




# Accelerating Time-To-Science in Geophysical Simulations

Ignacio Sarasua & Filippo Spiga


[isarasua@nvidia.com](mailto:isarasua@nvidia.com) [fspiga@nvidia.com](mailto:fspiga@nvidia.com)




THIS INFORMATION IS INTENDED TO OUTLINE OUR GENERAL PRODUCT DIRECTION. MANY OF THE PRODUCTS AND FEATURES DESCRIBED HEREIN REMAIN IN VARIOUS STAGES AND WILL BE OFFERED ON A WHEN-AND-IF-AVAILABLE BASIS. THIS ROADMAP DOES NOT CONSTITUTE A COMMITMENT, PROMISE, OR LEGAL OBLIGATION AND IS SUBJECT TO CHANGE AT THE SOLE DISCRETION OF NVIDIA. THE DEVELOPMENT, RELEASE, AND TIMING OF ANY FEATURES OR FUNCTIONALITIES DESCRIBED FOR OUR PRODUCTS REMAINS AT THE SOLE DISCRETION OF NVIDIA. NVIDIA WILL HAVE NO LIABILITY FOR FAILURE TO DELIVER OR DELAY IN THE DELIVERY OF ANY OF THE PRODUCTS, FEATURES, OR FUNCTIONS SET FORTH IN THIS DOCUMENT.

# HPC Enabling Scientific and Industrial Achievements

PARTICLE PHYSICS  
New Fundamental Particles




WEATHER & CLIMATE  
FourCastNet




GENOMICS  
Record DNA Sequencing




COSMOLOGY  
ID Dark Matter




RADIOLOGY  
GB: Laser Electron Accelerators




HIGH ENERGY PHYSICS  
Collision Reconstruction



ENERGY  
Plasma Modeling with FNO




VIRUS PATHOLOGY  
GB Finalist: CoVID variant classifier



ASTROPHYSICS  
DL Classifies Distant Galaxies



LIFE SCIENCES  
GB: Protein Similarity Search



# 2700+ GPU-Accelerated Applications

Plus many accelerated frameworks!

## ARTIFICIAL INTELLIGENCE

- PyTorch
- TensorFlow
- JAX

...

## CLIMATE & WEATHER

- Cosmos
- Gales
- WRF

...

## COMPUTATIONAL FINANCE

- O-Quant Options Pricing
- MUREX
- MISYS

...

## DATA SCIENCE & ANALYTICS

- Anaconda
- H2O
- OmniSci

...

## FEDERAL DEFENSE & OTHER

- ArcGIS Pro
- EVNI
- SocetGXP
- Cyllance
- FaceControl

...

## LIFE SCIENCES

- Amber
- LAMMPS
- GROMACS
- NAMD
- Relion
- VASP

...

## MANUFACTURING, CAD, & CAE

- Ansys Fluent
- Abaqus
- SIMULIA
- AutoCAD
- CST Studio Suite

...

## MEDIA & ENTERTAINMENT

- DaVinci Resolve
- Premiere Pro CC
- Redshift Renderer

...

## MEDICAL IMAGING

- aidoc
- PowerGrid
- RadiAnt

...

## OIL & GAS

- Echelon
- RTM
- SPECFEM3D

...

## RETAIL

- Everseen
- Deep North
- Third Eye Labs
- AWM
- Malong
- Clarifai
- Antuit

...

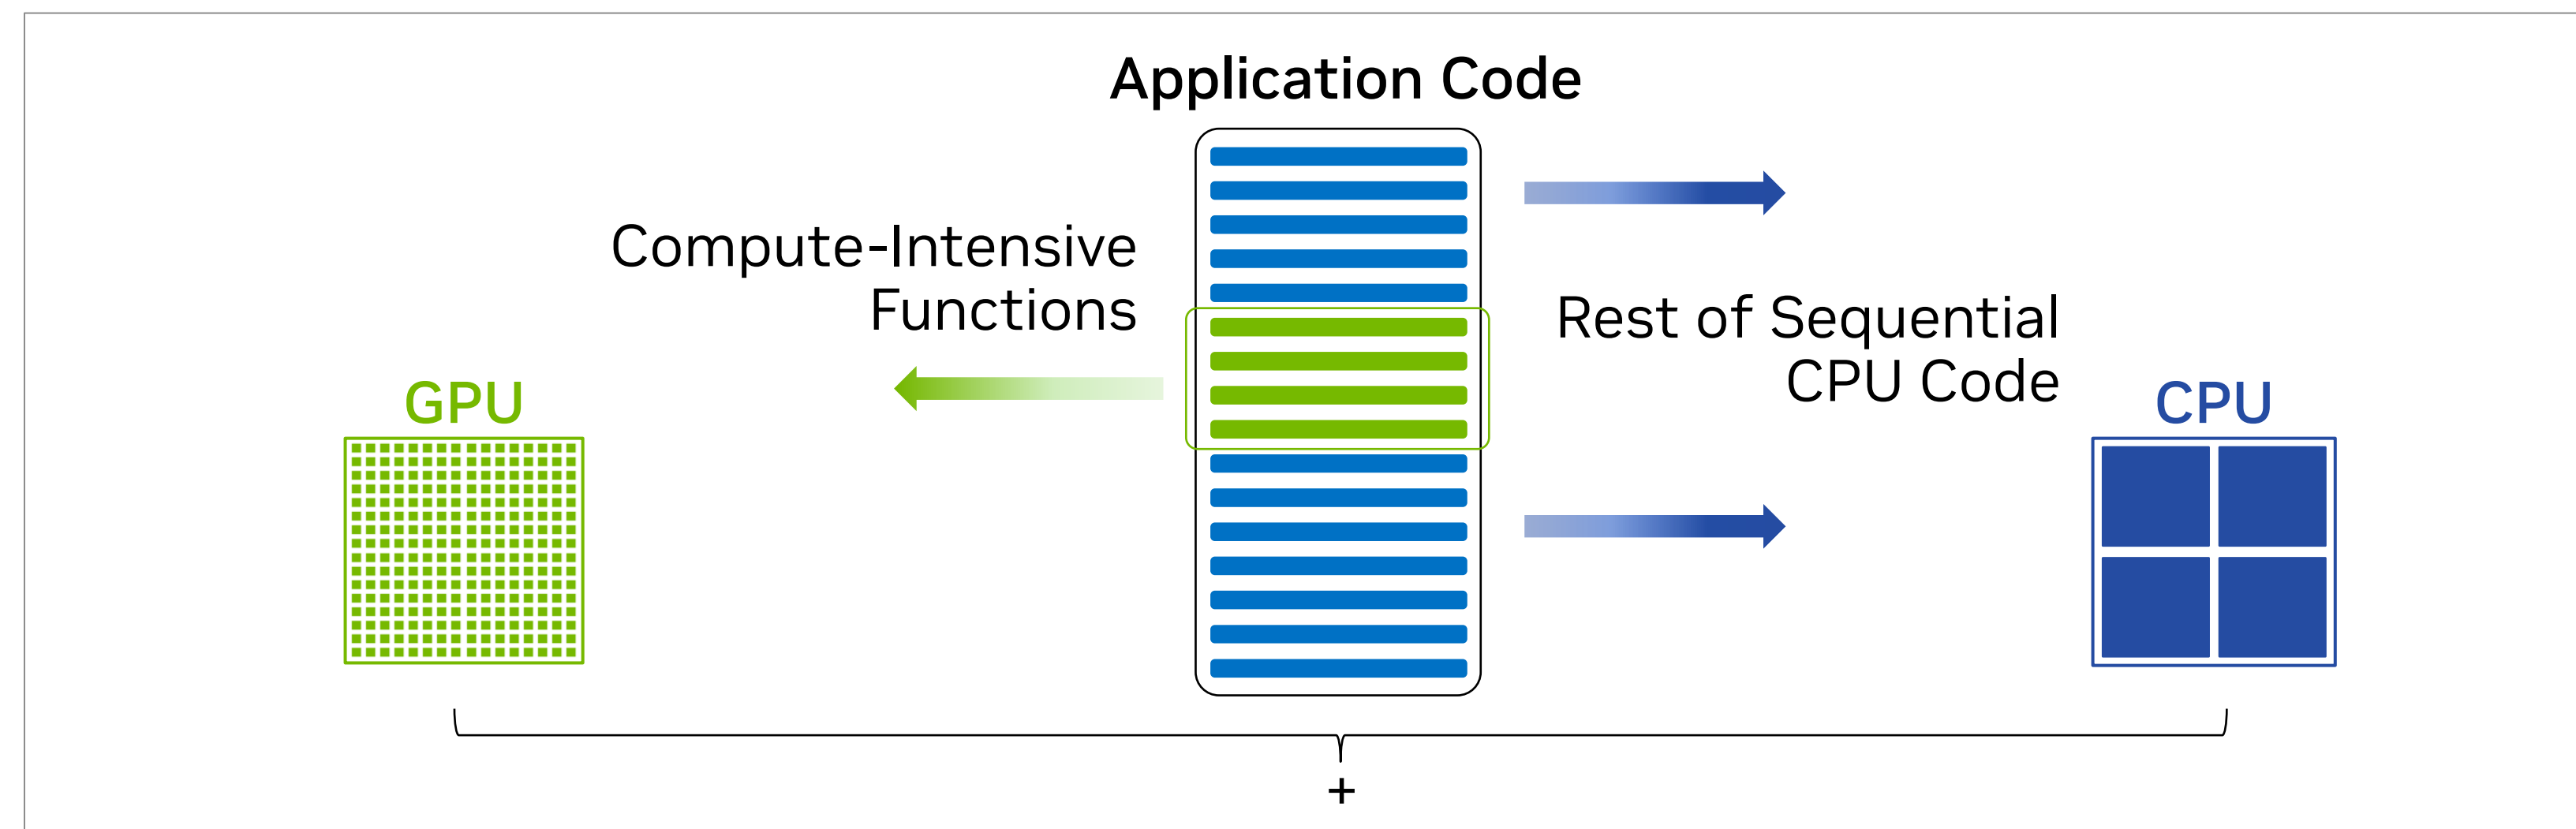
## SUPERCOMPUTING & HER

- Chroma
- GTC
- MILC
- QUDA
- XGC


...

# Why Accelerated Computing is Energy Efficient

Hardware, Software, & Networking to Optimize Performance & Efficiency




Moving Compute Intense Sections to the GPU



Consume Less Overall Energy

HPC applications gain across time



Center Panel: Geometric mean of application speedups vs. P100 | benchmark applications | Amber [PME-Cellulose\_NVE], Chroma [HMC], GROMACS [ADH Dodec], MILC [Apex Medium], NAMD [stmv\_nve\_cuda], PyTorch (BERT Large Fine Tuner), Quantum Espresso [AUSURF112-JR]; TensorFlow [ResNet-50], VASP 6 [Si Huge], [GPU node: with dual-socket CPUs with 4x P100, V100, or A100 GPUs. H100 values shown for 2022 projected performance subject to change

# NVIDIA H100

Unprecedented Performance, Scalability, and Security for Every Data Center

## Highest AI and HPC Performance

4PF FP8 (6X)| 2PF FP16 (3X)| 1PF TF32 (3X)| 67TF FP64 (3.4X)  
3.35TB/s (1.5X), 80GB HBM3 memory

## Transformer Model Optimizations


6X faster on largest transformer models

## Highest Utilization Efficiency and Security

7 Fully isolated & secured instances, guaranteed QoS  
2<sup>nd</sup> Gen MIG | Confidential Computing

## Fastest, Scalable Interconnect

900 GB/s GPU-2-GPU connectivity (1.5X)  
up to 256 GPUs with NVLink Switch | 128GB/s PCI Gen5



# Programming the NVIDIA platform

CPU, GPU, and Network

## ACCELERATED STANDARD LANGUAGES

ISO C++, ISO Fortran

```
std::transform(par, x, x+n, y, y,  
              [=] (float x, float y) { return y +  
a*x; }  
);
```

```
do concurrent (i = 1:n)  
  y(i) = y(i) + a*x(i)  
enddo
```

```
import cunumeric as np  
...  
def saxpy(a, x, y):  
  y[:] += a*x
```

## INCREMENTAL PORTABLE OPTIMIZATION

OpenACC, OpenMP

```
#pragma acc data copy(x,y) {  
...  
std::transform(par, x, x+n, y, y,  
              [=] (float x, float y) {  
                return y + a*x;  
              });  
...  
}  
  
#pragma omp target data map(x,y) {  
...  
std::transform(par, x, x+n, y, y,  
              [=] (float x, float y) {  
                return y + a*x;  
              });  
...  
}
```

## PLATFORM SPECIALIZATION

CUDA

```
__global__  
void saxpy(int n, float a,  
           float *x, float *y) {  
  int i = blockIdx.x*blockDim.x +  
          threadIdx.x;  
  if (i < n) y[i] += a*x[i];  
}  
  
int main(void) {  
  ...  
  cudaMemcpy(d_x, x, ...);  
  cudaMemcpy(d_y, y, ...);  
  
  saxpy<<<(N+255)/256,256>>>(...);  
  
  cudaMemcpy(y, d_y, ...);  
}
```

## ACCELERATION LIBRARIES

Core

Math

Communication

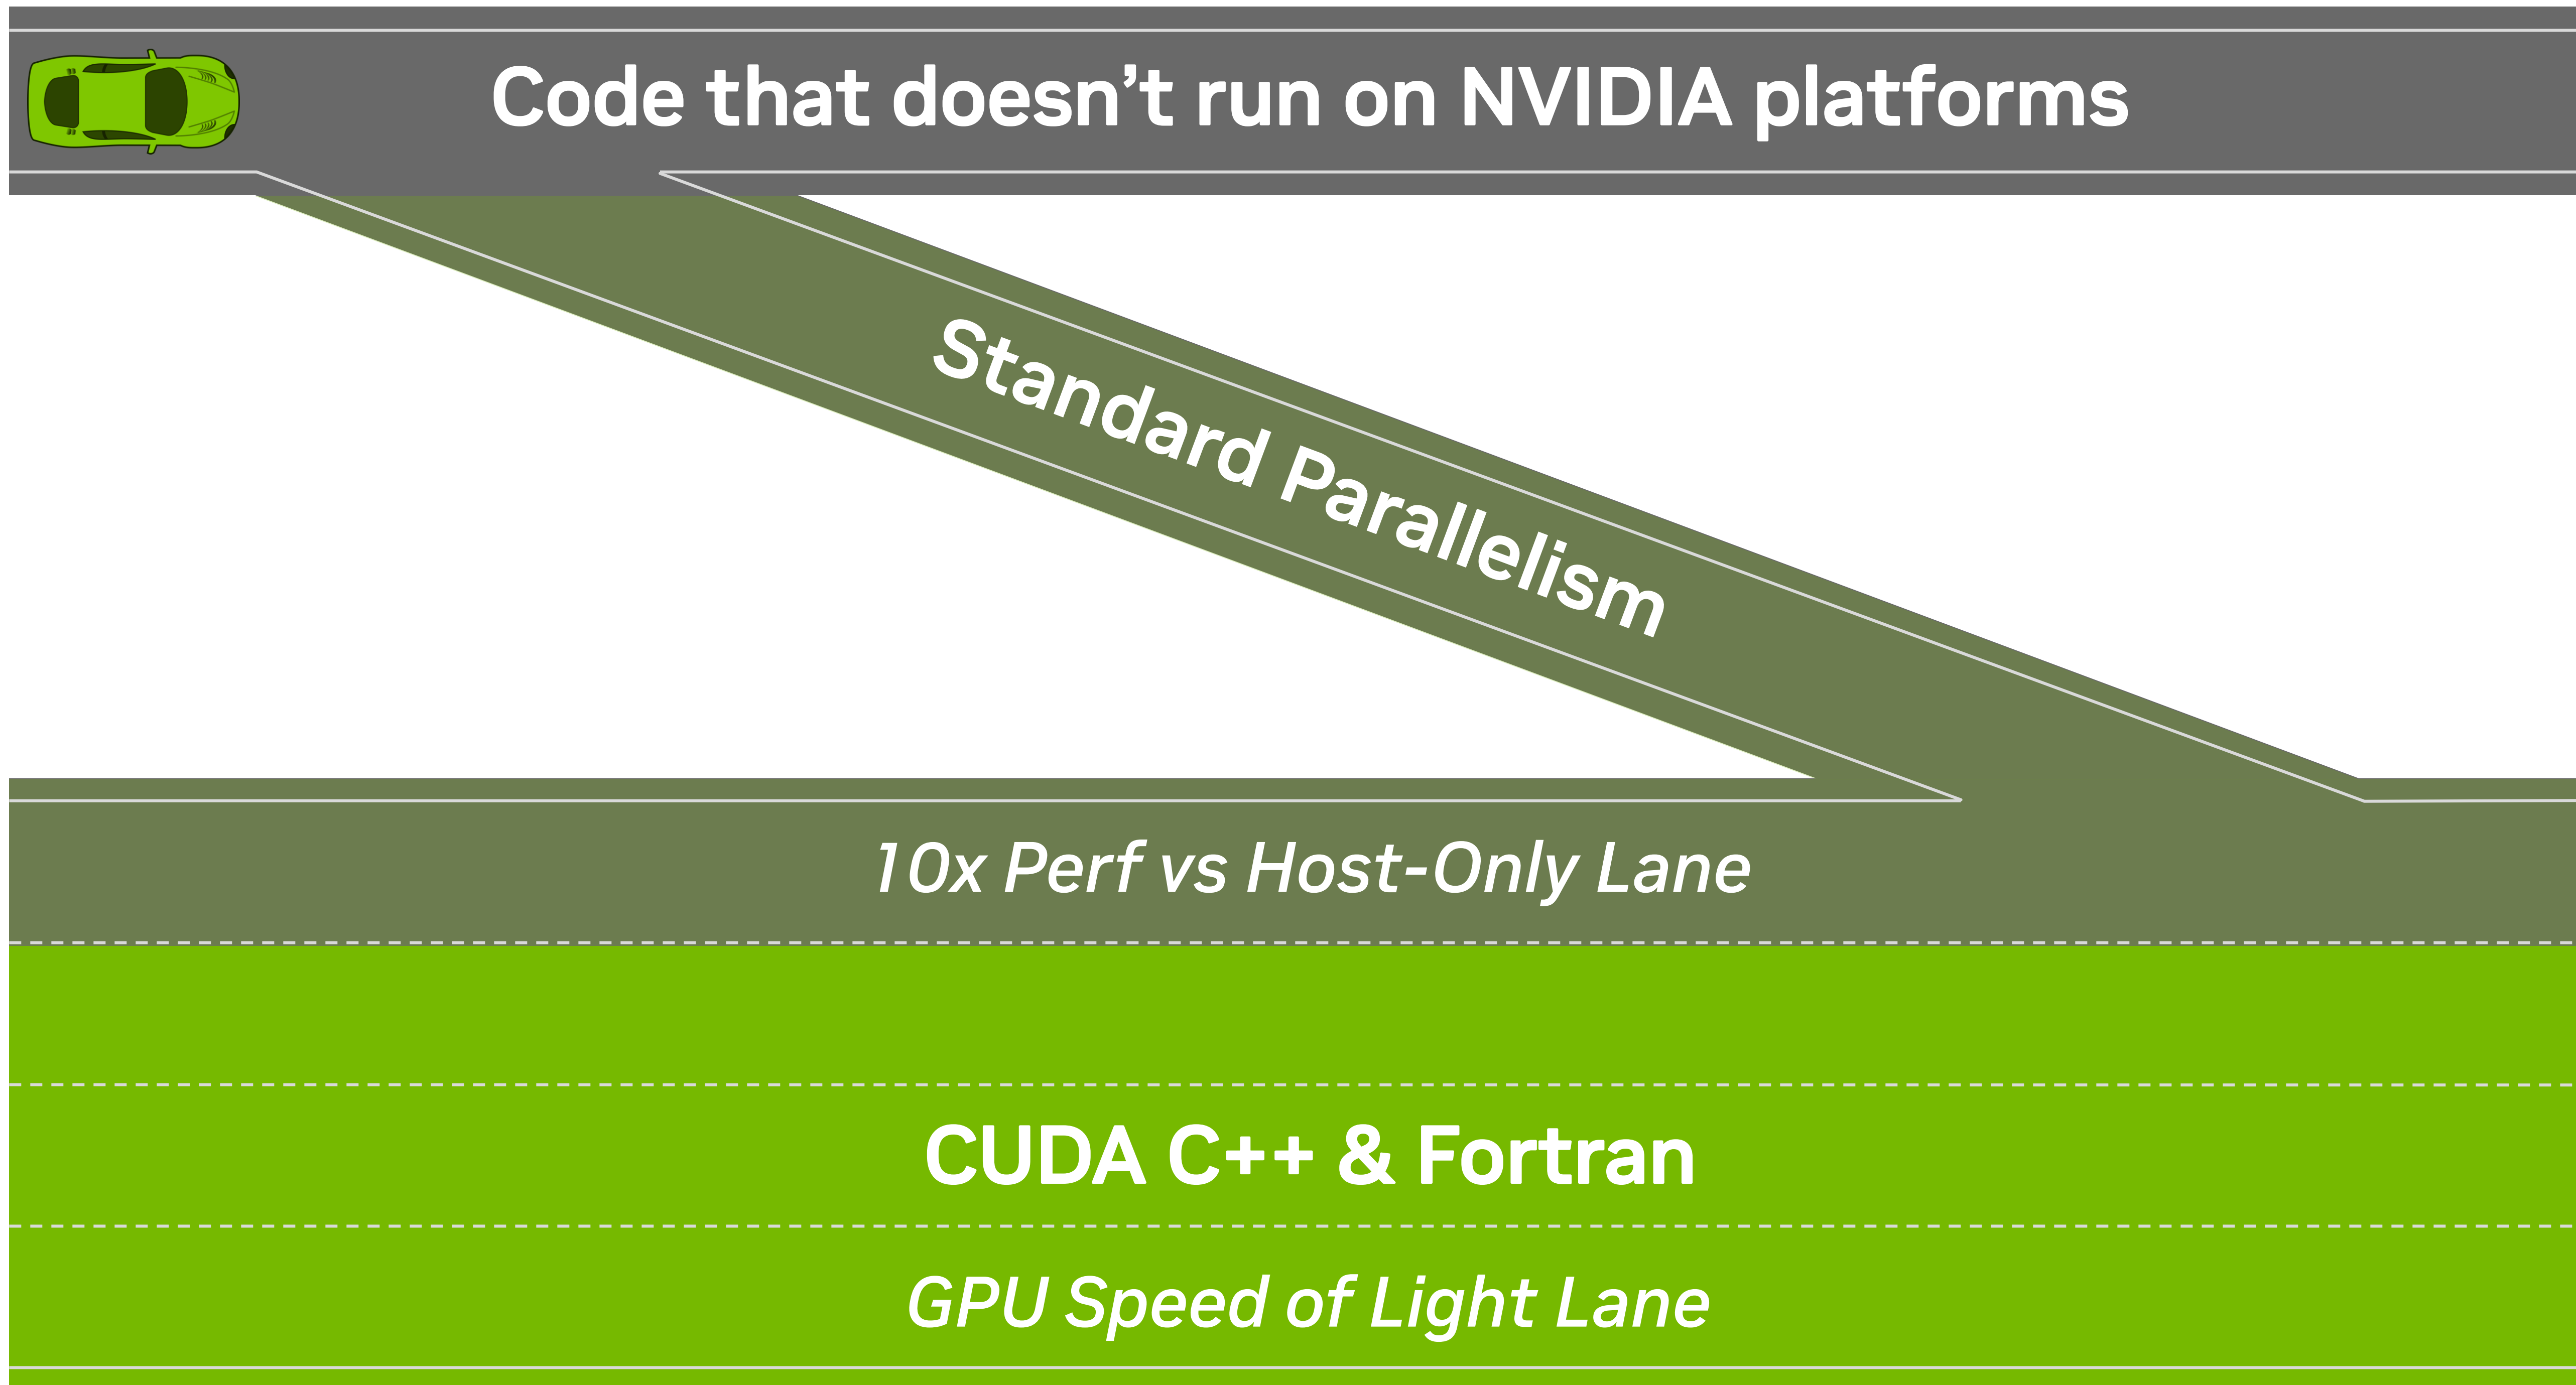
Data Analytics

AI

Quantum

# Scientists Need On-Ramps

Promote Parallelism, not Heterogeneity





# Accelerated Standard Languages

Parallel performance for wherever your code runs

## ISO C++

```
std::transform(par, x, x+n, y,  
              y, [=](float x, float y){  
                  return y + a*x;  
              })  
);
```

## ISO Fortran

```
do concurrent (i = 1:n)  
  y(i) = y(i) + a*x(i)  
enddo
```

## Python

```
import cunumeric as np  
...  
def saxpy(a, x, y):  
  y[:] += a*x
```

CPU

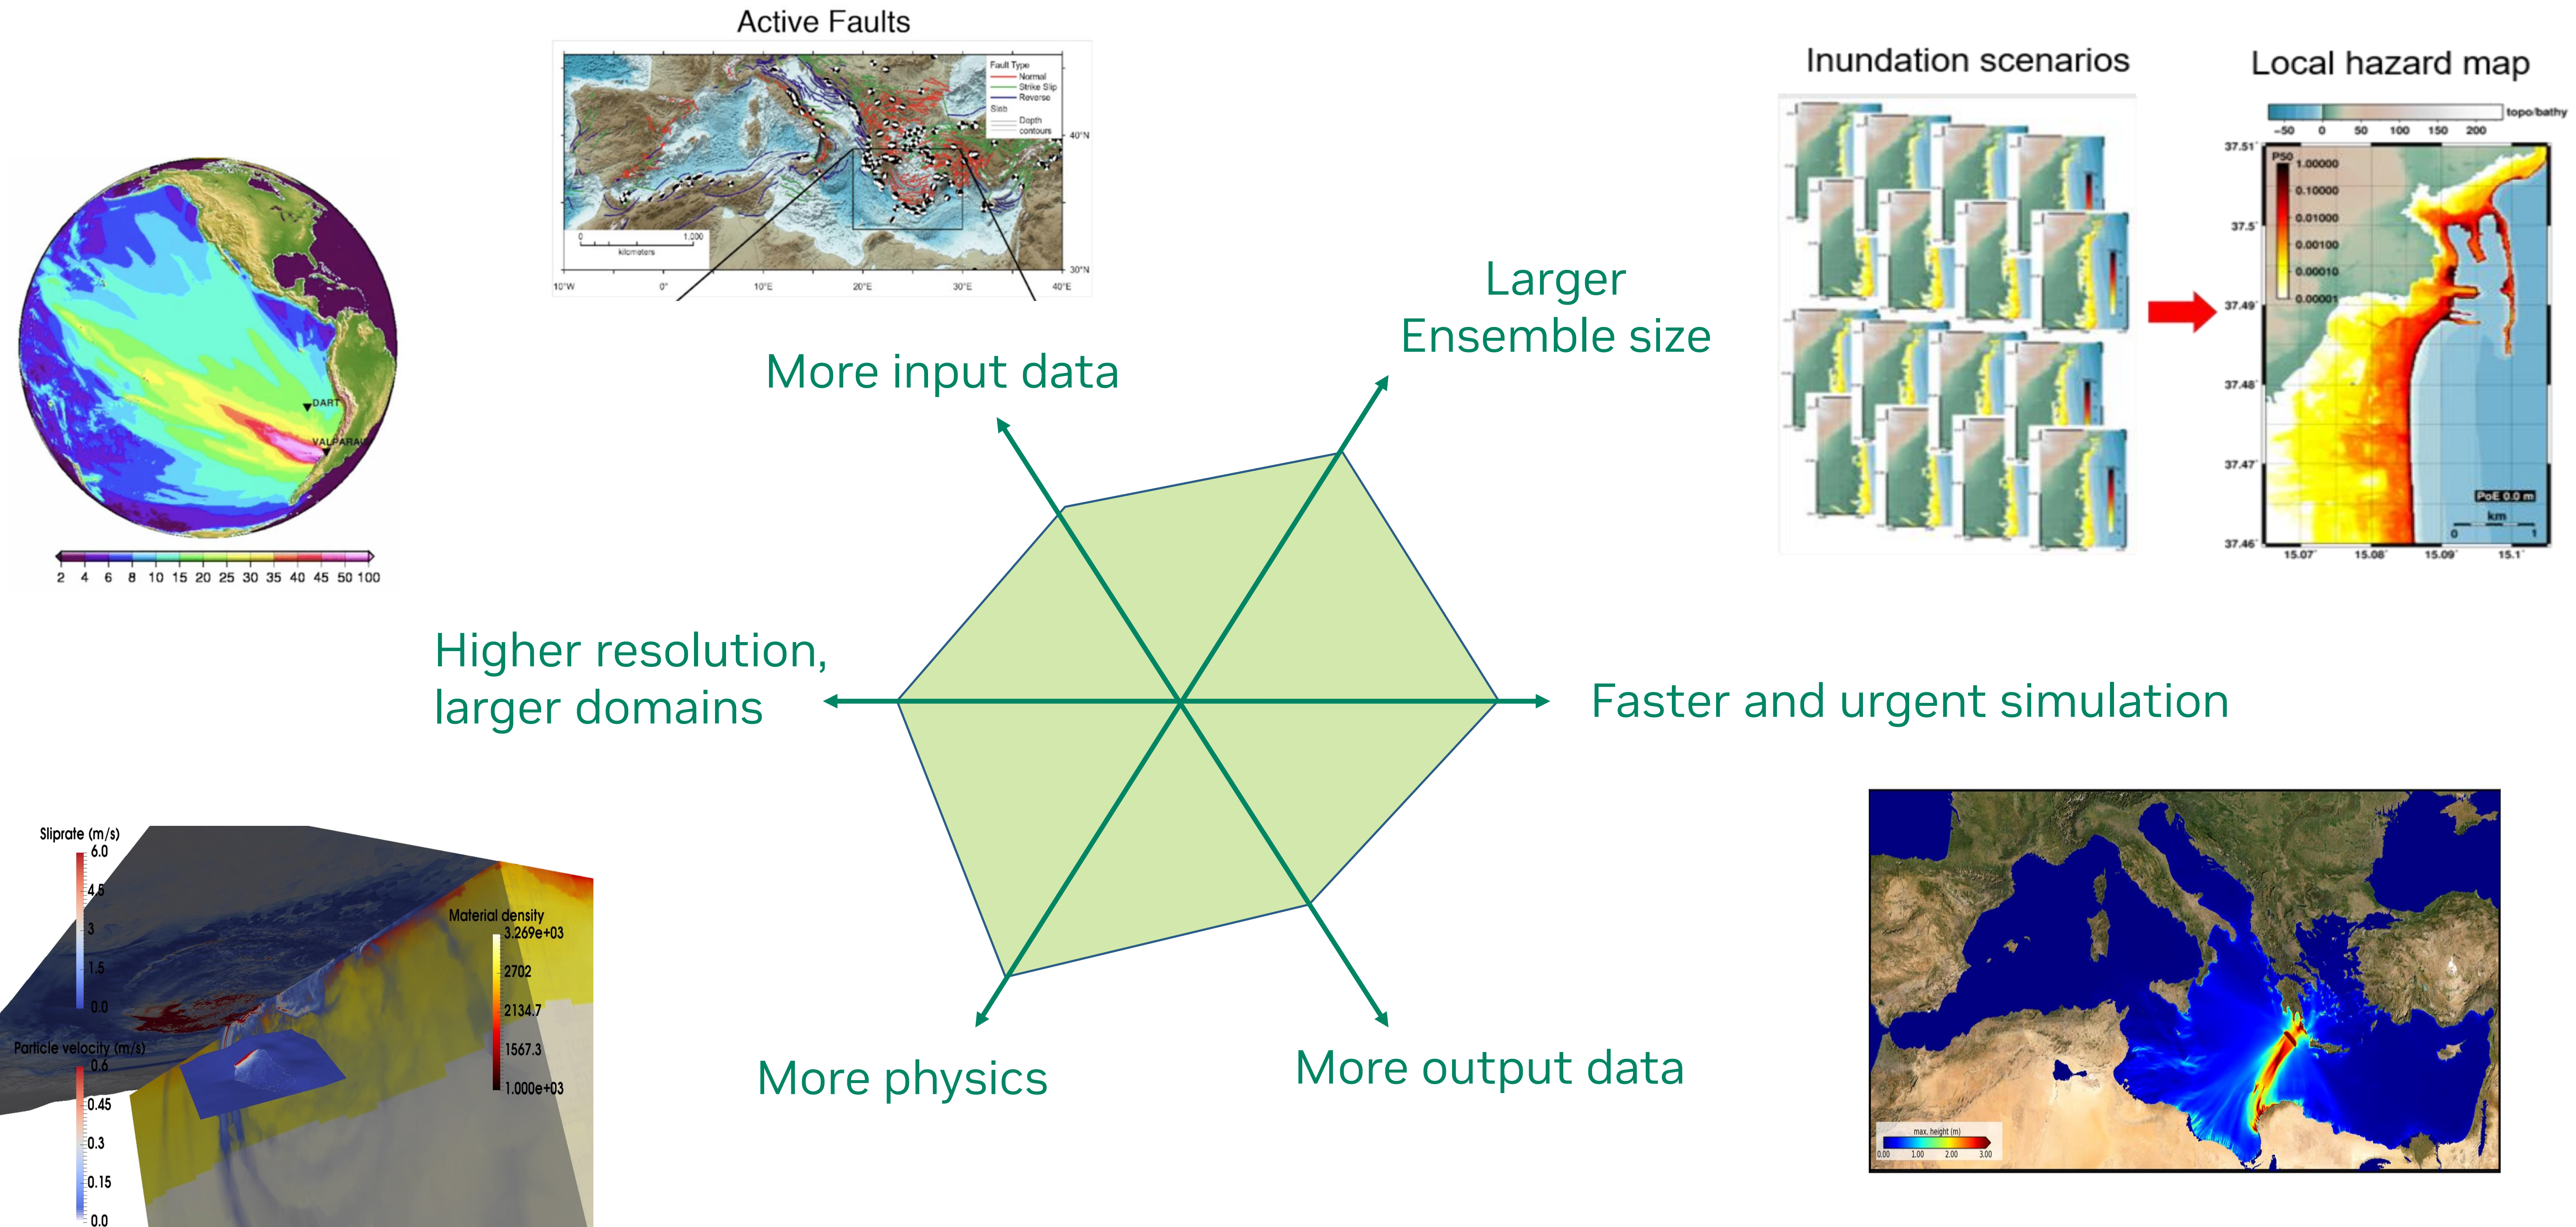
```
nvc++ -stdpar=multicore  
nvfortran -stdpar=multicore  
legate -cpus 16 saxpy.py
```

GPU

```
nvc++ -stdpar=gpu  
nvfortran -stdpar=gpu  
legate -gpus 1 saxpy.py
```

# GPU Acceleration Enables Breakthrough New Science

From **Petascale** ( $\sim 10^4$ - $10^5$  CPU) to **Exascale** ( $\times 1000!$ )



# GPU-accelerated ChEESA applications

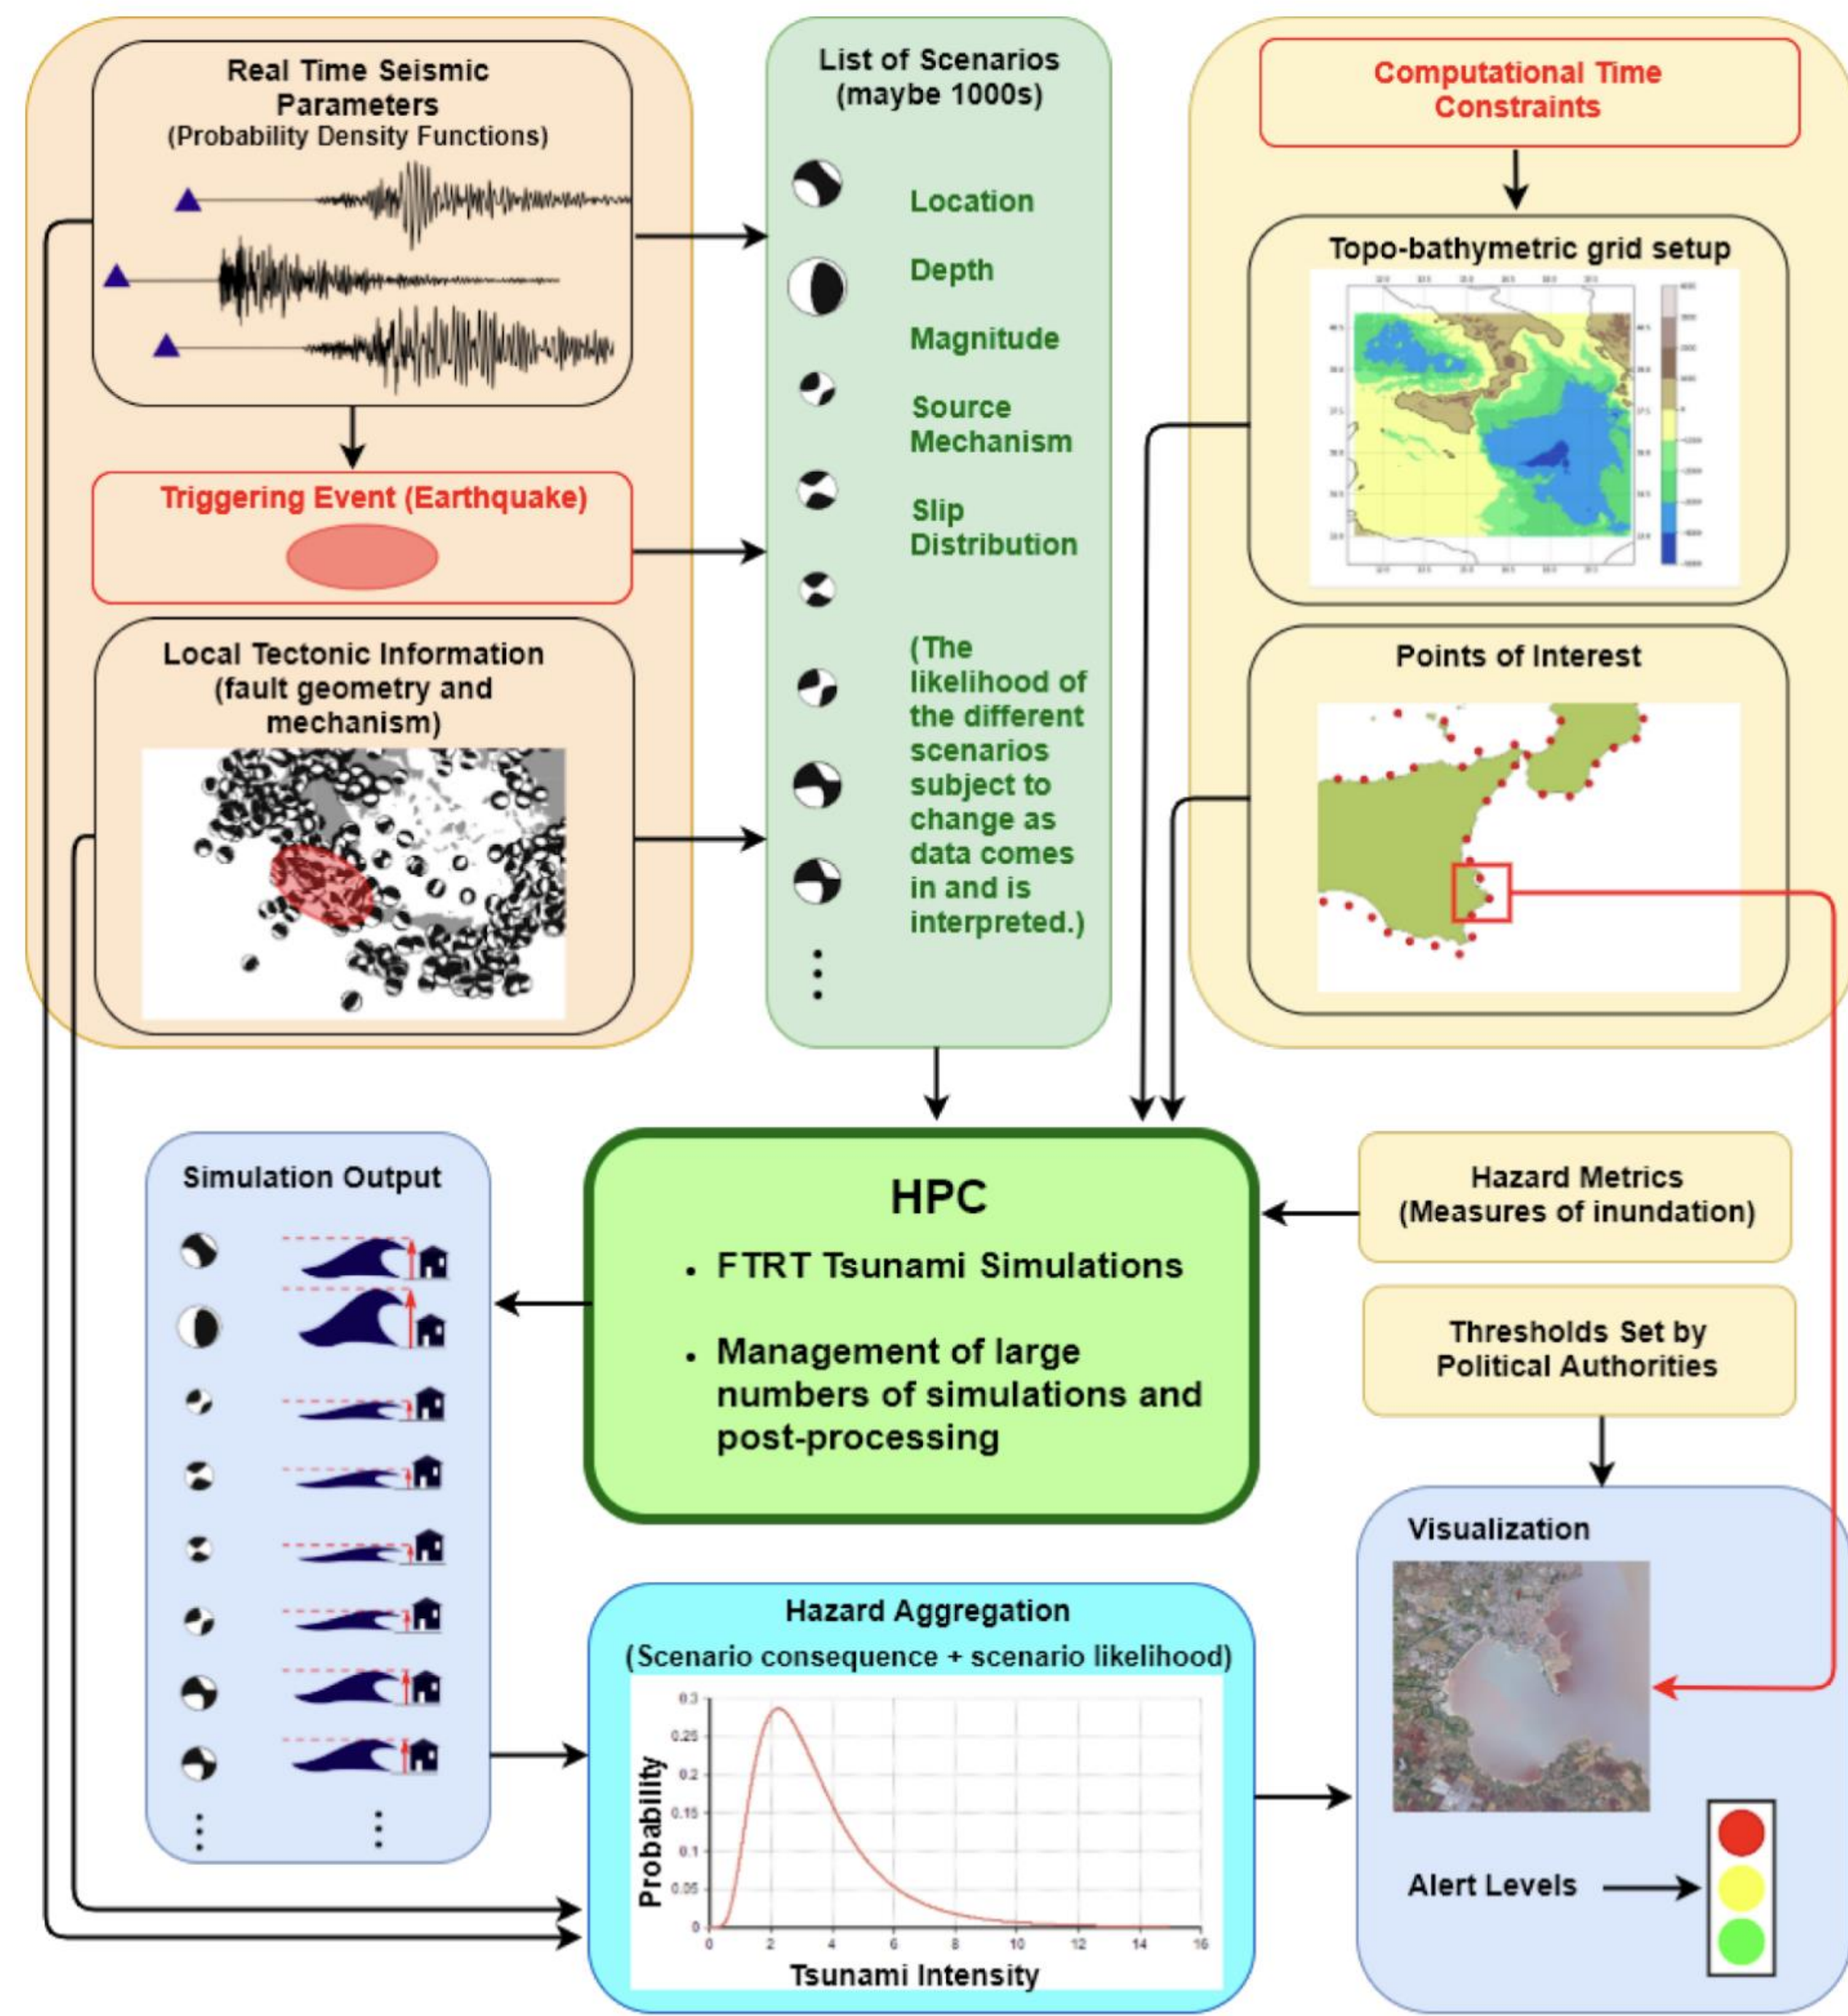
<https://cheese-coe.eu/>

Application	Programming Language	GPU	MPI+X
SeisSol	C/C++	Yes	CUDA
SPECFEM3D	Fortran	Yes	CUDA
ExaHyPE	C/C++	Partly	CUDA
xSHELLS	C/C++	Partly	CUDA
HySEA	C/C++	Yes	CUDA
FALL3D	Fortran	Yes	OpenACC



# GPU-enabled ChEESA workflow

The ChEESA PD8: Probabilistic Tsunami Forecasting (PTF) for early warning and rapid post event assessment



## Urgent computing workflow exercise on Marconi-100

Code: Tsunami-HySEA (TRL 6~7)

Scenario: On-the-fly simulations for POST-EVENT ASSESSMENT (tens of minutes) - NO Early Warning - the Samos Earthquake

- Simulation ensembles to be run from scratch on large enough HPC clusters in urgent computing mode.
- Provides exceedance probabilities for tsunami heights just off the coastline for almost equally spaced points of interest every 2 km in front of the coasts of the Mediterranean Sea

Outputs: ~38 000 scenario simulations

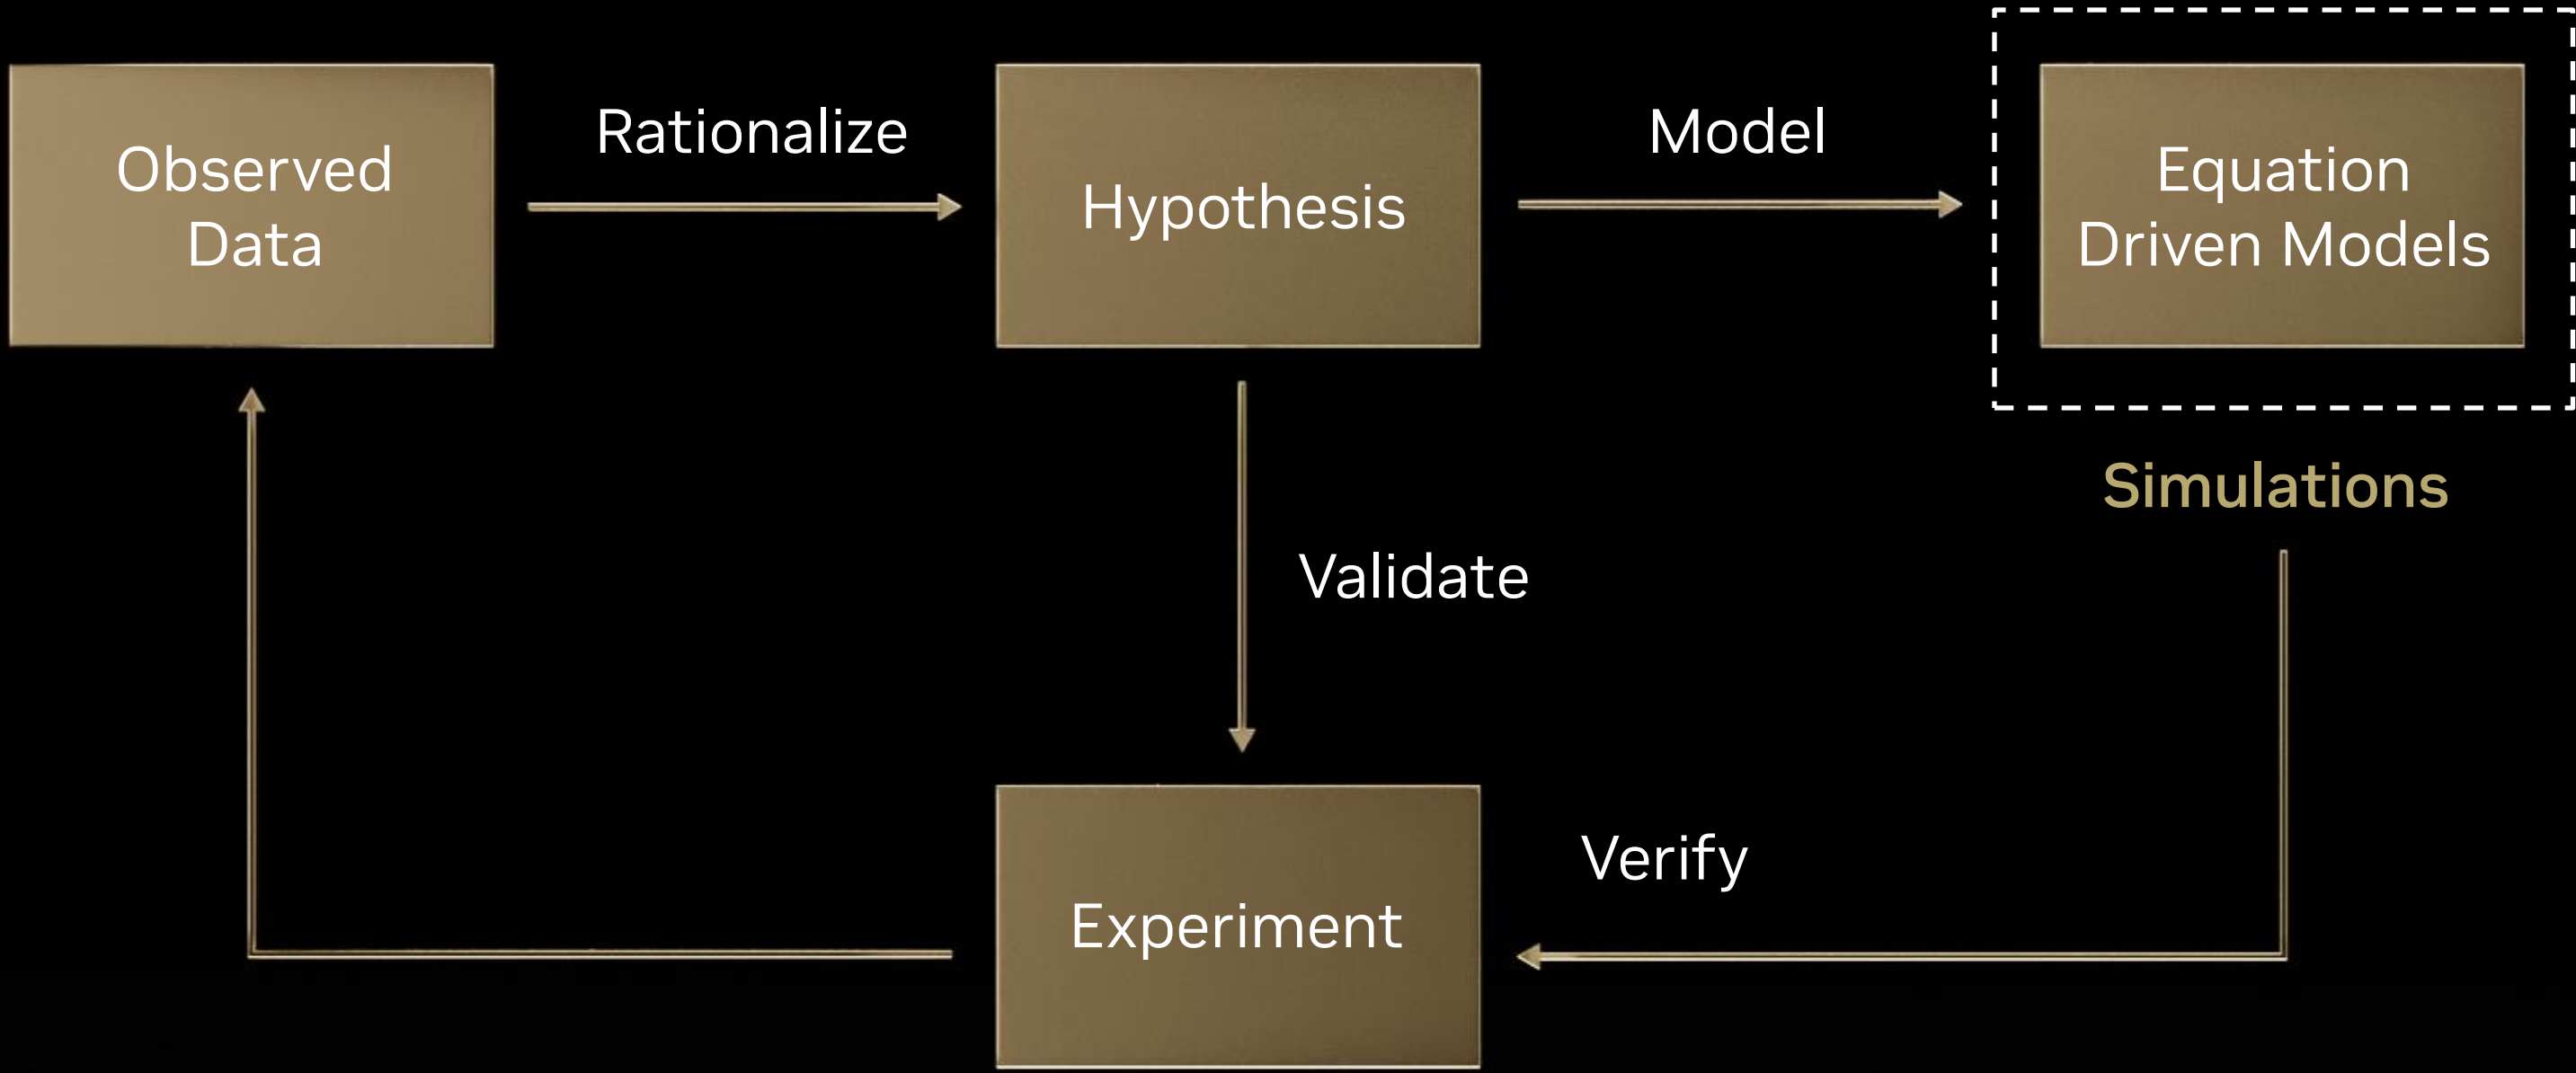
Resources: 805 nodes (3,220 V100 GPUs), ~25.7 PFlop/s



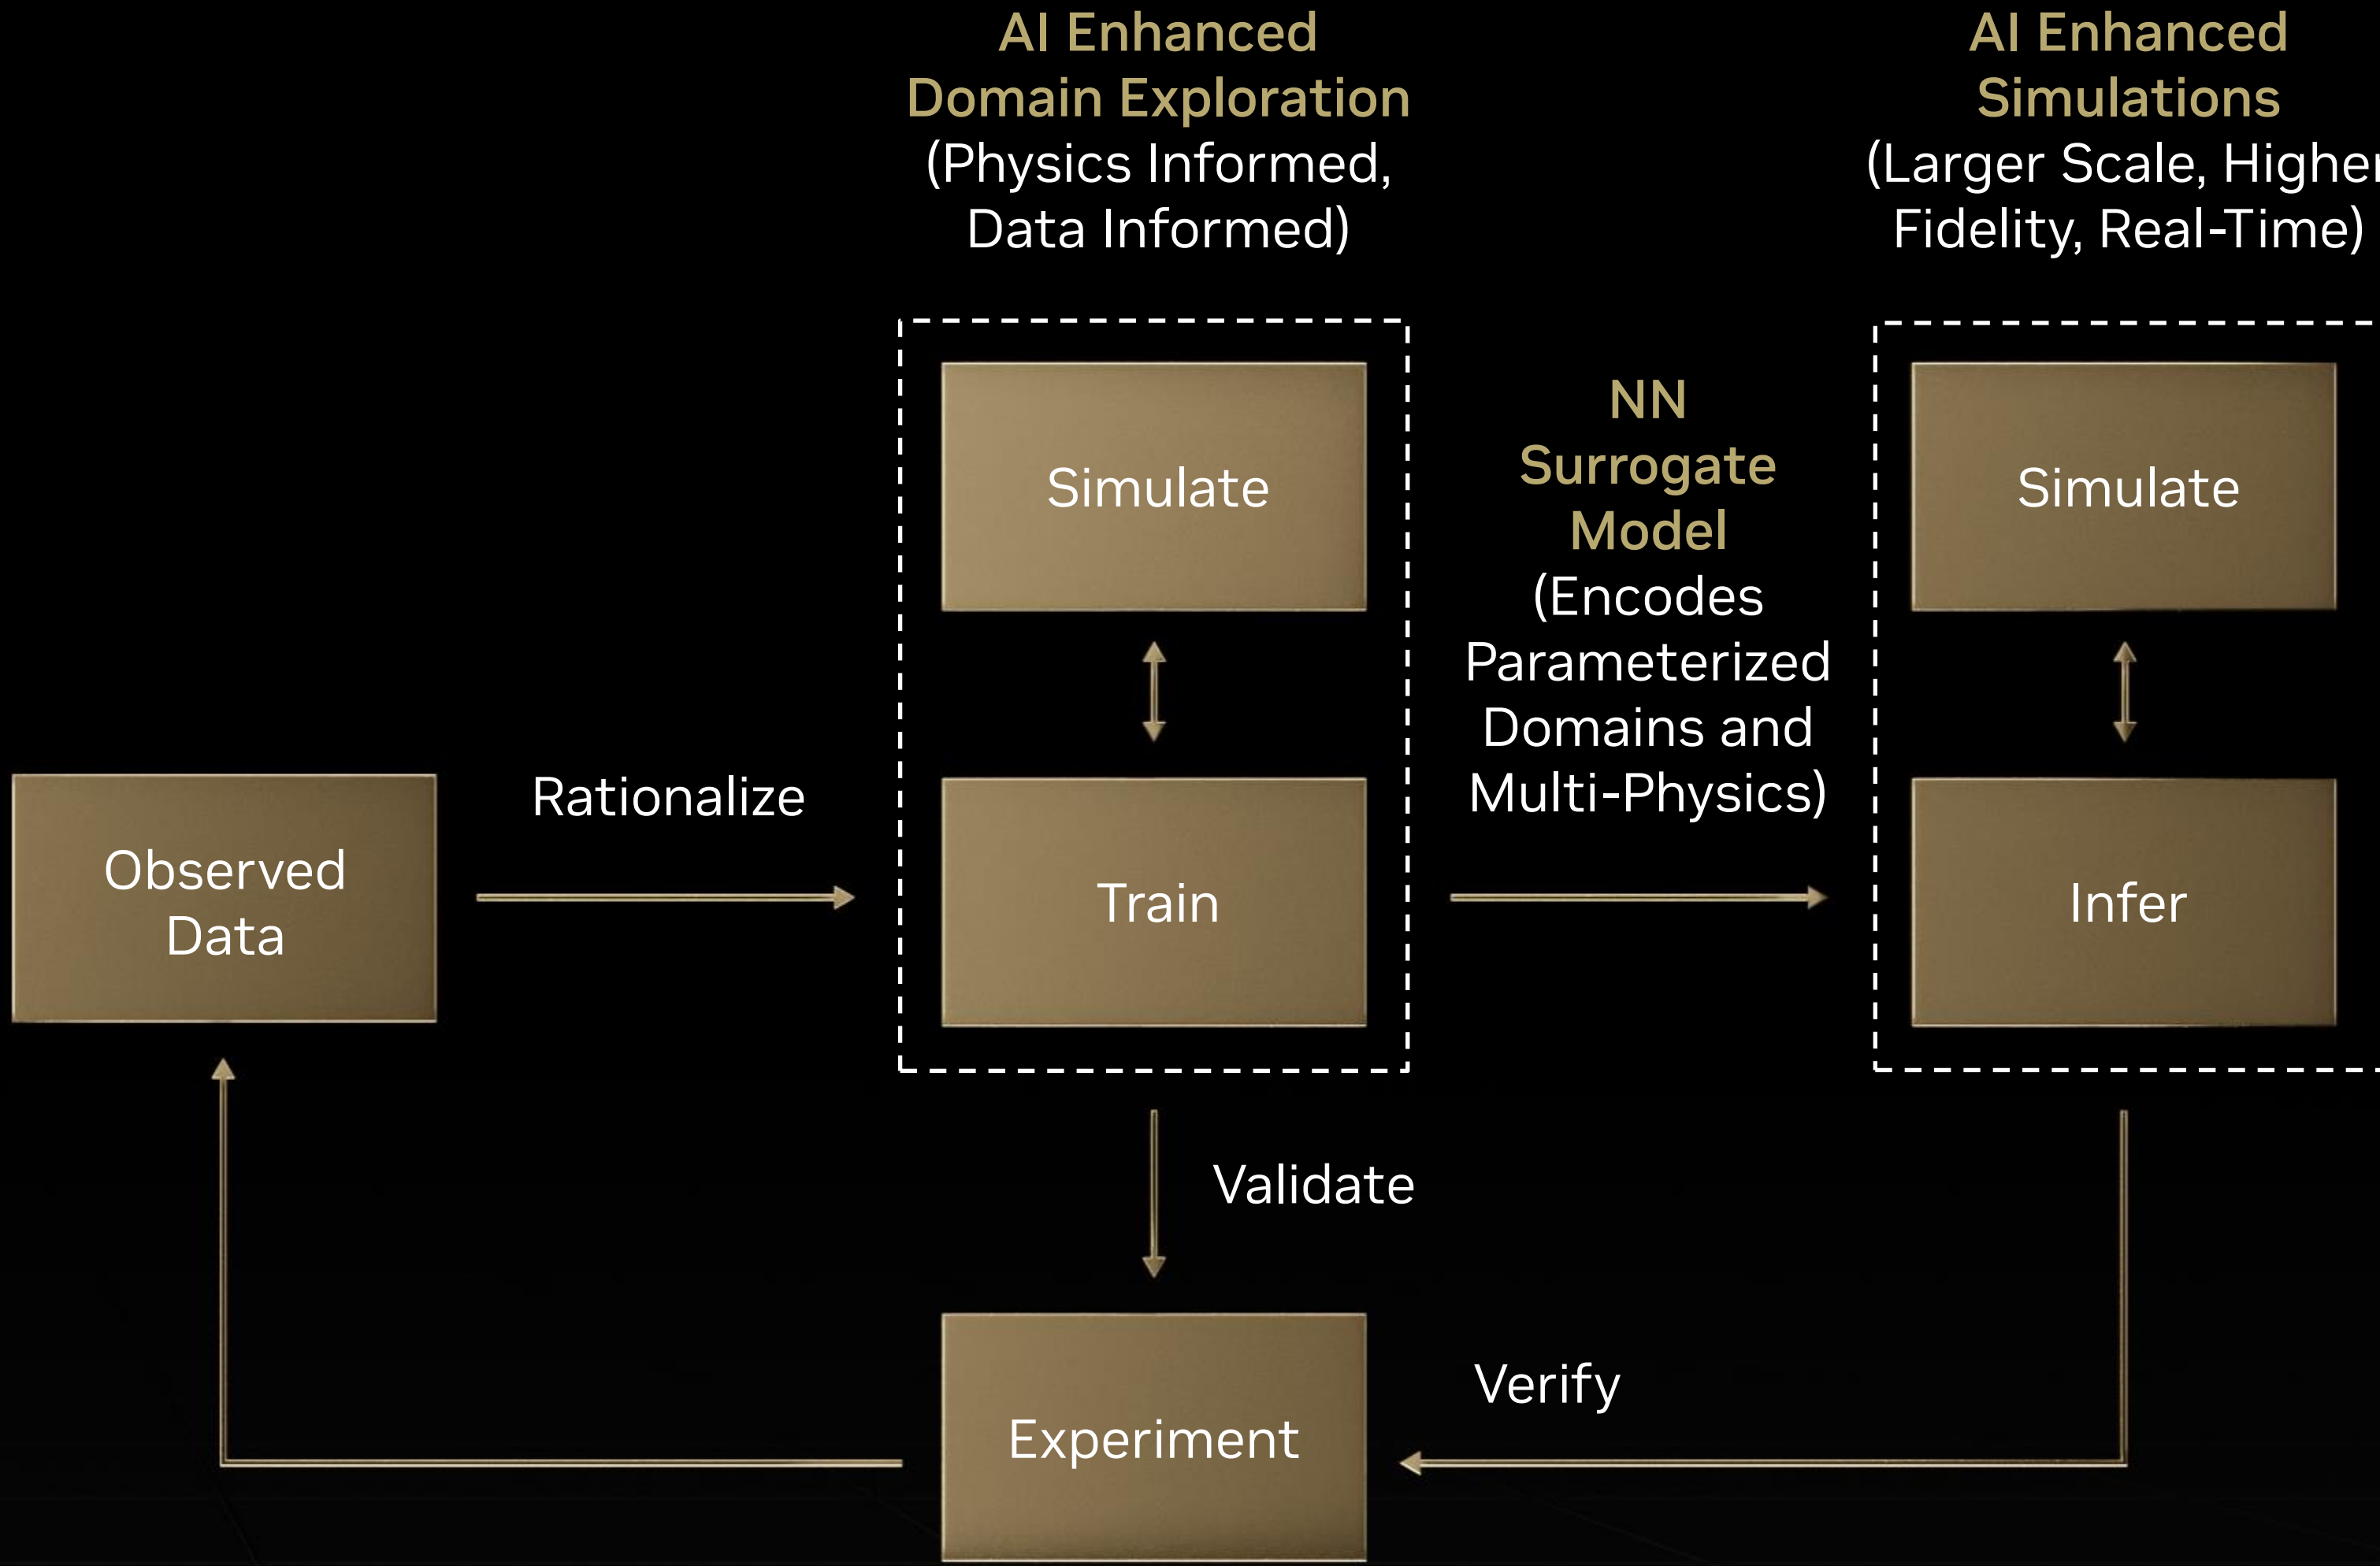
# Beyond HPC HPC+AI & Digital Twins

# AI IS THE 4<sup>TH</sup> PILLAR OF SCIENTIFIC DISCOVERY

## TRADITIONAL




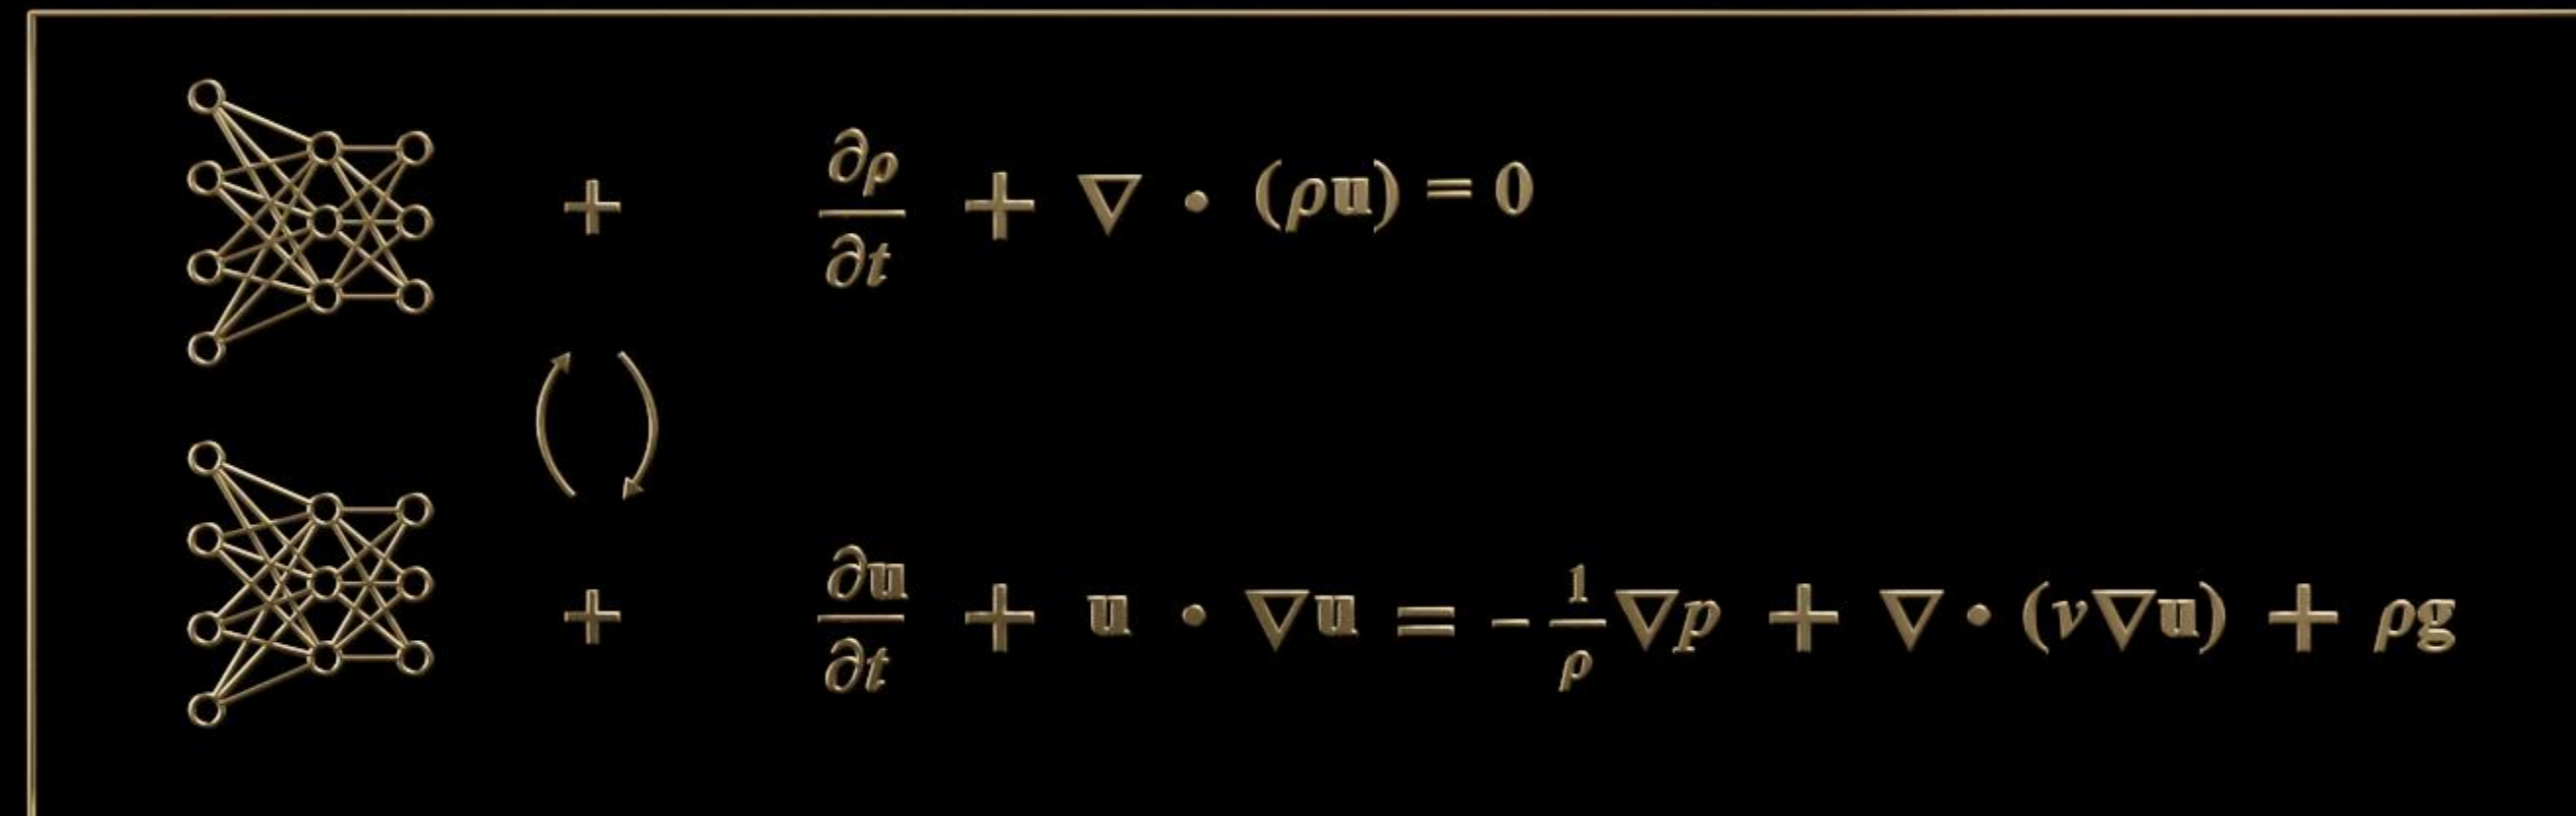

## DATA DRIVEN




# NVIDIA MODULUS

Shortcut to Surrogate Models for Interactive Simulation | Development Platform for AI Surrogate Models

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0$$
$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla p + \nabla \cdot (\nu \nabla \mathbf{u}) + \rho \mathbf{g}$$




WEATHER & CLIMATE MODELING  
FourCastNet



RENEWABLE ENERGY  
Siemens Gamesa Wind Farm



PRODUCT DEVELOPMENT  
Thermal Airflow Digital Twin




INDUSTRIAL DIGITAL TWIN  
Data Center CFD Acceleration



Now Available on NVIDIA LaunchPad and Major CSPs

# NVIDIA OMNIVERSE FOR SCIENTIFIC COMPUTING

Connecting Complex HPC 3D and Simulation Workflows



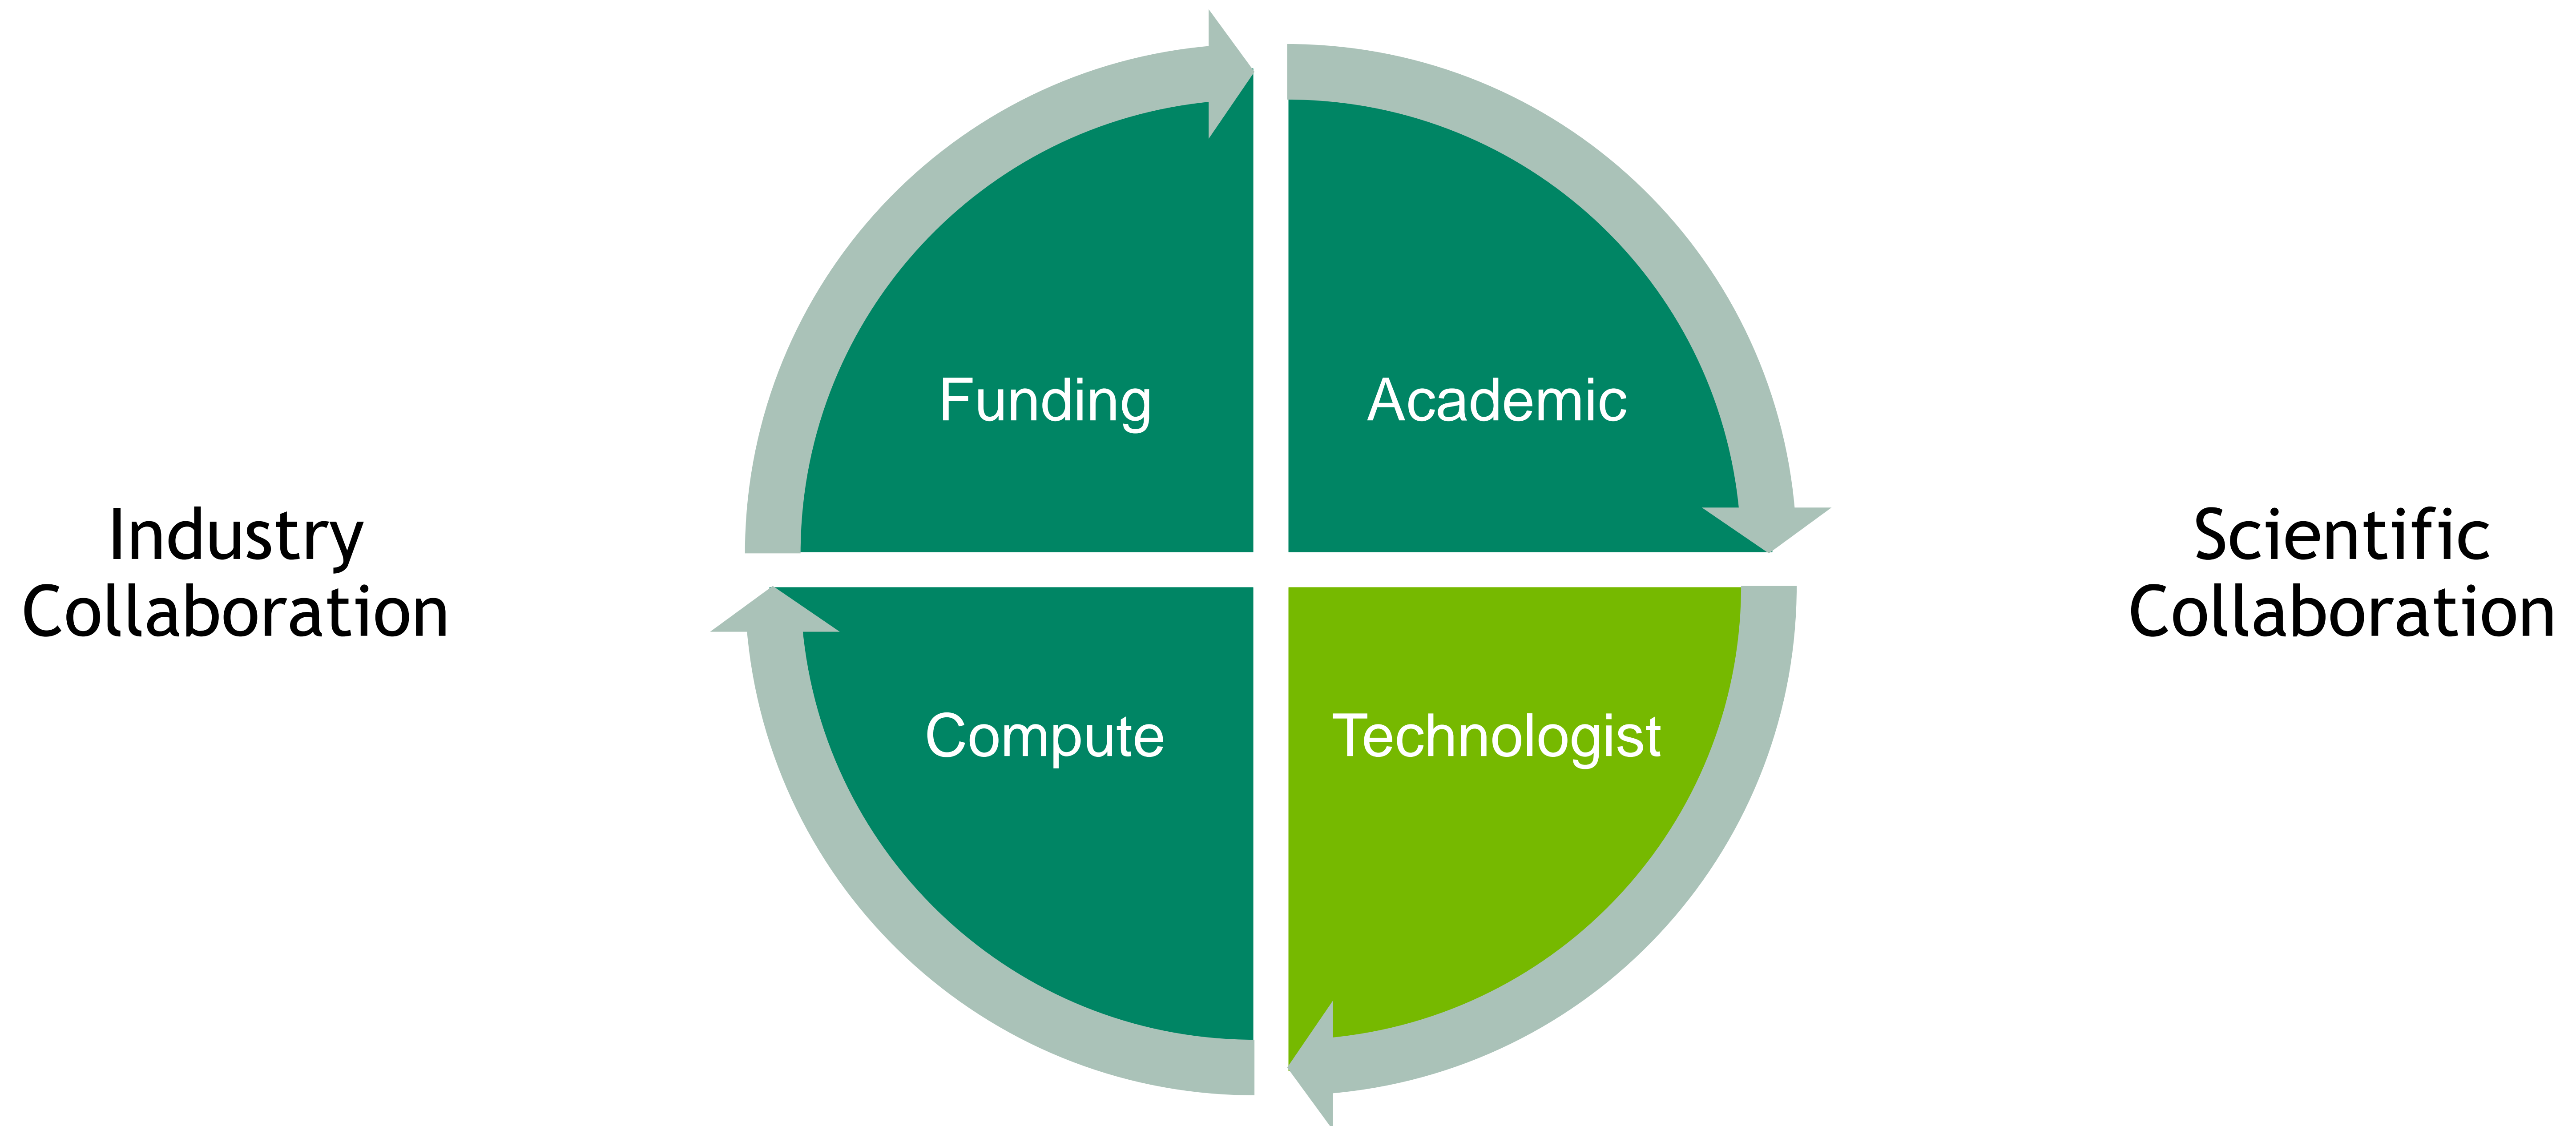




# ABOUT NVAITC

# RESEARCH COLLABORATIONS

Enabling AI research through efficient GPU Computing



## workstation



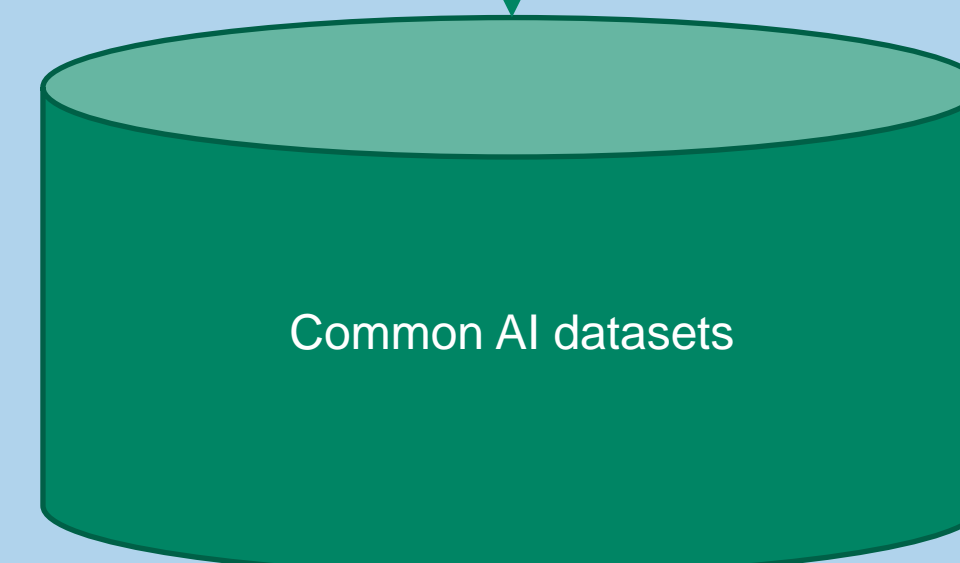

Each research has its own GPU for developing/debugging purposes. They all use the same software stack and push developed code to a private *github* organization.



Local or Remote development




## local cluster



~48 GPUs  
Researchers access to resources through SLURM. According to priorities, seniority, SLURM partitions access policy may vary. Senior researchers operate the local cluster. Containers are pulled from NVIDIA NGC.



## national facility



~5000 GPUs  
Support for singularity, enroot. DATASETS\_AI is a fileset containing the most common datasets (imagenet, epic-kitchens, meccano, synlidar, etc). Local conda channel (cineca-ai) to provide packages optimized for the local architecture. Containers are pulled from NVIDIA NGC.

