



Accelerating Earth system science knowledge and discovery through computational science

Thomas Hauser


thausen@ucar.edu

Director, Computational and Information Systems Lab

October 25th, 2024

What is the National Center for Atmospheric Research?


- NCAR was established by the National Science Foundation in 1960 to provide the university community with world-class facilities and services that were beyond the reach of any individual institution.
- NCAR's eight labs and programs collectively cover a breadth of research topics in Earth system science, from the effects of the Sun on Earth's atmosphere to the role of the ocean in weather and climate prediction, as well as supporting and training the next-generation of Earth system scientists.



ATMOSPHERIC CHEMISTRY OBSERVATIONS & MODELING (ACOM)

ACOM research focuses on advancing understanding and predictive capability for atmospheric composition and related processes.


ACOM



CLIMATE & GLOBAL DYNAMICS (CGD)

CGD develops understanding of the Earth system and is a leader of community-developed and community-owned models of the Earth system.

CGD



COMPUTATIONAL & INFORMATION SYSTEMS LABORATORY (CISL)

CISL is a world leader in supercomputing and cyberinfrastructure, providing services to NCAR, UCAR member universities, and the broader geosciences community.

CISL



EDUCATION, ENGAGEMENT & EARLY-CAREER DEVELOPMENT (EDEC)

EDEC is committed to inspiring, engaging and informing the public about the science going on at NCAR and in our broader scientific community.


EDEC



EARTH OBSERVING LABORATORY (EOL)

EOL provides state-of-the-art atmospheric observing systems and support services to the university-based research community for climate, weather, and related Earth system research.


EOL



HIGH ALTITUDE OBSERVATORY (HAO)

HAO conducts fundamental and applied research in solar-terrestrial physics using observational, theoretical, and numerical methods. Research at HAO extends from the solar core to the surface of the Earth.

HAO



MESOSCALE & MICROSCALE METEOROLOGY (MMM)

MMM advances the understanding of regional and fine-scale aspects of weather and climate and applies this knowledge to benefit society.

MMM



RESEARCH APPLICATIONS LABORATORY (RAL)

RAL conducts directed research and engineering to discover solutions to problems relevant to society and facilitates the transfer of our information, expertise, and technology to the public and private sectors.

RAL



NSF NCAR Facility Locations



Centergreen Campus



Foothills Lab Campus



NCAR-Wyoming Supercomputing Center



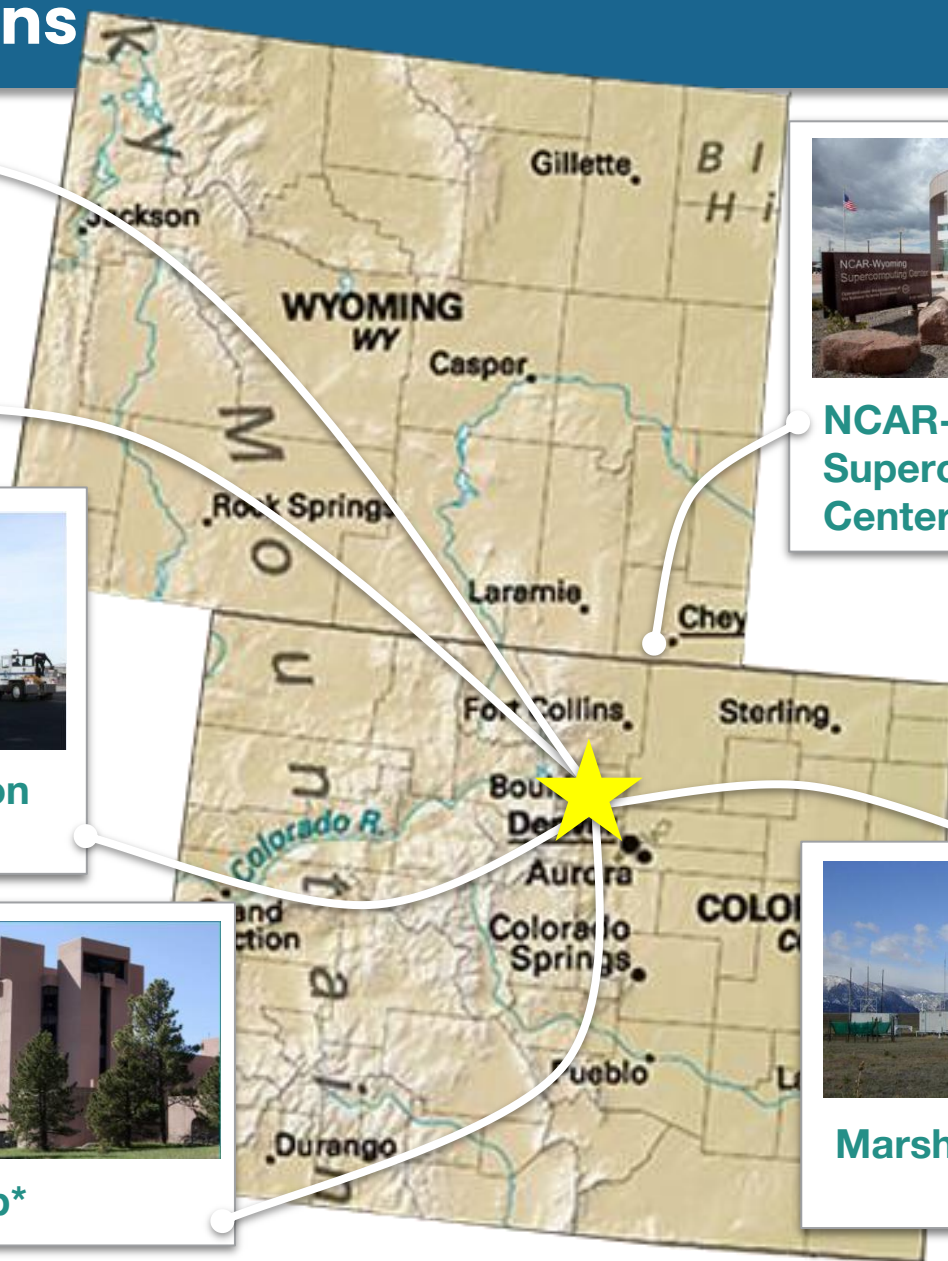
Research Aviation Facility*



Mesa Lab*



Marshall Field Site*



Mauna Loa Solar Observatory*

***NSF-titled properties**

Leading and enabling the community

Scientific Leadership & Expertise
+ Community Convener



Observational Facilities



Computational & Data Capabilities



Modeling Ecosystem



Conceptualize
problems

Apply models
& other facilities
to inform
concept

Followed by
deeper study
& analysis

New
knowledge

NCAR as a National Center

Accelerating Earth system science knowledge and discovery through computational science

Addressing disruptive trends

Societal Need

- Understanding local, extreme hazards and events
- Enable place-based science across natural and human systems
- Create climate resilience and mitigation solutions

- Easy of use and access
- Time to science
- Trust, robustness, security

Increasing Model Complexity

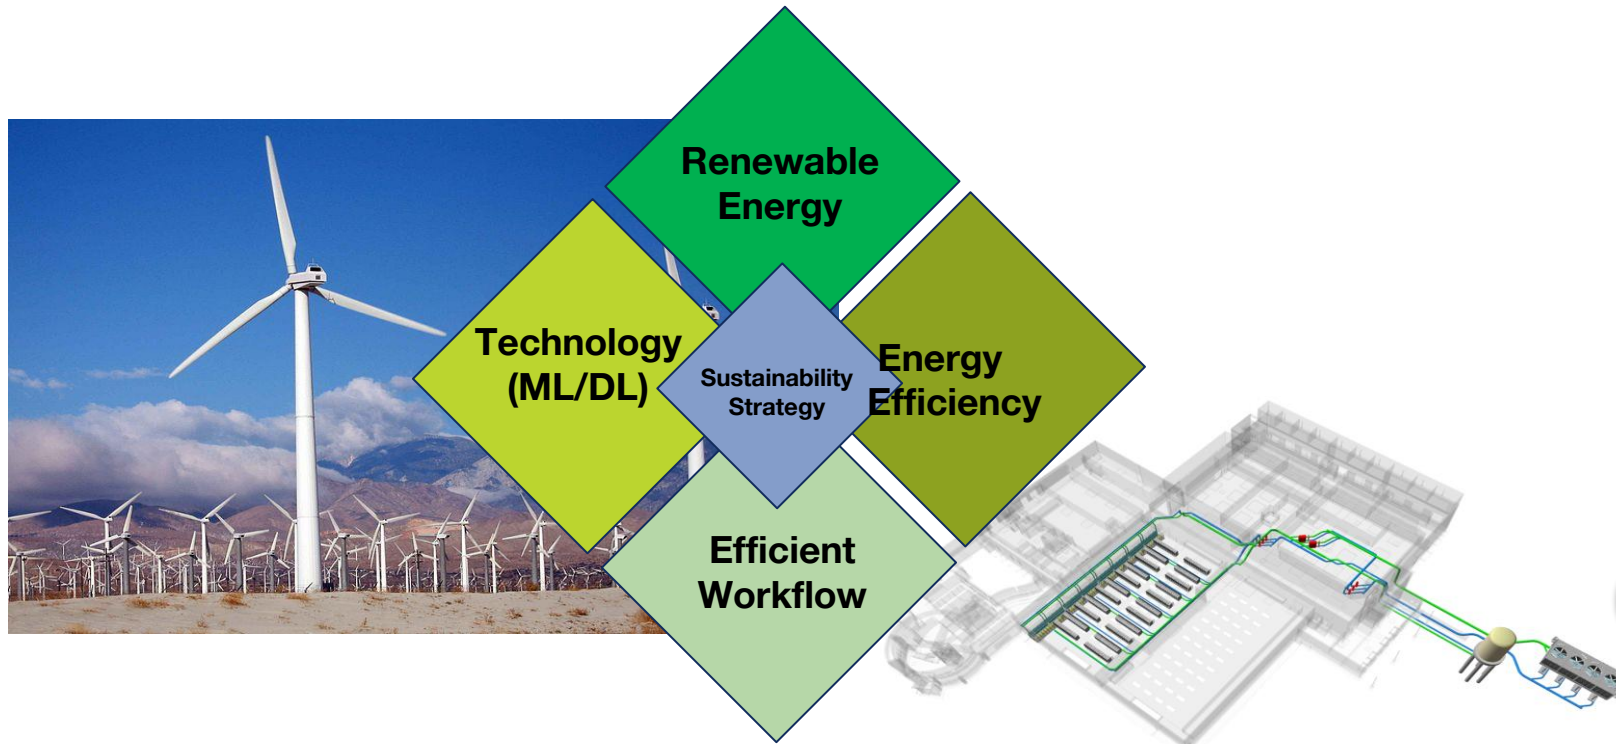
- Earth system coupled models across time and space scales
- Data assimilation across component models
- Finer spatial and temporal resolution simulations
- AI/ML revolution
- Software and its sustainability
 - Workforce

Rapid Changes in Compute & Data Technology

- Physical limits drive us towards heterogeneous hardware
- Energy efficiency
- Barriers of entry to advanced computing systems
- Data silos & inaccessibility of individual PI's data sets
- Cloud and edge computing

Integration and engagement across labs, community and stakeholders

NCAR-Wyoming Supercomputer Center (NWSC) – Sustainable Operations



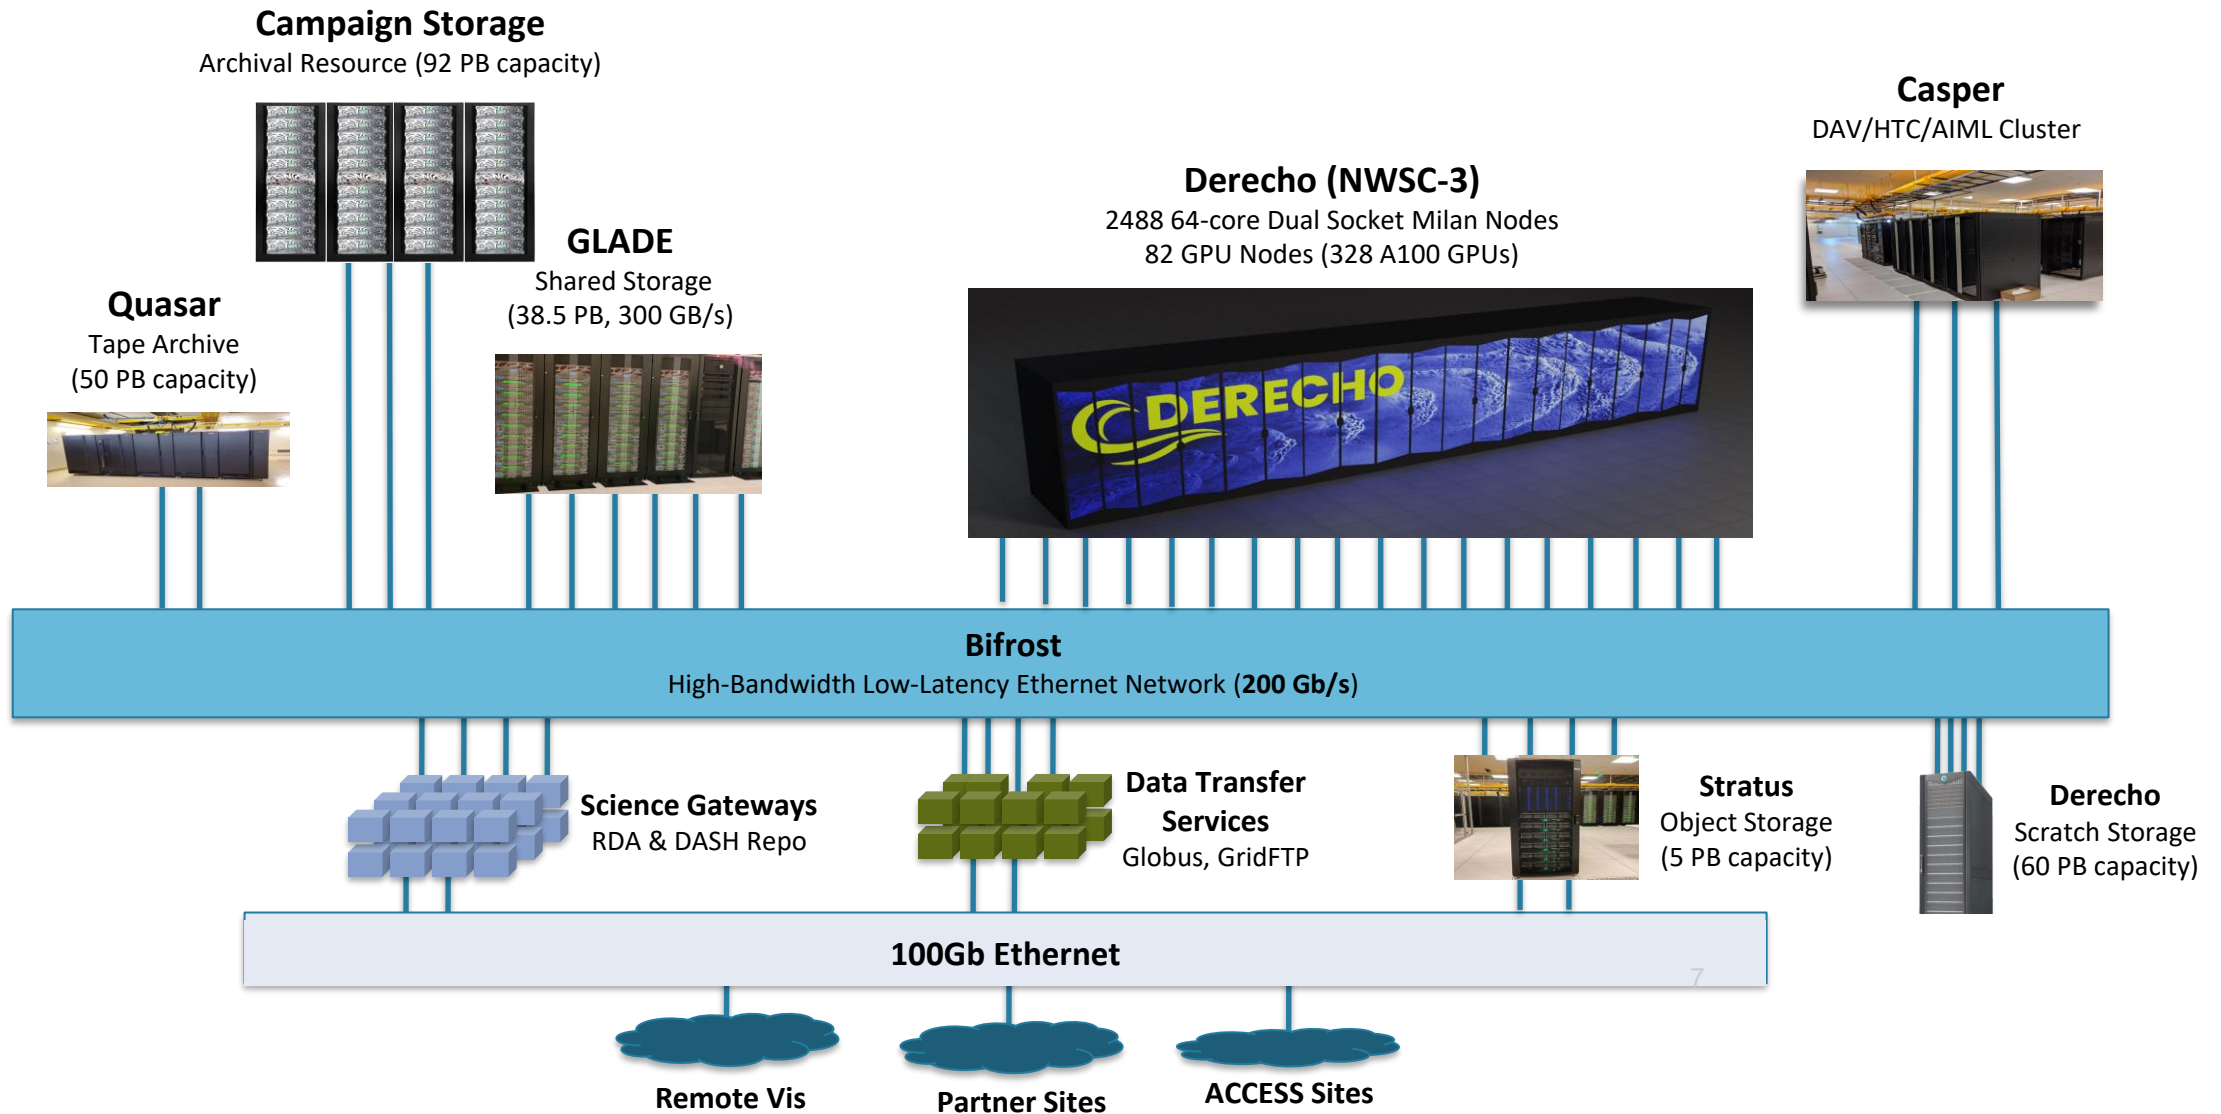
- Carbon-free energy at the NWSC
- Energy-efficient hardware
- Optimized end-to-end workflows
- AI/ML to reduce computation



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




NCAR's High-Performance Computing, Data, & Analysis Ecosystem



NSF NCAR's Supercomputer




- Derecho HPE/Cray XE
 - Includes HPC and PFS
 - Peak: **19.87** PetaFLOPS
 - 60 PB usable file system
- **3.51-fold improvement** over Cheyenne sustained Equivalent Performance (CSEP)
 - CPU – 2.84 CSEP ~80%
 - GPU – 0.67 CSEP ~20%
- Slingshot® Interconnect
- Fast scratch file system
 - Six HPE/Cray ClusterStor E1000 systems
 - 60 petabytes of usable file system space
 - 300 GB per second aggregate I/O bandwidth
 - 5,088 × 16-TB drives





CPU Node



CPU Cabinet


- 4 nodes per compute blade
- 1 slingshot injection
- 64 blades per cabinet
- 256 nodes per cabinet
- 210.7 kW
- 0.65 tons of mech cooling
- ~1.3 PFLOPS
- 0.29 CSEP

GPU Cabinet


- 2 nodes per compute blade
- 4 x GPU per node
- 4 Slingshot Injections
- 64 blades per cabinet
- 128 nodes per cabinet
- 190 kW
- 0.59 tons of mech cooling
- ~10.3 PFLOPS
- 1.04 CSEP



GPU Node



Derecho Community Portions



Community	Annual Portion
University	981 million core-hours 1 million GPU-hours
NCAR	825 million core-hours 850,000 GPU-hours
CESM	451 million core-hours 465,000 GPU-hours
Wyoming	345 million core-hours 355,000 GPU-hours
AMPS	53 million core-hours 55,000 GPU-hours

Enabling large scale computing in the earth and related geospace community

Enable breakthrough science through co-design with the scientific community

CM1, MURaM, MUSICA

Create portable, scalable, energy efficient code for modern architectures

CM1, MURaM, MUSICA, CESM, MPAS-A, SAMURAI, EarthWorks, 3Dynamo

Optimize input/output workflows





MURaM, MUSICA, EarthWorks, CESM

Provide software verification and reproducibility techniques


CESM, MPAS-A, PyCECT

Develop techniques to reduce data volume

CESM, MURaM, MPAS-A, LDCCPy



MURaM - Advancing Solar Physics with GPUs

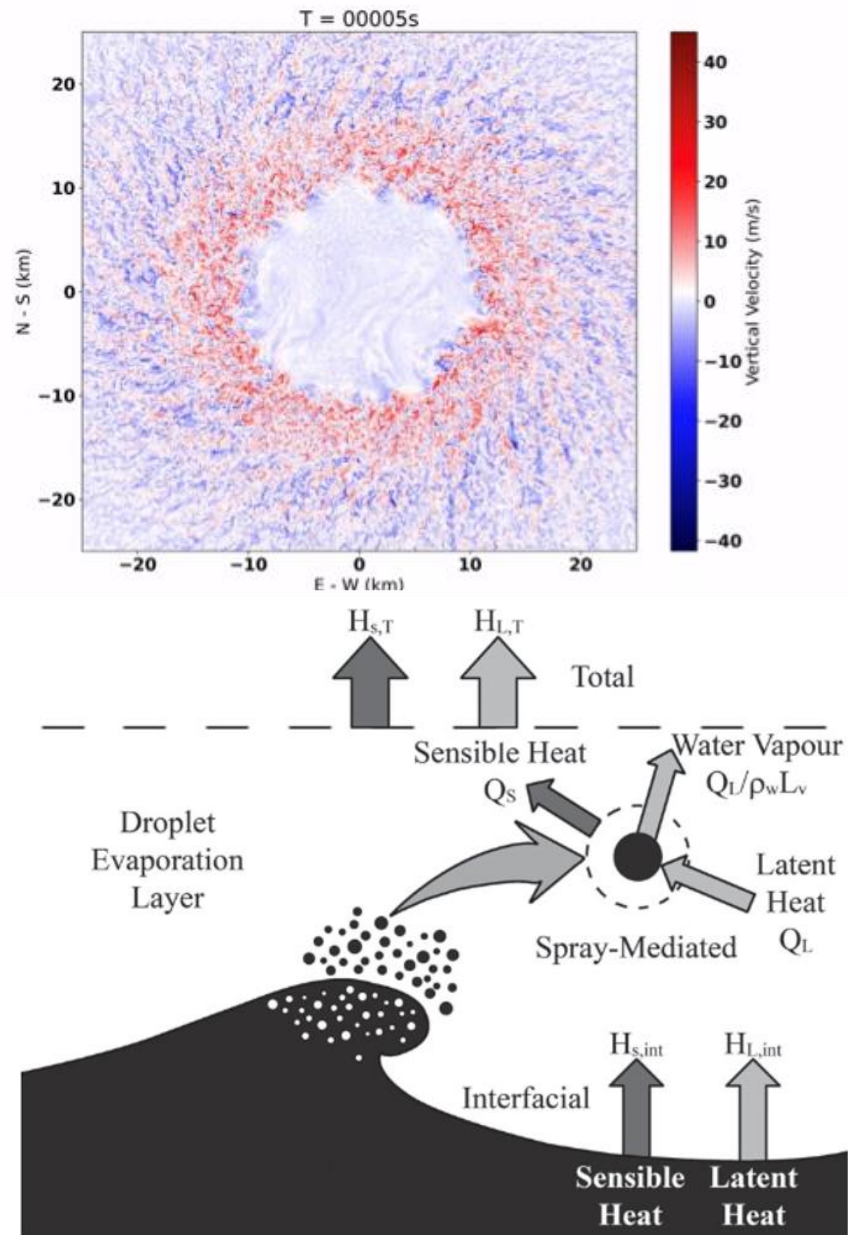


“Quiet” Sun Photosphere
2km resolution

- Simulated ~60 hrs of photospheric magneto-convection on the sun
- Lower resolution case (~1.8B gridpoints) ran on up to 320 A100 GPUs on NCAR’s Derecho cluster
- ≈ 65k CPU cores running same case for comparison
- **Running these simulations on GPUs prevented the emission of ~60 tons of carbon or about 4 person-years of emissions**

	GPU	CPU
Total emission amounts for these simulations	14.2 tons	73.4 tons

Porting the Cloud Model 1 (CM1) onto GPUs



- 1024x1024x512 grid
- Up to $O(10^{10})$ particles
- 12 hours on 128 GPU on NCAR's Derecho cluster
- <1500 gpu-hours per simulation
- **Running these simulations on GPUs prevented the emission of 47.5 tons of carbon or about 3 person-years of emissions**

	GPU	CPU
Total emission amounts for these simulations	4.05 tons	51.59 tons

Image credit: David Richter

Ensemble consistency tests for model development

Goal: Ensure that changes during the model development life cycle do not adversely affect the results!

Challenge: Many changes during the model development cycle are not bit-for-bit:

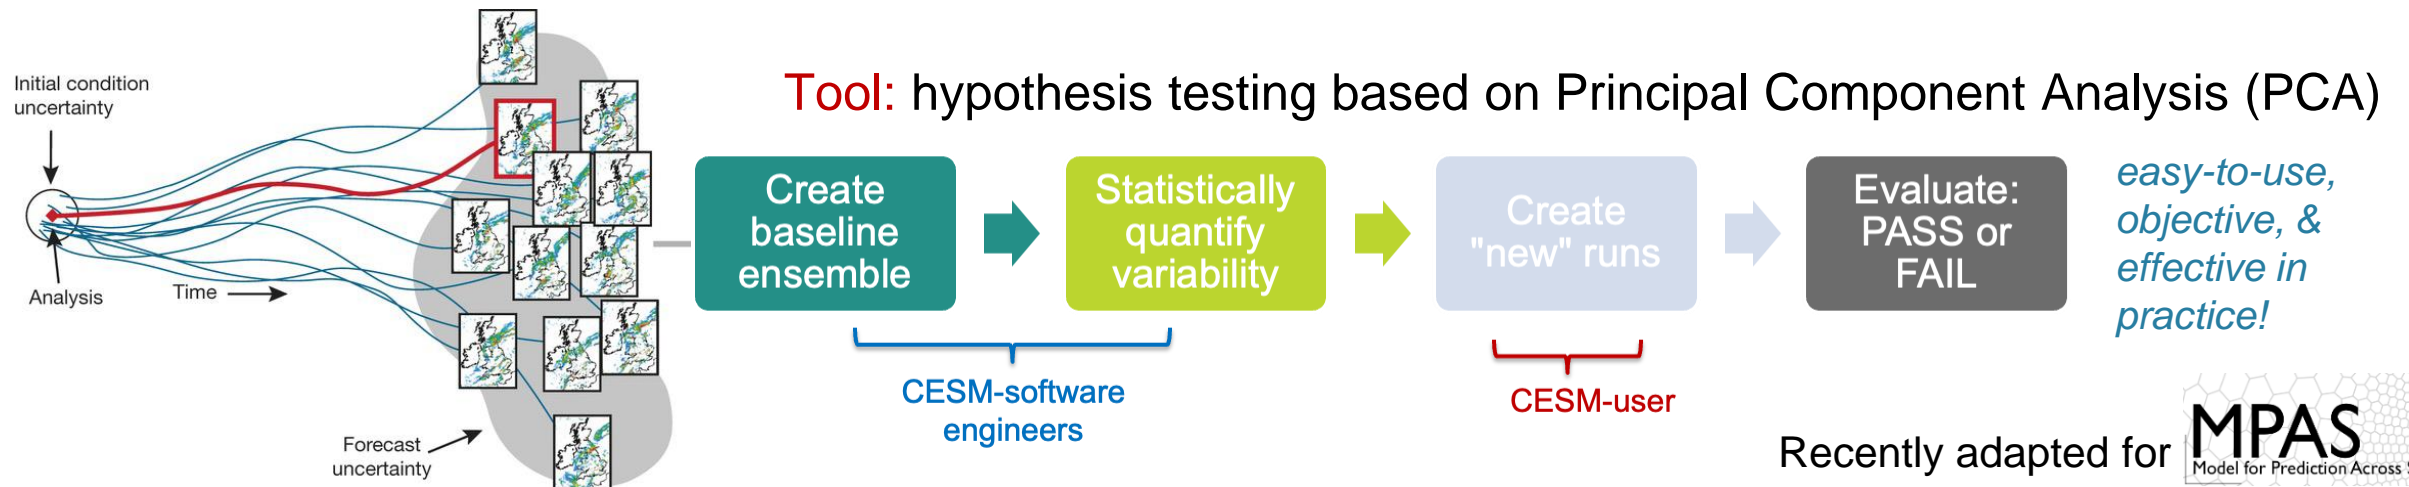
- port to new environment
- code modifications (e.g., optimizations)
- compiler changes
- heterogeneous computing (e.g., GPUs)



How do we assess whether the new output is still “correct” (when not bit-for-bit identical)?

Reframe “correctness” : Is the new data **statistically distinguishable** from the original data?

Our approach: evaluate in the context of the climate model’s variability (an **ensemble** of “accepted” runs)



Lossy compression for Earth system models

Data generation trends at NCAR are unsustainable!

Goal: Use lossy compression to reduce model storage ...without (negatively) impacting science results!

Challenges for climate data:

- spatial & temporal dependencies
- output variables have very different characteristics
- typical compression metrics don't have much meaning
- parallel I/O


Research focus areas:

(1) How to evaluate the effects of lossy compression on model data?



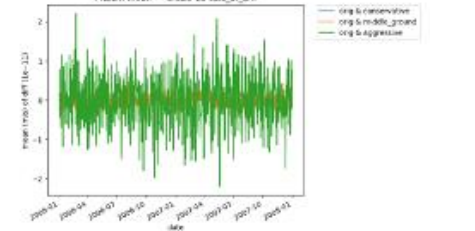
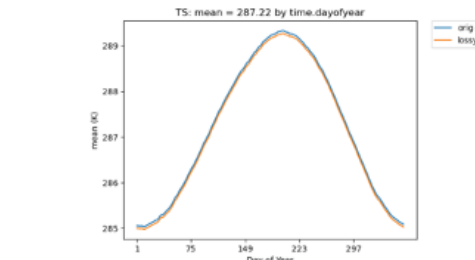
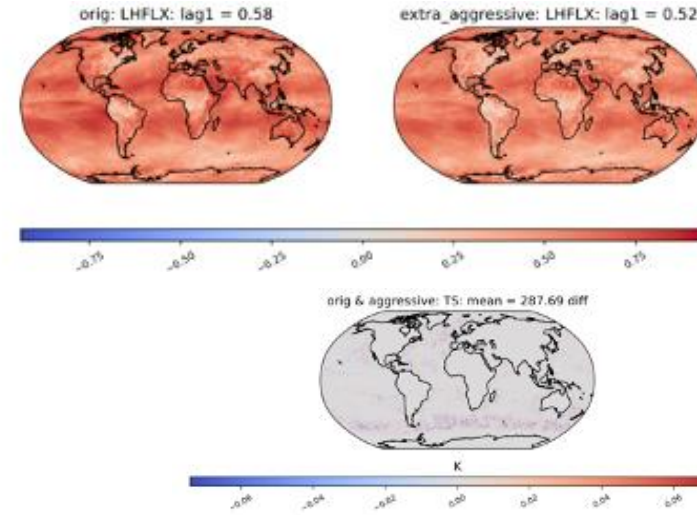
- quantifiable metrics (and thresholds) for derived quantities and statistics
- a Python package (LDCPy) for analysis of differences in large spatio-temporal datasets

(1) From the data's features, determine what type of (and how much) compression to use.

- an automated tool to determine the optimal algorithm and compression rate for each variable!



Storage resources are limiting climate science objectives!



Storage Platforms for Scientific Workflows



- Procured with HPC system
- High bandwidth & IOPS
- 3 – 6 month retention

- Procured annually
- Medium bandwidth
- Allocated access
- 5-yr retention

- Med/Low bandwidth
- Controlled access
- Long-term preservation (with review)


Enabling Earth System Data Science



<https://raijin.ucar.edu>

Extensible, Scalable Community
Geoscience Analysis Tools for
Unstructured Grids

Open, scalable, community-driven
data analysis and visualization tools
for Earth System Science



PROJECT PYTHIA

The go to resource for learning the *Scientific Python Ecosystem*



NCAR's Data Repository Landscape

Research
Data
Archive
(RDA)

*Observations,
Reanalyses,
mesoscale models -
not necessarily
archival copy*

Large, Complex Obs and Model Data

Climate
Data
Gateway
(CDG)

*Climate models,
mostly CMIP -
some limited
access*

Geoscience
Data
Exchange
(GDEX)

*NCAR data, some
external - all open
access, archival.*

**Long Tail and
Selected NSF Facility
(CIF) Data**

EOL Field
Data
Archive
(FDA)

Field Project Data

**EOL/NSF LAOF
Field Campaign
Data**

HAO
Repository

*Solar & Space
Weather Data*

**MLSO and CSAC
Solar Facility Data**

NCAR
UCAR

DASH
Digital Asset Services Hub



NCAR is sponsored by
National Science Foundation

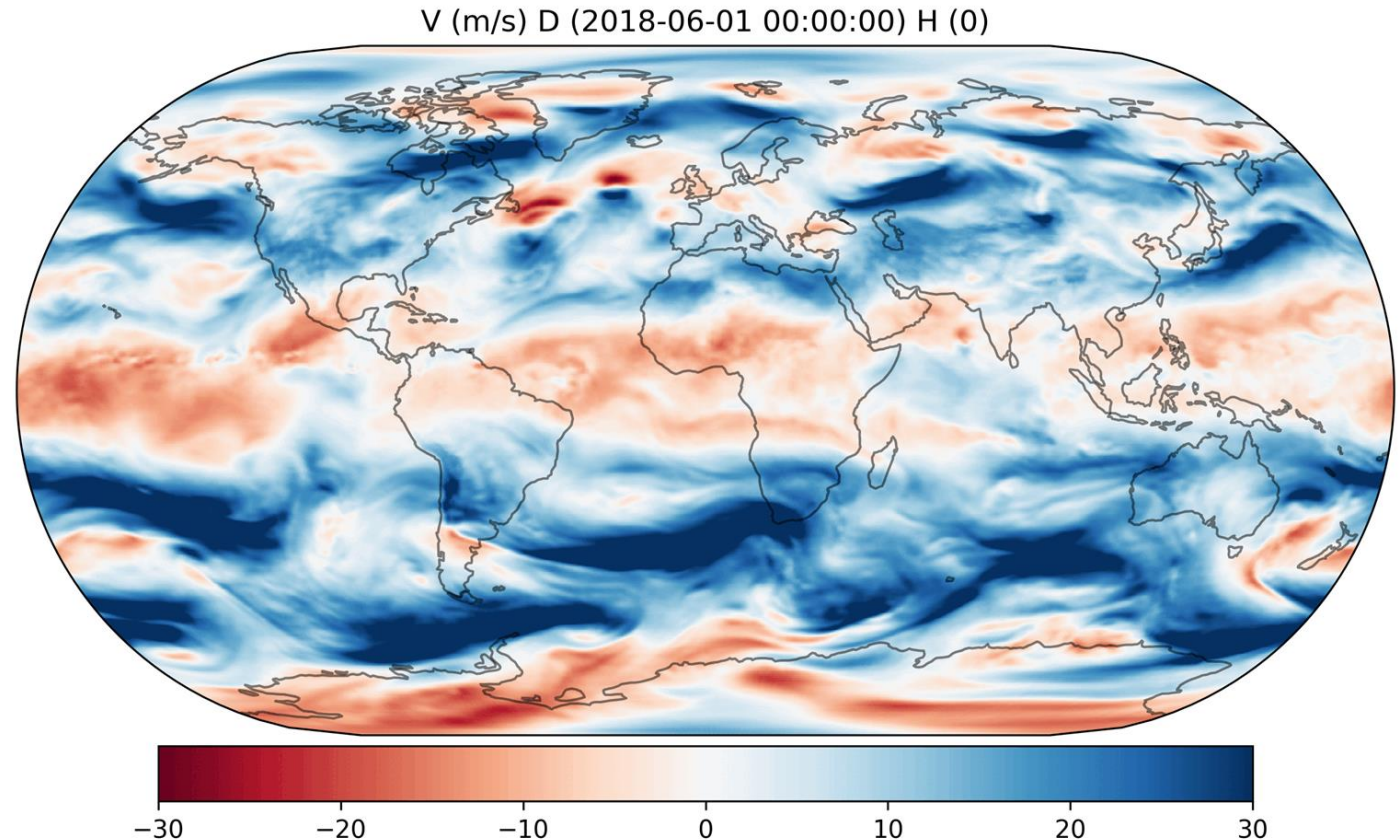


Search Data, Software, Models and Publications

Search...

CREDIT: Community Runnable Earth Digital Intelligence Twins

- Current AI weather prediction groups have released model weights but not the code to train them at scale
- NSF NCAR developing open framework to *train* and *run* global and regional AI weather models
- Goals
 - Usability
 - Integration with community datasets and models
 - Support Scientific Investigation
 - Accelerate deployment across many use cases



Team: John Schreck (CISL), Will Chapman (CGD),
DJ Gagne (CISL), Dhamma Kimpara (CISL), Kyle Sha (MMM)


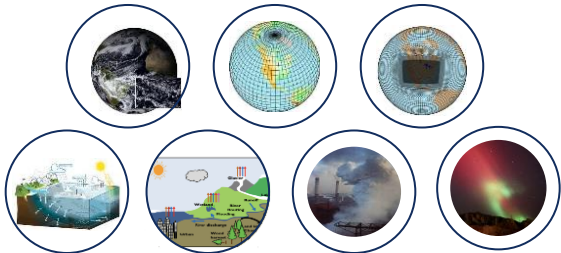
NSF NCAR's computational science strategy



- **Enabling and supporting the community through a digital Earth System Science lab**
- **Democratizing Research Capabilities in support of Earth System Science**
- **Sustainability**




NSF NCAR Community Software Facility




NCAR Compute Infrastructure



Translational Computational Science



NSF NCAR Research Data Commons




The Future of NSF NCAR's Computing Infrastructure

- Plans

- Carbon Free power at the NWSC
- Integrate and federate with the NSF CISE-funded computing ecosystem
- Focus on end-to-end workflows
- Expand our resources through (cross-agency) partnerships
- Data-intensive architecture
 - AI
 - Data assimilation
- What's the architecture of the next supercomputer at NCAR?
 - Co-design with the Community Software Facility


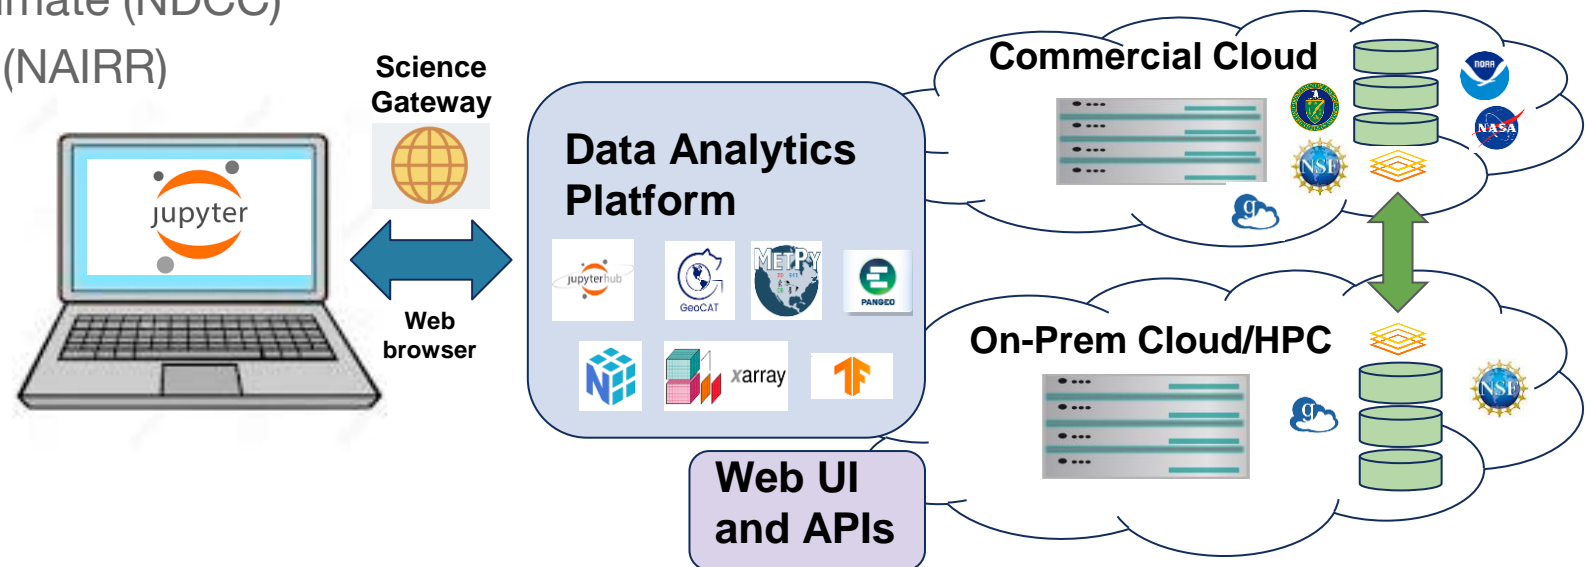
- Early wins

- CISL/COO task force for carbon-free electricity at the NWSC
 - Long-term contract for electricity will be cost-neutral
- Funded to integrate our data infrastructures with the Open Science Data Federation
- Allocation of some of our resources through the NSF ACCESS program
- CONTINENTS international project in collaboration with EPCC and NCAS



A unified NSF NCAR data infrastructure - Data Commons named GDEX

- **Unified platform for NSF NCAR and community data**
 - **Highest priority for CISL**
- Sustainable business model (engagement with NSF)
- Support the Geo community with data management services
- **Cloud platform services for data analysis, AI/ML, and visualization**
- Large Language Models (How Large Language Models Will Disrupt Data Management <https://doi.org/10.14778/3611479.3611527>)
- NSF Opportunities
 - National Discovery Cloud for Climate (NDCC)
 - National AI Research Resource (NAIRR)



Cloud Technologies for Data Intensive Workflows

Data Commons analytics platform and prototyping environment

- Data proximate analysis from anywhere with a network connection
- Access analysis-ready, cloud-optimized data
- Lightweight access could replace Data Analysis allocations

Interactive web visualizations and dashboards


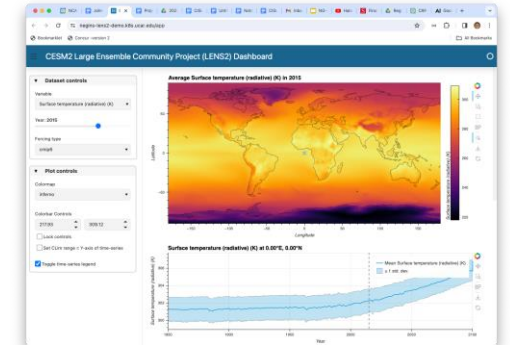

- Interactive dashboards created by scientists, shared with colleagues
- Explanatory web visualizations for the public, policymakers, and more

Training

- Ephemeral, custom JupyterHub instances
- Containerized apps (e.g. WRF Cloud)
- Lightweight access could reduce the need for Classroom allocations

Continuous Integration / Continuous Delivery (CI/CD) support for code development

- Target for GitHub *self-hosted runners*
- Automate software testing with custom hardware, access to NCAR data
- Support for Community Software Facility



Applied Computer & Computational Science Research

- Machine Learning must evolve into Scientific ML
 - Progress must happen via models | algorithms | software
 - HPC empowers ML – as it has empowered Simulation before
- Develop a comprehensive NCAR AI/ML strategy
- Support the software modernization effort through applied computer science research
 - Domain Specific Abstractions
 - Compiler technologies
 - Mixed Precision, compression, ...
 - Verification methods
- Extending our capability through connections and partnerships
- Student and PostDoc programs
- Joint project with Edinburgh Computing Center
 - Energy efficiency from the facility to the software
 - Domain specific abstractions
 - Investment in compiler technologies
 - Source to source translations
- Participation in Chameleon Cloud
 - Translational computer science
 - Partnership with TACC and University of Chicago
- HLRS
- DOE
- USGS
- Oghers



**Translational
Computational Science**

Modernizing NSF NCAR's Model Ecosystem – Community Software Facility



- What is the Community Software Facility?
 - Resourced to manage best practices and governance of our “core” community software
 - Resourced to support the use of the software across leadership class hardware systems
 - Resourced to maintain software while avoiding toxic technical debt
- Project to create the facility “Pathfinder project for the CSF”
 - Review standard of current software and bring into the facility
 - Deliver interoperability of our codes
 - Modelling systems within our facility that have common components
 - Deliver traceability between weather and climate
 - Improve “quality” of code within
 - Readable/good standards
 - Modular
 - Portable
 - Efficient
- **It's a people change activity**
 - There will be changes in working practices and organizational structure needed
 - It needs to be embedded in our working practices



Leading in a rapidly changing technology environment




**NSF NCAR Community
Software Facility**



NCAR Compute Infrastructure



**Translational
Computational Science**



**NSF NCAR Research Data
Commons**

