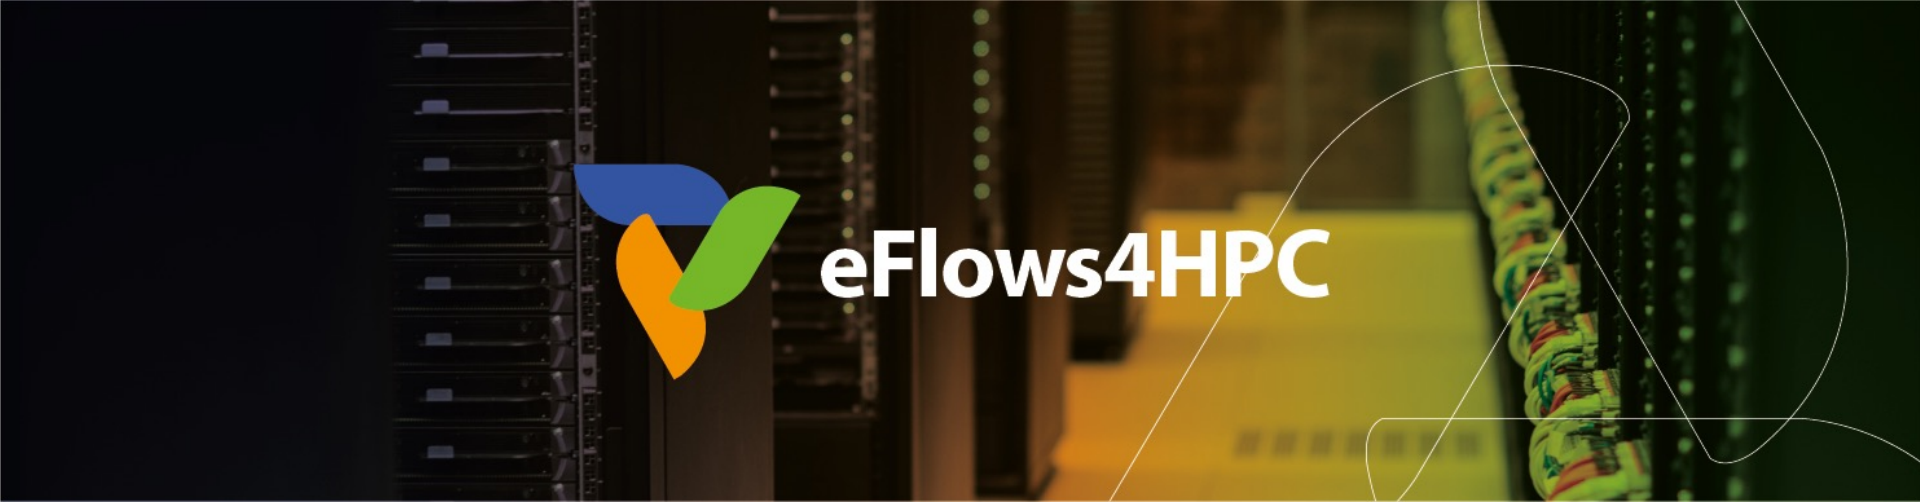

Dynamic and Intelligent workflows with eFlows4HPC

Rosa M Badia– BSC

ScaDL 2021

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955558. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Germany, France, Italy, Poland, Switzerland, Norway.

Complex workflows and complex infrastructures

- Different methodologies and software stacks
  - HPC modeling and simulation
  - Big data
  - Artificial Intelligence
- There is a need for integration of the different aspects in a single application workflow
- Other aspects:
  - Specific hardware
  - Compute continuum

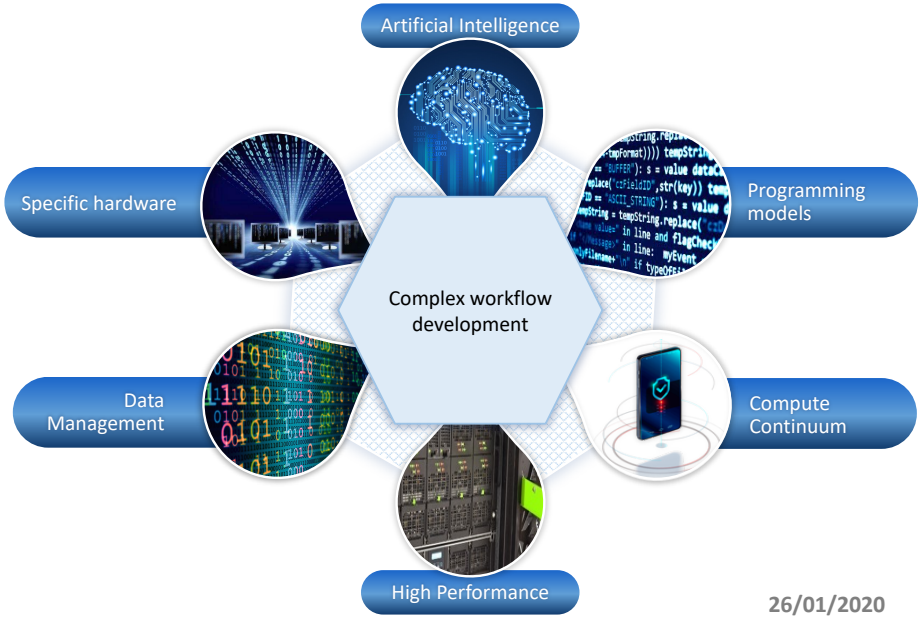

Challenges that we aim to tackle

- Openness, transparency, re-usability, reproducibility and accessibility of the workflows and their results
  - Can we adopt methods similar to those used in the cloud?
- Simplify the development of complex workflows while keeping their capabilities and performance
  - Simple and intuitive
  - Bridging the gap between HPC, AI and Big Data
- Workflow dynamicity
  - React to changes, events
  - Urgency, real-time
- Data management
  - Offer new layer that deals with data managemetn, new storage devices

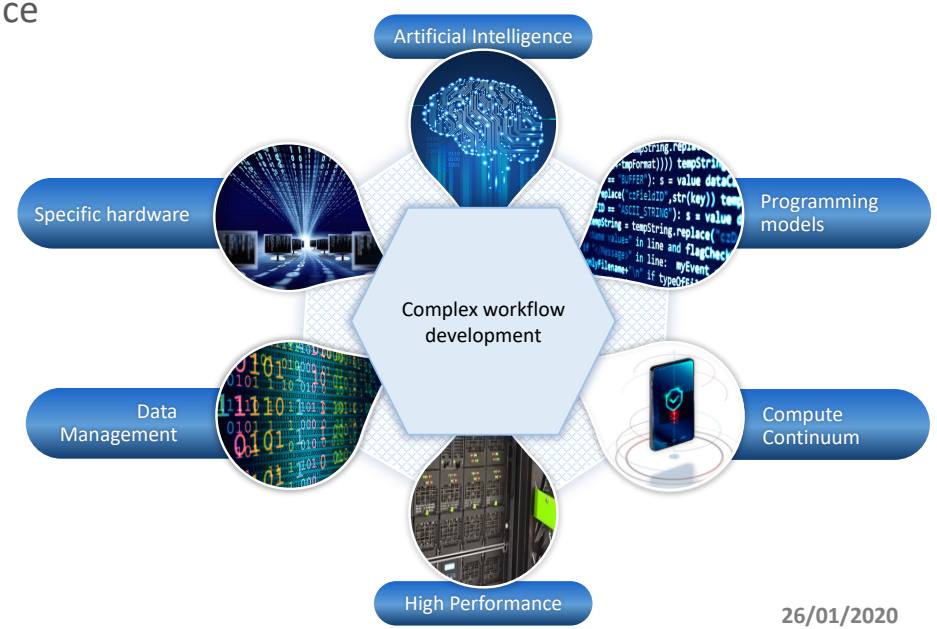

EuroHPC and its projects

- The European High Performance Computing Joint Undertaking (EuroHPC JU) is a joint initiative between the EU, European countries and private partners to develop a World Class Supercomputing Ecosystem in Europe
  - Procuring and deploying 3 pre-exascale and 5 petascale systems in Europe
  - Two additional Exascale systems
- Supporting research and innovation activities
  - Call on Jan 2020: EuroHPC-02-2019: HPC and data-centric environments and application platforms
  - High Performance Computing (HPC) and data driven HPC software environments and application oriented platforms

ACROSS

HEROES

MICROCARD

REGALE

eFlows4HPC

Outline

- The eFlows4HPC project
  - Pillars' applications
  - General objectives
  - Initial architecture
- Conclusions

Users' Communities

Pillar I: Digital twins

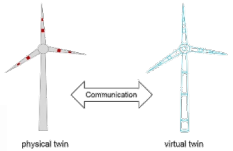

EXCELLERAT

Pillar II: Climate

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Pillar III: Urgent Computing

ChEESE

use

HPC Workflow as a Service

eFlows4HPC Software Stack

Architectural optimizations

Federated HPC Infrastructure

Cloud Infrastructure

Pillar I: Manufacturing

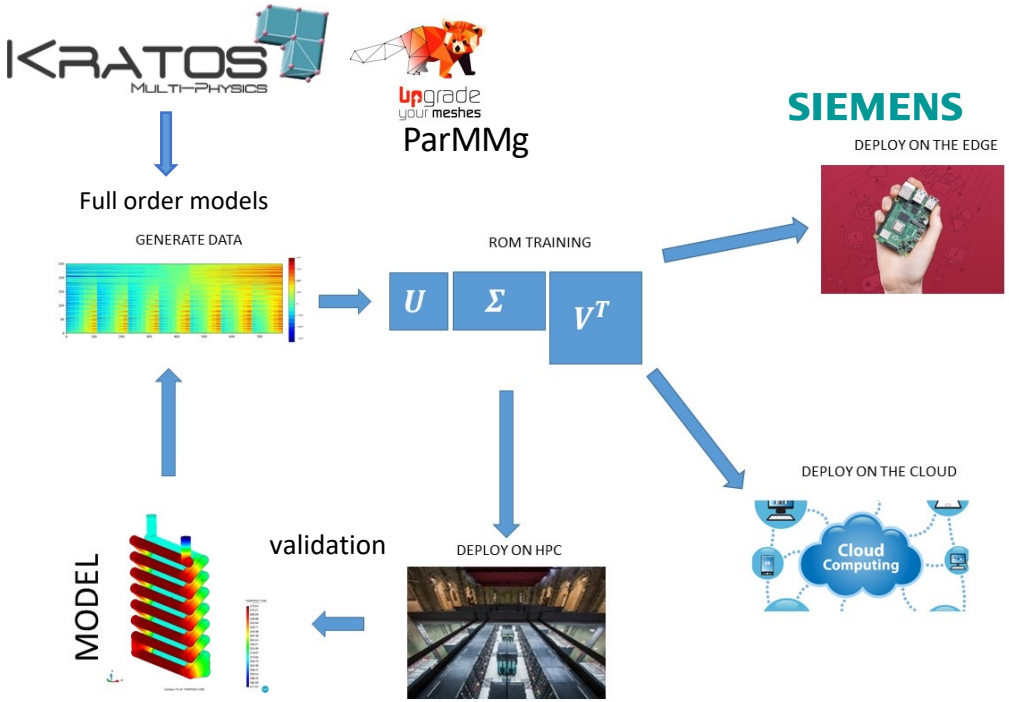

Pillar I focuses on the construction of DigitalTwins for the prototyping of complex manufactured objects:

- Integrating state-of-the-art adaptive solvers with machine learning and data-mining
- Contributing to the Industry 4.0 vision

Pillar II: Climate

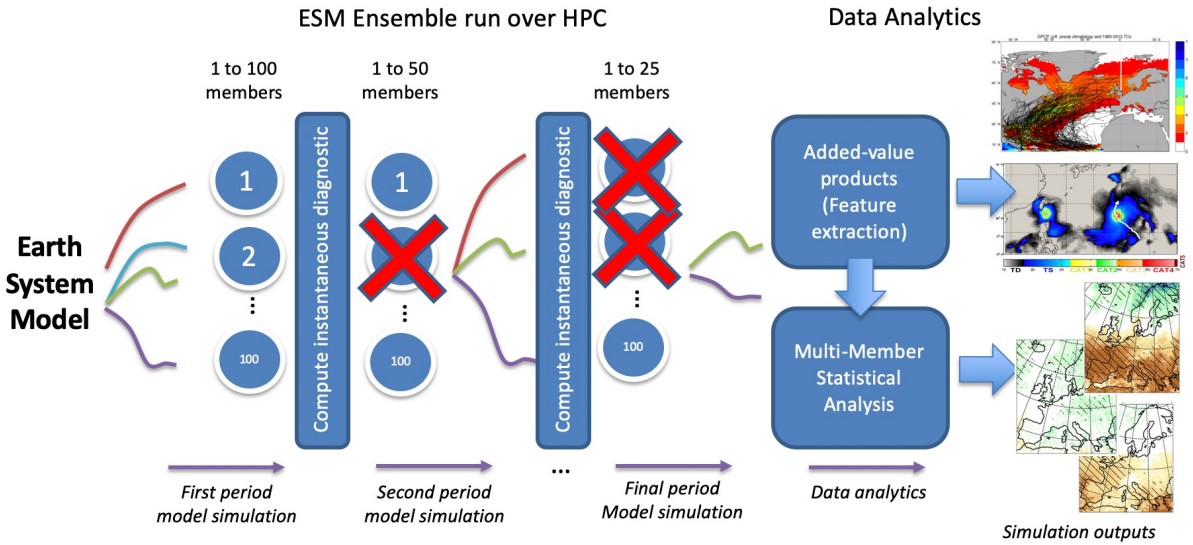

FESOM2
Finite volumE
Sea ice-Ocean Model

Pillar II develops innovative adaptive workflows for climate and for the study of Tropical Cyclones (TC):

- in the context of the CMIP6 experiment
- including in-situ analytics

Pillar III: Urgent computing for natural hazards

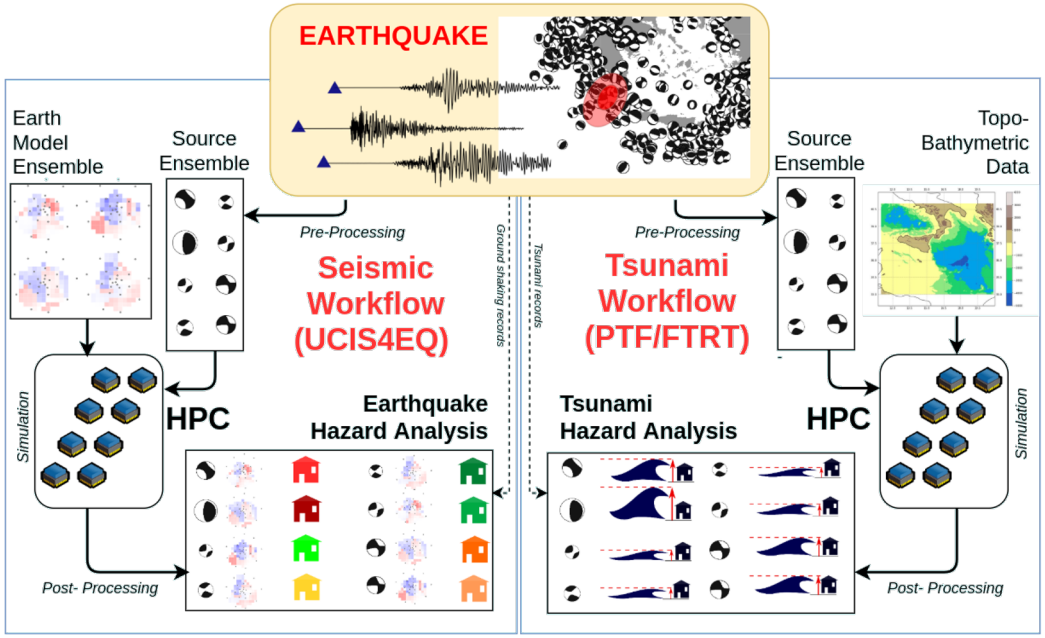

Pillar III explores the modelling of natural catastrophes:

- Earthquakes and their associated tsunamis shortly after such an event is recorded
- Usage of urgent computing on emergencies
- Use of AI to estimate intensity maps
- Use of DA and AI tools to enhance event diagnostics

Tsunami-HySEA GPU-based code

Main objectives

- To deliver an **open European workflows software stack** (eFlows4HPC) for the development and management of complex workflows, by integrating and extending existing European components and enabling their energy-efficient execution
- To enable the definition of complex workflows integrating **HPC simulation and modelling** with **high-performance data analytics** and **machine learning** in **scientific and industrial applications**
- To define an **open methodology** that widens and eases the use of workflows by existing and **new HPC communities and users** — HPC Workflows as a Service (**HPCWaaS**).
- To provide **means to increase openness, transparency, reusability, and reproducibility** of computation results by means of providing a catalogs, repositories and registries that store data sets and software components, including whole workflow instances
- To **optimize specific application kernels** from use cases to efficiently accommodate the **EPI** and **other emerging platforms and architectures** (GPUs, FPGAs, and custom accelerators), addressing **energy efficiency** goals as well

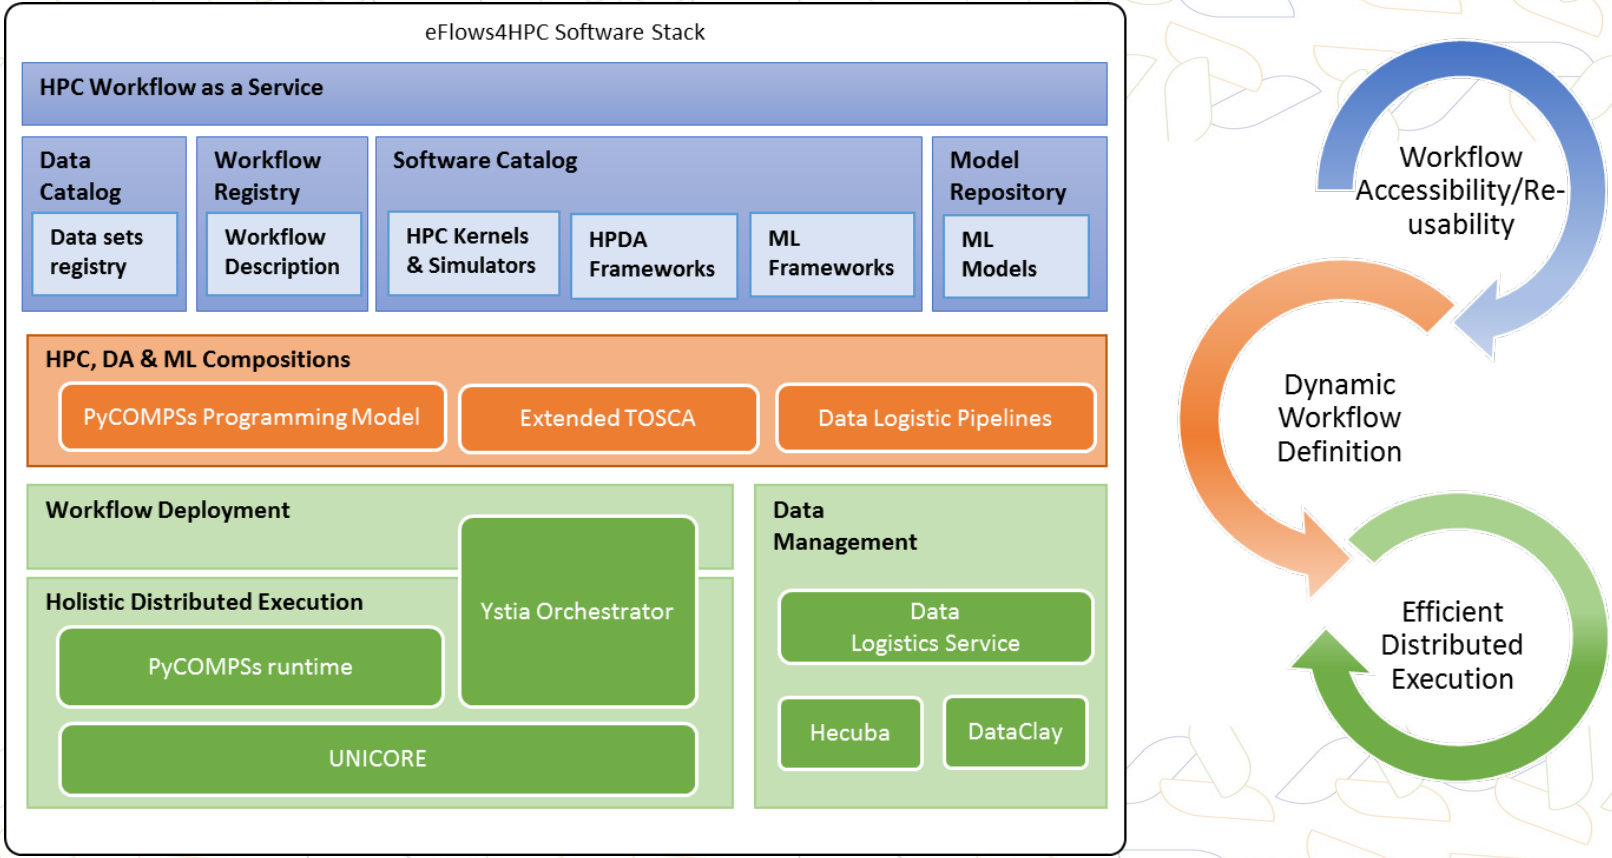
eFlows4HPC Software Stack
HPC Workflow as a Service
Data Catalog
Data sets registry
Workflow Registry
Workflow Description
Software Catalog
HPC Kernels & Simulators
HPDA Frameworks
ML Frameworks
Model Repository
ML Models
HPC, DA & ML Compositions
PyCOMPSs Programming Model
Extended TOSCA
Data Logistic Pipelines
Workflow Deployment
Ystia Orchestrator
Holistic Distributed Execution
PyCOMPSs runtime
UNICORE
Data Management
Data Logistics Service
Hecuba
DataClay

Workflow Accessibility/Re-usability
Dynamic Workflow Definition
Efficient Distributed Execution

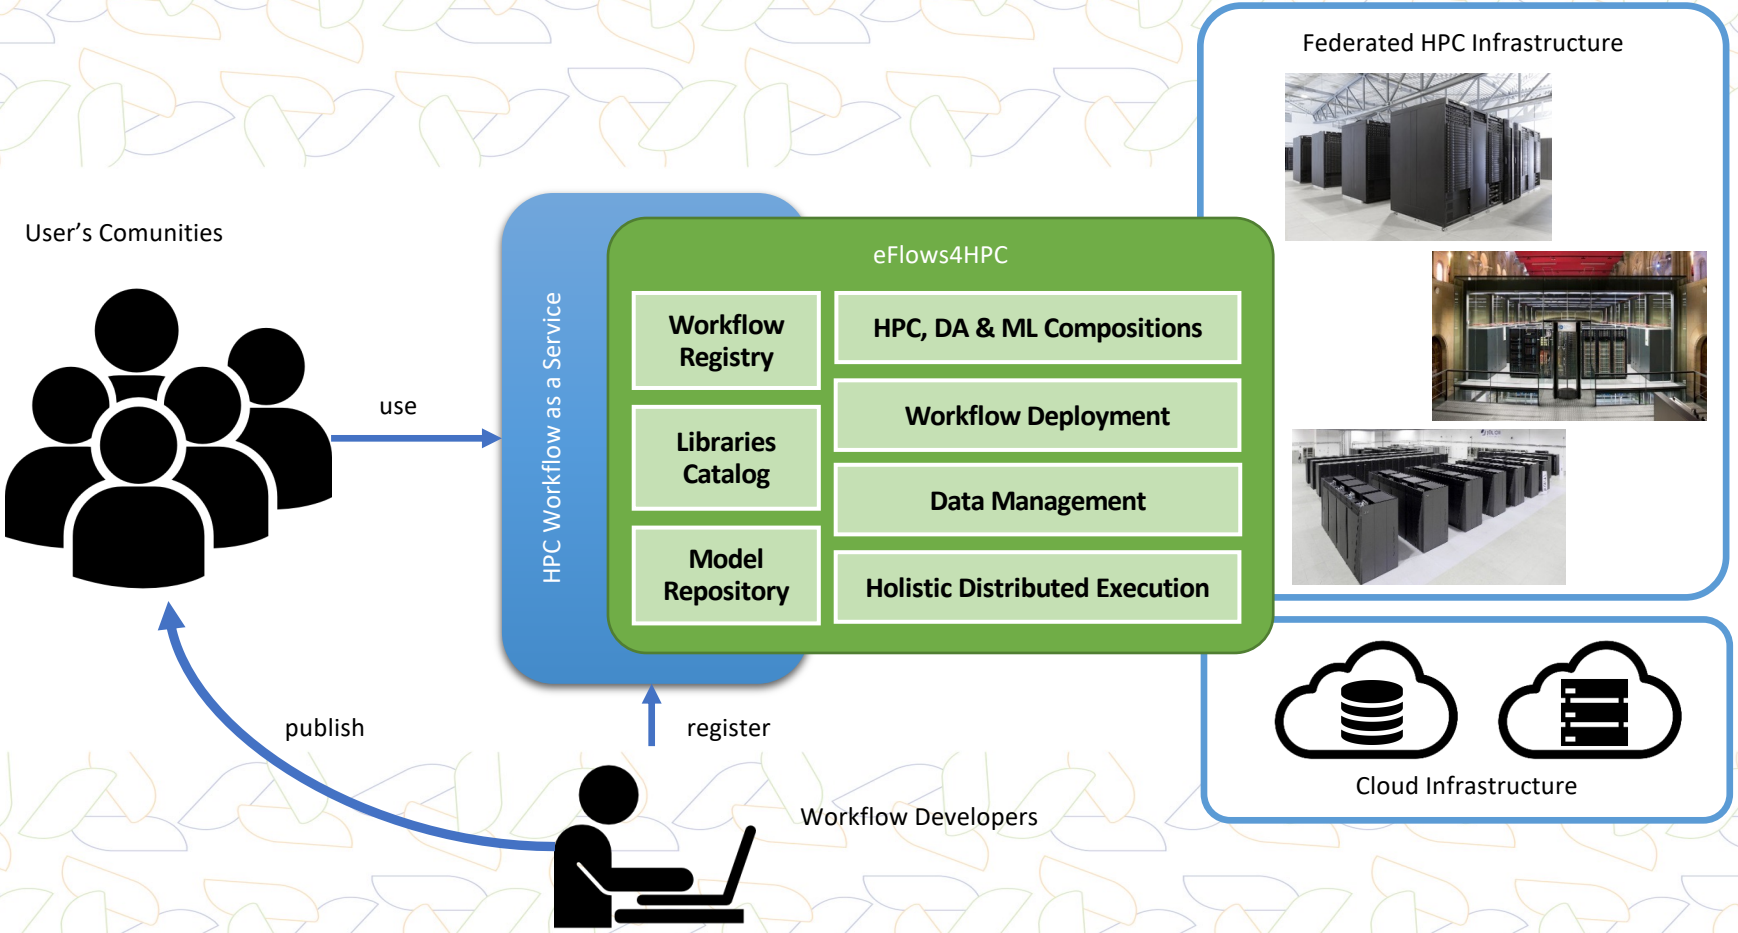
Federated HPC Infrastructure
User's Comunities
use
HPC Workflow as a Service
eFlows4HPC
Workflow Registry
HPC, DA & ML Compositions
Libraries Catalog
Workflow Deployment
Data Management
Model Repository
Holistic Distributed Execution
publish
register
Workflow Developers
Cloud Infrastructure

Development

- Different catalogues, repositories and registries will be source for workflow components
- TOSCA description:
  - Required software and services
  - Deployment and configuration
- PyCOMPSs:
  - Logic of the dynamic workflow
- Data logistics pipelines:
  - Data acquisition, data movement, and data storage
- Workflow description registered in the workflow registry with the HPCWaaS interface
  - Service endpoint available for later invocation of the workflow

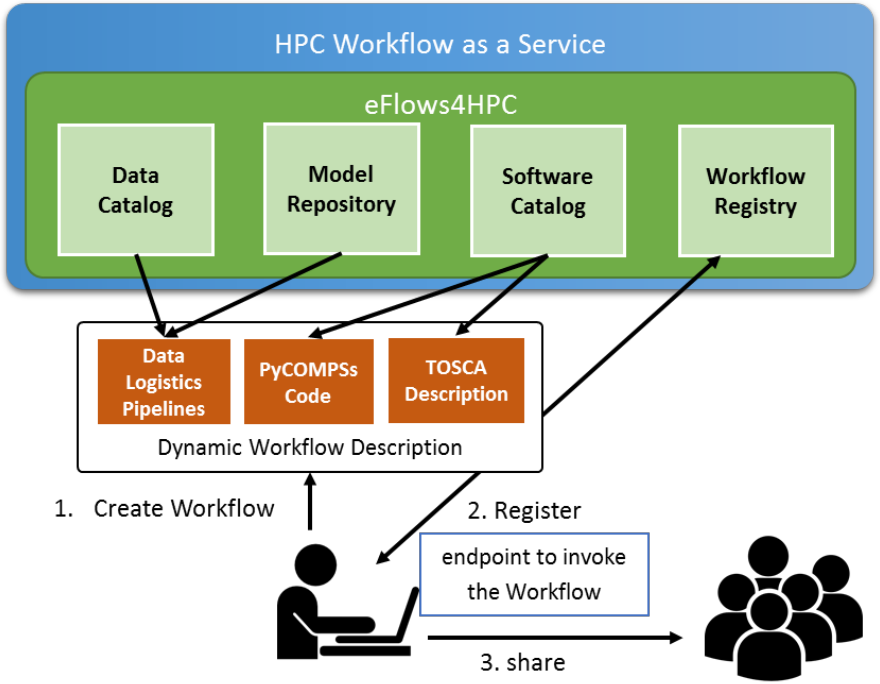

Programming with PyCOMPSs/COMPSs

- Sequential programming, parallel execution
- General purpose programming language + annotations/hints
  - To identify tasks and directionality of data
- Builds a task graph at runtime that express potential concurrency
- Tasks can be sequential, parallel (threaded or MPI)
- Offers a shared memory illusion to applications in a distributed system
  - The application can address larger data storage space: support for Big Data apps
  - Support for persistent storage
- Agnostic of computing platform
  - Enabled by the runtime for clusters, clouds and container managed clusters

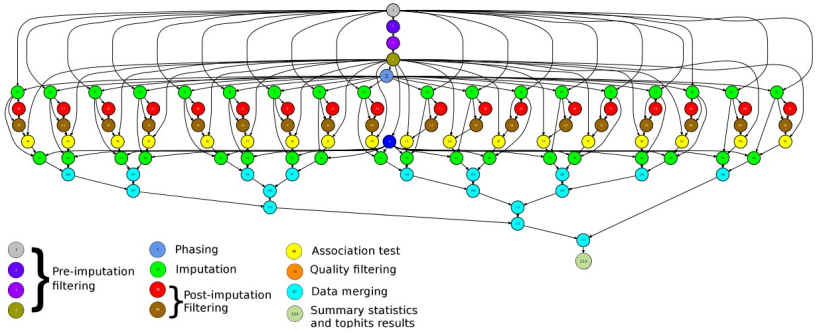

Machine learning libraries and tools

- Dislib
  - Machine learning library parallelized in PyCOMPSs – scikit learn syntax
  - Runs in distributed computing – ds-array
- EDDL
  - Neural networks training and inference
  - Distributed training
  - CPU, GPU, FPGA support
- HeAT
  - Distributed (MPI) n-dimensional tensor library exploiting GPUs for performance
  - NumPy compatibility
  - ML and DL
- FastML
  - Model Training Management focused on HPC deployment
  - Support of TensorFlow, Keras, Pytorch, Scikit-Learn
  - Model tuning through Hyper Parameter Optimisation features
  - Distributed Training through Horovod
- Multiple integration activities planed during the project

Codex AI Suite

FastML

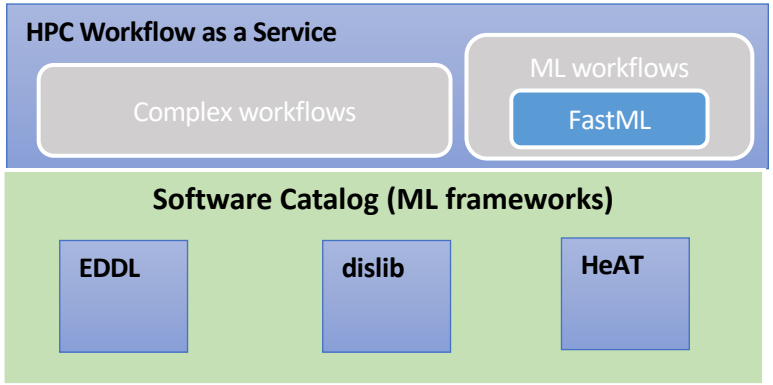

EDDL

HEAT
Helmholtz Analytics Toolkit

dislib

Deployment

- Software, models and services deployment
- Configuration
- Data stage-in

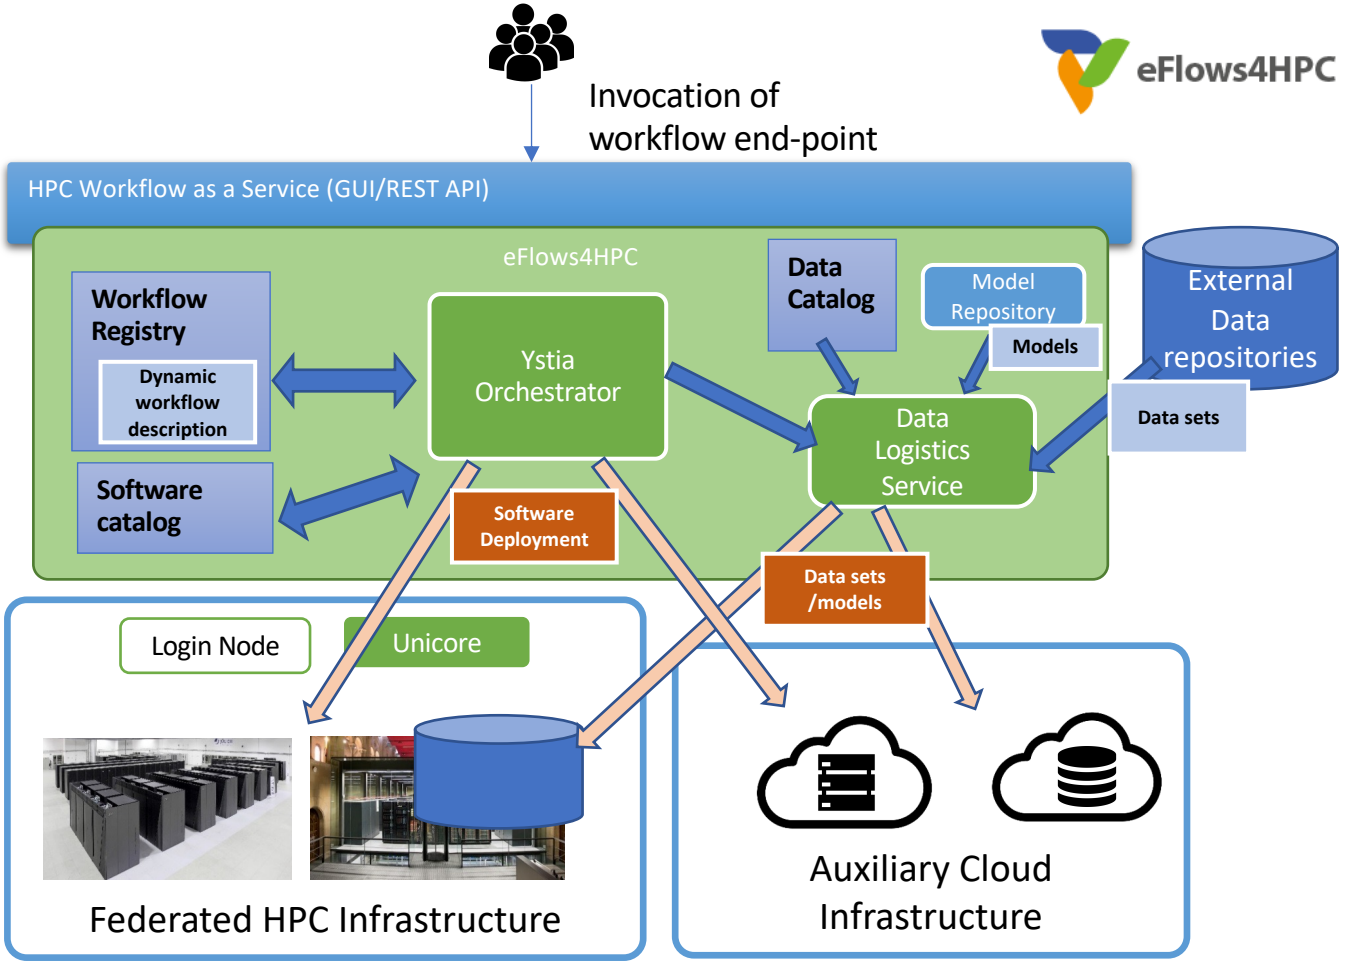

Operation- Workflow Execution

- Support to different types of tasks
  - MPI, ML, DA
- Dynamicity
  - Runtime task-graph
  - Task-level FT
  - Exceptions
- Data management
  - Persistent stg
- Optimized kernels
  - EPI, GPU, FPGA

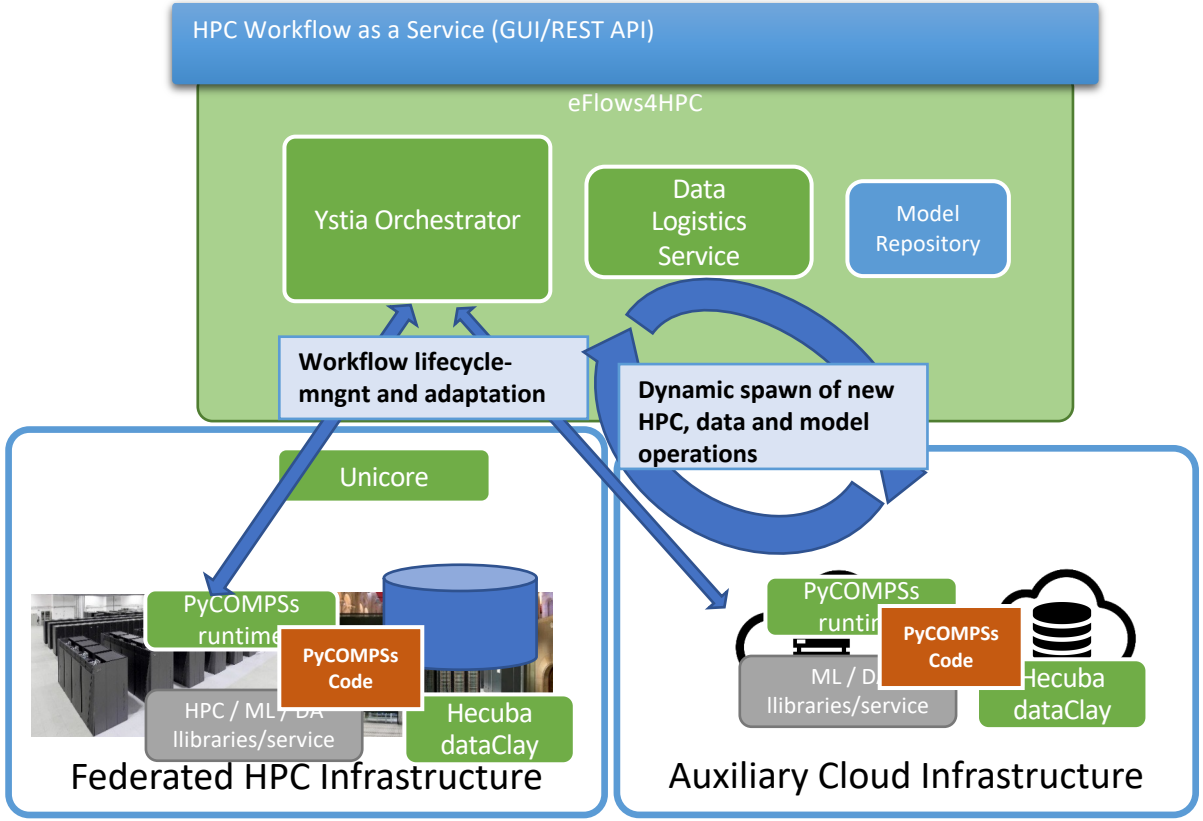

Project WP Structure and project status

- Started January 2021
- Requirement elicitation
  - From pillars' applications
  - From HPC centers
- Software architecture
  - Components interaction, interoperability
  - Design of the architecture
- First results expected end of year

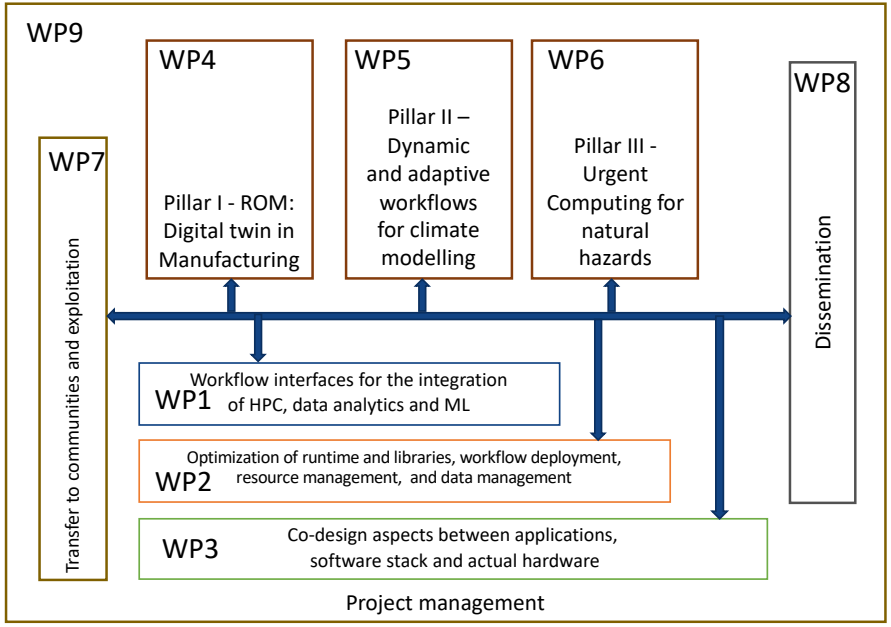

Conclusions

- There is a need for providing tools for the development of complex workflows that include HPC modeling and simulation, artificial intelligence components and big data
- eFlows4HPC aims at providing a software stack that supports the development, deployment and execution of complex and dynamic workflows
- The HPCWaaS aims to provide a functionality similar for FaaS in cloud for complex workflows in HPC

Project partners

eFlows4HPC

eFlows4HPC

Enabling dynamic and Intelligent workflows
in the future EuroHPC ecosystem

www.eFlows4HPC.eu

@eFlows4HPC eFlows4HPC Project

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955558. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Germany, France, Italy, Poland, Switzerland, Norway.