# Creating a multi-PB Weather dataset through GPU-Accelerated Parallel Processing

Jeff Adie

Distinguished Engineer, NVAITC

NVIDIA Singapore

# Introduction

## Our Body of Work

NVIDIA is the world leader in accelerated computing. We are fundamentally changing how computing works and what computers can do. The next industrial revolution has begun.

# NVIDIA HPC & AI PLATFORM

# CUDA-X Accelerates Industries

**PhysicsNeMo**
**Warp**
AI Physics

**cuDF**
**cuML**
Data Science and Processing

**cuDSS**
**cuSPARSE**
**cuFFT**
**AmgX**
Computer Aided Engineering

**TRT-LLM**
**Megatron**
**NCCL**
**cuDNN**
**CUTLASS**
**cuBLAS**
Deep Learning

**cuEquivariance**
**cuTensor**
Quantum Chemistry

**cuQuantum**
**CUDA-Q**
Quantum Computing

**Earth-2**
Weather Analytics

**MONAI**
Medical Imaging

**Parabricks**
Gene Sequencing

**cuOpt**
Decision Optimization

**Aerial**
**Sionna**
5G/6G Signal Processing

**cuLitho**
Computational Lithography

**cuPyNumeric**
Numerical Computing

# The need for high-fidelity Weather Forecasting

# Weather Prediction is Integral to Modern Society

Simulations and forecasts drive planning and decision-making

Agriculture & Forestry

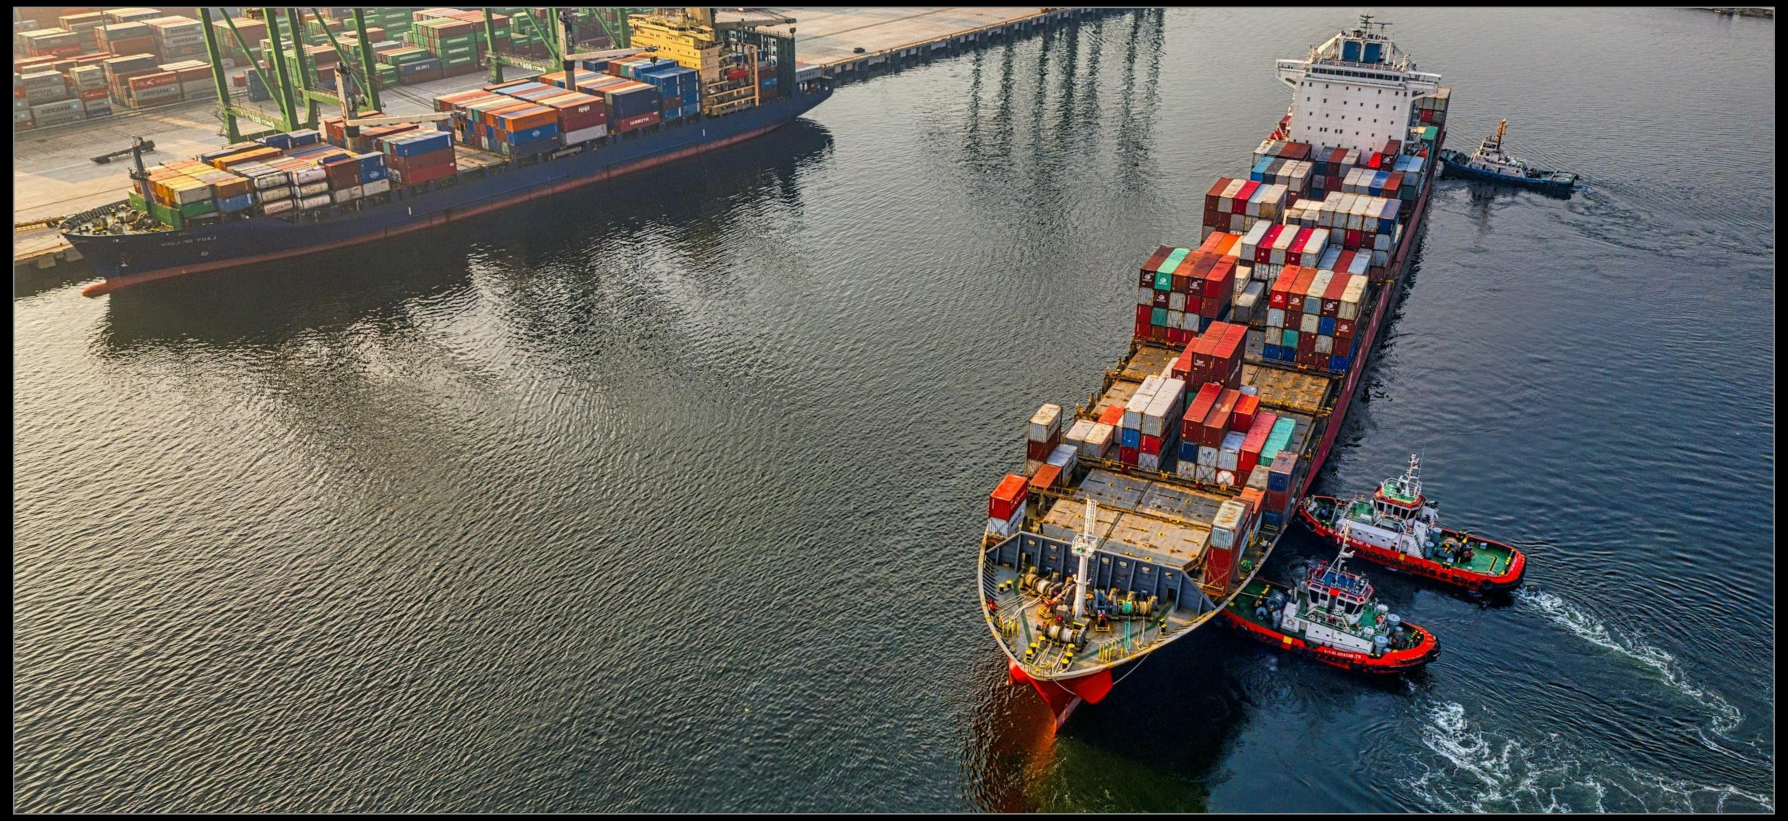

Transport & Logistics

Energy

Risk management

Land development

Recreation

# Climate and Weather Simulation is a Massive Computation Challenge

AI models are skilled enough to drive more accessible and efficient forecasting


Figure adapted from: Schneider, T., Teixeira, J., Bretherton, C. *et al.* "*Climate goals and computing the future of clouds*". *Nature Climate Change* 7, 3–5 (2017)

# Advantages of AI Weather Models

Compelling skill, resource requirements, and accessibility

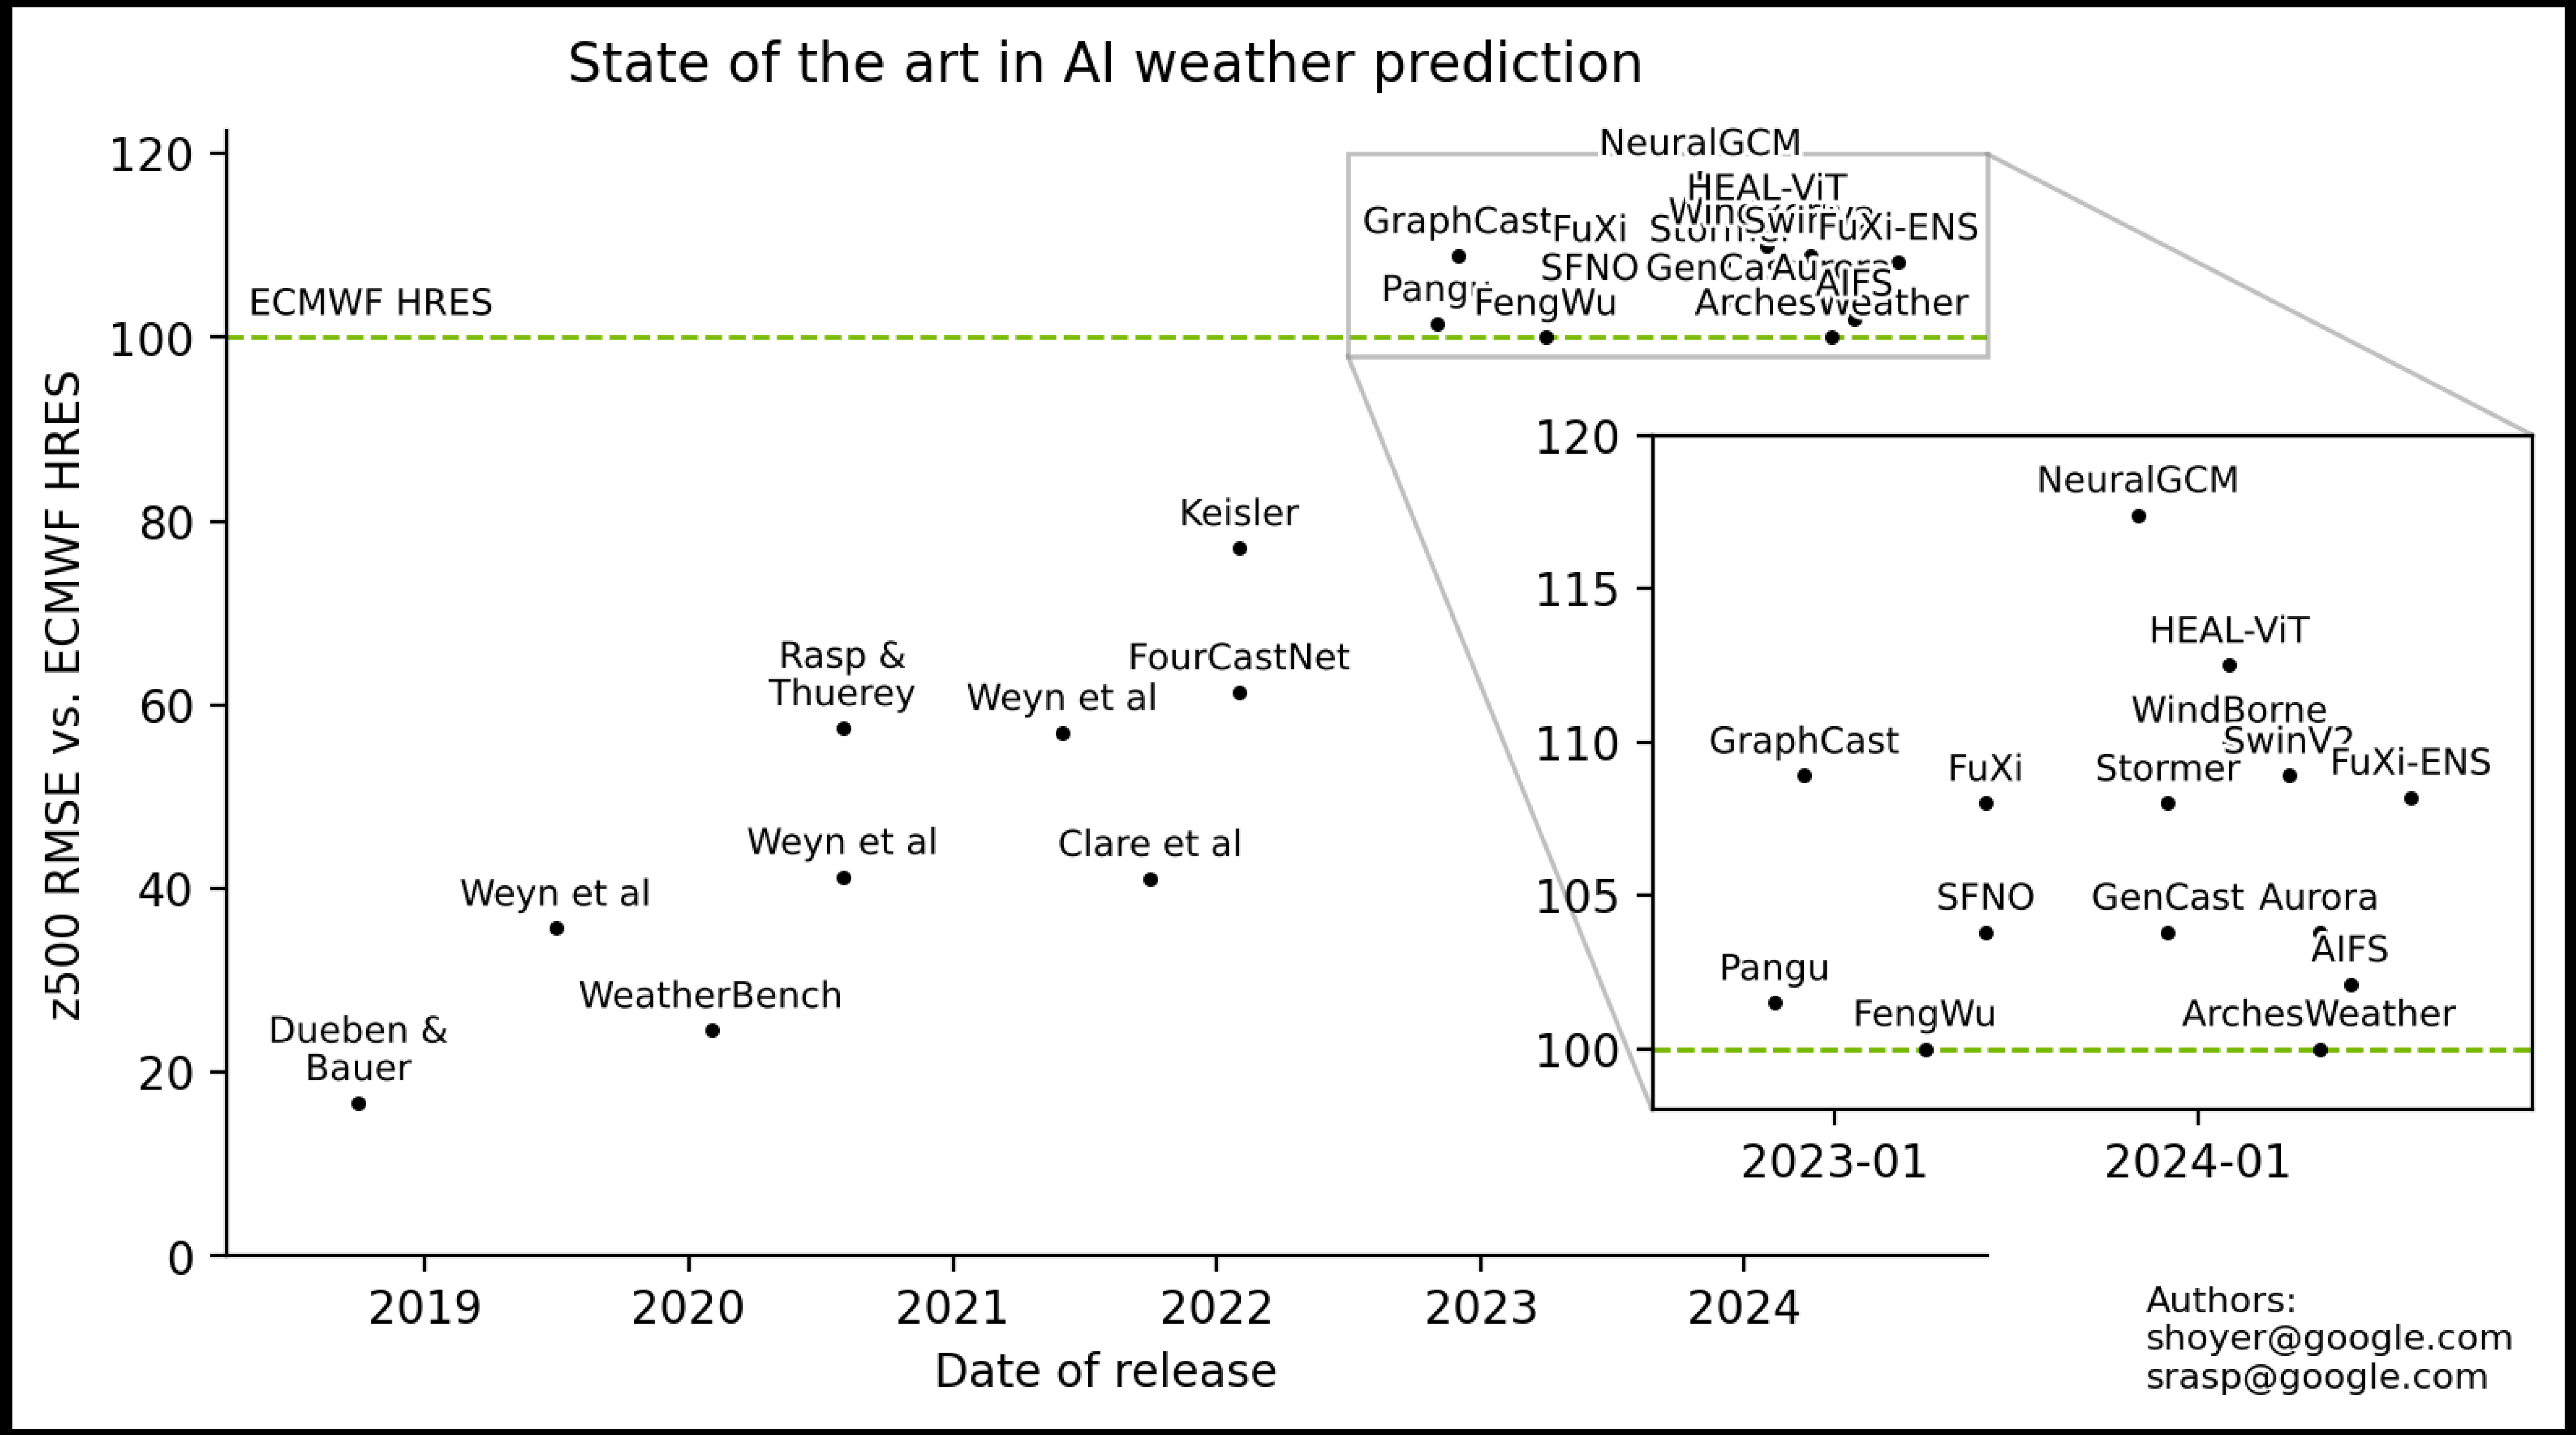


Rasp, Stephan (2024). AI-Weather SotA vs. Time.
https://doi.org/10.6084/m9.figshare.28083515.v1

**High skill**

**Less hardware**

**Faster inference**

**More accessible**

Custom weather applications are now possible

# Two-Phase Life Cycle of AI Weather Models

Once trained, AI models make rapid predictions

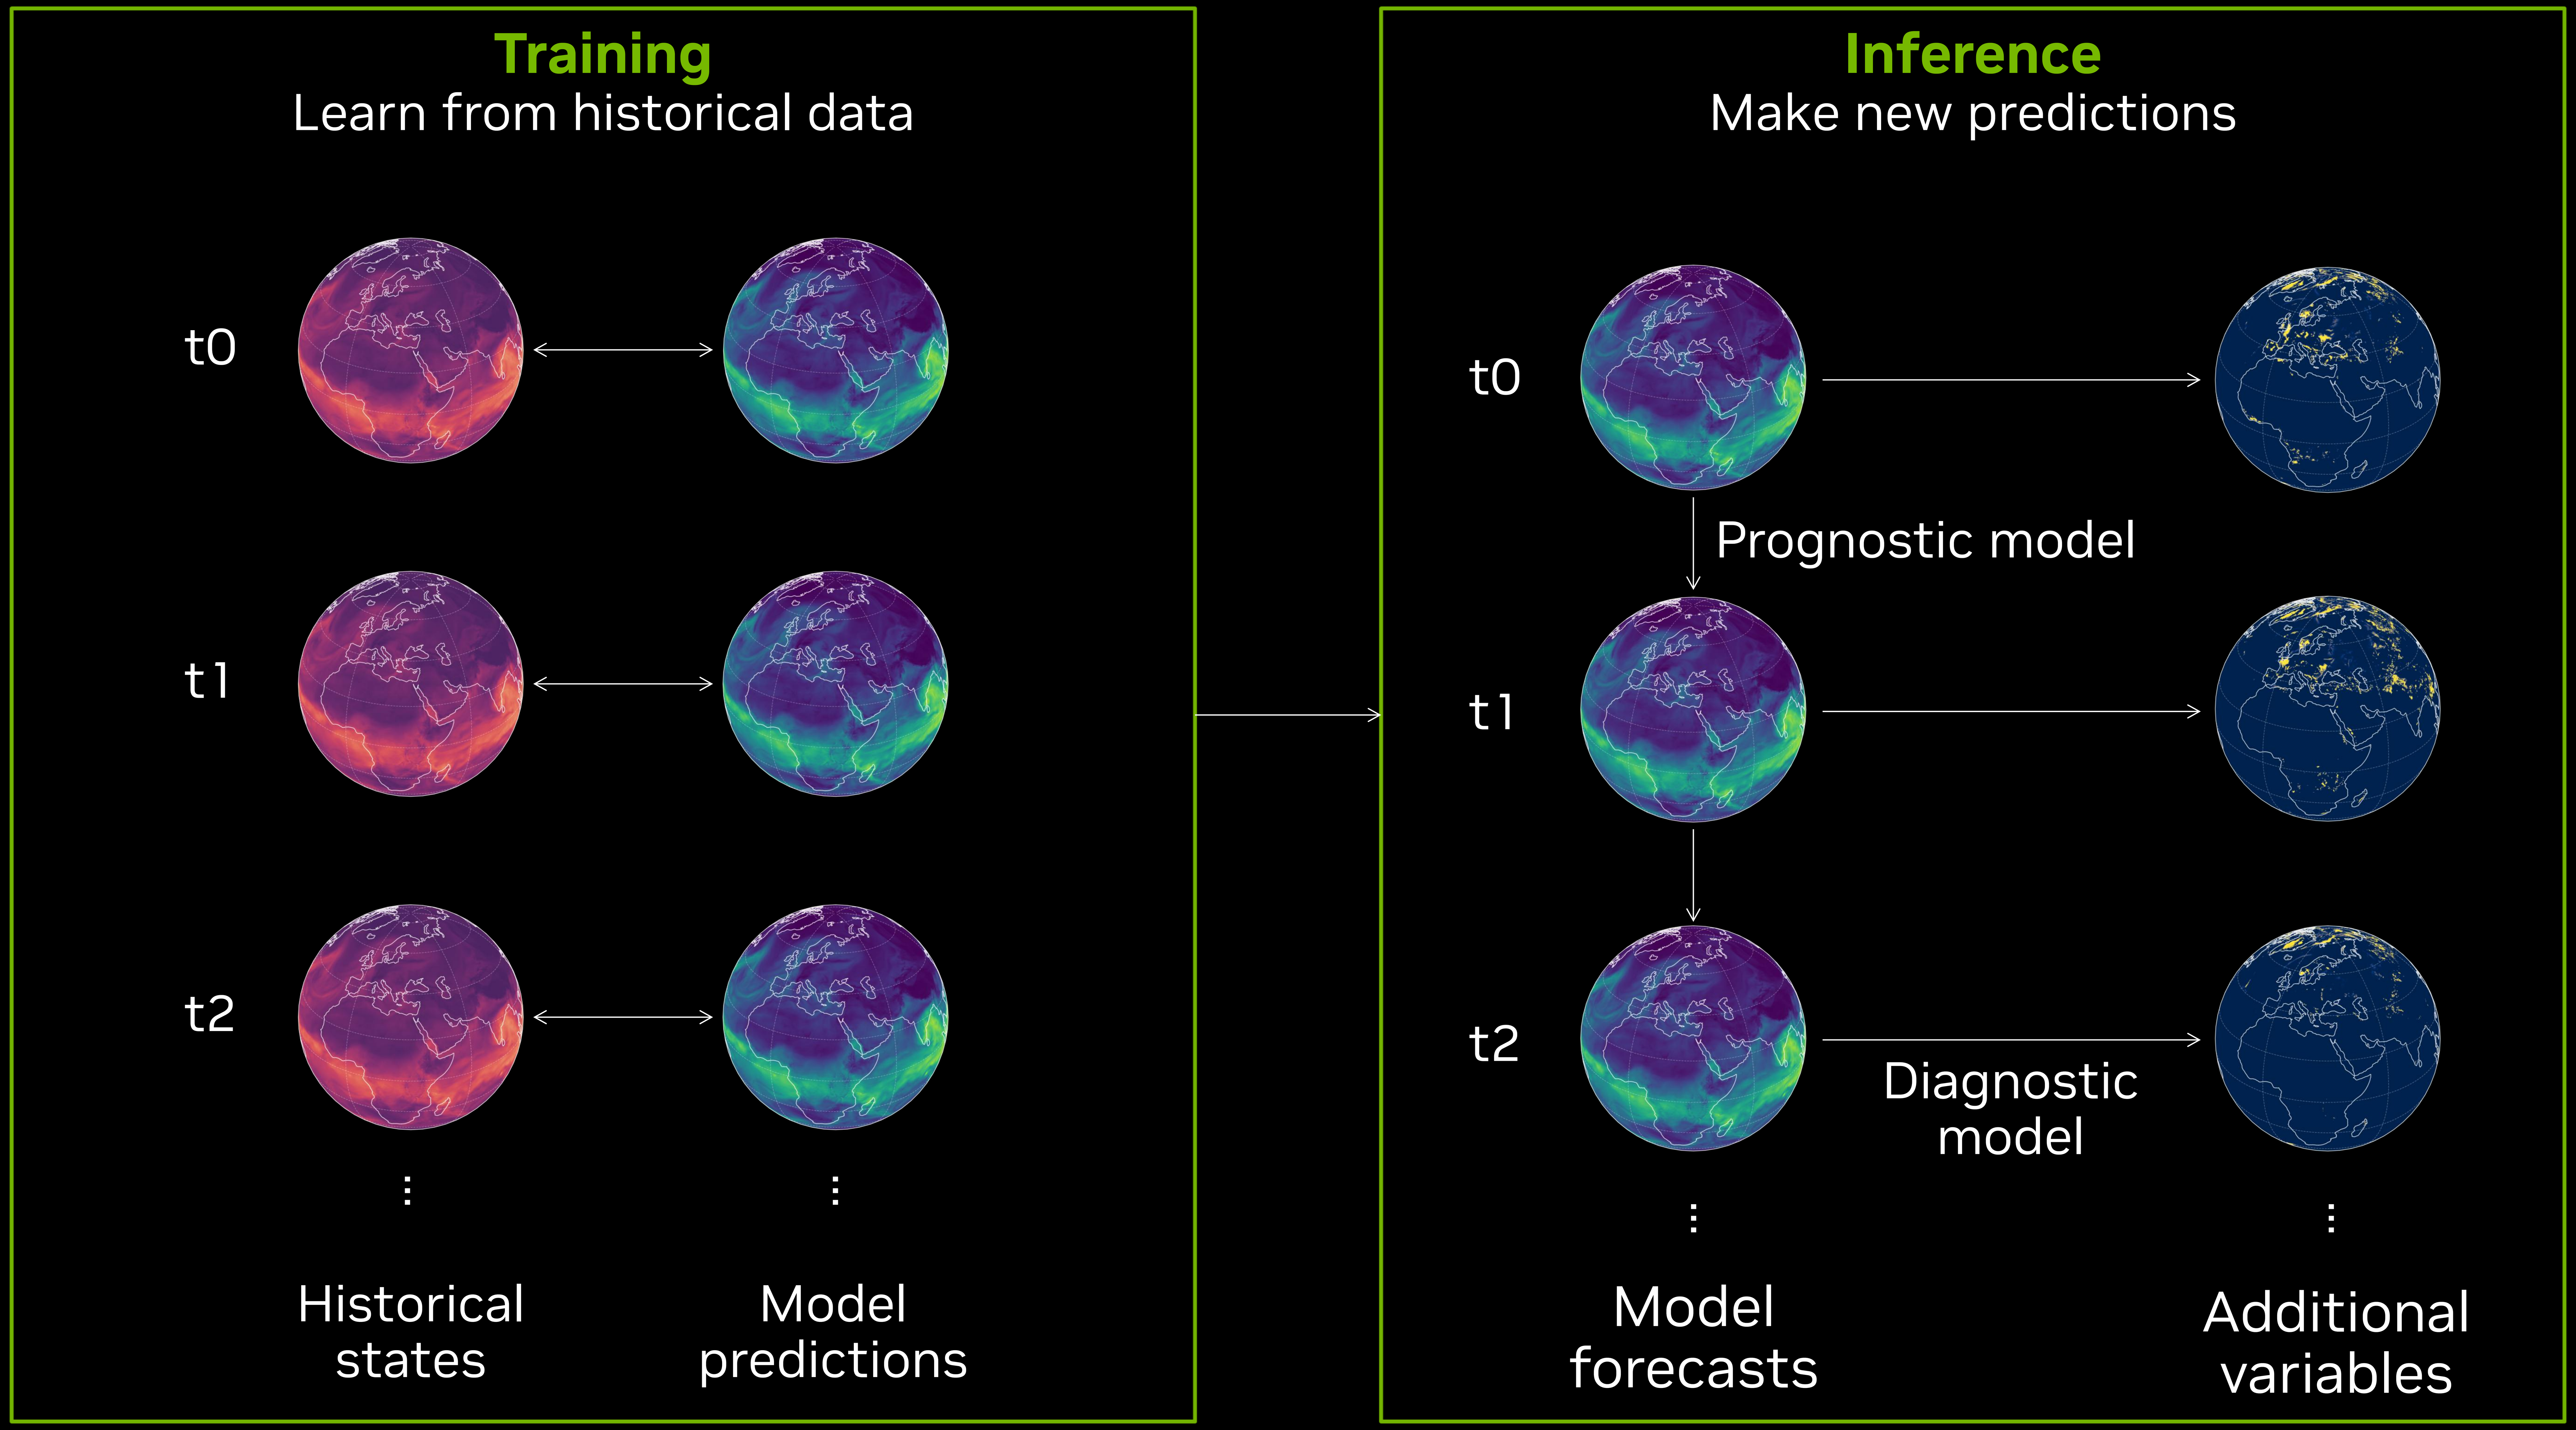

# Creating the Weather-1K Dataset

# “WEATHER-1K”

# Forecasting at 1-kilometer scale.

*MITRE × The Weather Company (April 2026): a strategic collaboration advancing AI-driven, high-resolution global weather forecasting — built on Weather 1K, a public 1-kilometer-resolution atmospheric dataset developed in MITRE’s Federal AI Sandbox.*

## AI Supercomputing

Trained in MITRE’s Federal AI Sandbox on NVIDIA DGX infrastructure — building the next generation of meteorological models.

## The Dataset

A public, high-resolution continental U.S. training dataset mapping wind, temperature, pressure, and reflectivity — powered by TempoQuest’s AceCAST modeling.

## Operational Impact

Sharper predictions for high-impact severe weather — hurricanes, wildfires, flooding — improving enterprise and emergency decisions.

**Model. Forecast. Act.**

MITRE





NVIDIA

# Federal AI Sandbox – NVIDIA DGX SuperPOD

## Federal AI Sandbox Technical Specs

1 SU NVIDIA H100 DGX SuperPOD AI supercomputer

**248 GPUs in 31 DGX nodes** each containing:
- 8 x H100 GPUs (640 Tensor cores; 80GB)
- 2 x Intel Xeon CPUs (112 cores)
- 900 GB/s InfiniBand optical interconnect
- 350 kW power consumption

**9 PB of VAST high-performance storage**

**Delivers 1 ExaFLOP** (1018 8-bit floating point operations per second) of AI compute supporting the training of **multi-billion parameter foundational AI models**

NVIDIA

TempoQuest

MITRE

# Weather 1k Overview

- AceCast and HRRR have same computational engine (WRF)
  - Worked with NVIDIA and TempoQuest on model setup.
  - Settings mimic the HRRR.
- WRF not optimized for GPUs
- Increased spatial and temporal resolution
  - 1-km vs 3-km
  - 10-min vs 1-hour
- Continuously ran 6 simulations across 4 nodes (32 GPUs, 192 total)
  - Produced data for 2022 – 2023
  - 24-hour forecasts at 10-minute time increments
  - Each forecast took approximately 10 hours
- Created a ~5200 x 3300-pixel domain over CONUS (HRRR ~ 1800 x 1100)

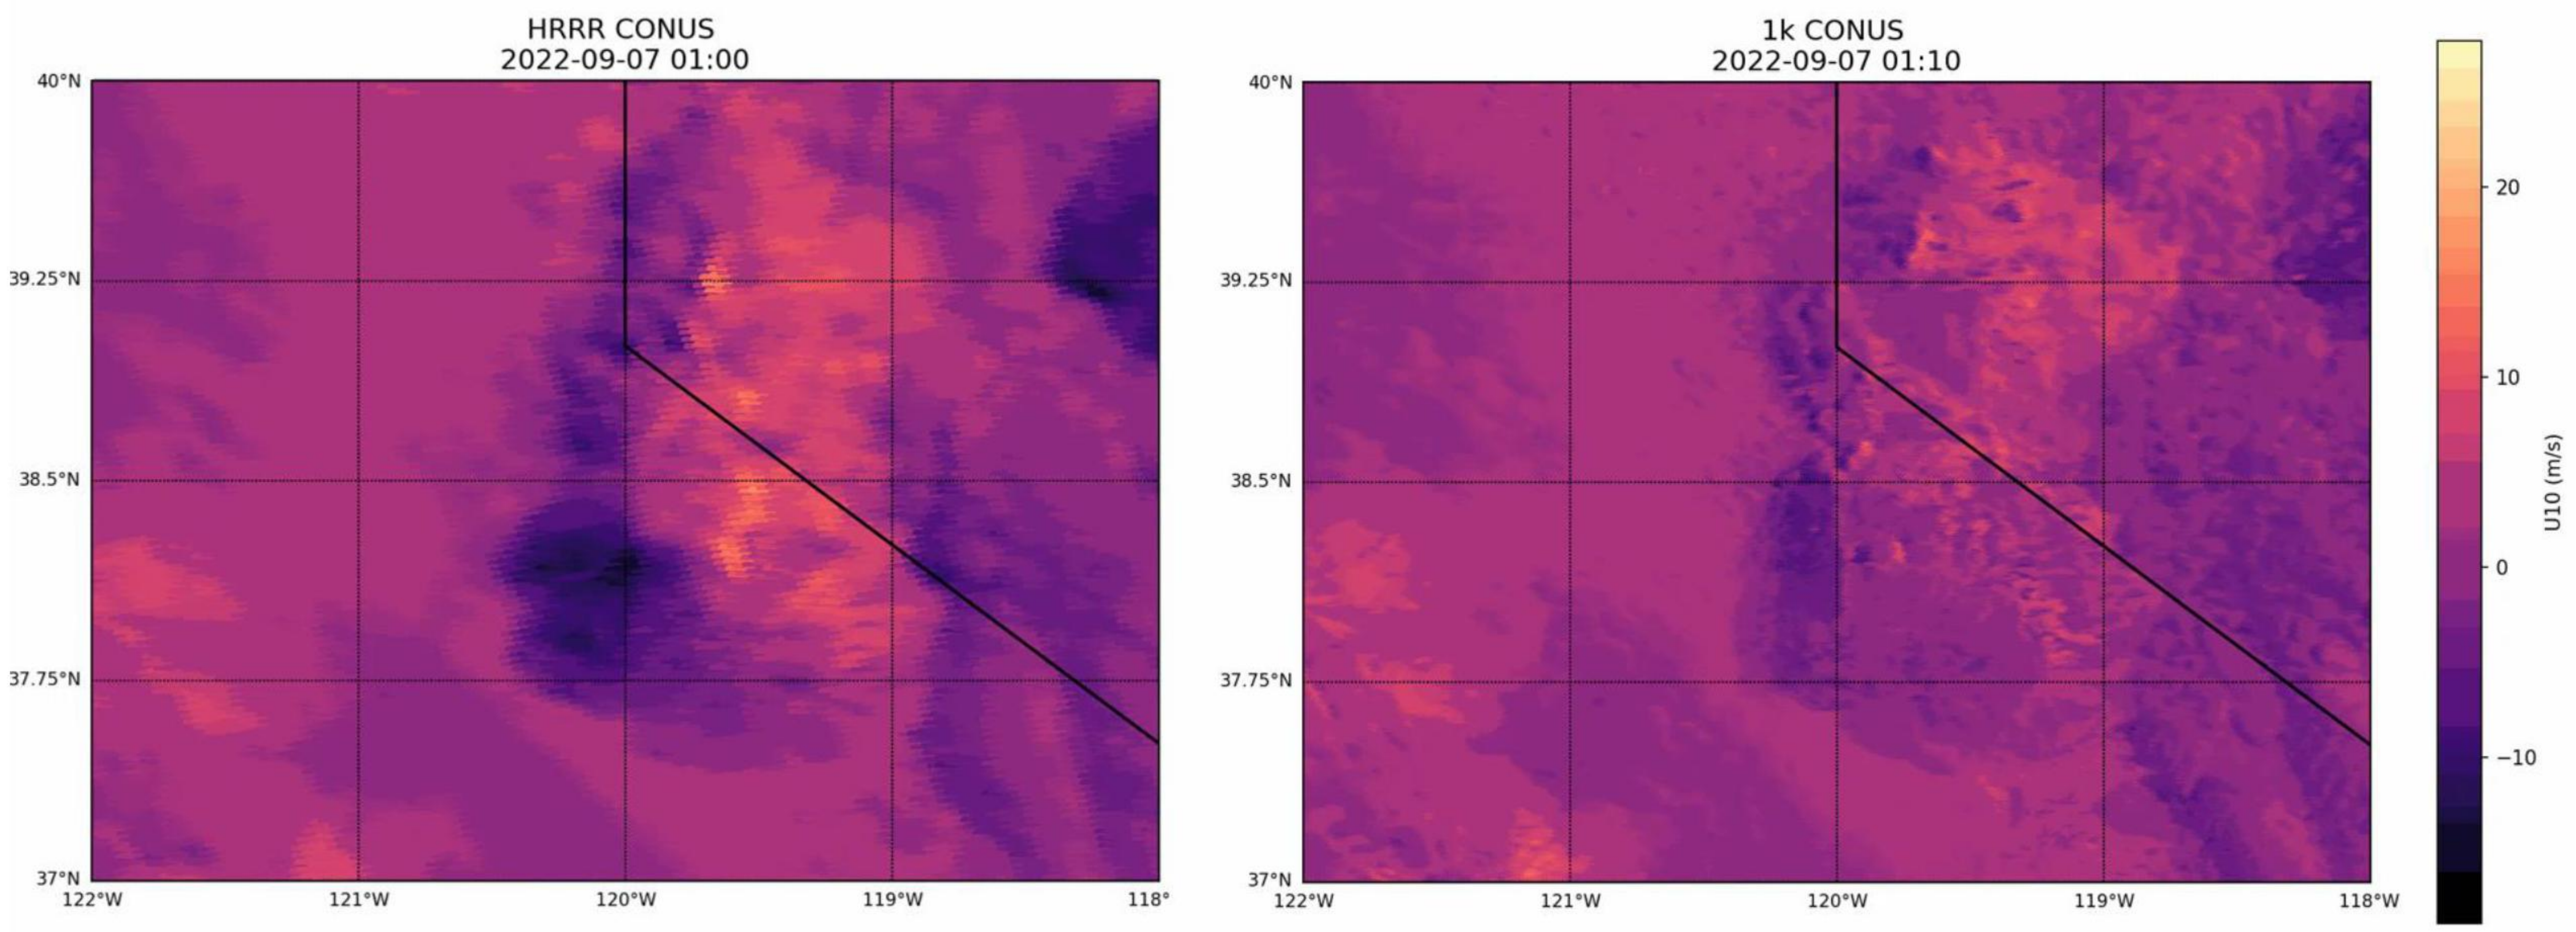


Comparison of HRRR output (left) to Weather-1k output (right). HRRR model output is at 3-km spatial resolution and hourly timesteps. Weather-1k output is at 1-km spatial resolution and 10-minute timesteps.

# Compute Challenges

Creating a high-fidelity training dataset with 100k samples

- To create sufficient training samples we needed 2 years of weather prediction at 10 minute intervals:
  - 365 x 2 x 24 x 6 = 105,120 samples
- Each 24 hour forecast ran independently:
  - Used 32 GPUs per forecast
  - We ran 6 forecasts in parallel across 192 GPs on 24 nodes
  - Approximately 10 hours wallclock for each forecast
  - Total 350,000 GPU hours
- Training for a 3km -> 1km diffusion model:
  - Ran on 64 GPUs (8 nodes)
  - Total of 17,000 GPU hours for training

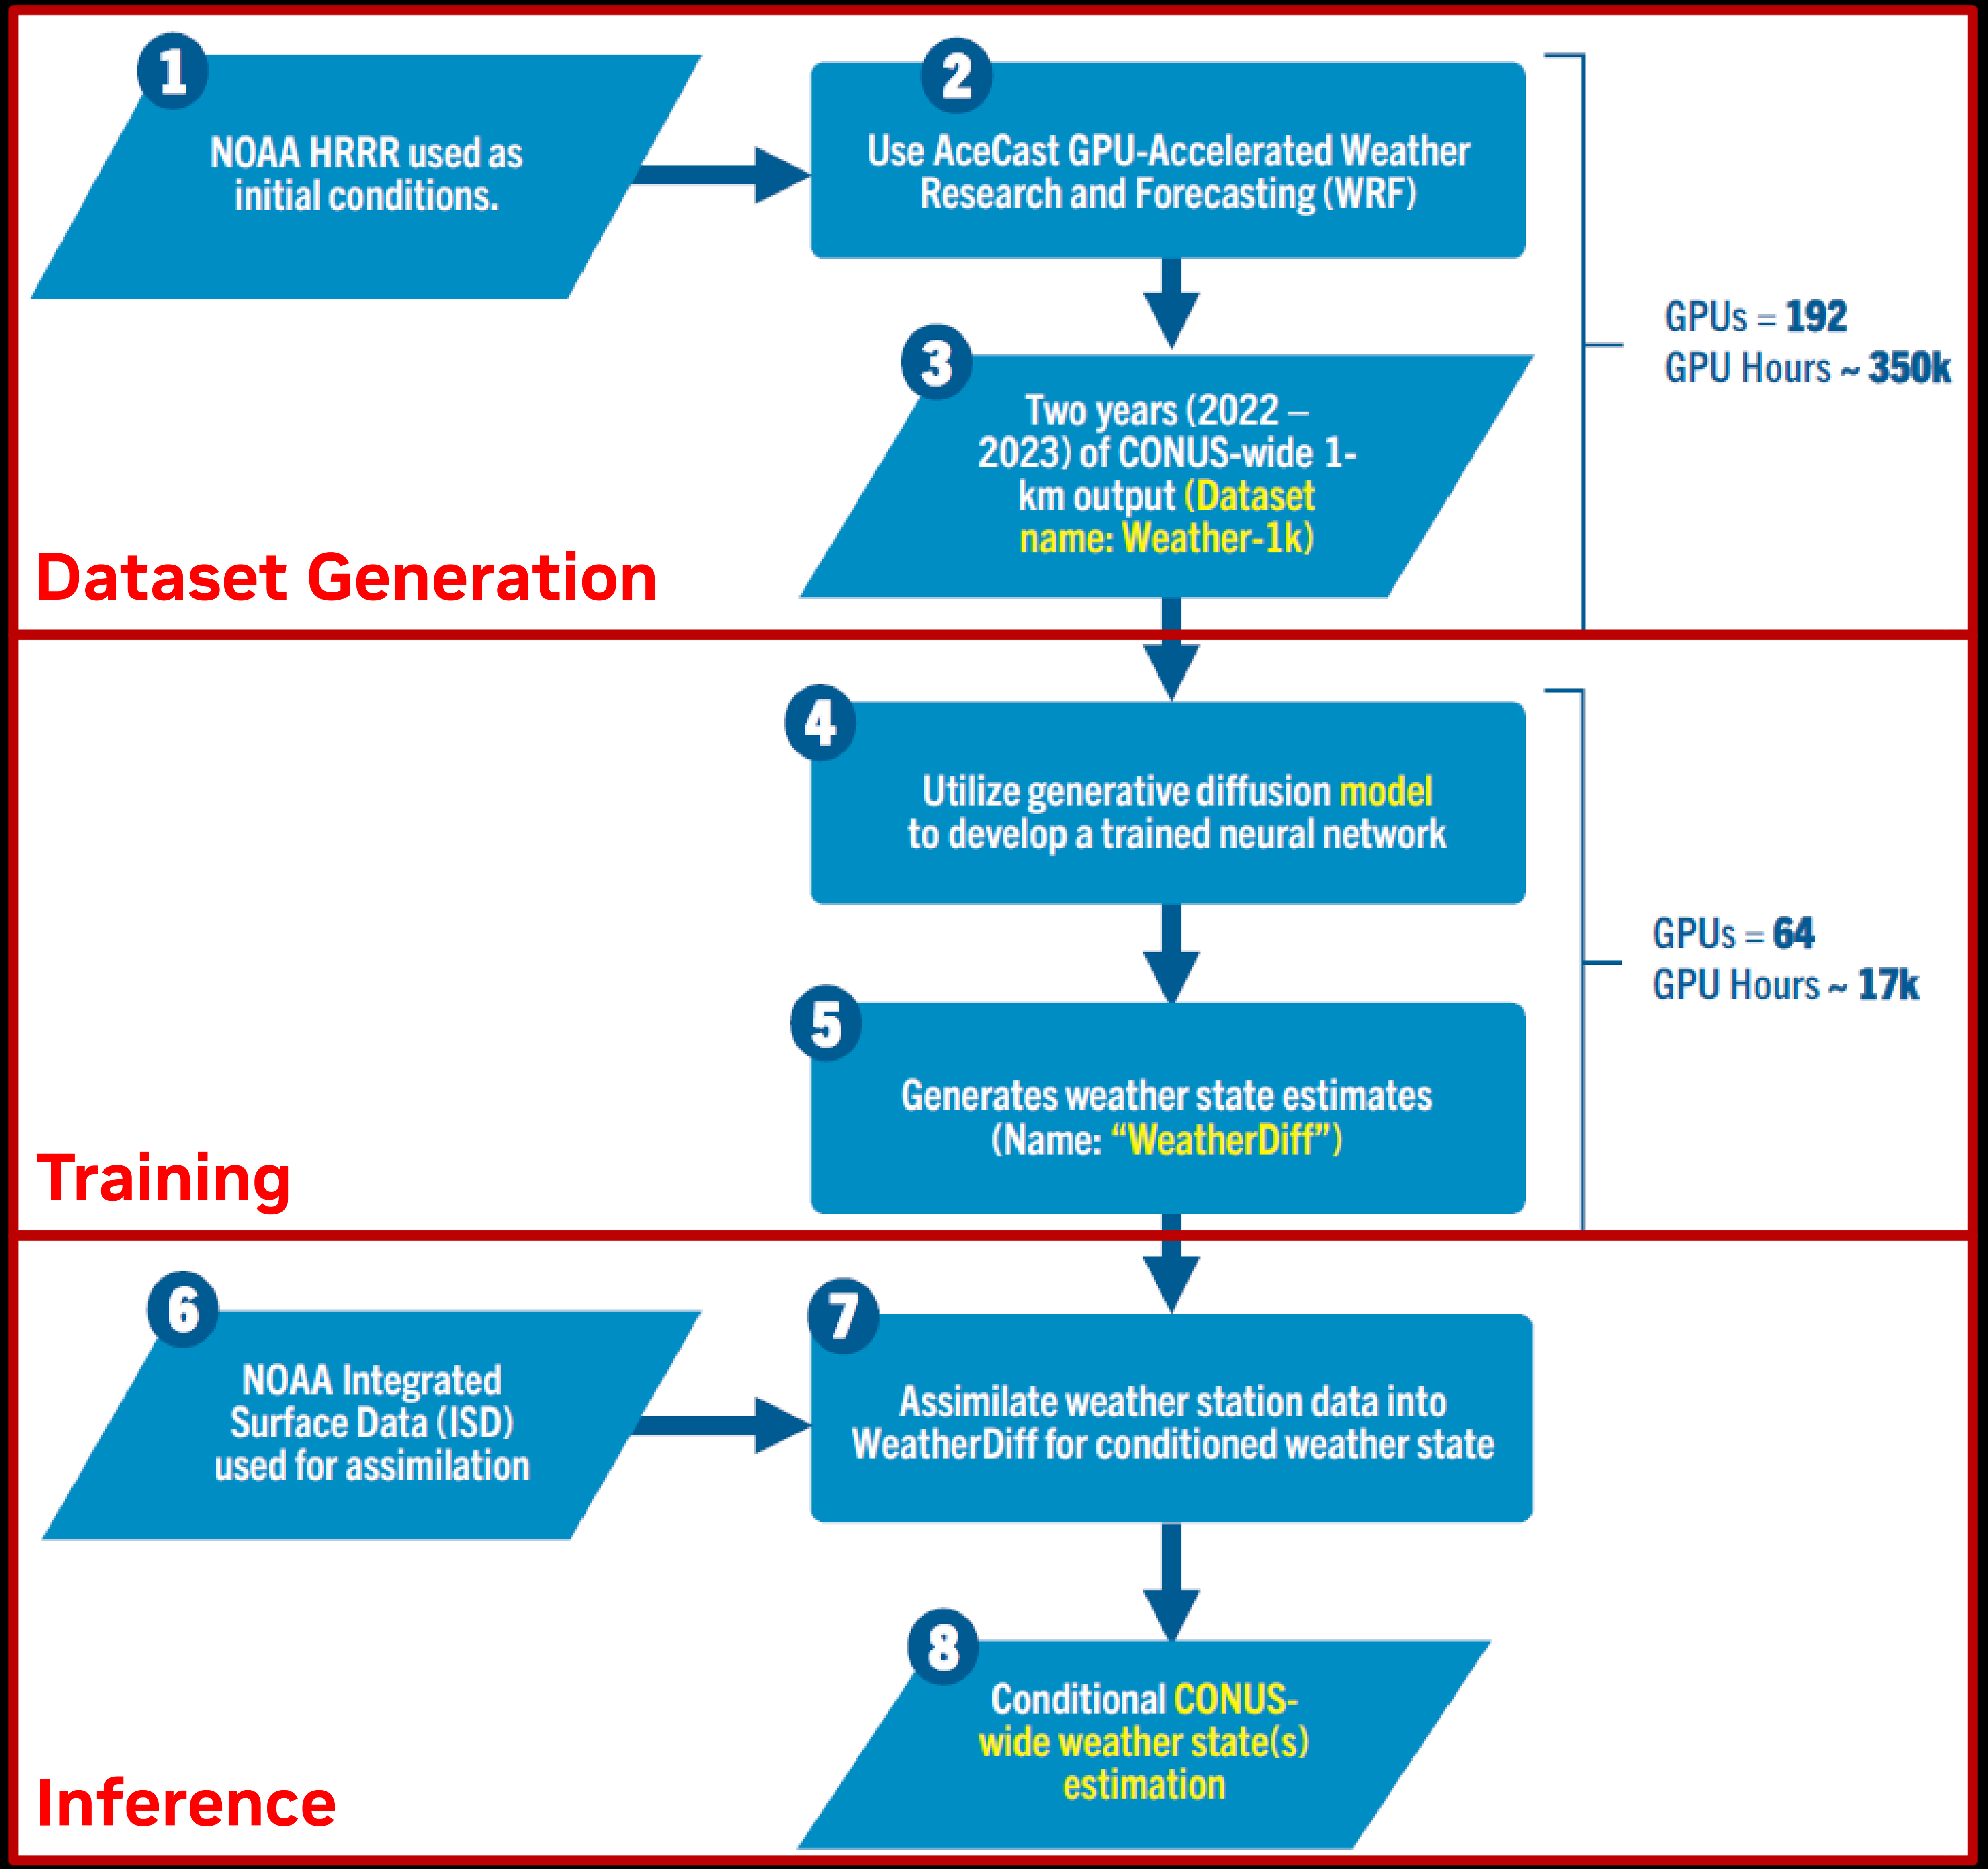

# Compute Challenges

## Parallelising the Forecasts

- WRF supports a hybrid scheme for 2D (i,j) domain decomposition
- The third dimension (k) is always handled within a single thread.
- The domain is broken into MPI ranks as patches, which can be further subdivided into OMP threads as tiles.
- We use MPI only, with the domain divided into 32 patches
- Optimal patch dimensions were 1x32 (1D-decomposition) for two reasons:
  - Simplifies the communications patterns as halos have coalesced memory access – much faster to pack/unpack N/S halos vs E/W
  - Improve I/O timings when writing out the patches (writes large, contiguous chunks rather than lots of small chunks).

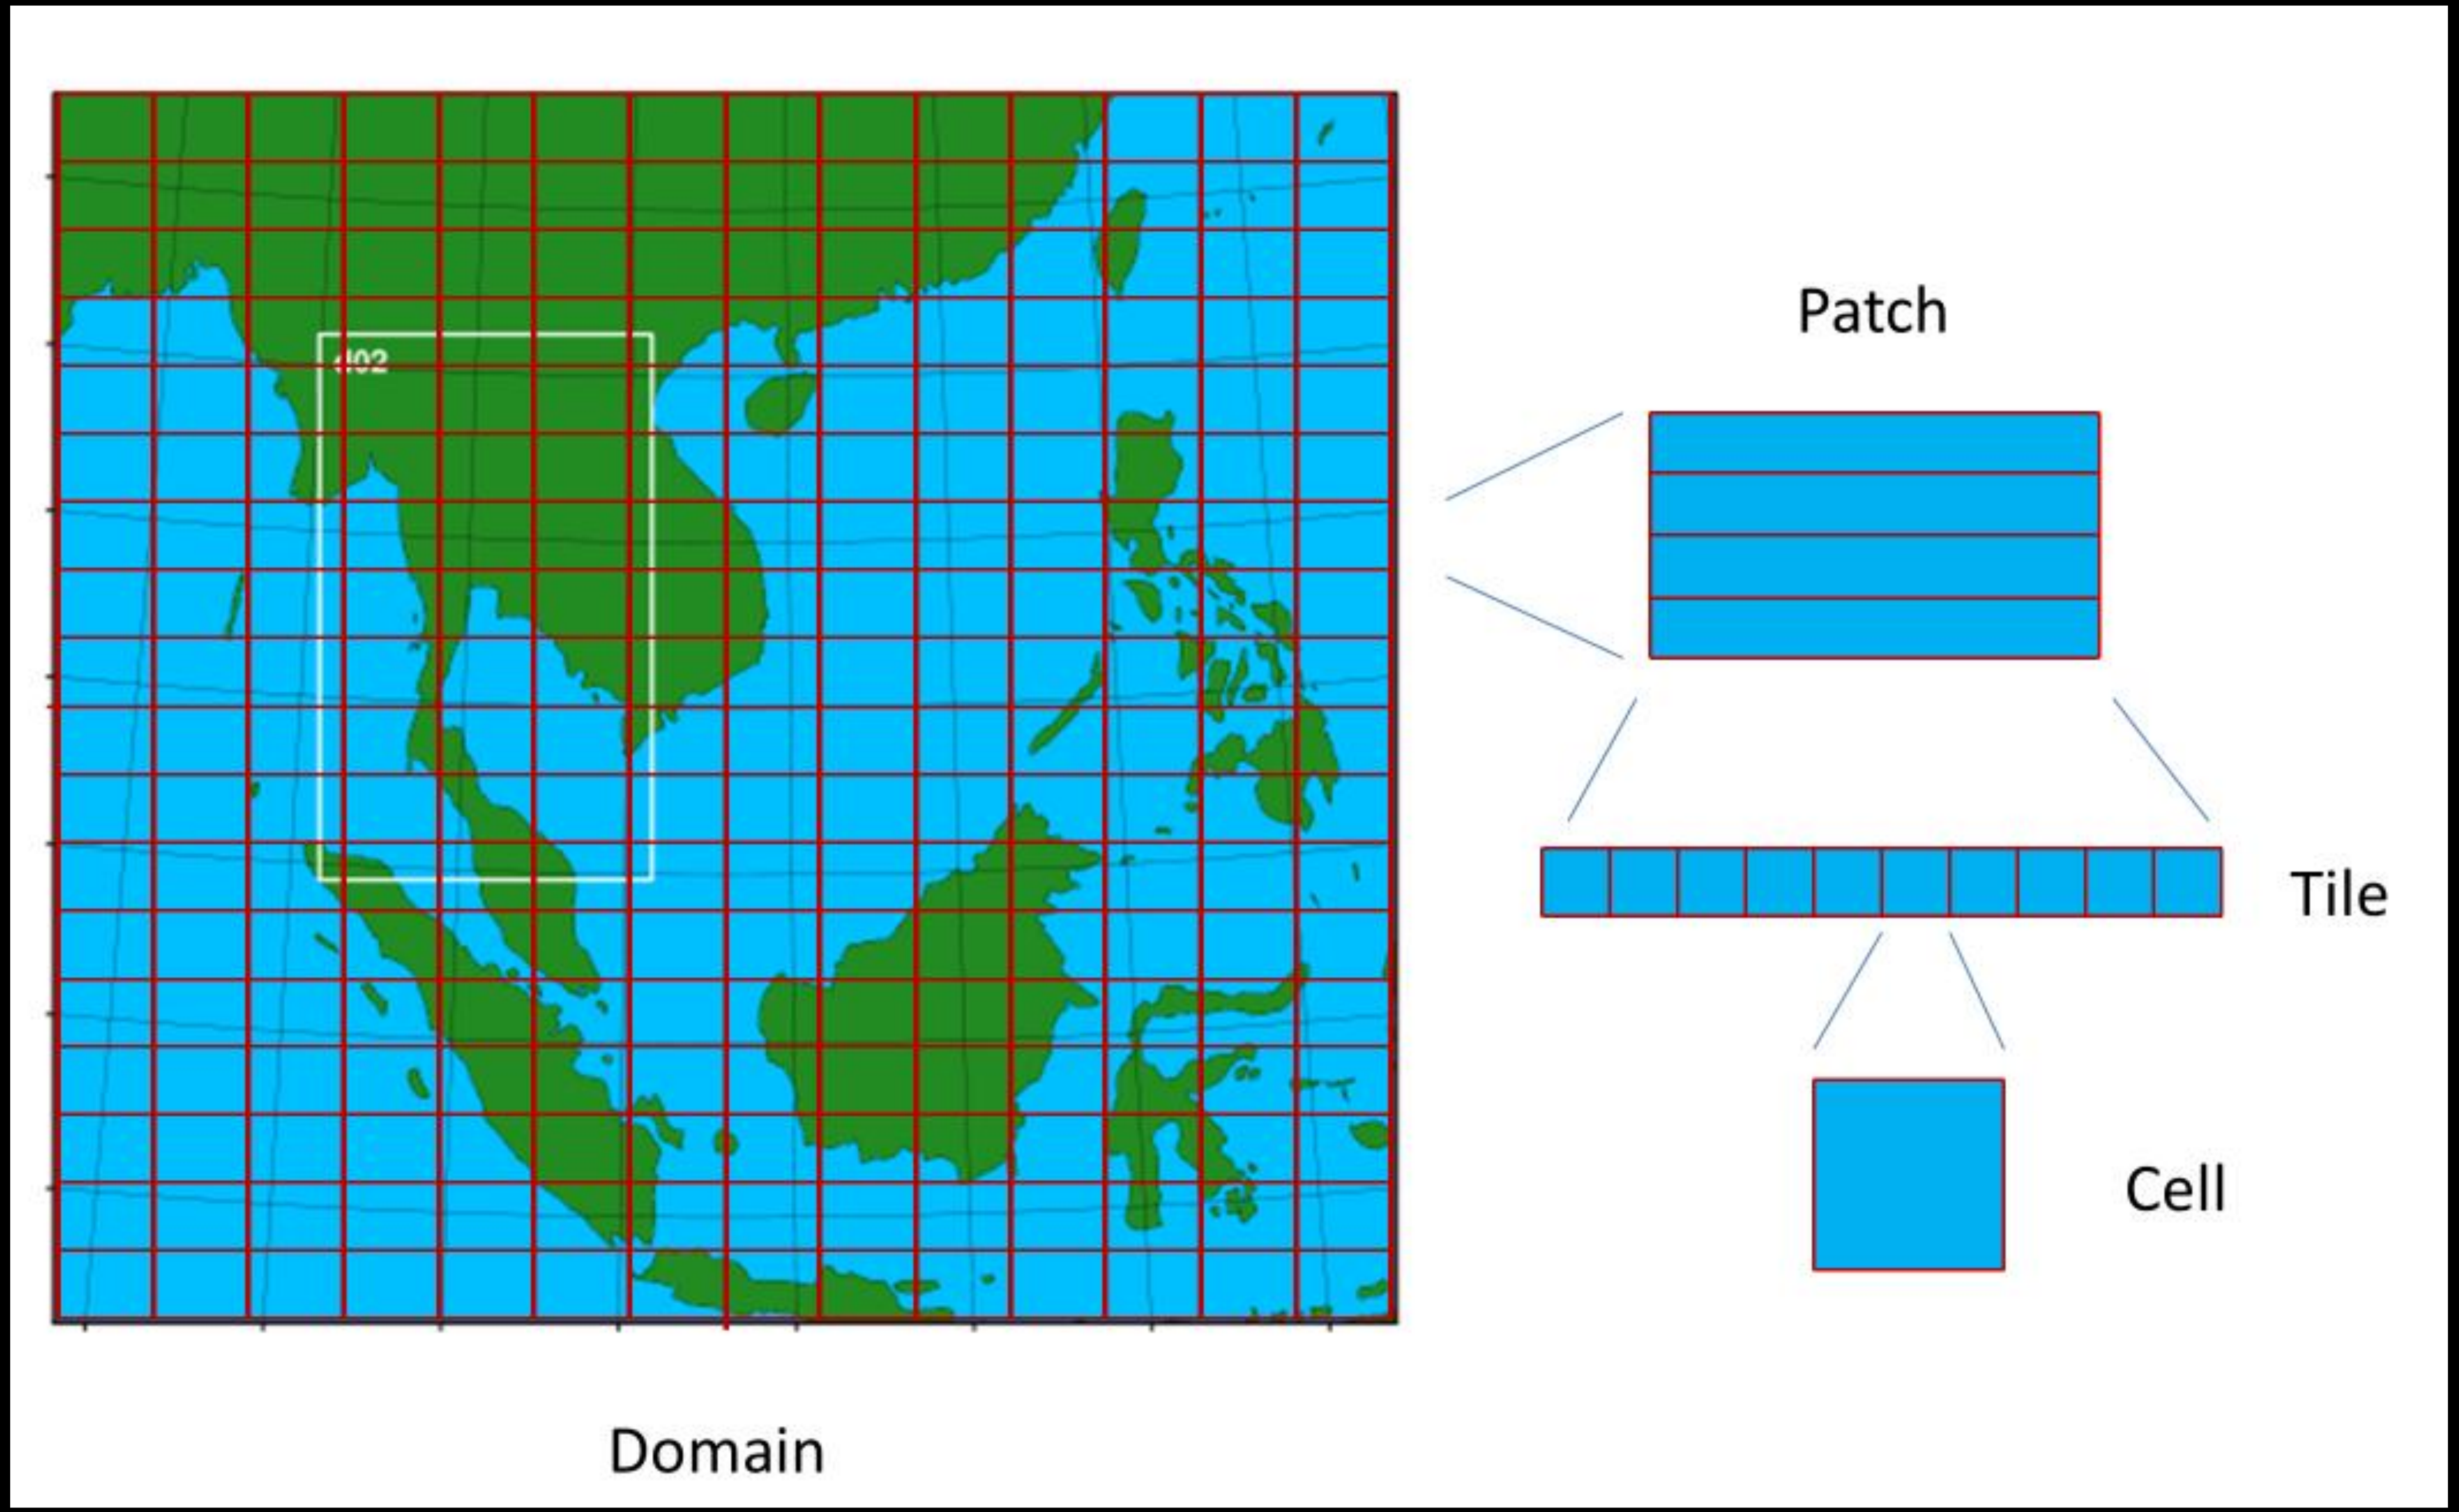

# Compute Challenges

## Profiling the runs

- Added custom timers using NVTX Markers

| Name | Time (s) | Time (%) |
|---|---|---|
| WRF Total | 58953.342068 | 100.00 |
| Initialization | 67.944473 | 0.12 |
| Allocate | 15.950707 | 0.03 |
| I/O (Read) | 40.807770 | 0.07 |
| I/O (Write) | 0.000000 | 0.00 |
| HALO/Nesting (MPI) | 0.328014 | 0.00 |
| HALO/Nesting (non-MPI) | 0.078611 | 0.00 |
| Compute/Other | 10.779371 | 0.02 |
| Integration | 58885.397560 | 99.88 |
| I/O (Read) | 101.059078 | 0.17 |
| I/O (Write) | 22916.229178 | 38.87 |
| HALO/Nesting (MPI) | 3545.977074 | 6.01 |
| HALO/Nesting (non-MPI) | 108.324041 | 0.18 |
| Compute/Other | 32213.808190 | 54.64 |
| LW Radiation | 142.295548 | 0.24 |
| SW Radiation | 249.819088 | 0.42 |
| Surface Layer | 725.574731 | 1.23 |
| Land Surface | 2994.578241 | 5.08 |
| PBL | 8966.006227 | 15.21 |
| Cumulus | 0.000000 | 0.00 |
| Microphysics | 12520.291206 | 21.24 |

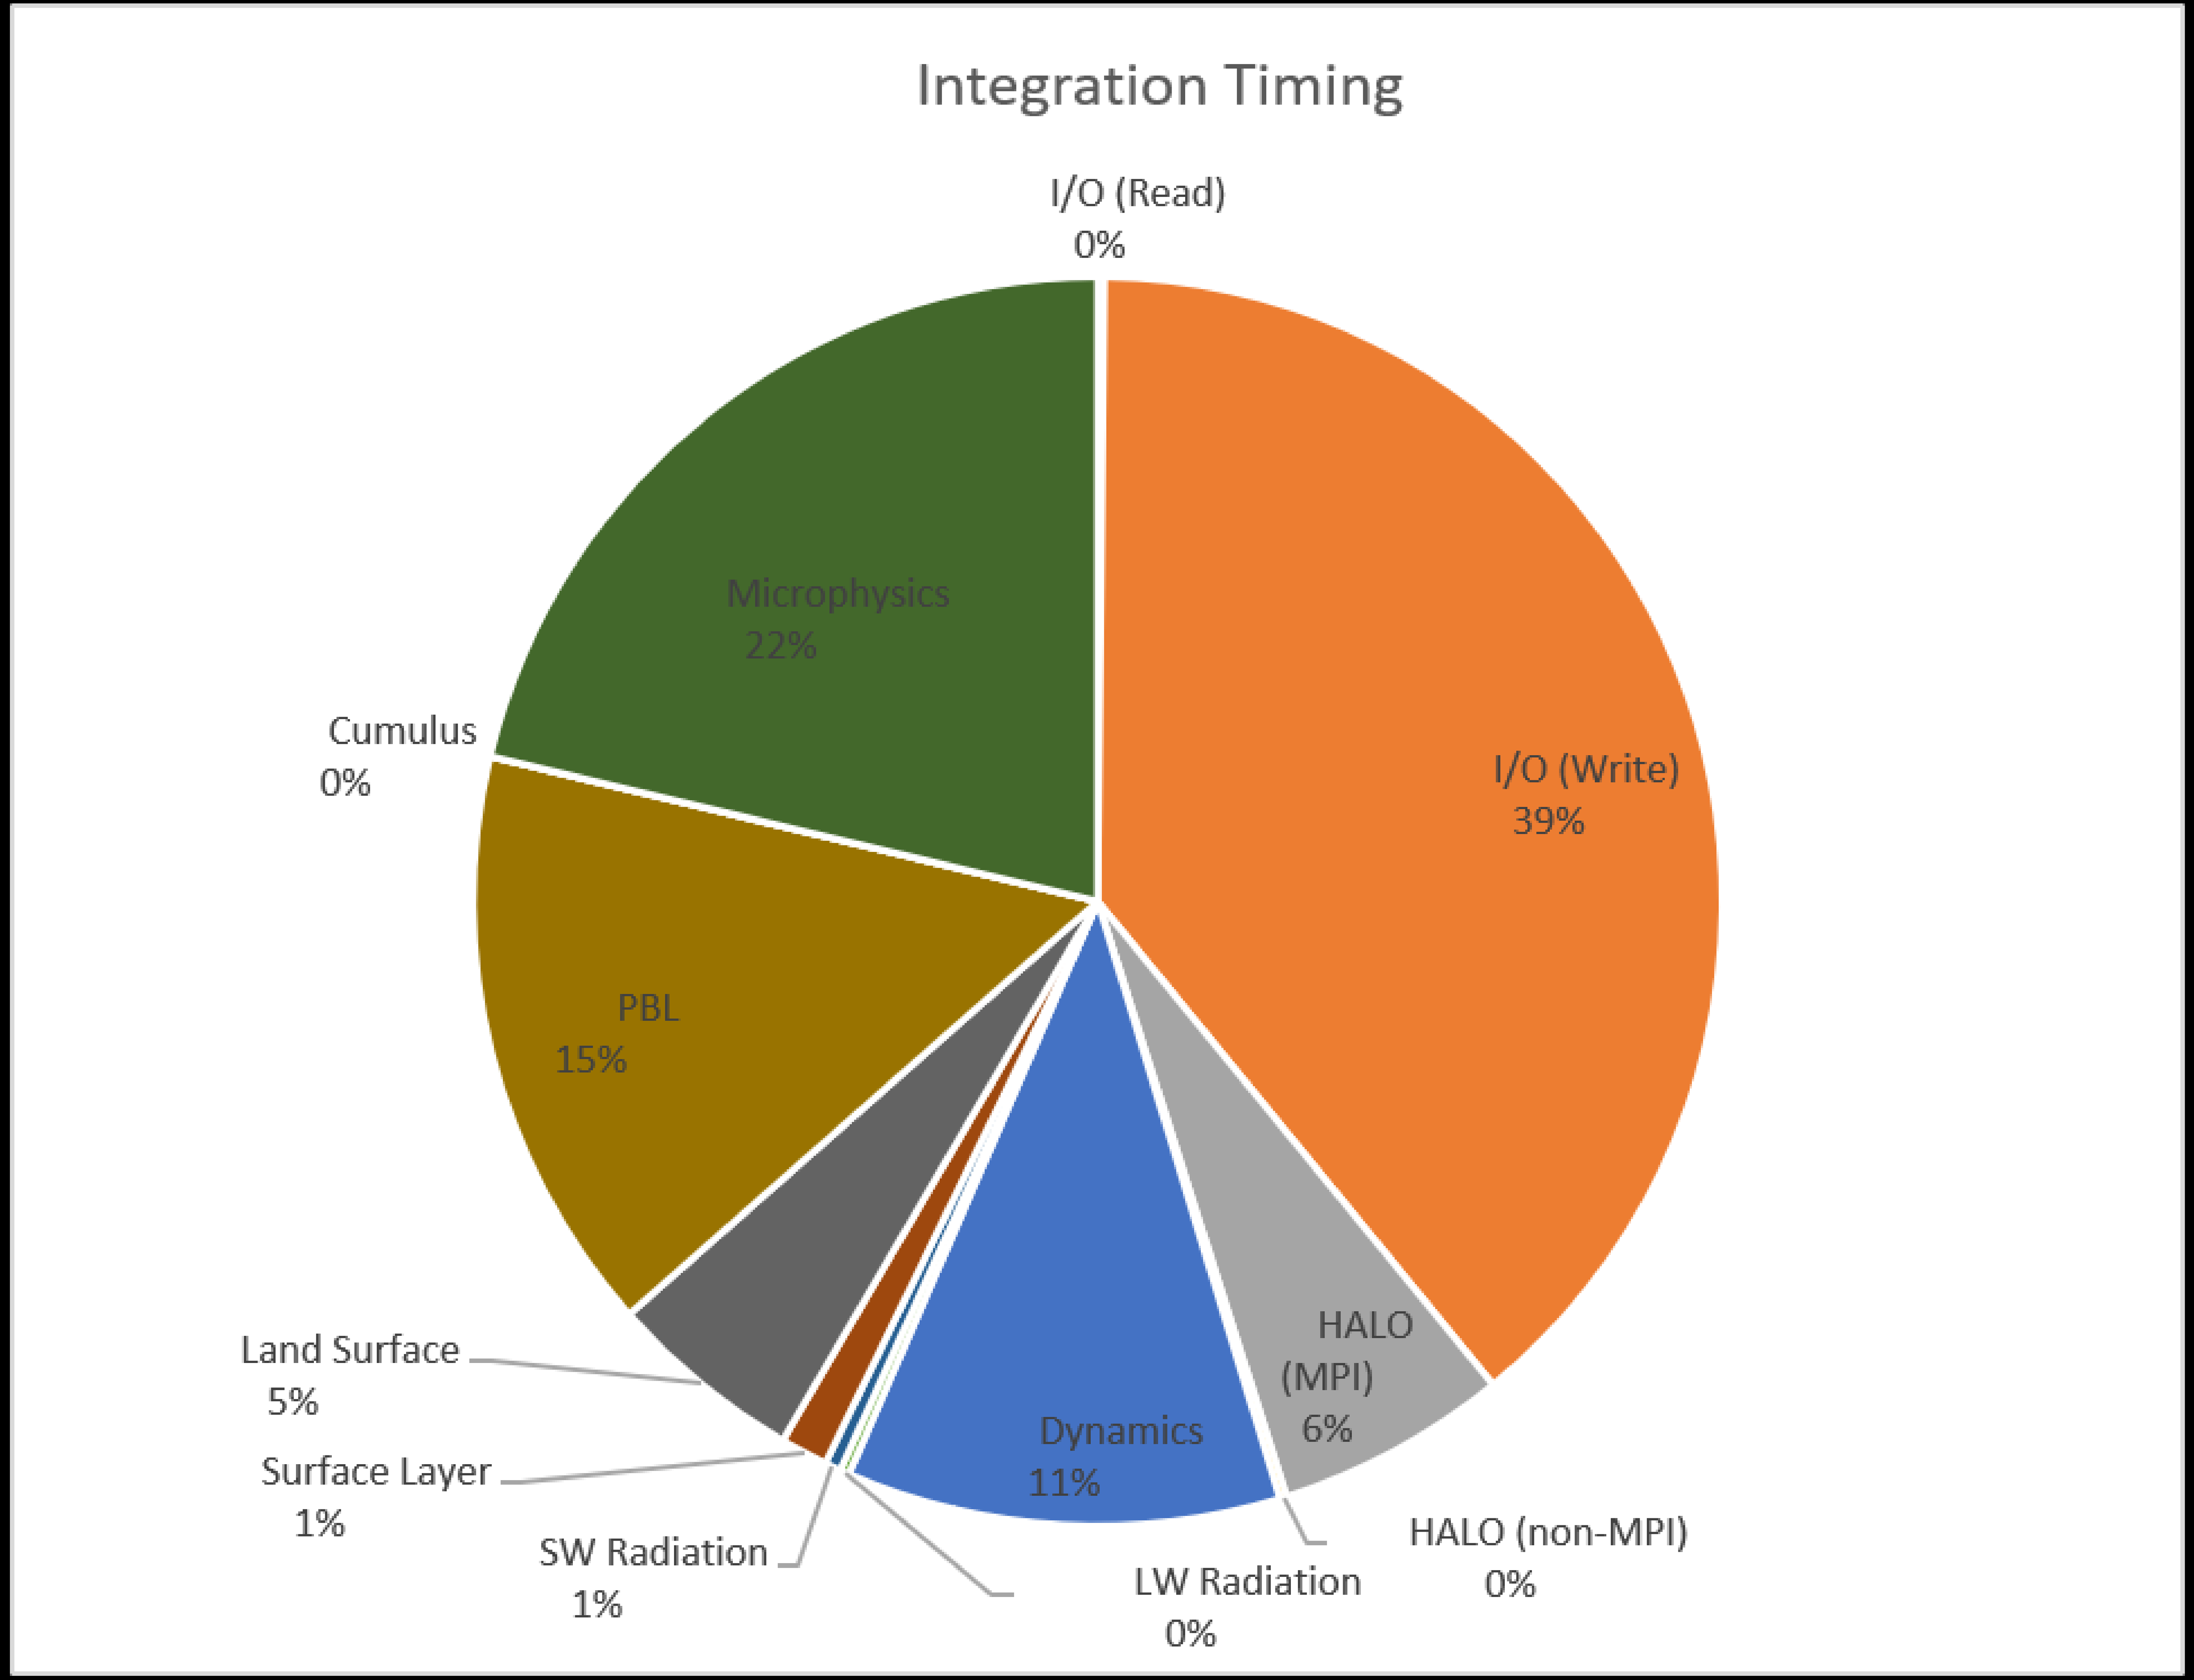

*“A supercomputer is a device for turning compute-bound problems into I/O-bound problems”* – Ken Batcher

# Data Challenges

## 8PB of data on a 9PB Storage system!

- The training dataset is huge:
  - Produced 730 forecasts (2 years from 2022-2023)
  - Each forecast is 24 hours, saved at 10 minute intervals (144 samples)
  - 5200x3300x51 = 875 Million gridpoints, 144 timesteps per forecast
  - Approximately 11 TB of raw data per forecast
- Initially used Parallel I/O to write to Lustre filesystem
  - I/O was 40% of the total Forecast generation time
  - Experimented with I/O quilting, but discovered per-rank splitfiles to be 43% faster than parallel-IO
  - Cost is the (asynchronous) recombination of the patch splitfiles
- Developed an I/O postprocessing pipeline
  - Take NetCDF files from WRF as they are output to disk
  - Interpolate vertical axis from model levels to pressure levels
  - Compress data with lossless compression scheme
  - Output to AI-Friendly zarr format
  - Final output: ~4PB

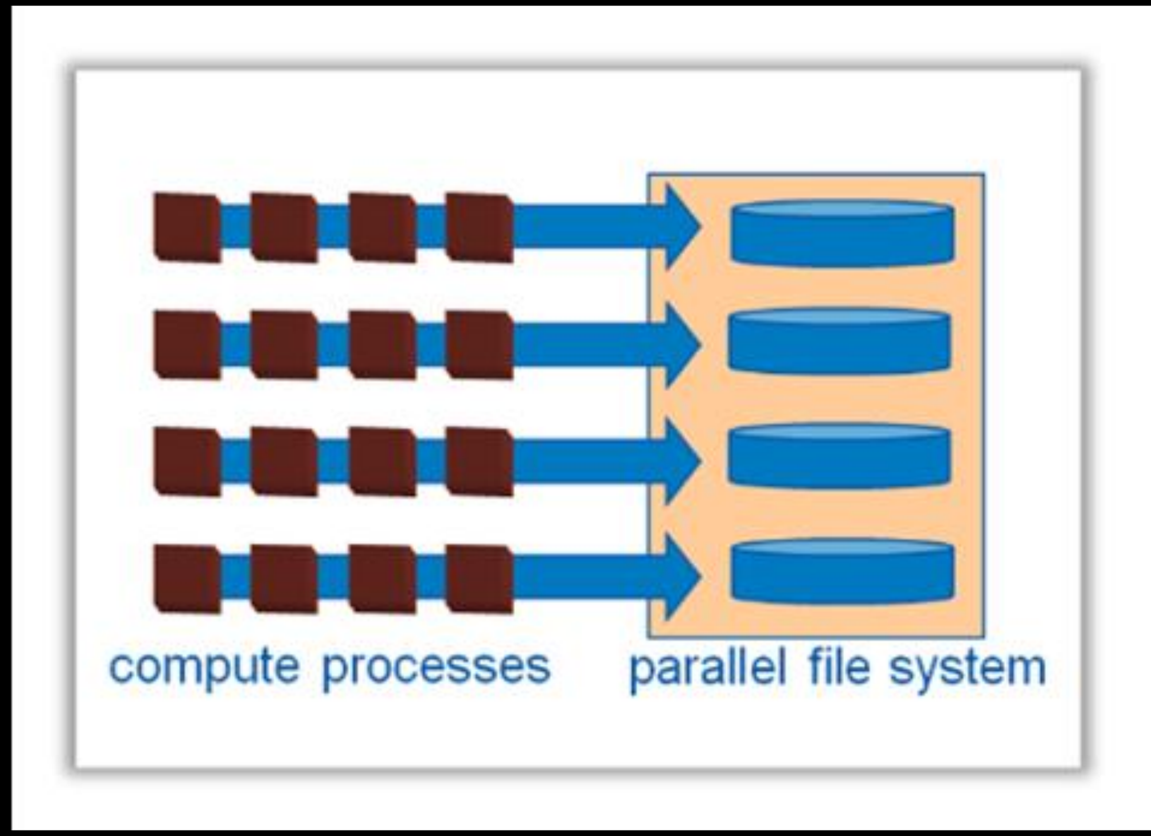


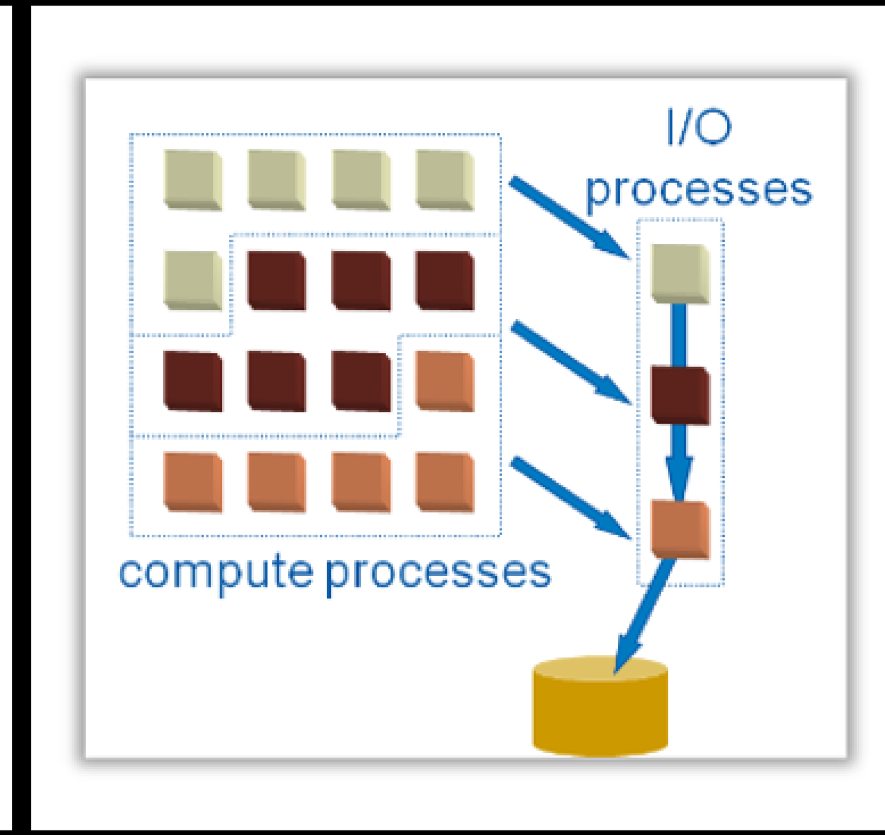


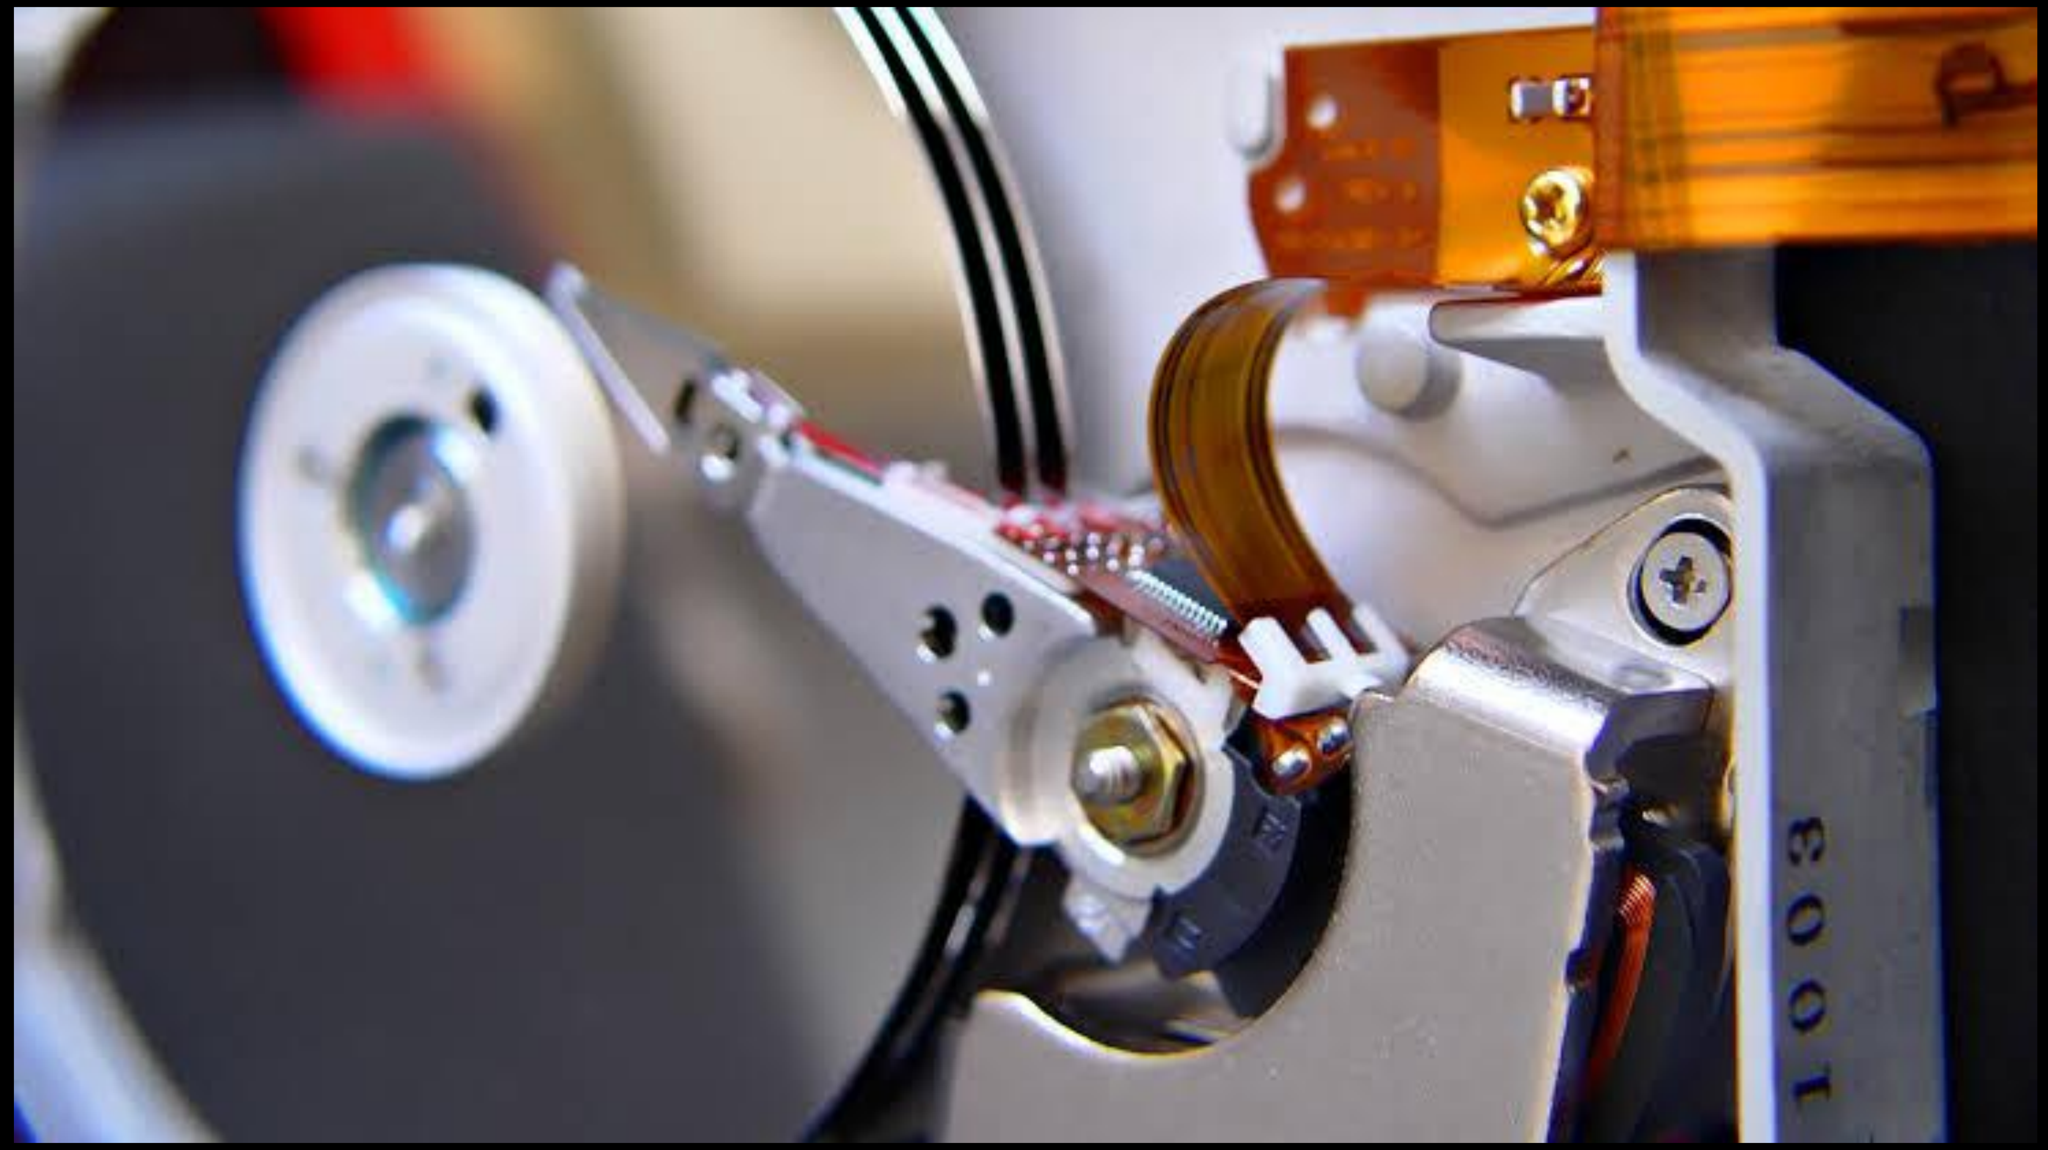

# Results

# MITRE – USA

Corrdiff – High Resolution 1km HRRR

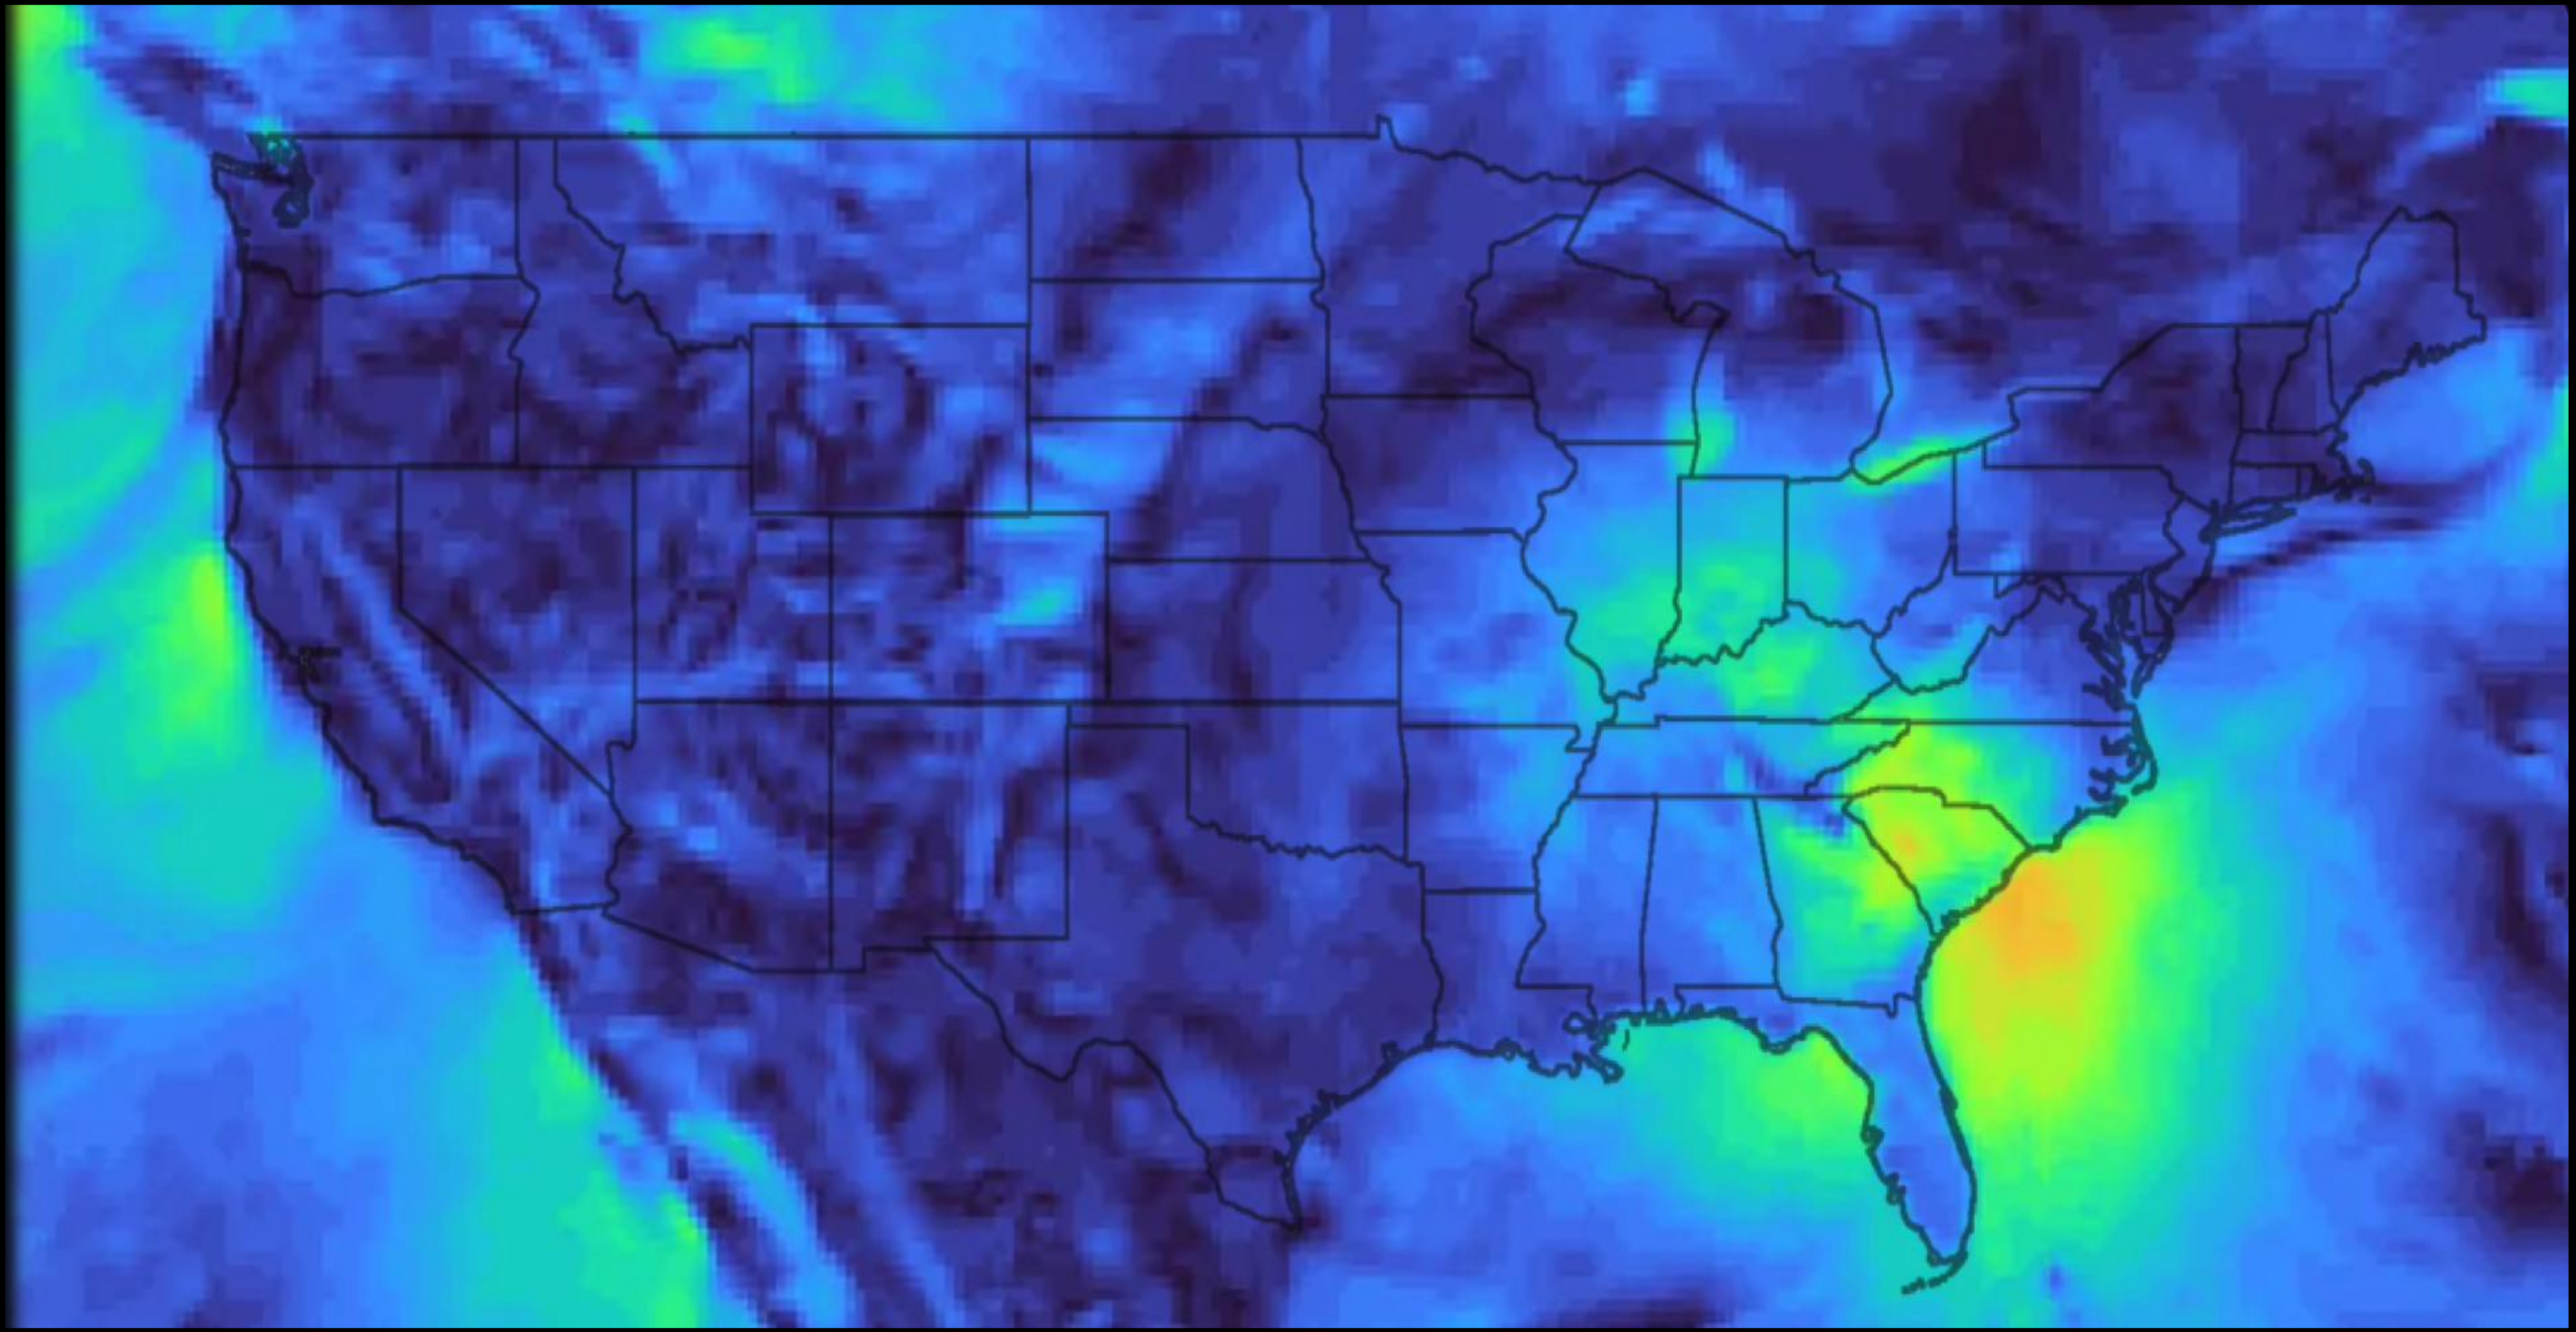

# MITRE – USA

## Corrdiff – High Resolution 1km HRRR prediction of HRRR precipitation type from GEFS 25-km inputs

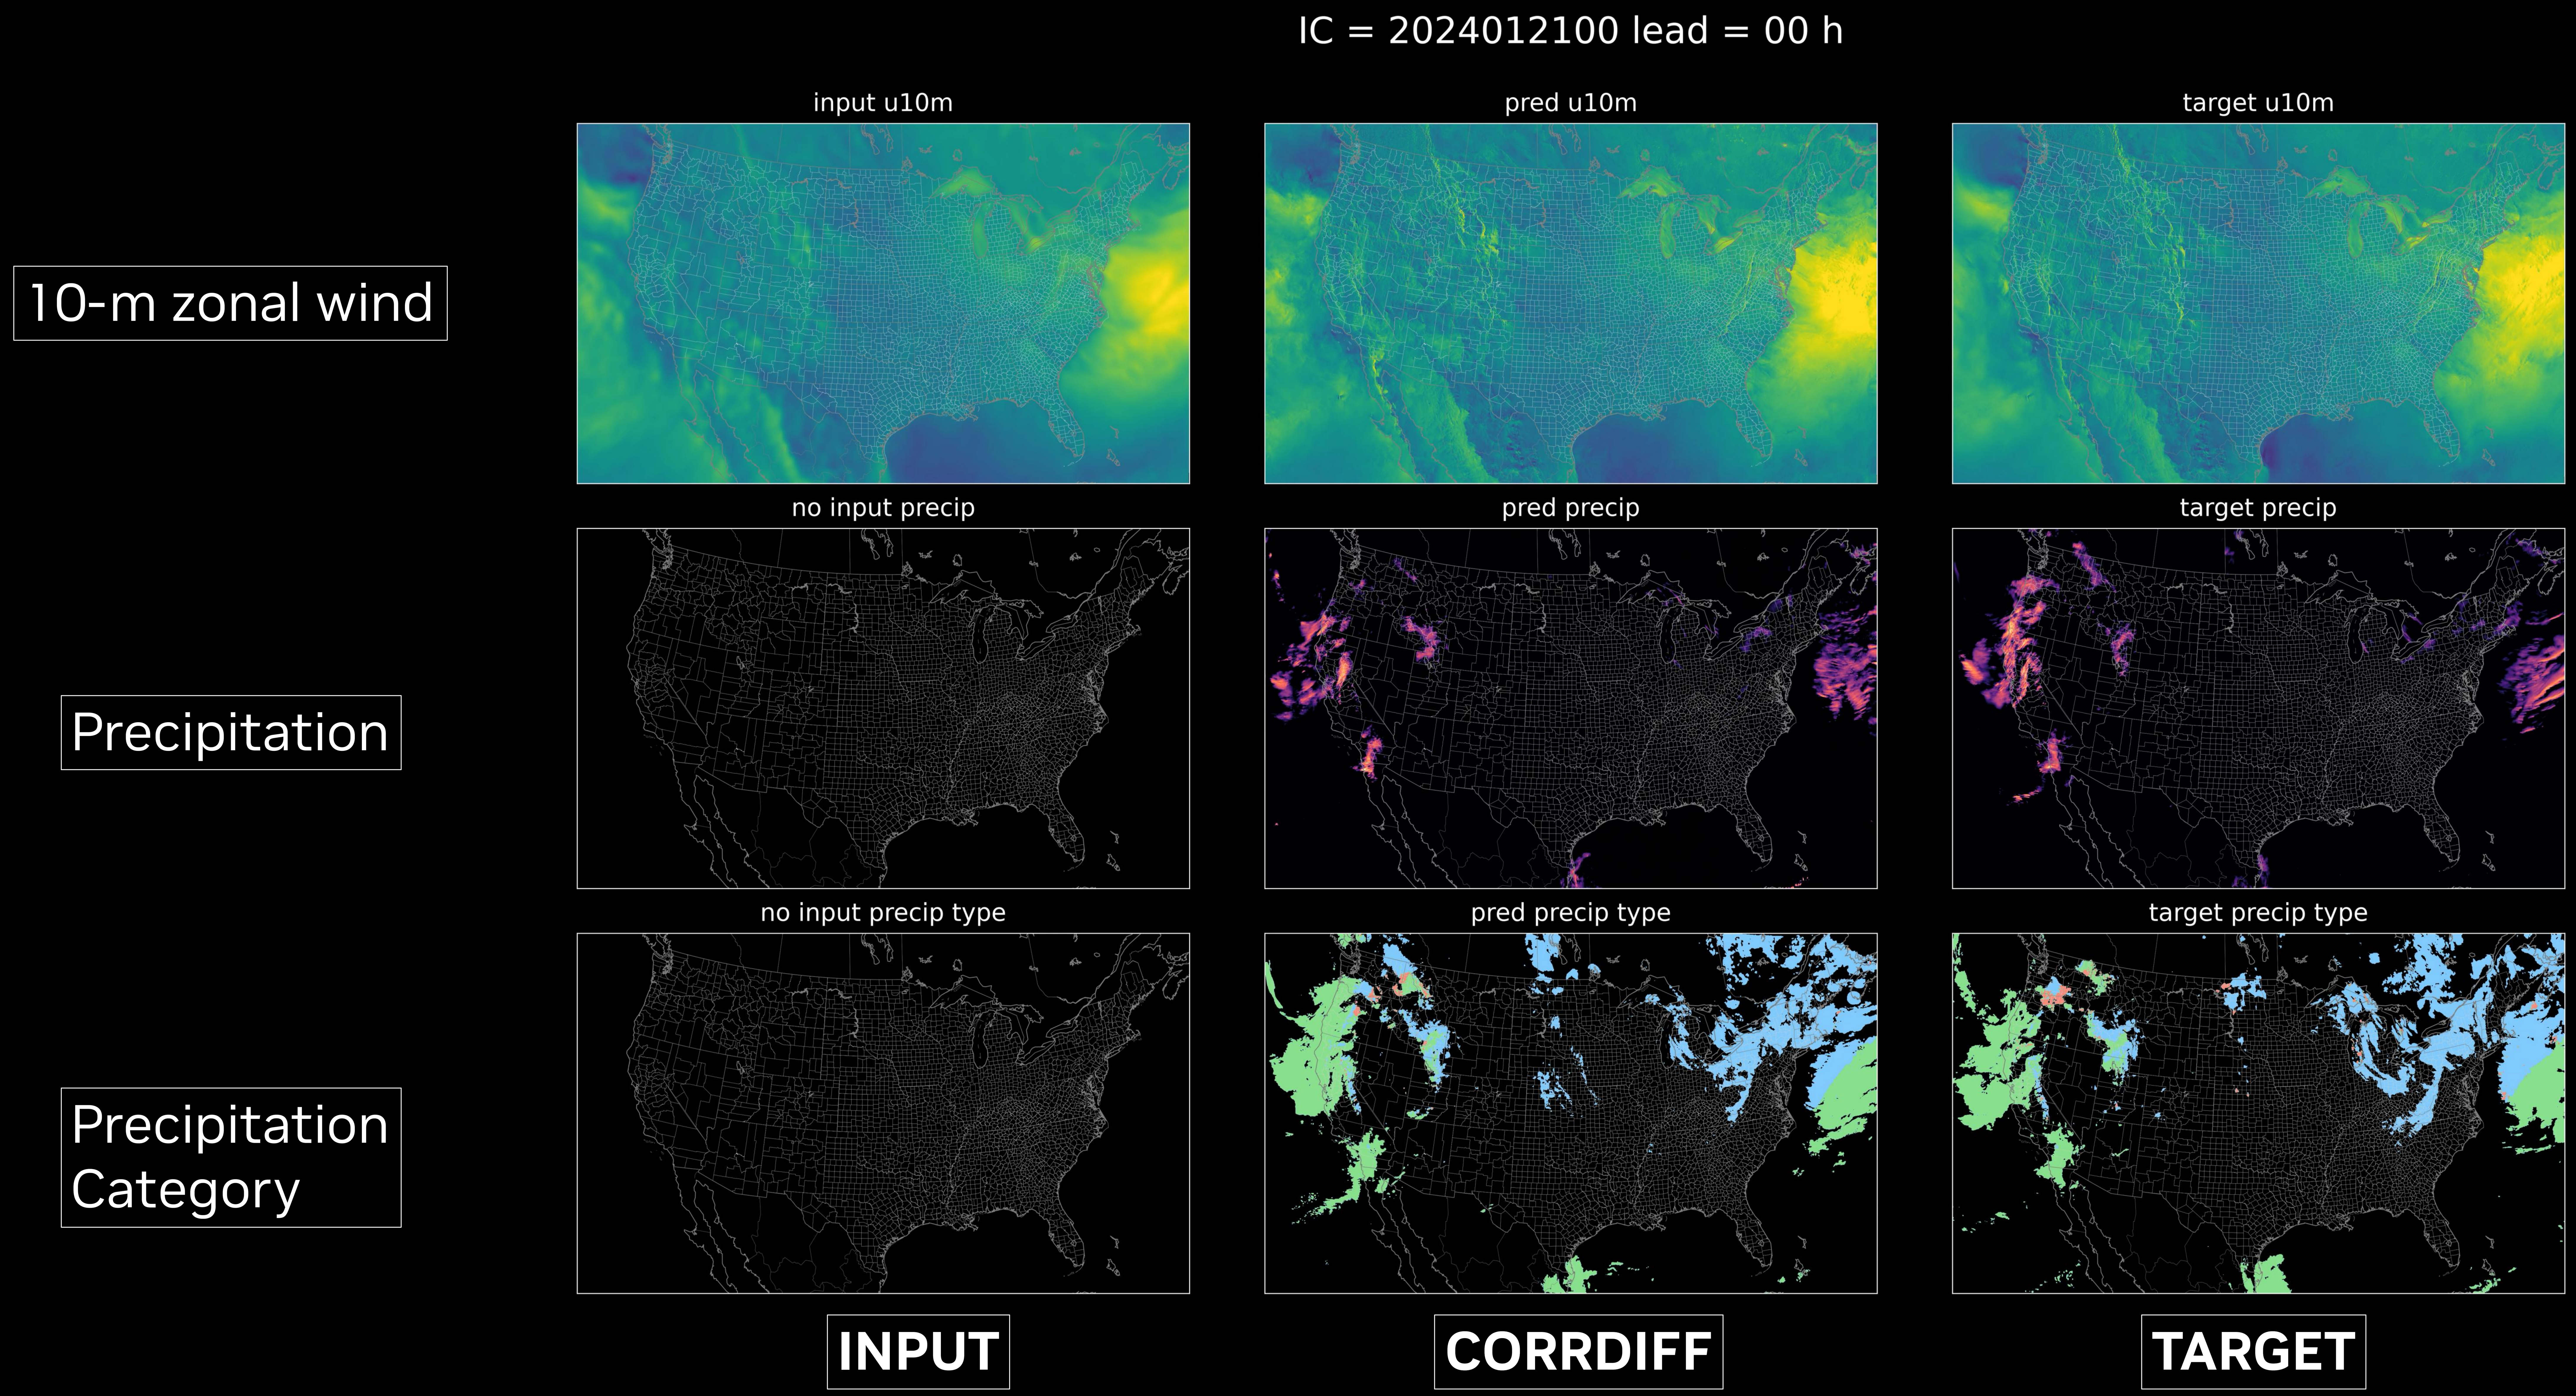

# Fire Weather: 3-km to ~100-m Resolution AI weather forecast pipeline

NVIDIA

CorrDiff

HRRRCast-1km

GSL

MITRE Weather1k

CFD Emulator

*Deterministic 3-km forecasts of all variables*

- Collaboration with Dr. Isidora Jankov (GSL)

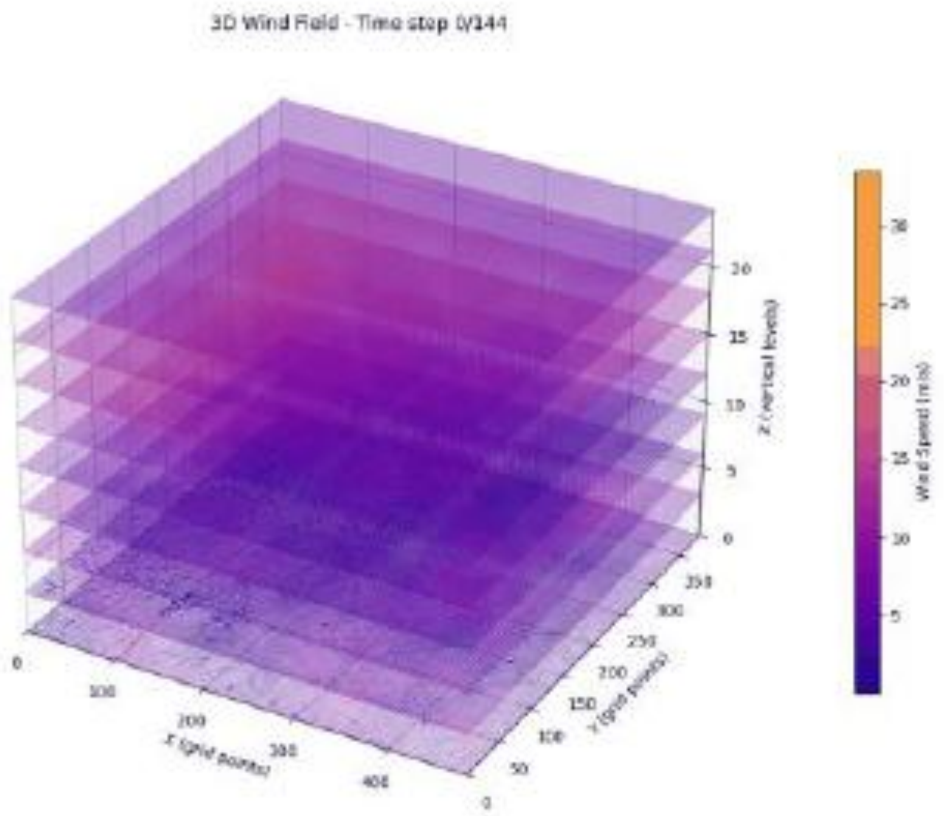

*Probabilistic 1-km forecasts of relevant fire weather variables*

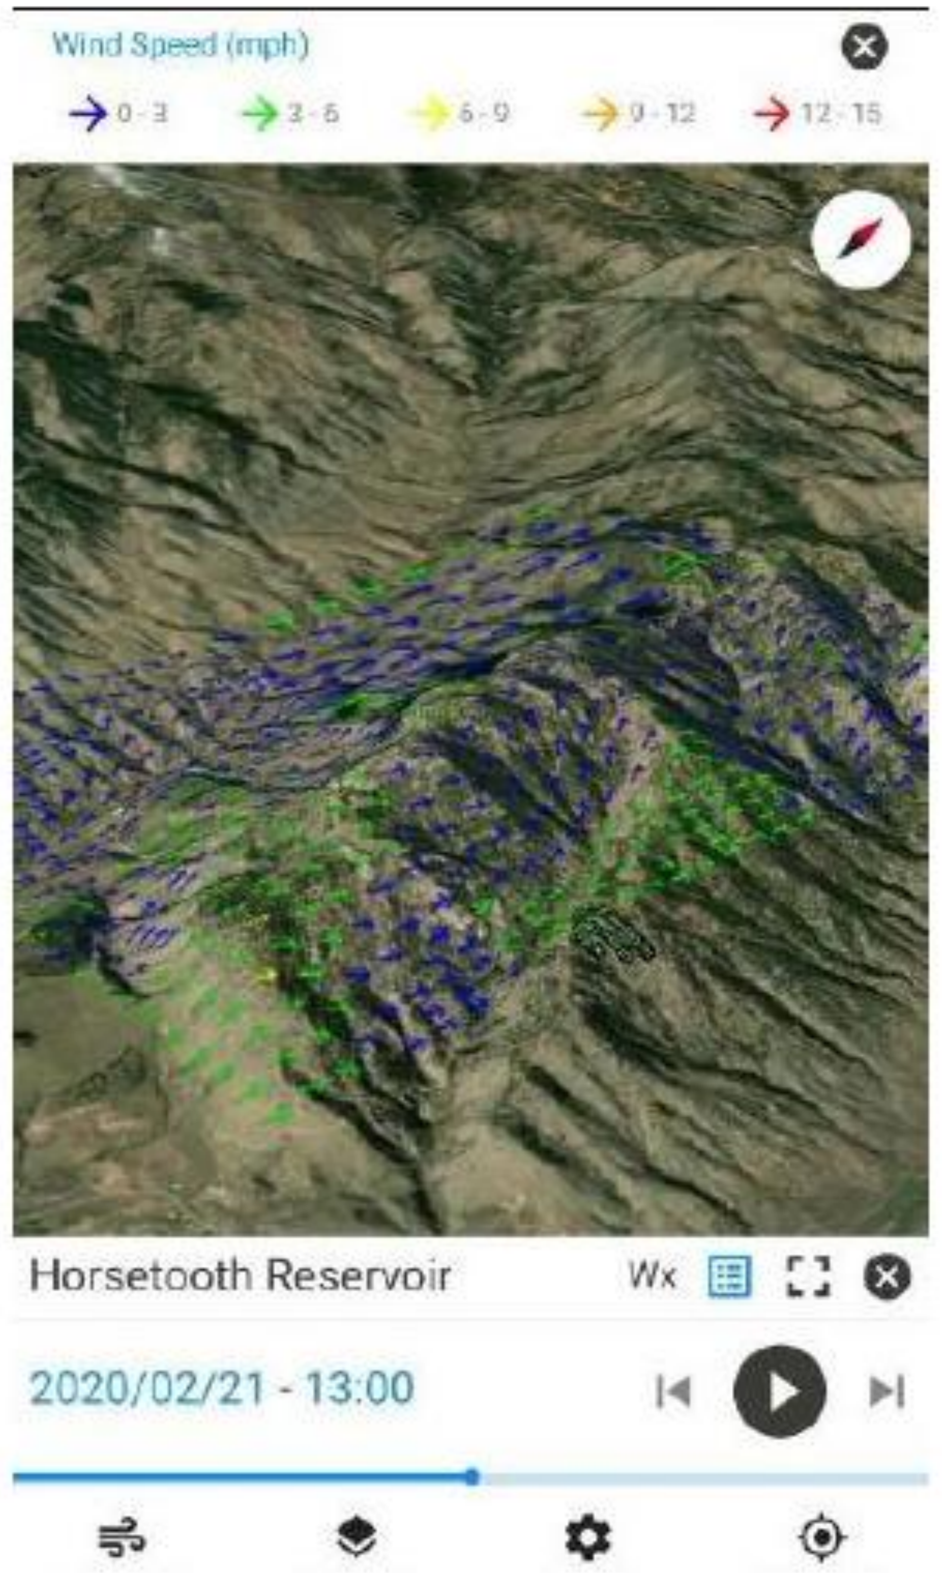


*Deterministic ML wind forecasts at ~100-m*

- Collaboration with USFS

# National Pilot Program for Unified Fire Warning System

- Leverage foundational science needed to support a unified fire warning system in the United States
- Develop actionable implementation pathway WFOs can adopt and pilot at select locations
- Collaboration with the University of Oklahoma

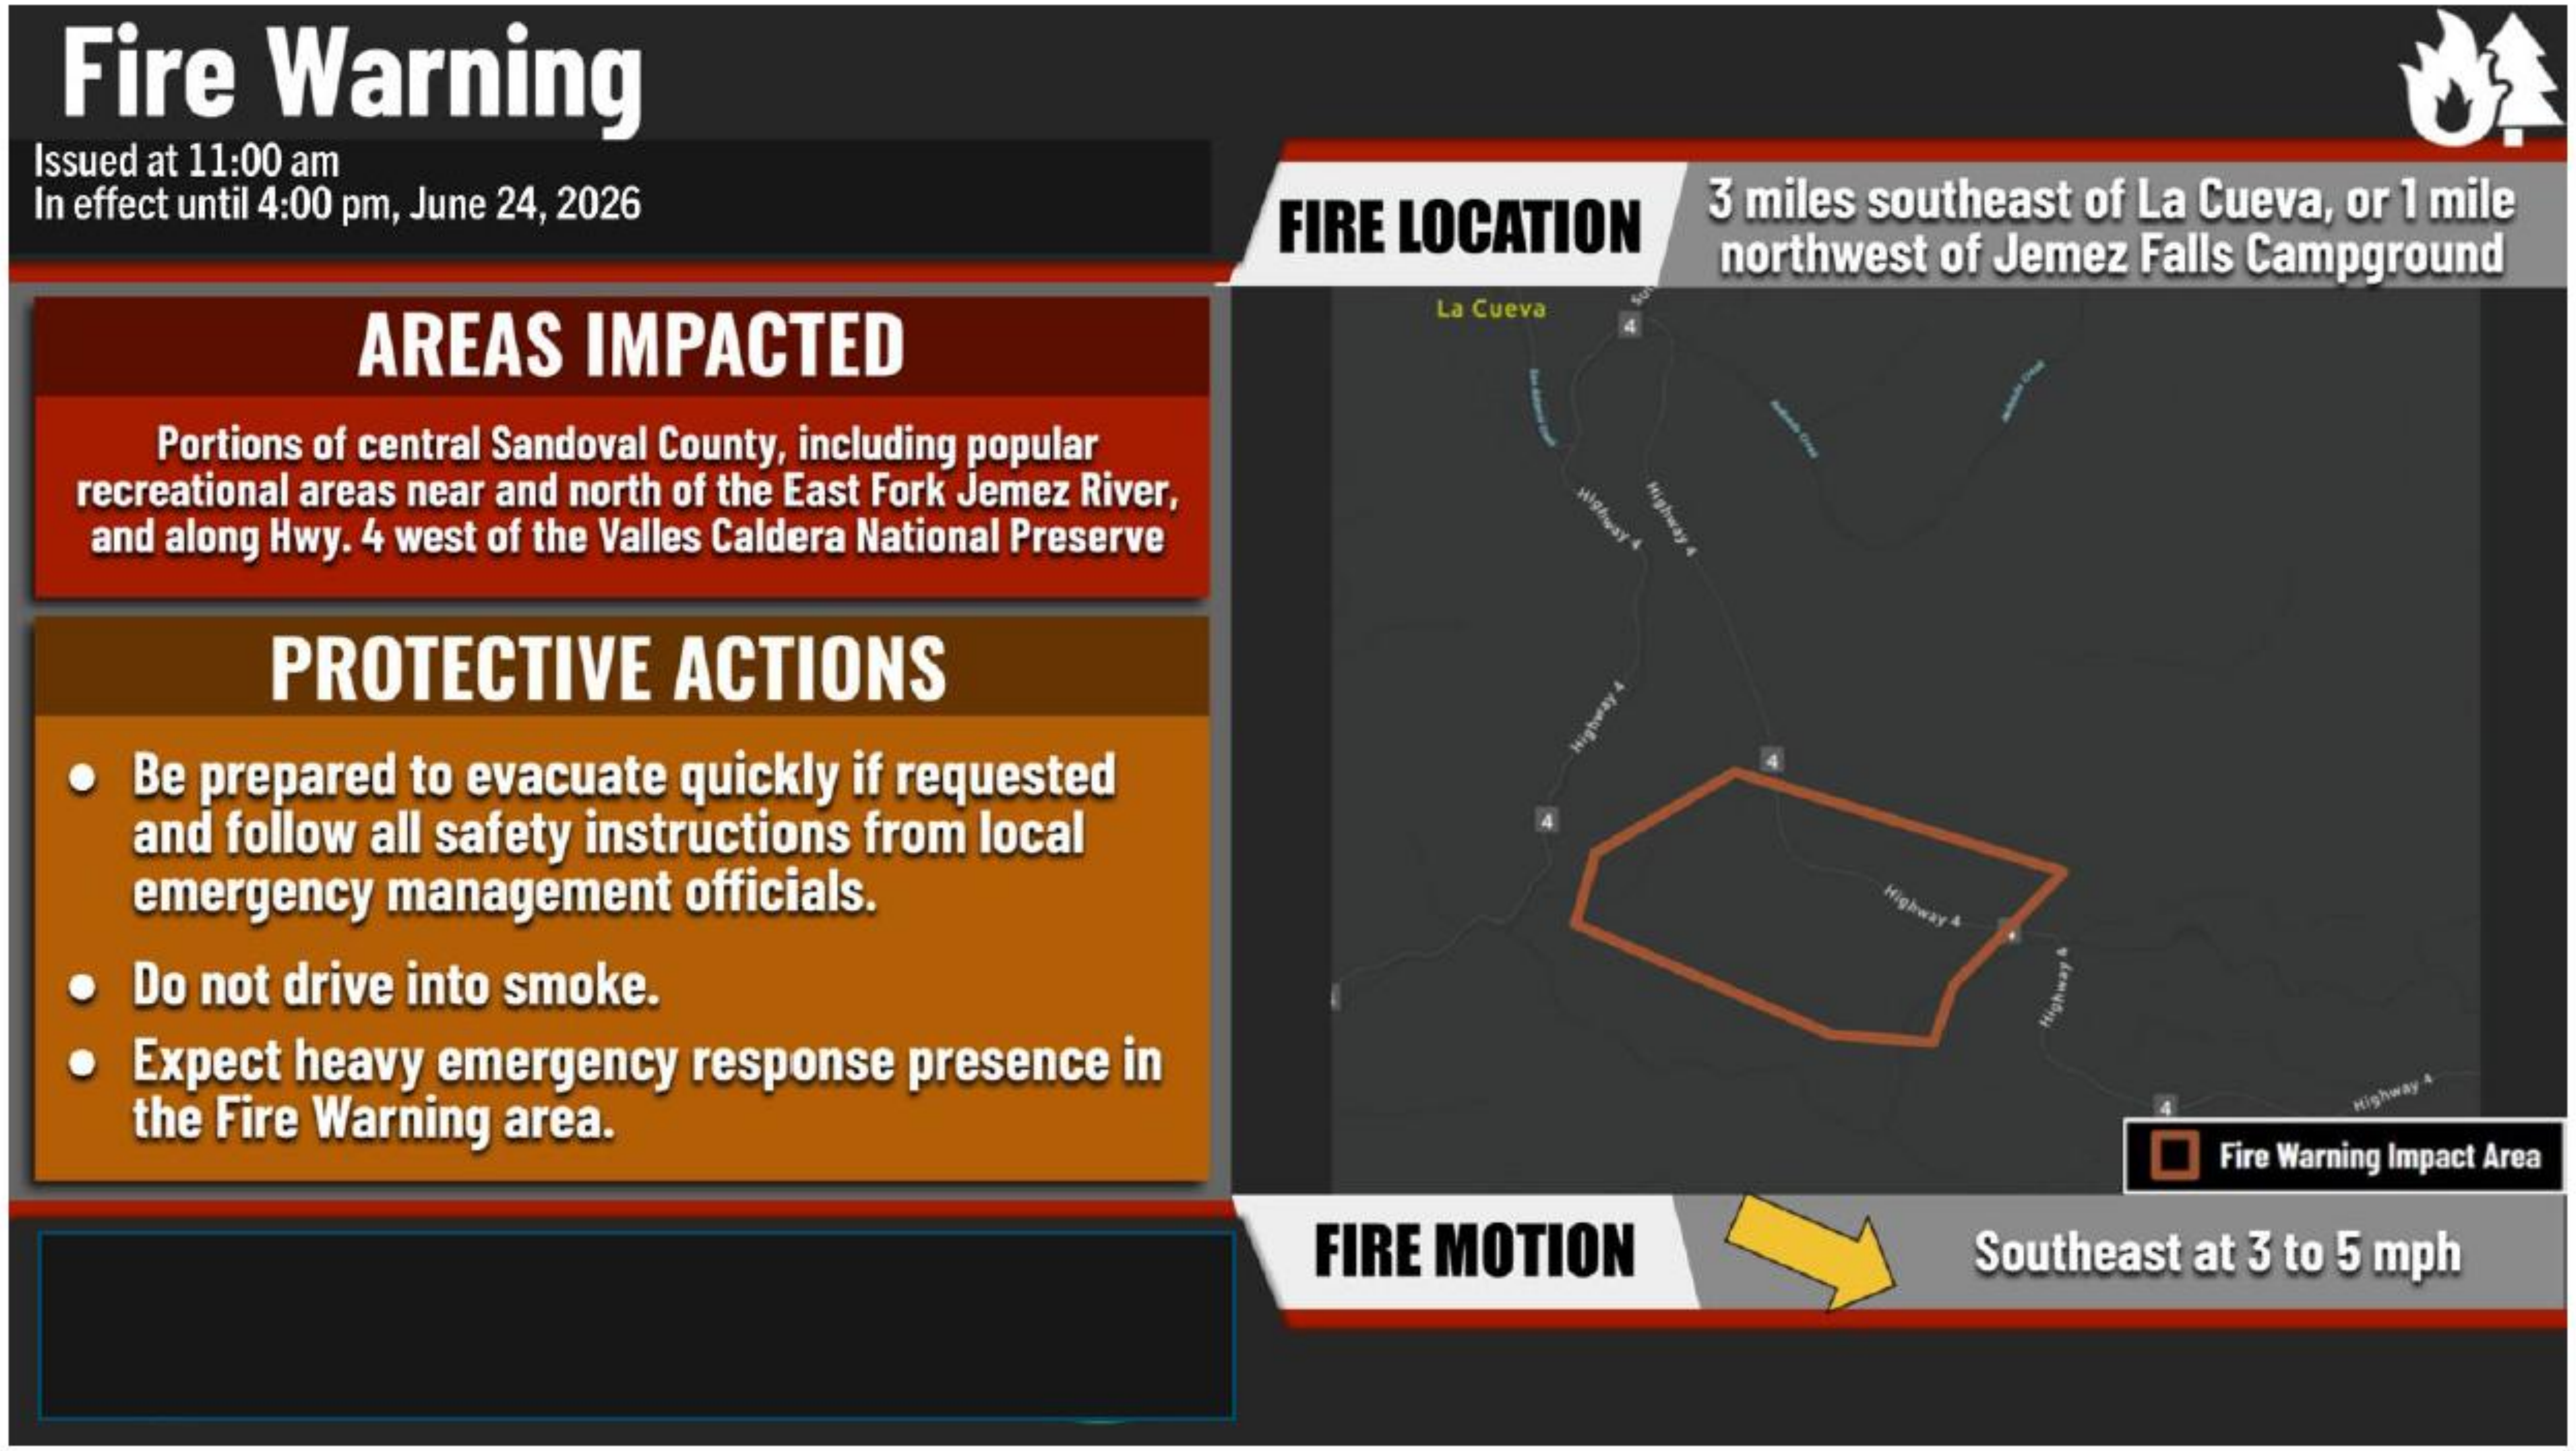


Example fire warning based on McCauley Springs Fire on June 24, 2026 in New Mexico. The image above is being used as an example illustration and was not issued in this format by any jurisdiction.

MITRE

jadie@nvidia.com

Thank You!

`