
CERN
openlab

AI workflow lifecycle on Digital Twins for multi-sciences

Matteo Bunino (matteo.bunino@cern.ch), Alexander Zoechbauer, Kalliopi Tsolaki, Rakesh Sarma, Ilaira Luise, Maria Girone, Sofia Vallecorsa

interTwin - Digital Twin Engine for science

Co-design and implement the prototype of an **interdisciplinary Digital Twin Engine**

Open-source platform based on **open standards**

Large spectrum of **diverse use cases** from **physics** and **earth observation** sciences

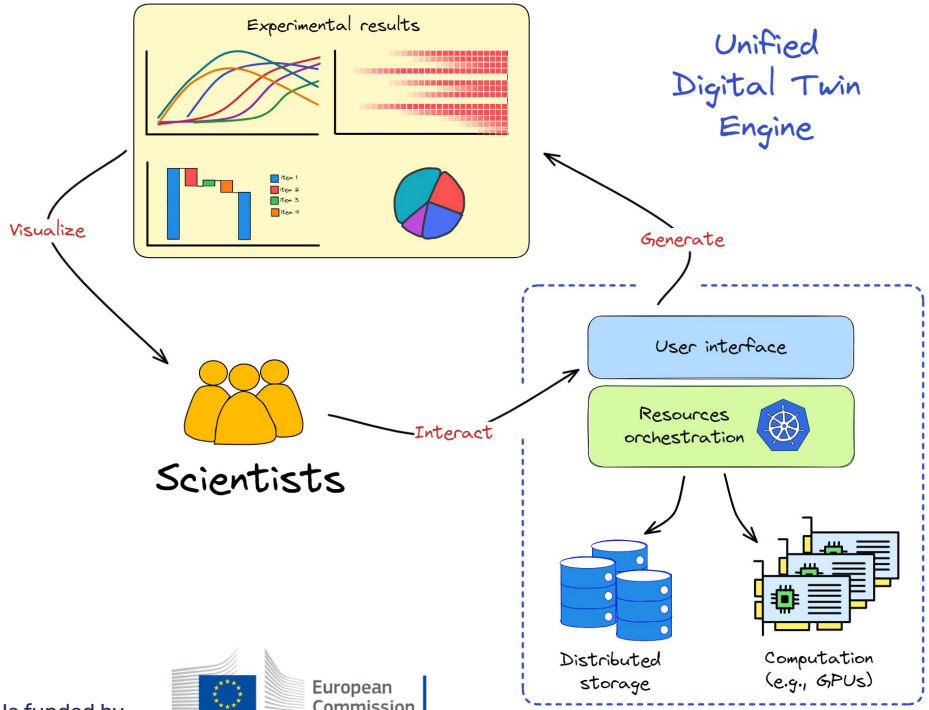

Website: https://www.intertwin.eu/

interTwin

Is funded by

European Commission

interTwin - Digital Twin Engine for science

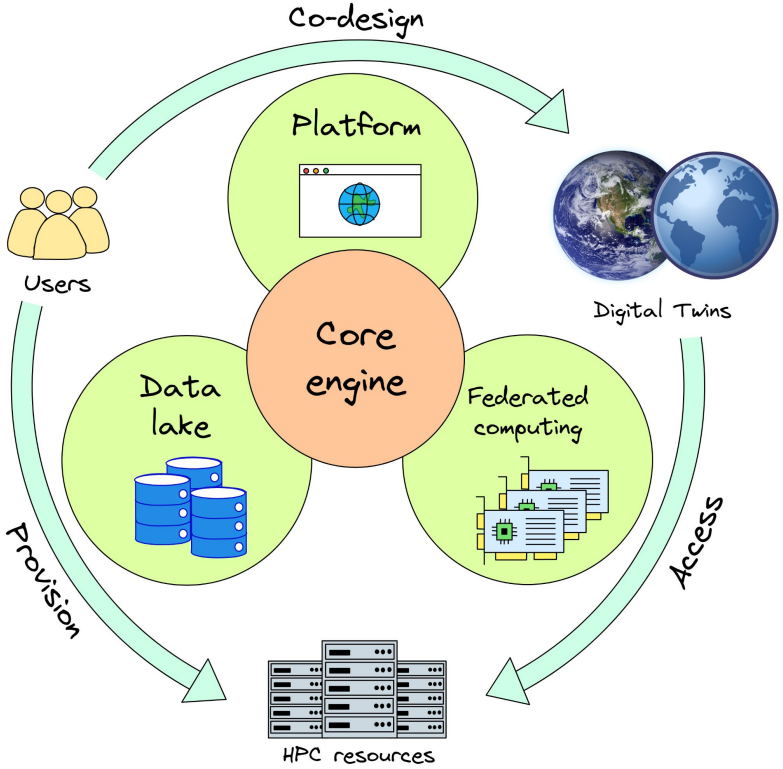

interTwin

interTwin - Digital Twin Engine for science

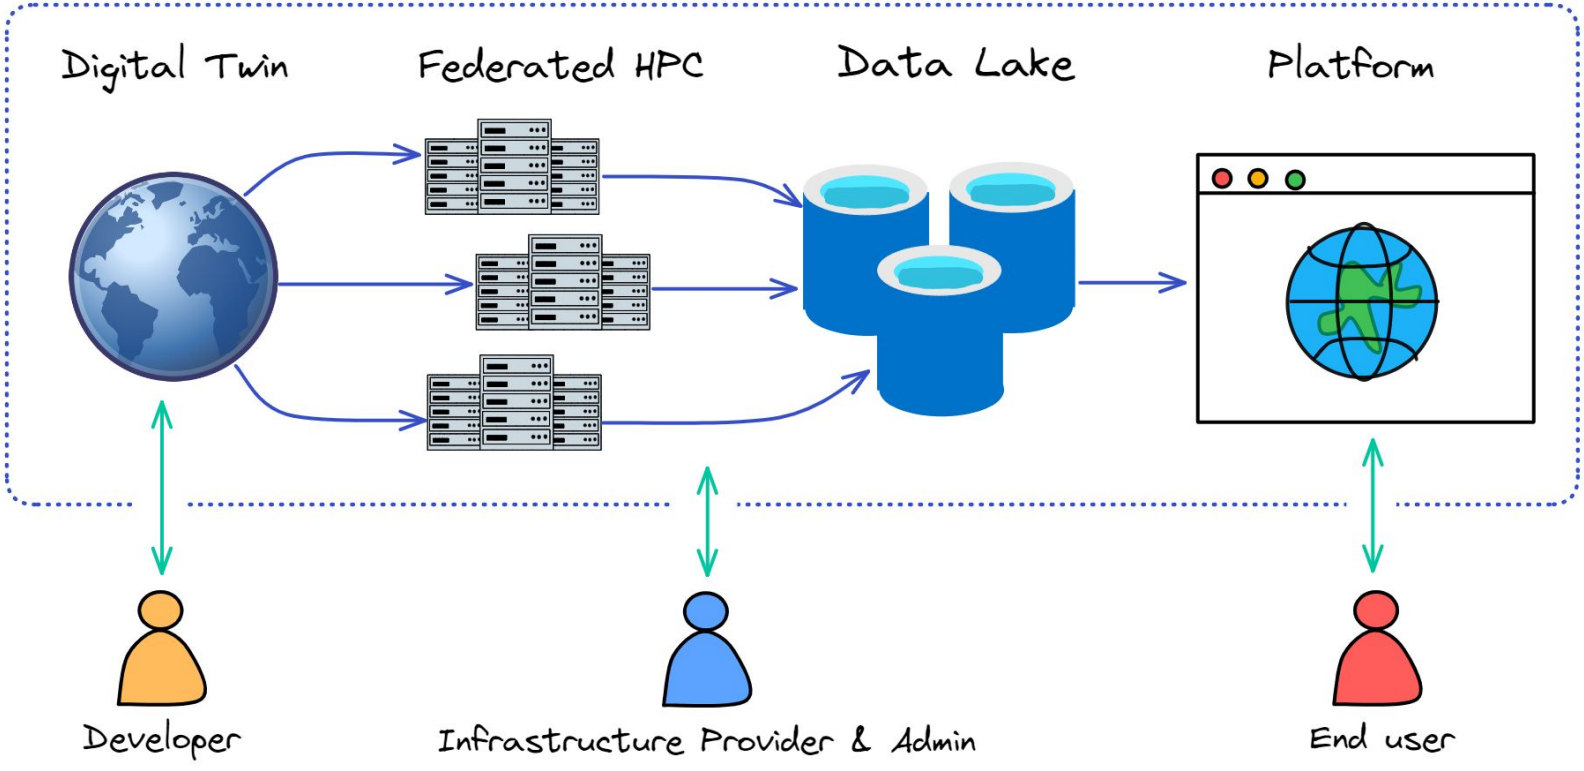

interTwin

Multi-sciences DTs

interTwin

interTwin use cases

interTwin

Earth Observation Use Cases in InterTwin

Cyclone Detection
CMCC, CNRS, Univ. of Torino

Fire Hazard Map Generation
CMCC, CNRS, Univ. of Torino

Early Flood Warnings
Deltares, EURAC, Technical Univ. of Vienna

Drought Prediction
CERFACS, EURAC, Deltares

interTwin

Climate Change Future Projections of Extreme Events

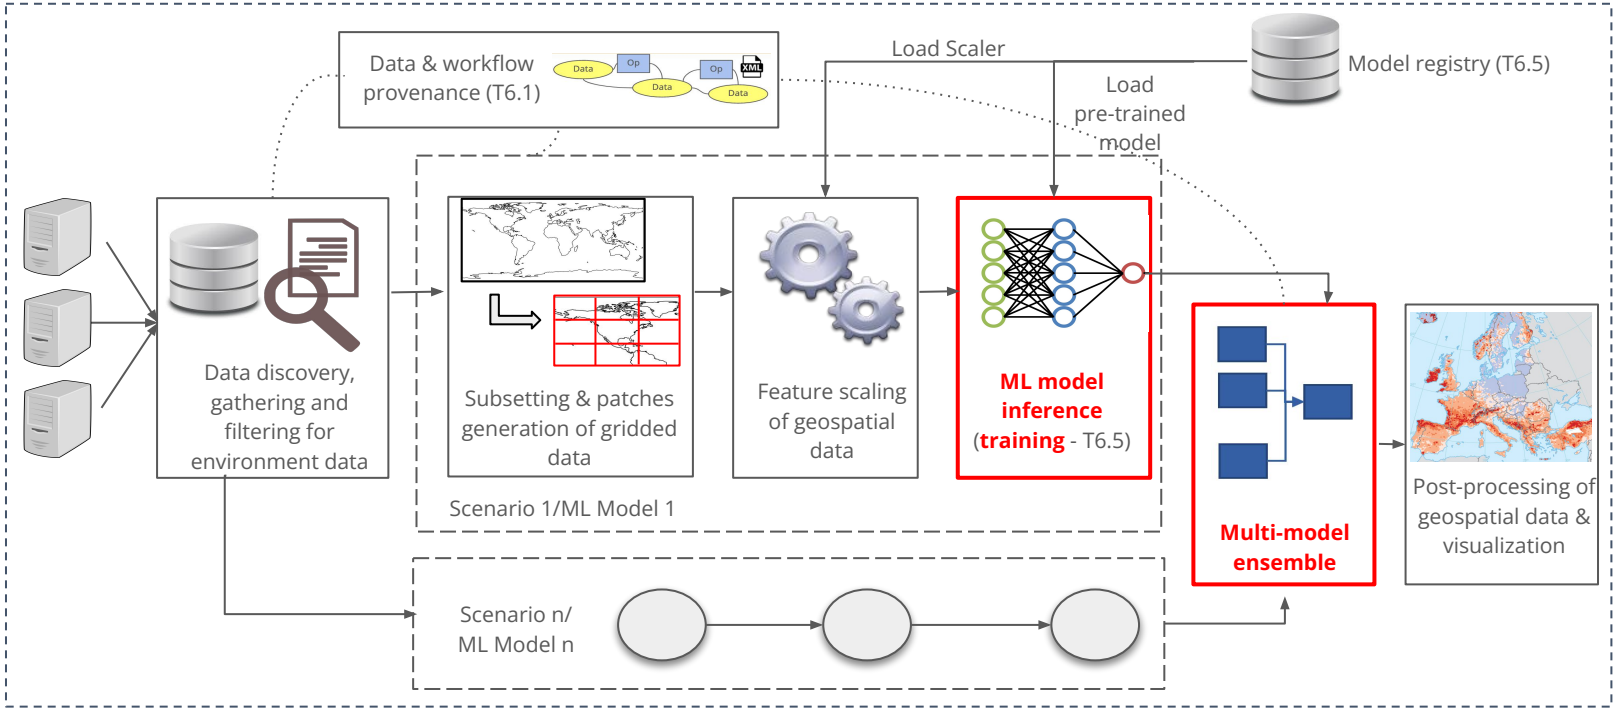

Image courtesy of Donatello Elia (CMCC)

Physics Use Cases in InterTwin

Radio Astronomy
Univ. of Heidelberg,
Max Planck Society

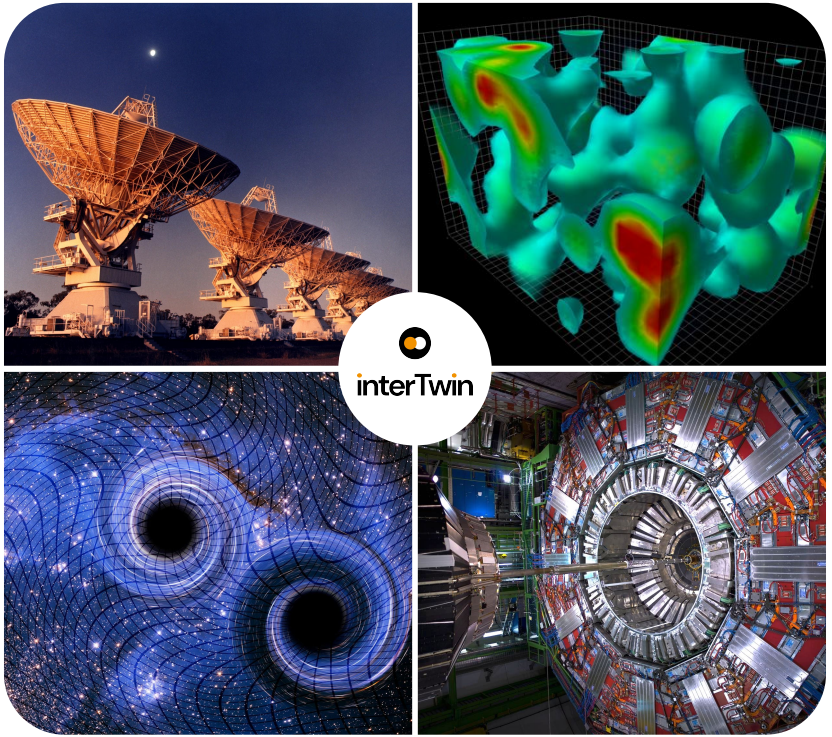

Lattice QCD
CSIC, CNRS

Gravitational Wave Astronomy
INFN

High Energy Physics
CERN, CNRS

interTwin

DT of Particle Detector

Detector Prototyping & Optimization

Build data-driven tool that **simulates detector response** and integrates operation conditions from experimental setups (test-beams).

Online ML for Detectors

Adapt **real-time** detector and/or data acquisition configuration with respect to run conditions

Quality verification & Validation frameworks

Model convergence and accuracy of the generated data should be monitored.

Development of sample-based validation framework in collaboration with HEP community.

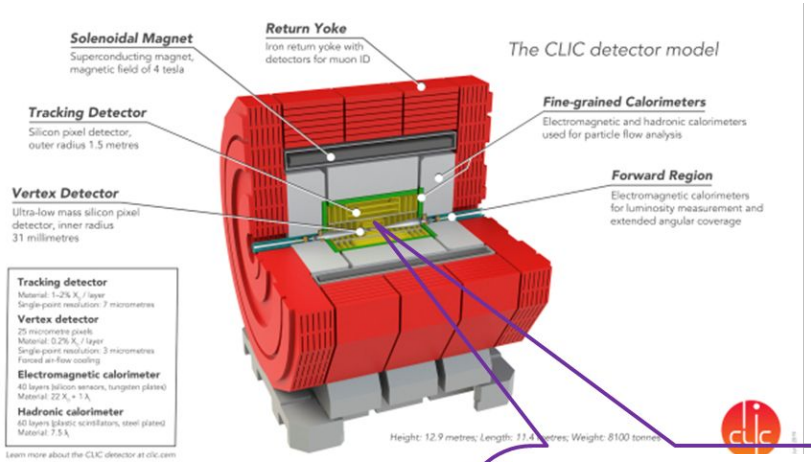

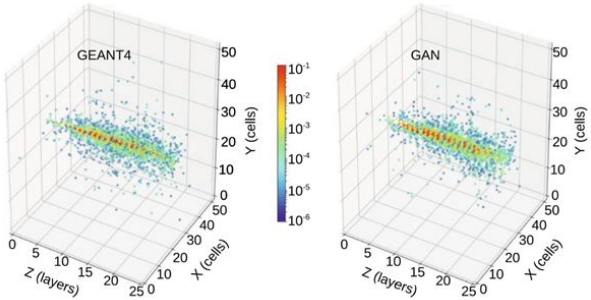

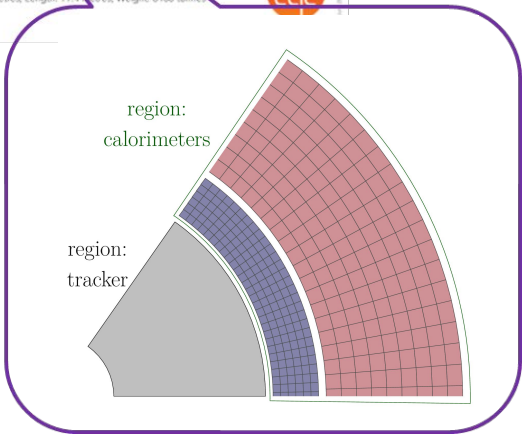

Fast Simulation of a High Granularity Calorimeter by Generative Adversarial Networks. Gul Rukh Khattak et al. https://arxiv.org/abs/2109.07388 DOI: https://doi.org/10.48550/arXiv.2109.07388

interTwin

Benefits of interdisciplinary approach

- **Collaboration**: Increase in cross-community development efforts and unification of frameworks used - “breaking down silos”.
- **Portability**: Run DT workflows infrastructure agnostic across multiple HPC centers in Europe.
- **Extensibility**: Easy addition of new use cases.
- **Modularity**: Customizable according to specific use case’s needs.

DT Core Engine - ML capabilities

Our contribution to the Core Engine

Support AI-based digital twin applications in science:

- **Reproducibility, Reusability, and Modularity**
- **Framework-independent** (e.g., PyTorch, TensorFlow, XGBoost, MLFlow, WandB)
- **UX/UI**: user-friendly GUI (e.g., JupyterLab)
- Off-the-shelf AI tooling:
  - **Hyper-parameters optimization**
  - Scalability (e.g., **distributed ML**)
  - State of the Art **models repository**
- **Seamless access to infrastructure** (cloud and HPC resources)

Our contribution to the Core Engine

Unified distributed training (informal representation)

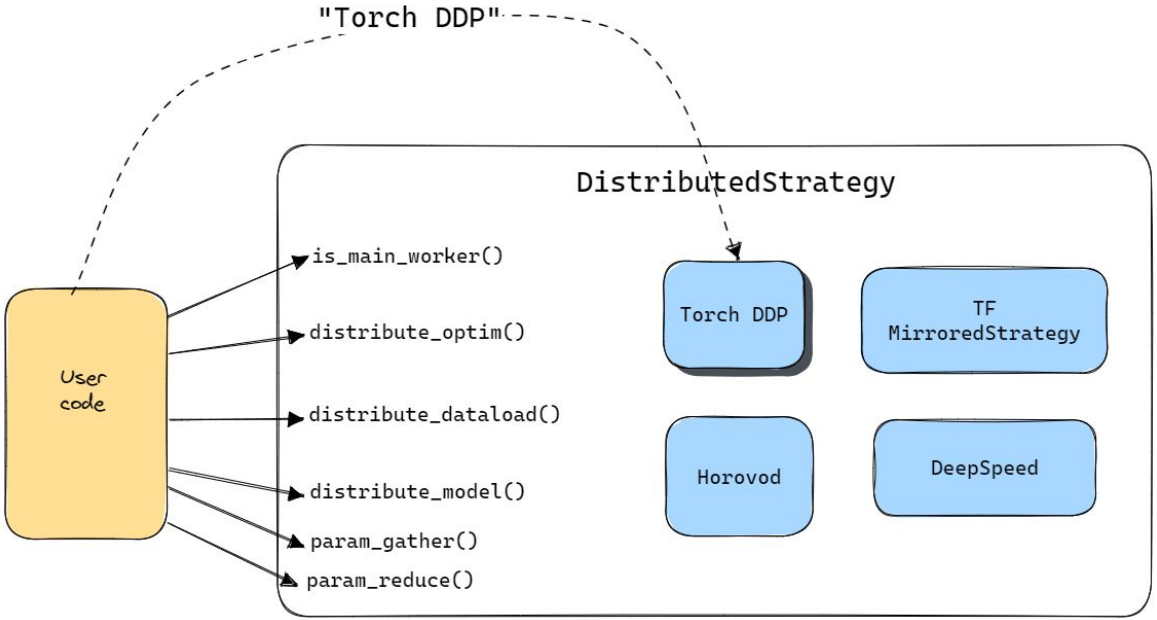

Our contribution to the Core Engine

Generic logger (informal representation)

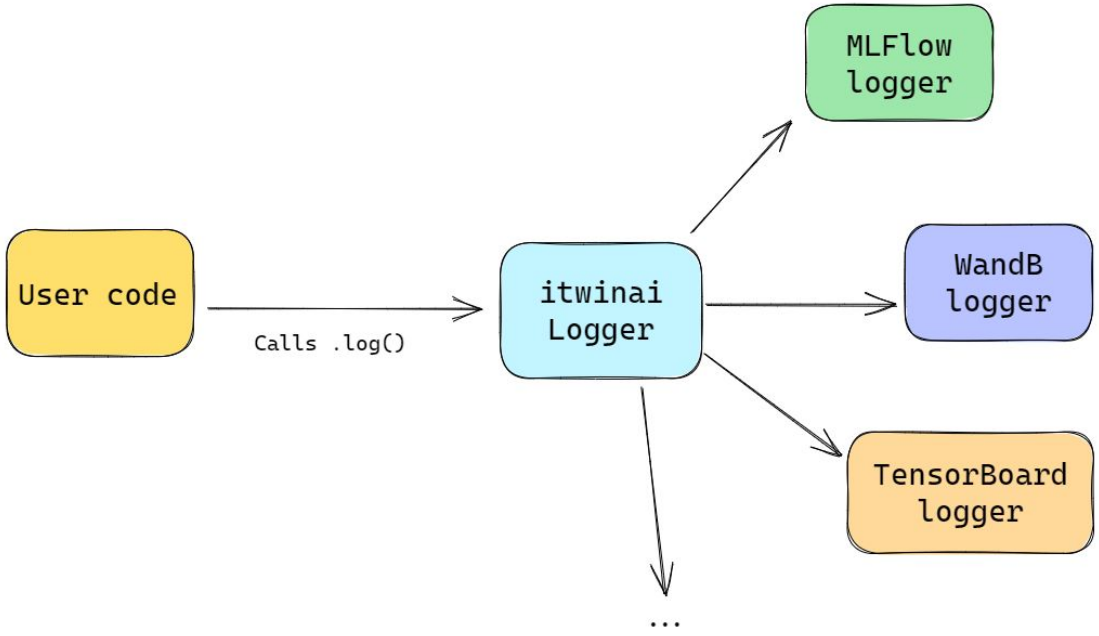

Itwinai - ML tooling for DT applicaitons

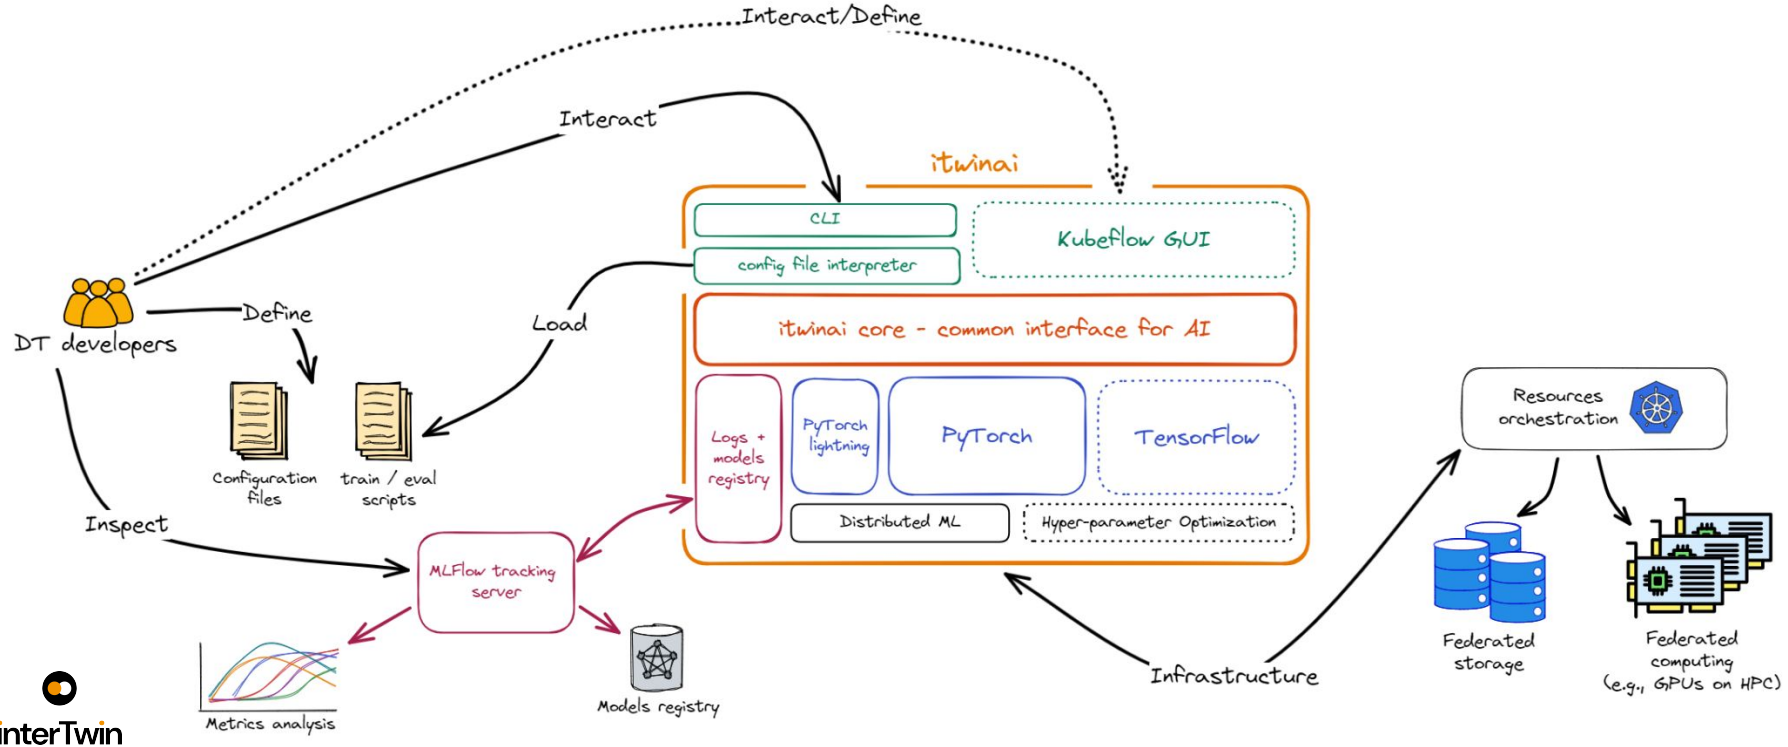

interTwin

UX/UI: KubeFlow

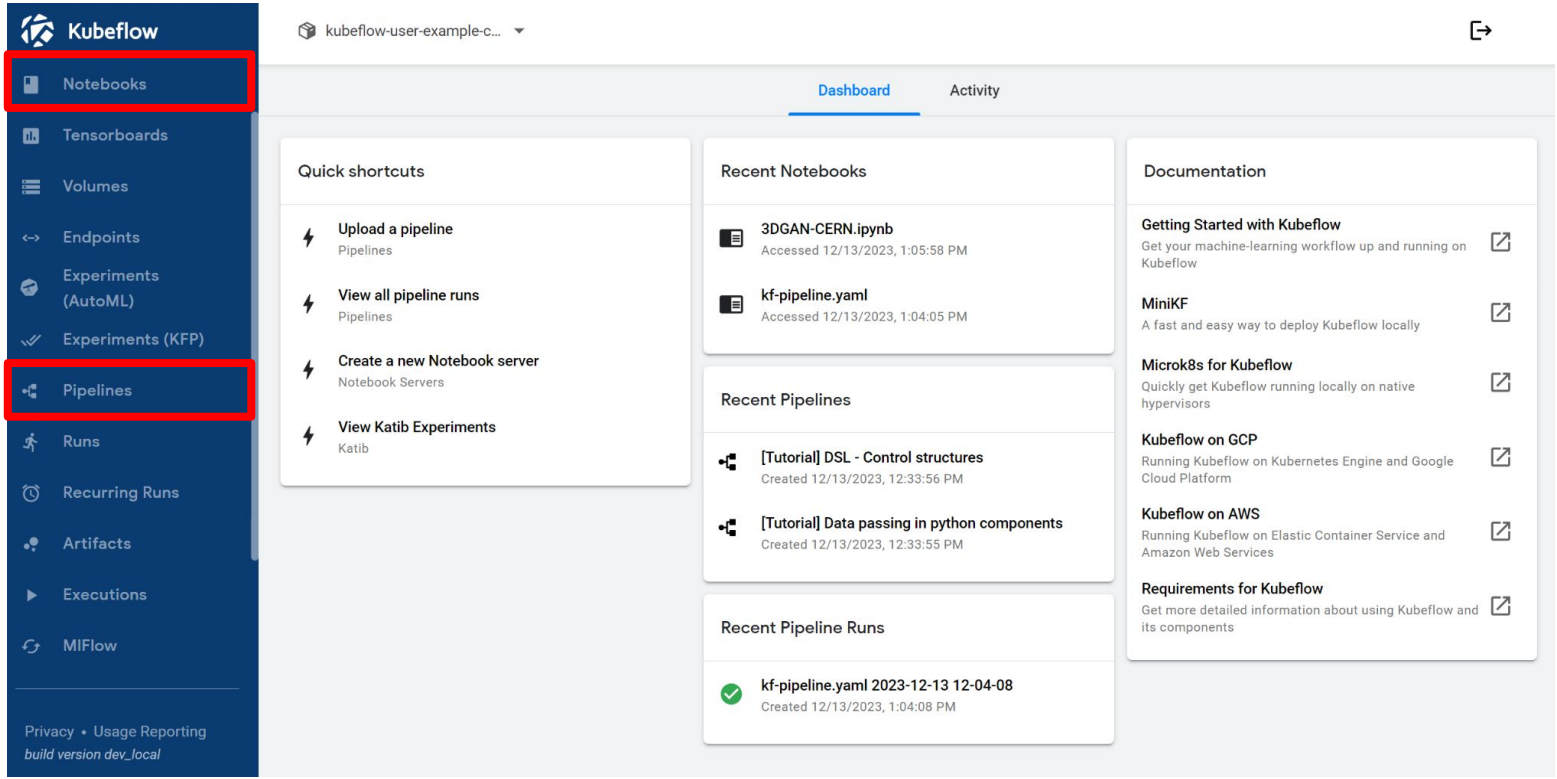

UX/UI: KubeFlow and JupyterLab

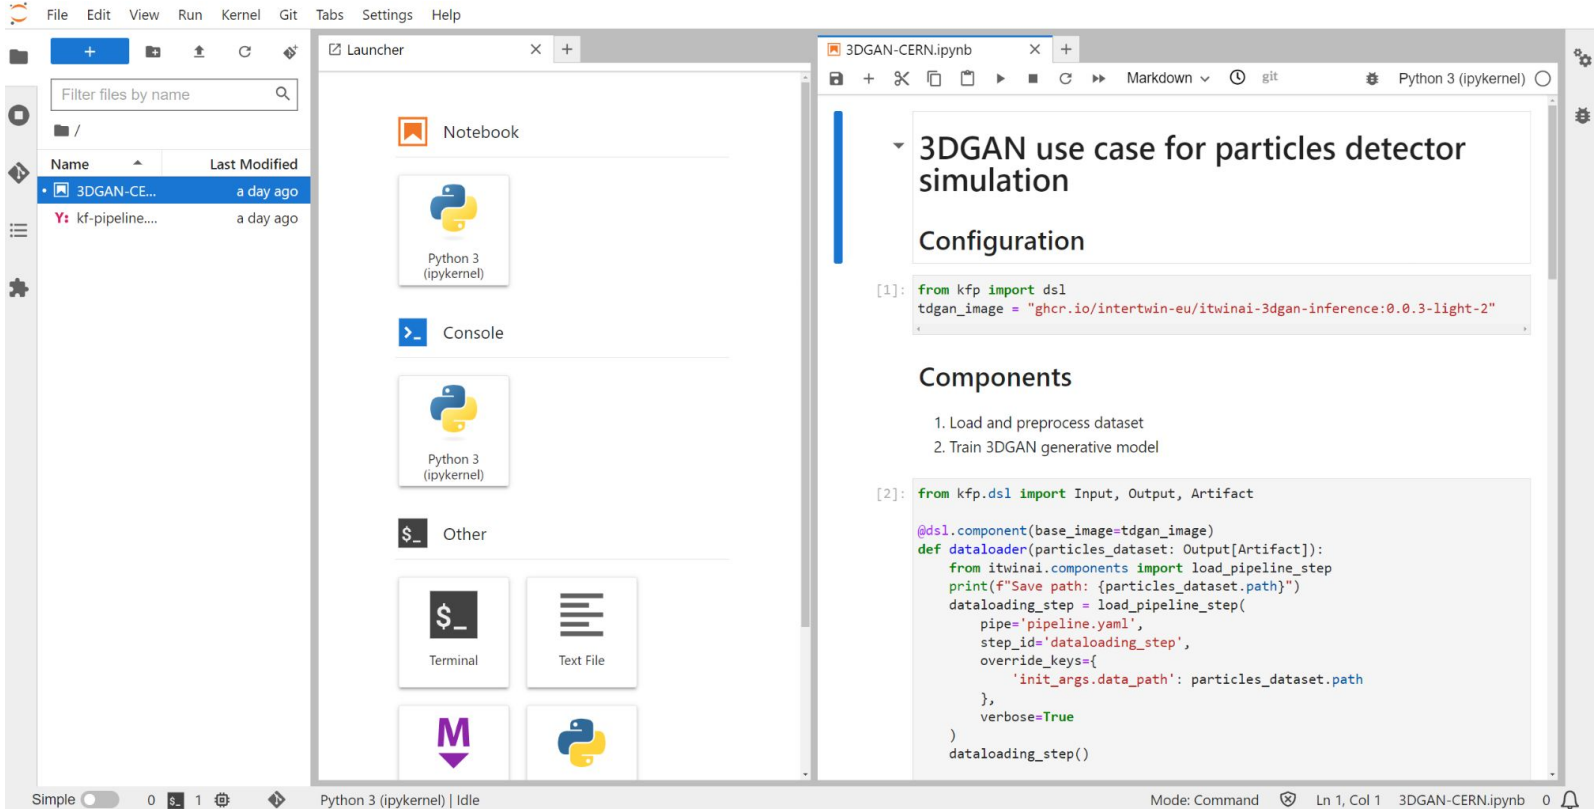

interTwin

UX/UI: KubeFlow pipelines

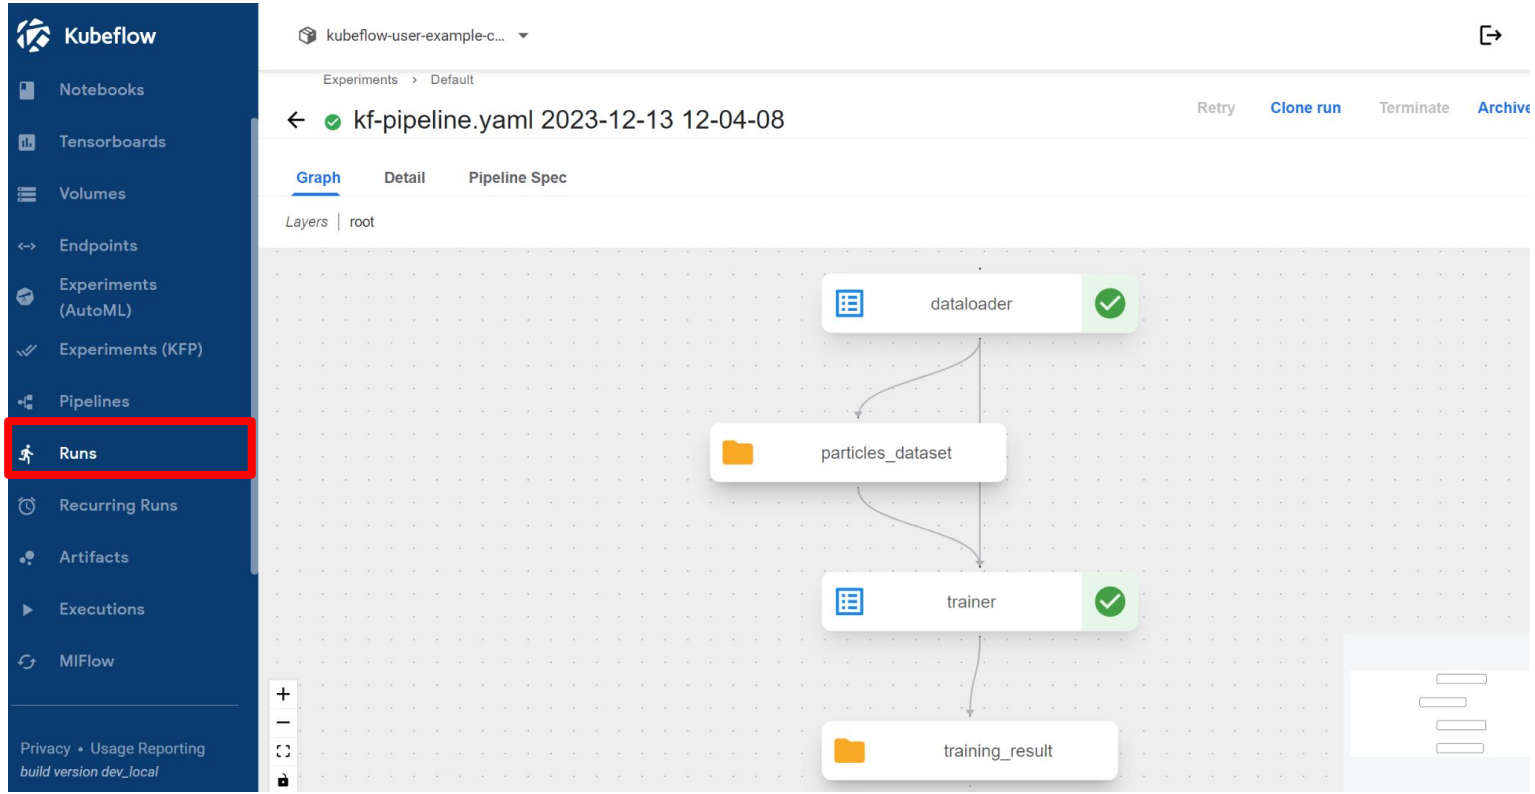

UX/UI: KubeFlow and MLFlow

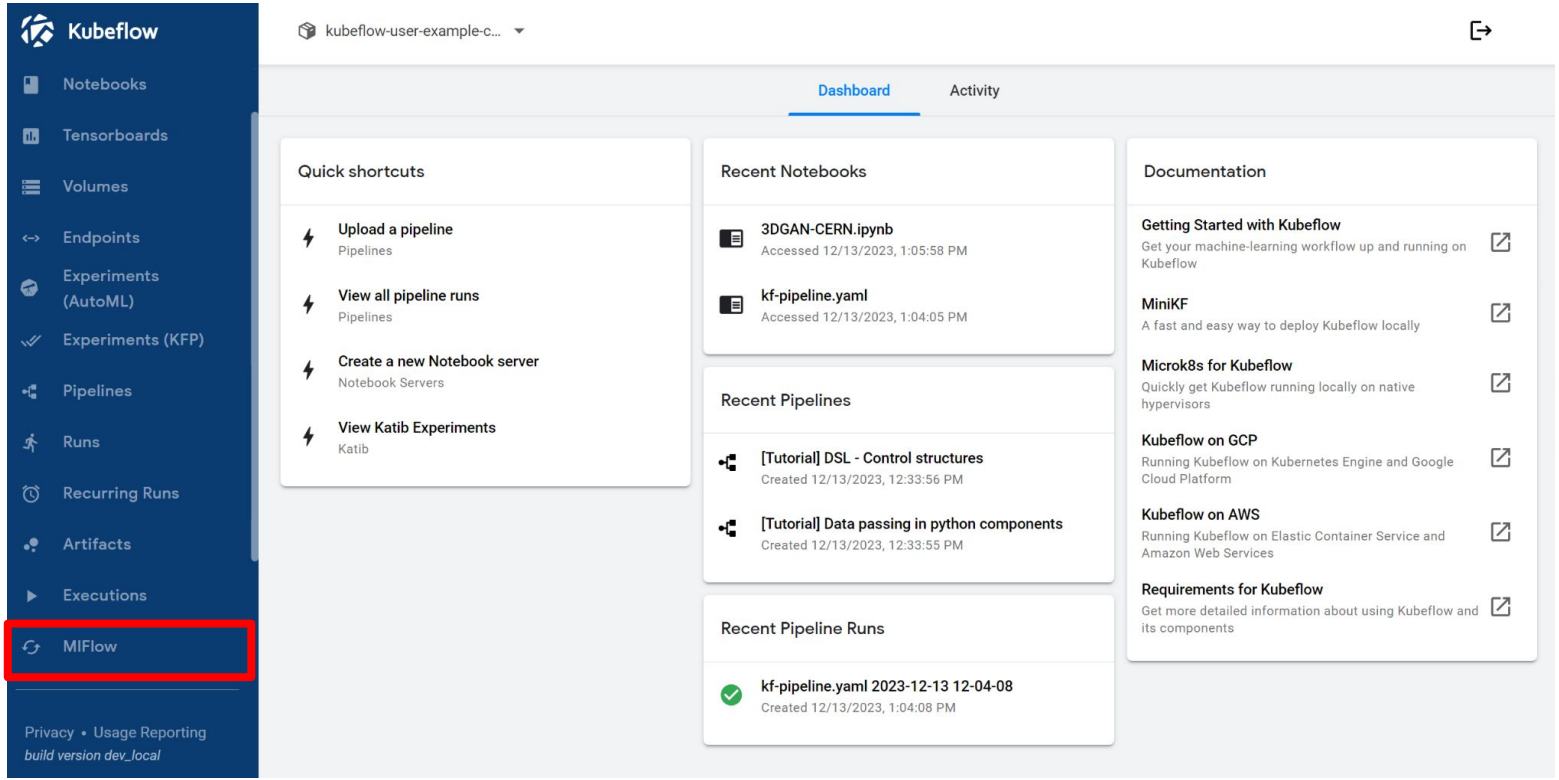

interTwin

UX/UI: MLFlow logs

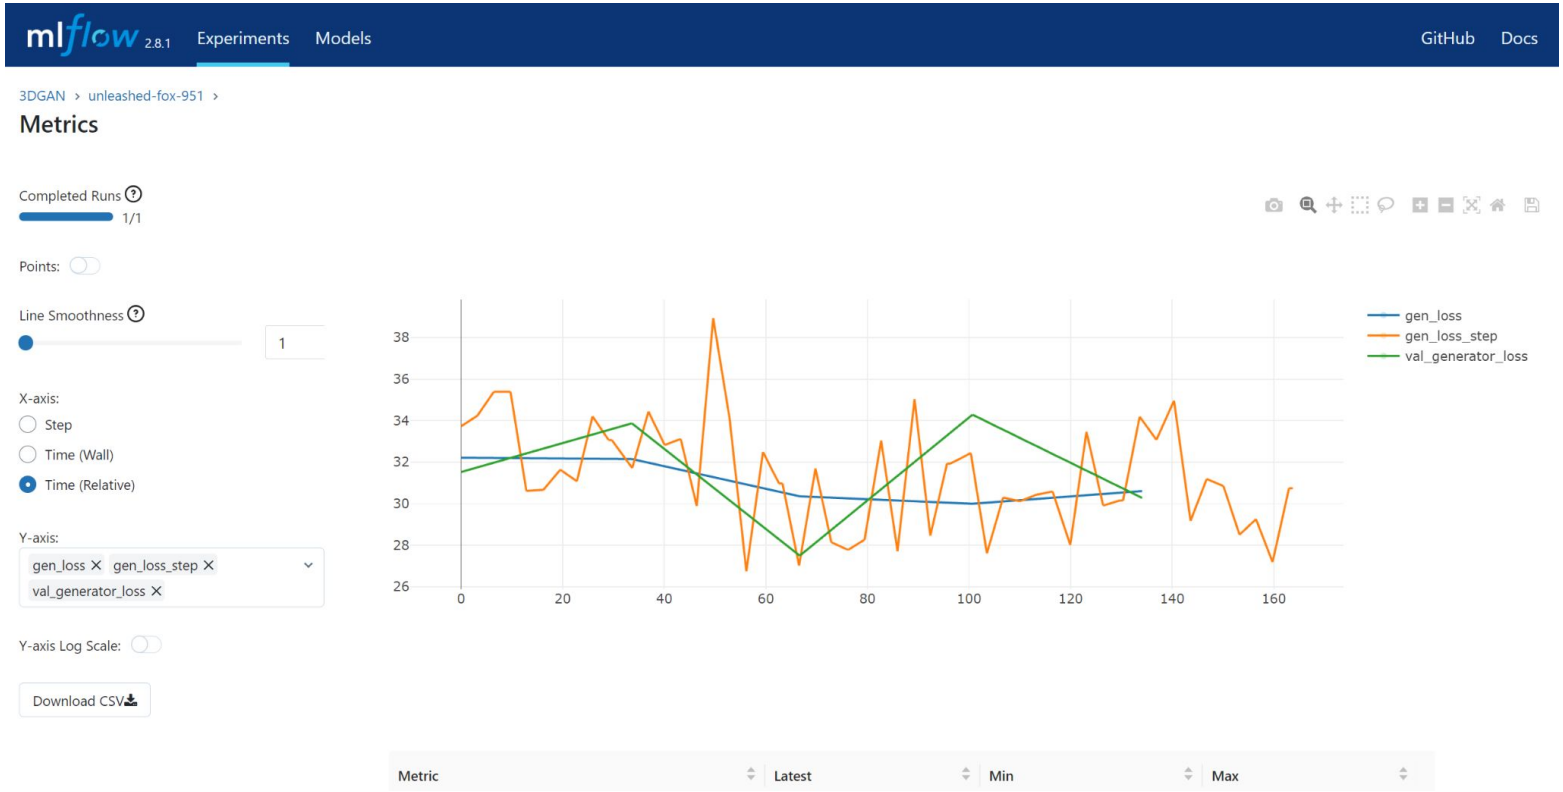

UX/UI: MLFlow models registry

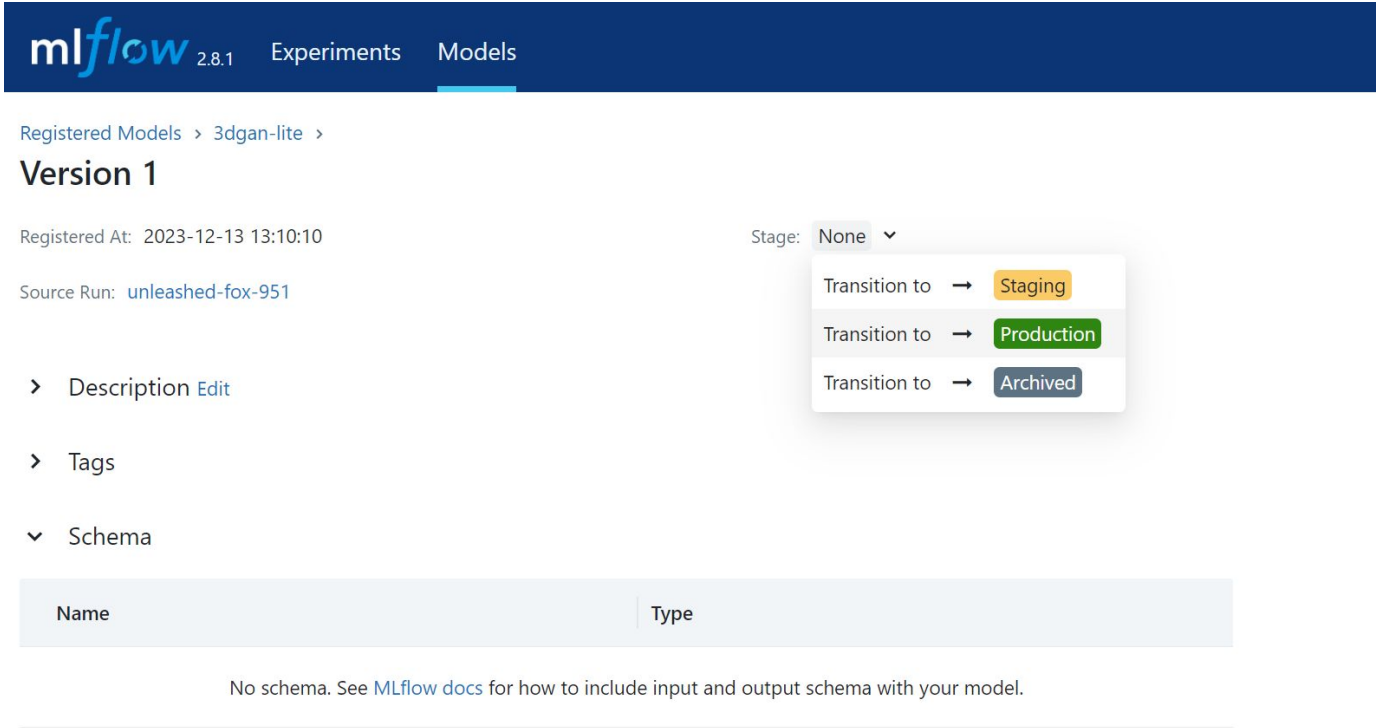

Latest news

Development status of itwinai library:

- Support for **PyTorch**, investigation towards **TensorFlow**
- **AI workflows**: exploration of **KubeFlow** Pipelines.
- **Distributed ML**: integration of existing strategies (e.g., DDP, Horovod, DeepSpeed)..
- **Link with the infrastructure**: Docker/Singularity container, offloaded through WP5's **interLink** on cloud/HPC systems.
- **ML logs and models**: MLFlow tracking

Demo time

Demo video: https://www.youtube.com/watch?v=NoVCfSxwtX0

GitHub repository: https://github.com/interTwin-eu/itwinai

- For the moment, please refer to the “dev” branch: https://github.com/interTwin-eu/itwinai/tree/dev
- Some **tutorials** available, **more to come**. Check “tutorials” folder: https://github.com/interTwin-eu/itwinai/tree/dev/tutorials
- Additional examples are under “use-cases” folder: https://github.com/interTwin-eu/itwinai/tree/dev/use-cases

Other DT initiatives

EMP^2: Environmental Modelling and prediction platform

First proof-of-concept of a machine-learning based global environmental model trained on terabytes of observational data

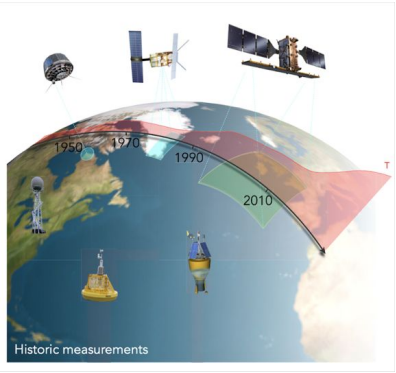

Observations

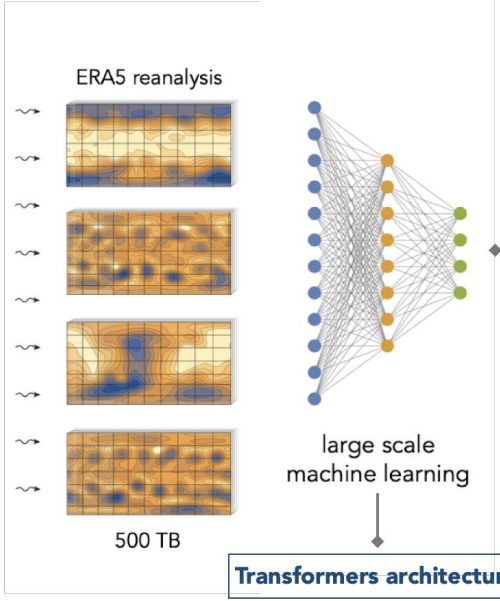

Spatio-temporal representation of atmospheric dynamics

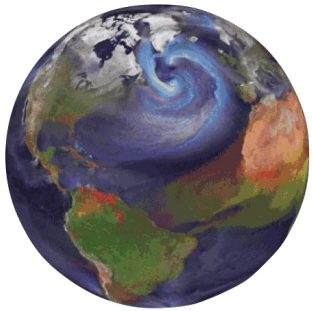

Model
given by the trained neural network

R&D at Juelich SSC:
6x10^5 GPU hours granted in 2022

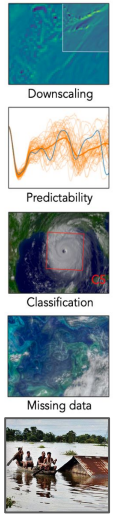

Applications

Visual interface

interTwin

OTTO VON GUERICKE UNIVERSITÄT MAGDEBURG INSTITUT FÜR SIMULATION UND GRAFIK

JÜLICH Forschungszentrum

CERN

CIPEA CERN Innovation Programme on Environmental Applications

Slide courtesy of Ilaria Luise

Nvidia Omniverse (OV)

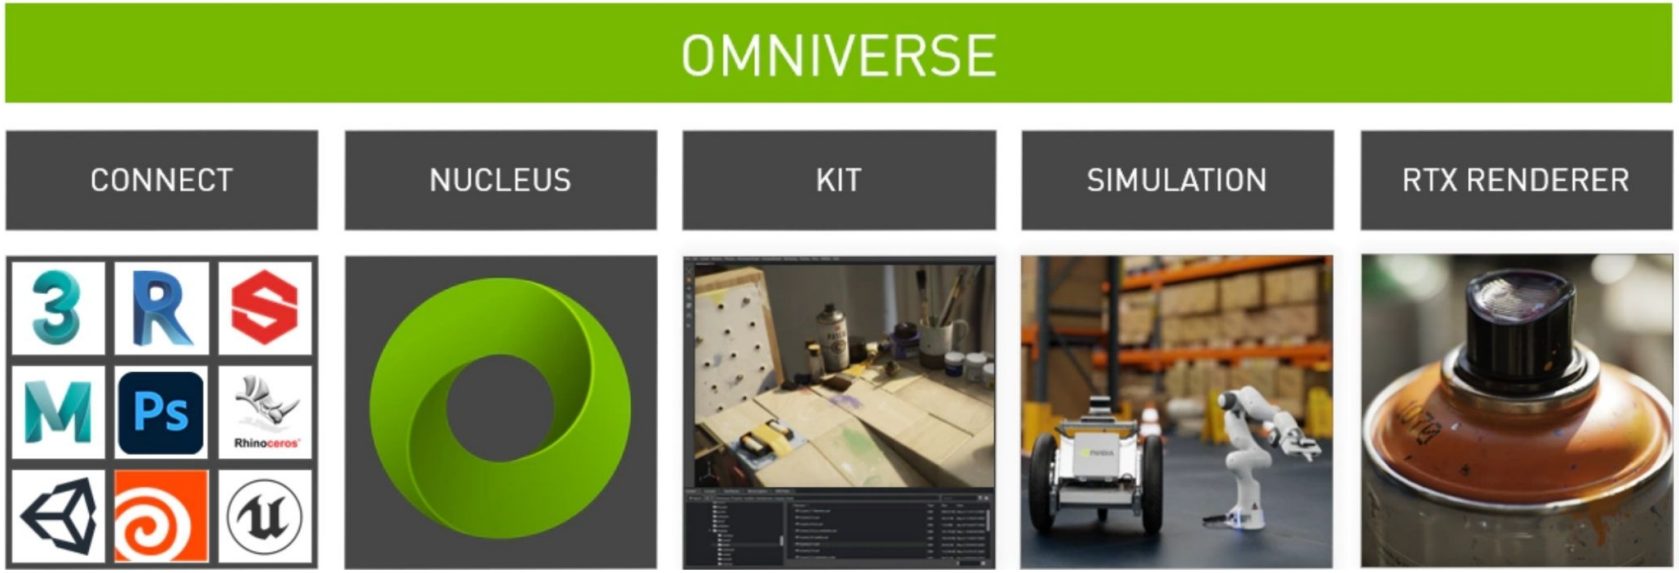

OV for DTs of the accelerators complex

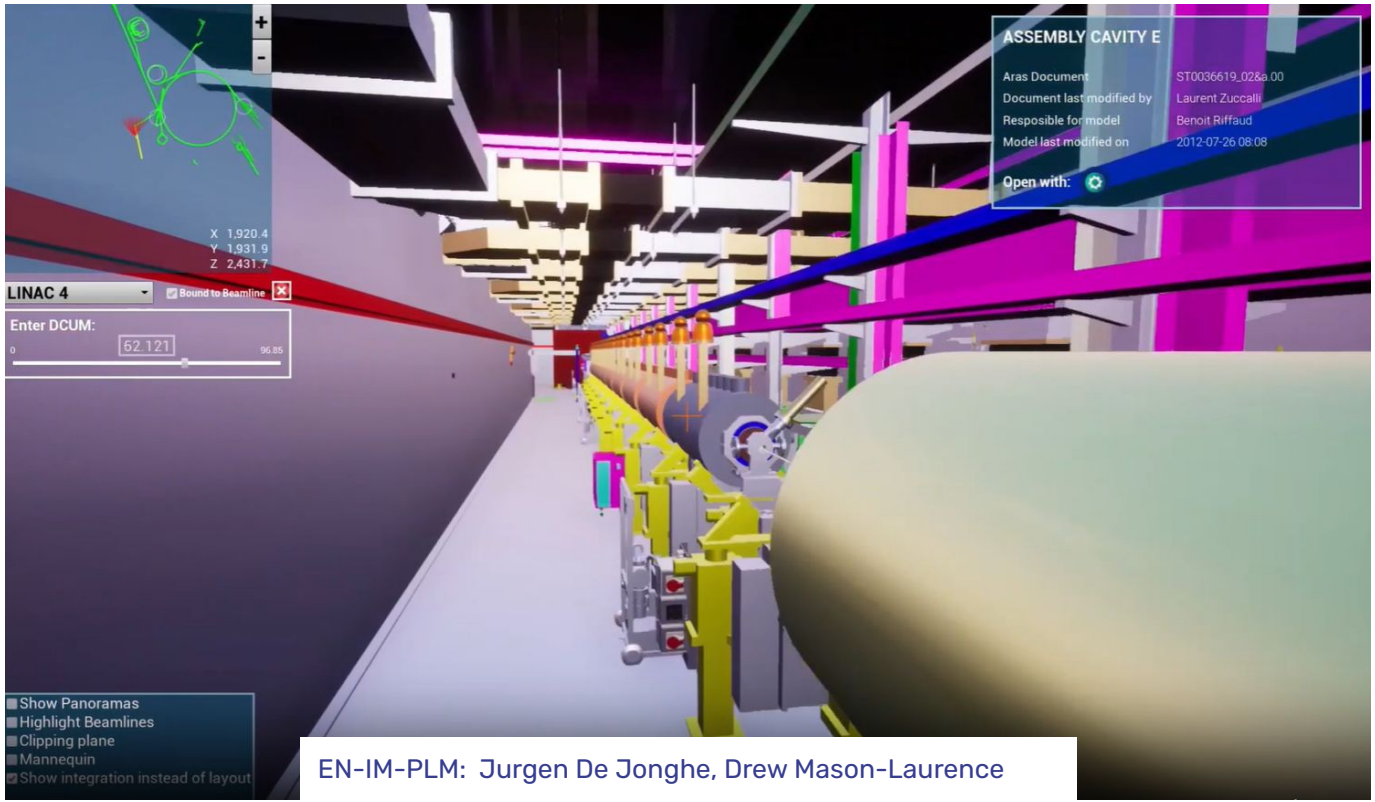

EN-IM-PLM: Jurgen De Jonghe, Drew Mason-Laurence

Nvidia Omniverse - Opportunities

- Detector assembly simulation

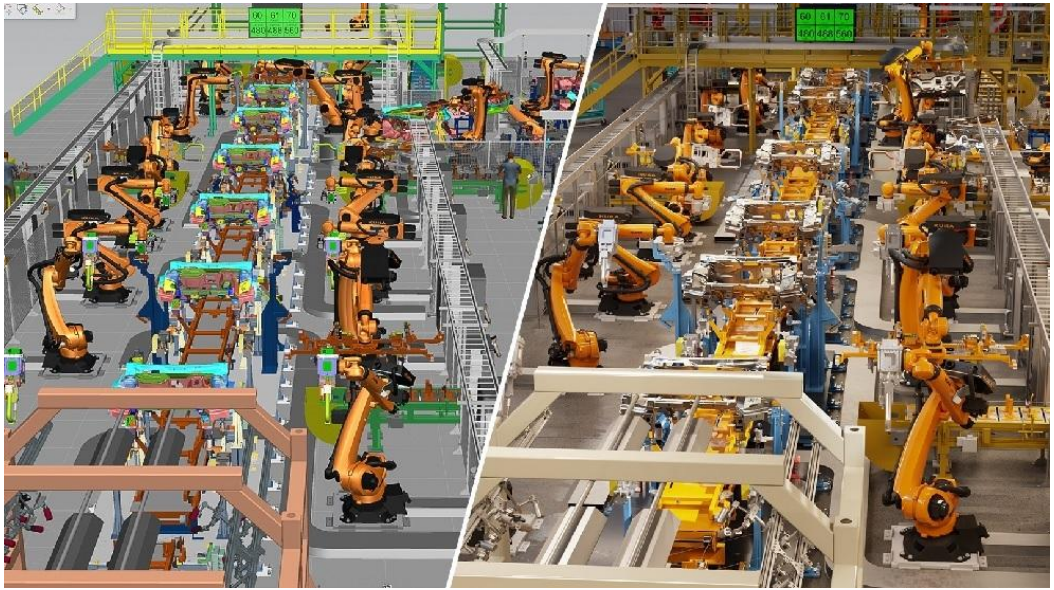

Image credits: https://blogs.sw.siemens.com/thought-leadership/2022/06/29/creating-the-industrial-metaverse-siemens-xcelerator-nvidia-omniverse/

Nvidia Omniverse - Opportunities

- Robotic simulation. IsaacGym can also train Reinforcement learning agents.

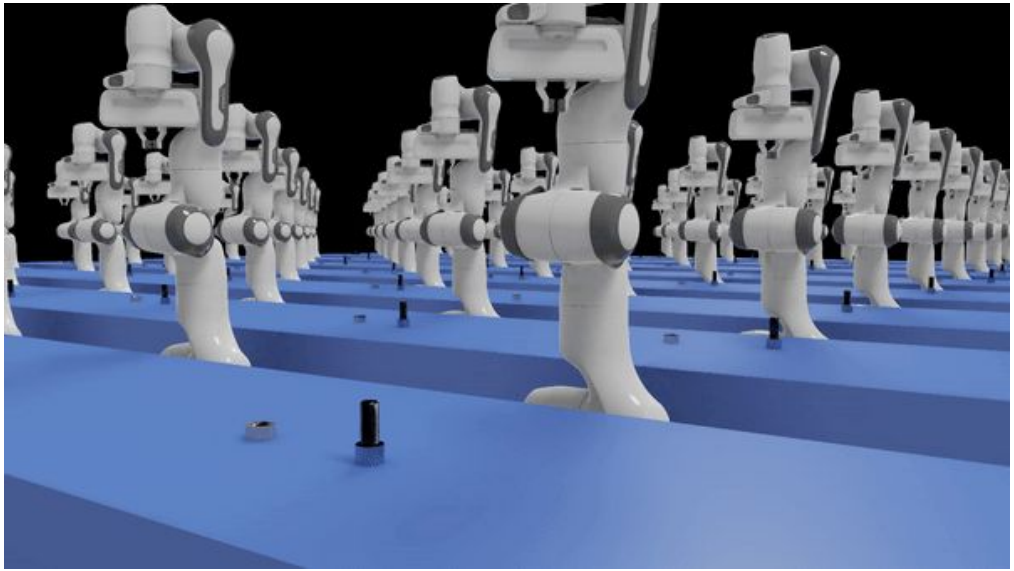

Image credits: https://developer.nvidia.com/blog/advancing-robotic-assembly-with-a-novel-simulation-approach-using-nvidia-isaac/

See also: https://github.com/NVIDIA-Omniverse/OmniIsaacGymEnvs

Nvidia Omniverse - Opportunities

- Visualizations for HEP with Geant4

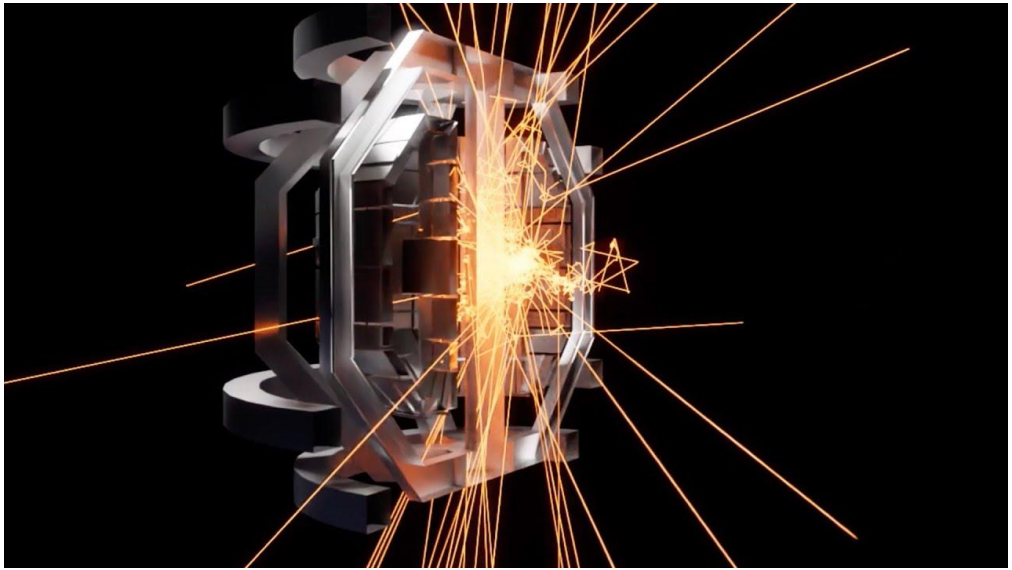

Image credits: https://blogs.nvidia.com/blog/ukaea-digital-twins-omniverse/


CERN
openlab

Digital Twin Engine

A focus on the architecture

interTwin consortium

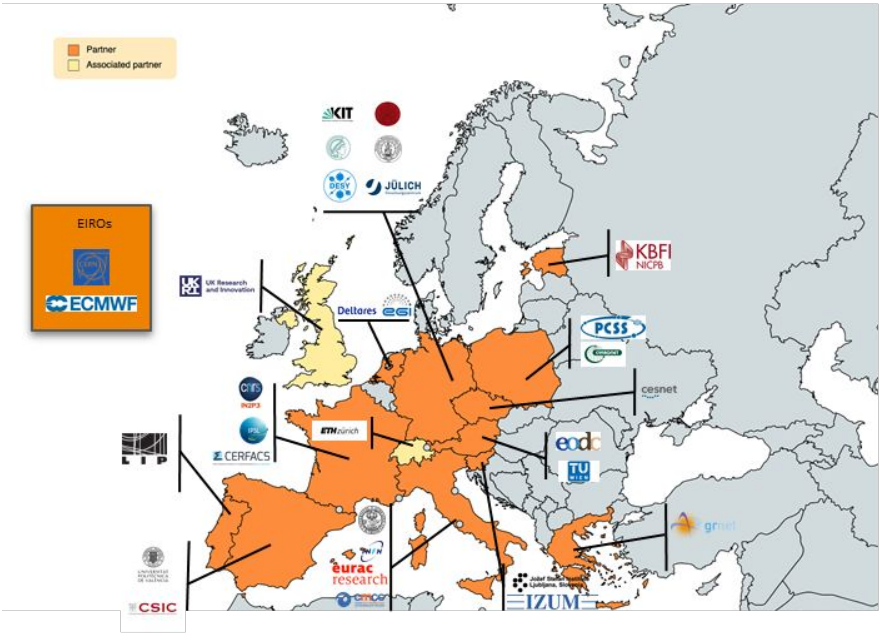

EGI Foundation as coordinator

29 **Participants**, including 1 affiliated entity and 2 associated partners

Consortium at a glance

10	11	14
Providers	Technology providers	Community representants
cloud, HTC , HPC resources and access to Quantum systems	delivering the DTE infrastructure and horizontal capabilities	from 5 scientific areas; requirements and developing DT applications and thematic modules

interTwin

interTwin - DT Engine stack

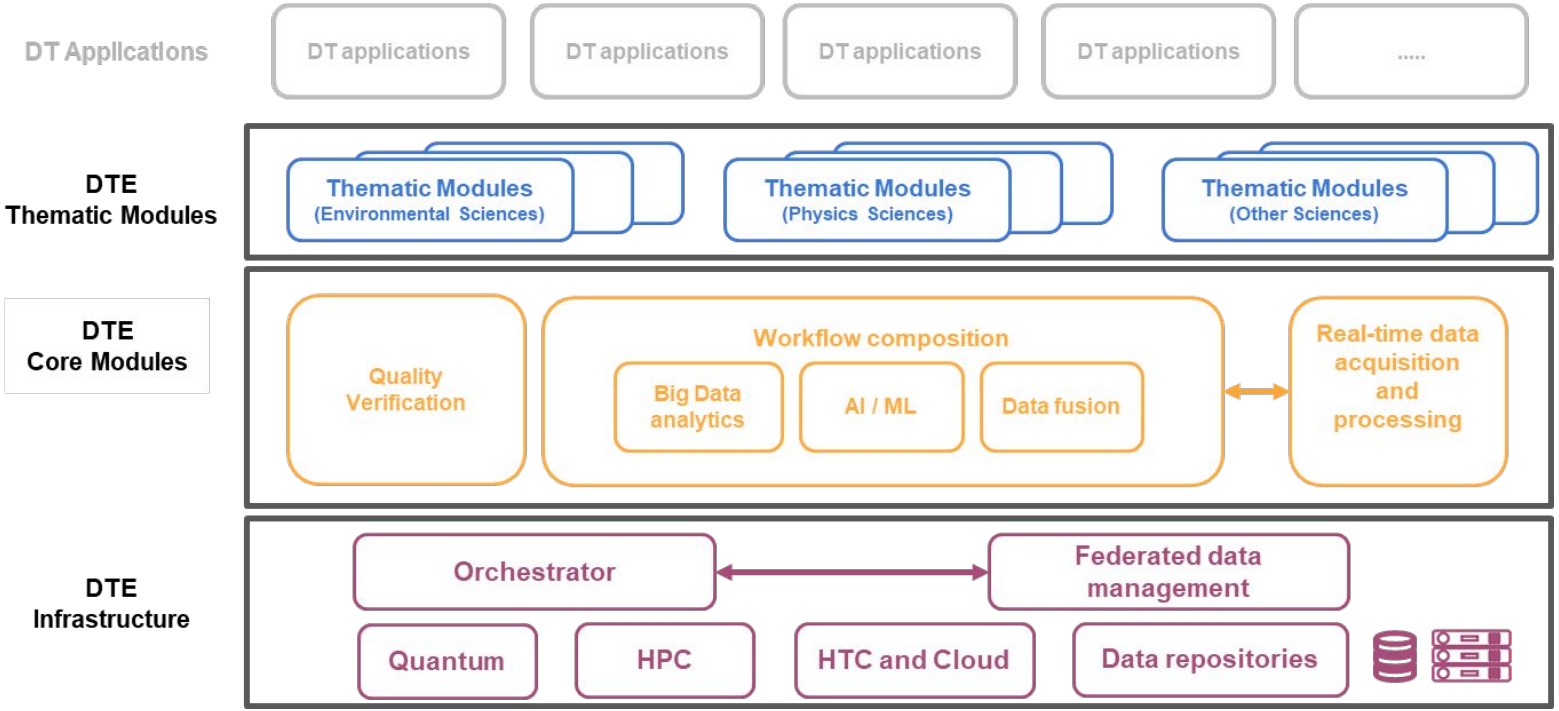

interTwin

DT Engine for science

Today

- DTs developed in isolation
- Community-specific technologies and standards
- **Great overheads** (i.e., reinventing the wheel)

Boils down to… **need for improvement.**

Tomorrow

- Unified DTE framework
- Standard DT lifecycle management, thanks to co-design
- **Low overheads** (engineering)

Results into… **accelerated science!**

interTwin

Digital twin engine - System context

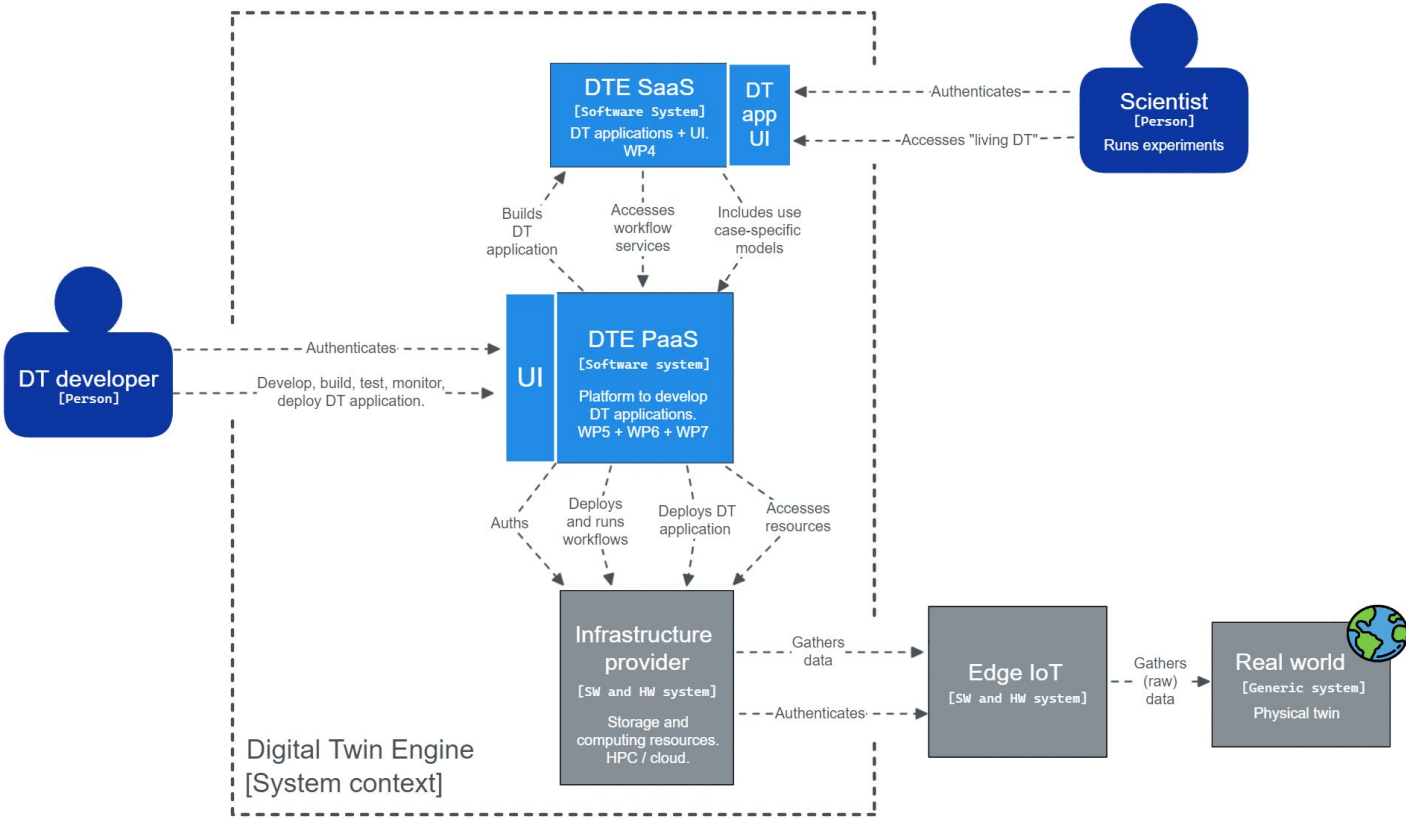

Digital twin engine - DT workflows

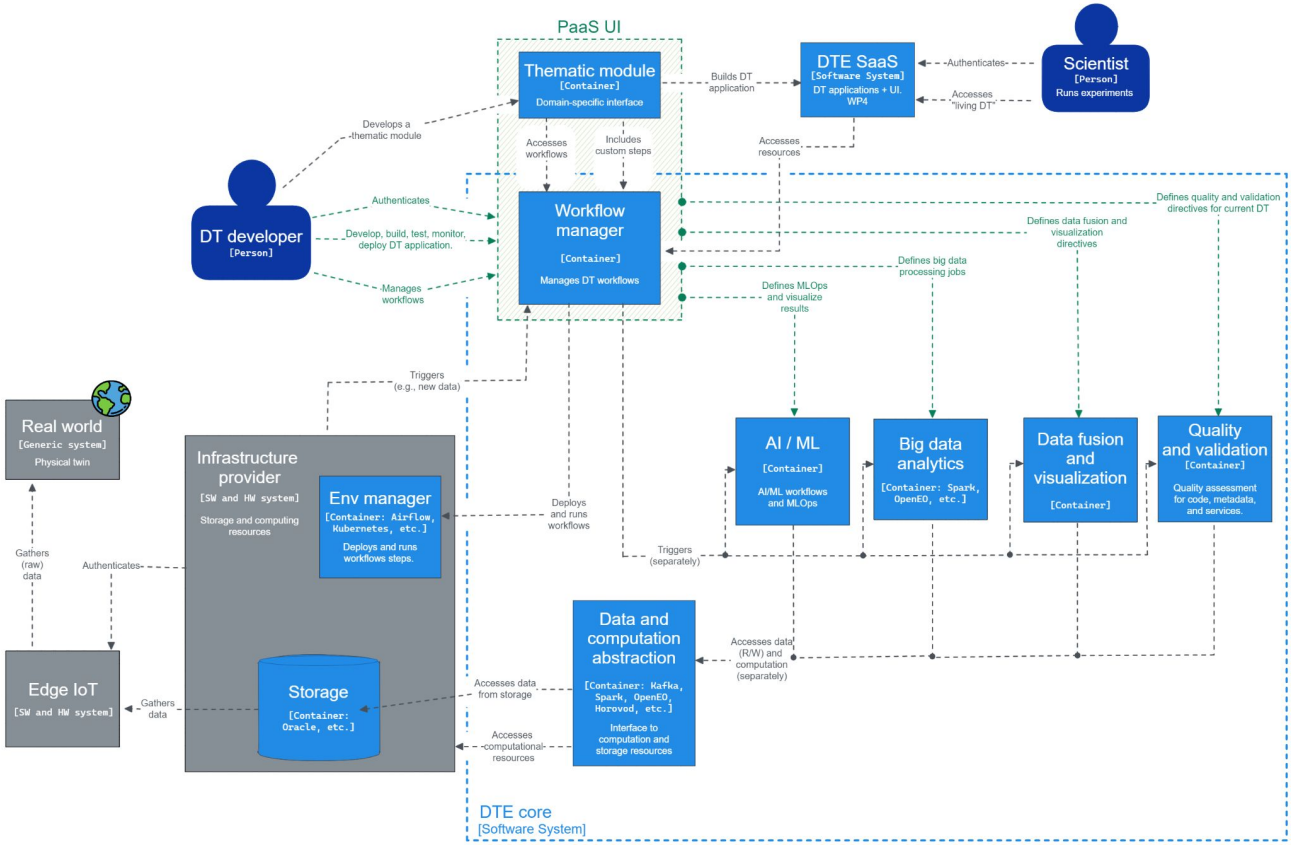

Use case-specific components

Reuse pre-existing components and workflows from sciences.