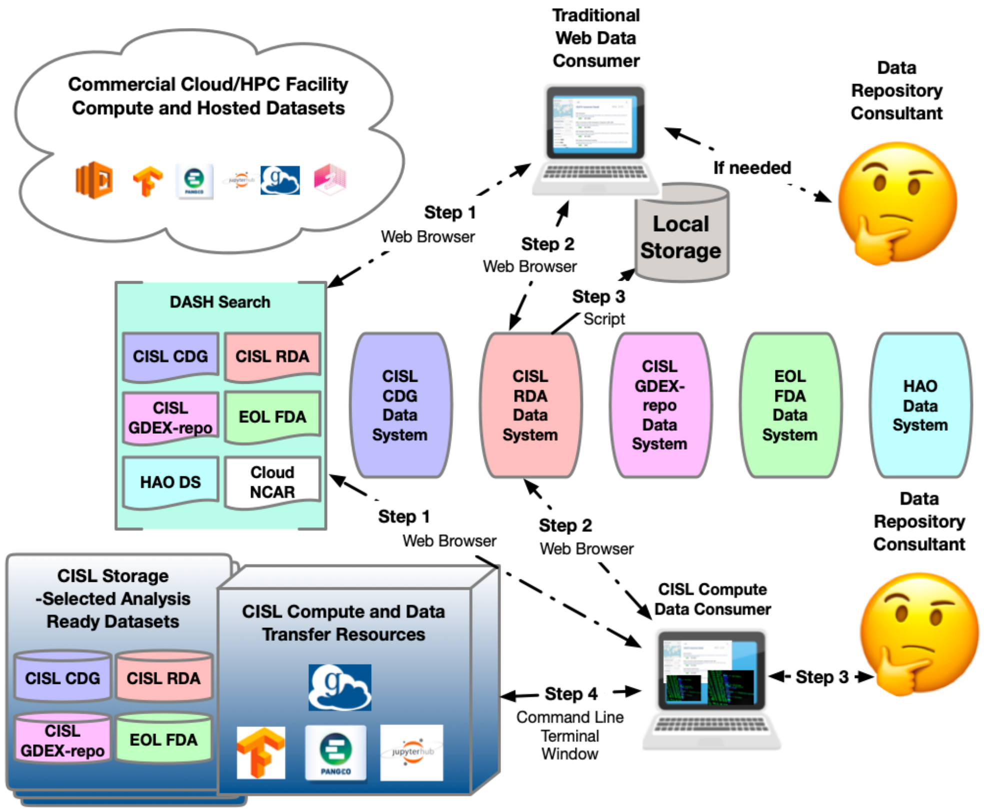

NSF NCAR GDEX - Integrated Research Data Commons

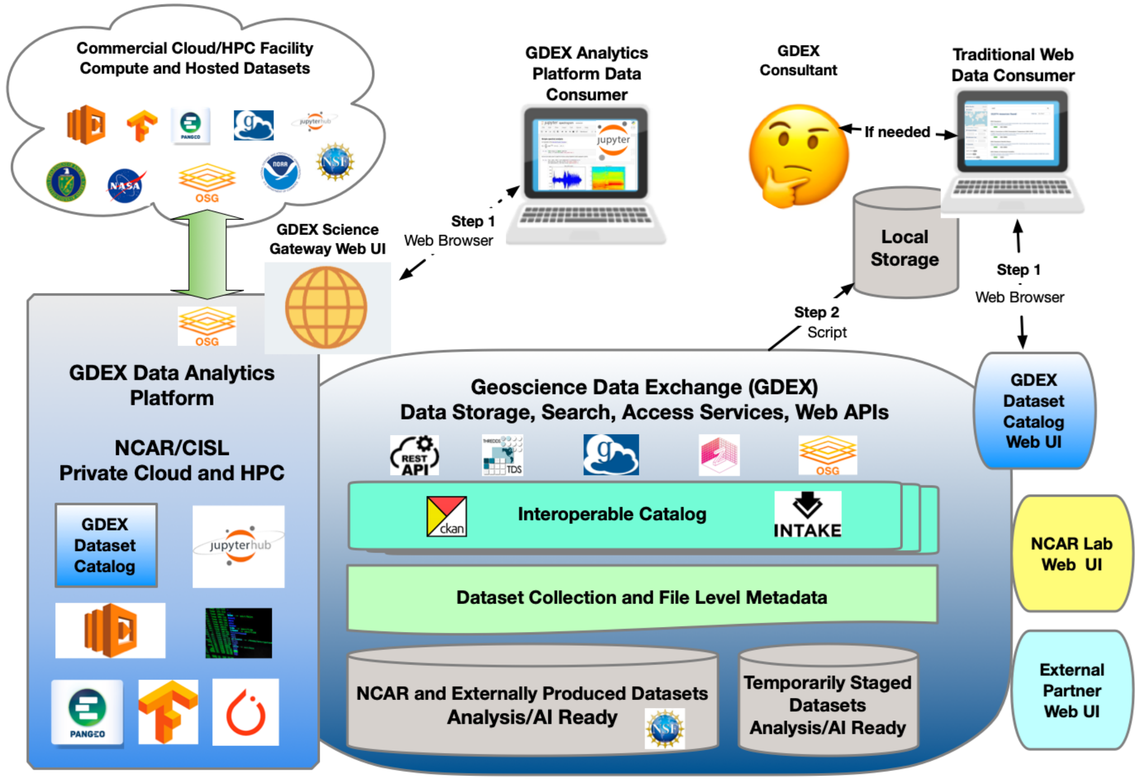

Grossman, R.L. Ten lessons for data sharing with a data commons. *Sci Data* 10, 120 (2023). https://doi.org/10.1038/s41597-023-02029-x

- Risk of doing nothing
- **Culture change is difficult**
  - Team Science
  - Software engineering practices
  - Workforce composition
- Partnerships

Questions?

Thank You !

Contributions:
Jadwiga (Yaga) Richter, NCAR DO
Glen Romaine, NCAR DO
Tricia O'Keefe, NCAR DO
David John Gagne, CISL
Doug Schuster, CISL