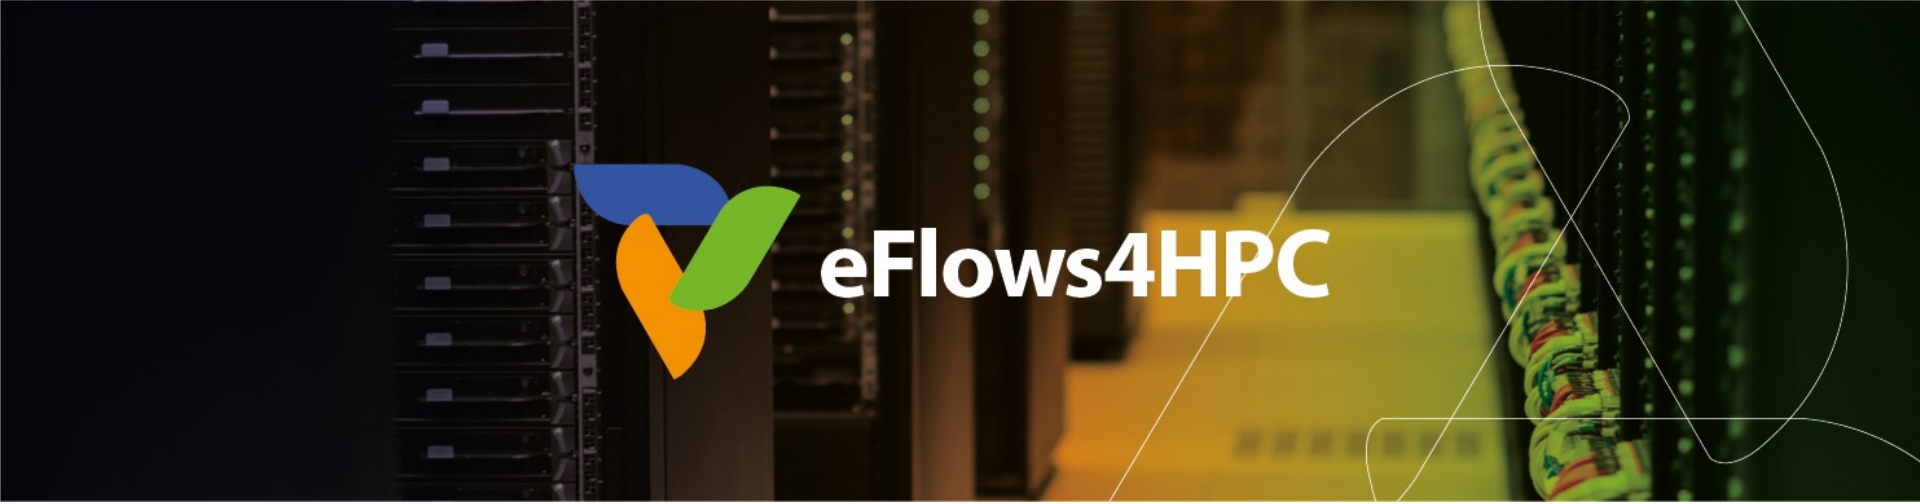

HPC workflows for scientific applications

Rosa M Badia (BSC)

Innovative HPC workflows for industry
Barcelona Supercomputing Center - 10th January 2024

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. MCIN/AEI/10.13039/501100011033 and the European Union NextGenerationEU/PRTR (PCI2021-121957)

Agenda (I)

9:00 - 9:30	Arrival and registration	
9:30 – 10:00	eFlows4HPC project overview	Rosa M Badia (BSC)
10:00 – 10:30	Demo about HPCWaaS	Jorge Ejarque(BSC)
10:30 – 11:00	Coffee break	
11:00 – 11:30	Presentation about provenance with PyCOMPSs	Raül Sirvent (BSC)
11:30 – 12:00	Services for industry and academia from the EuroCC network of Competence Centres in HPC	Oriol Pineda (BSC) Carlos Teijeiro (NCC Netherlands)
12:00 – 12:30	HPC-enabling technologies for decarbonization of the power and transportation sectors: the Center of Excellence in Combustion	Daniel Mira (BSC)
12:30 - 13:10	Distributed SKA science-driven workflows at extreme scales : lessons from SKA precursors/pathfinders and next SKA challenges	Susana Sanchez (IAA) Damien Gratadour (CNRS) Jean Pierre Vilotte (CNRS-INSU)
13:10 - 14:00	Lunch break	

Agenda (II)

14:00 - 14:30	Bridging AI and HPC in the Center of Excellence RAISE	Andreas Lintermann (Juelich)
14:30 – 15:00	Workflow Developments in EXCELLERAT P2 for the European HPC Strategy in Engineering	Gregor Weiß (HLRS)
15:00 – 15:30	Towards a framework to integrating CFD and ML in heterogeneous supercomputers	Oriol Lehmkuhl (BSC)
15:30 – 16:00	Discussion	
16:00 - 16:30	Conclusions	Rosa M. Badia (BSC)

Visitors

- Please **always wear your credentials visible** during your stay at the BSC – Repsol building.
- Visitors cannot be unattended and should be with BSC staff during their stay.
- Remember to **return your credential in the plastic baskets** next to reception when leaving the building.

WIFI

Username:

guest100026

Password:

@aNB1$Om

Logistics: WC and catering

WC

When going out of the auditorium room, please **turn left**, open glass door and you will see the WC on your right front side

Catering

Due to the weather forecast, we will have our breaks indoors:

- **Coffee break**: in front of the WCs, outside the auditorium turning left
- **Lunch break**: foyer Torre Girona

Photos

- For project dissemination purposes, photos will be taken during the event to appear on eFlows4HPC and BSC social media accounts and websites and project reports.

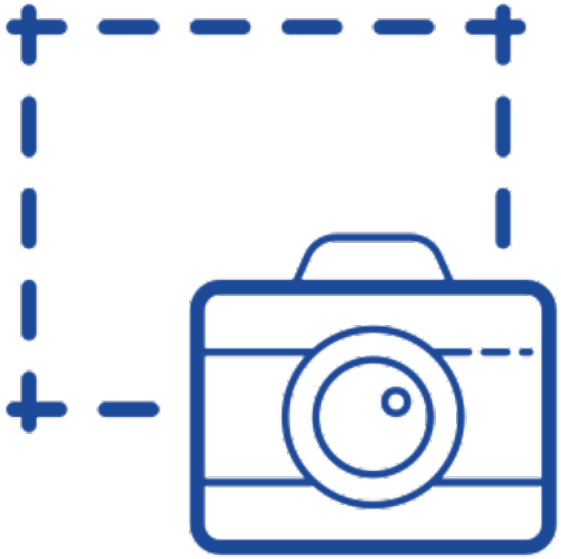

EFLOWS4HPC OVERVIEW

Complex workflows and complex infrastructures

- EuroHPC aims at developing a World Class Supercomputing Ecosystem in Europe
  - Procuring and deploying pre-exascale and petascale systems in Europe
- These systems will be capable of running large and complex applications
- Applications demand the composition of HPC, artificial intelligence and data analytics
- The development, installation, execution and of workflows is manual and error prone:
  - New tools and methodologies are needed

eFlows4HPC in a nutshell

- **Software tools stack that makes easier the development and management of complex workflows:**
  - Combine different aspects
    - HPC, AI, data analytics
  - Reactive and dynamic workflows
    - Autonomous workflow steering
  - Full lifecycle management
    - Not just execution
    - Data logistics and Deployment
- **HPC Workflows as a Service:**
  - Mechanisms to make easier the use and reuse of HPC by wider communities
- **Architectural Optimizations:**
  - Selected HPC – AI Kernels Optimized for GPUs, FPGA, EPI
- **Validation Pillar's**
  - End-user workflows linked to CoEs

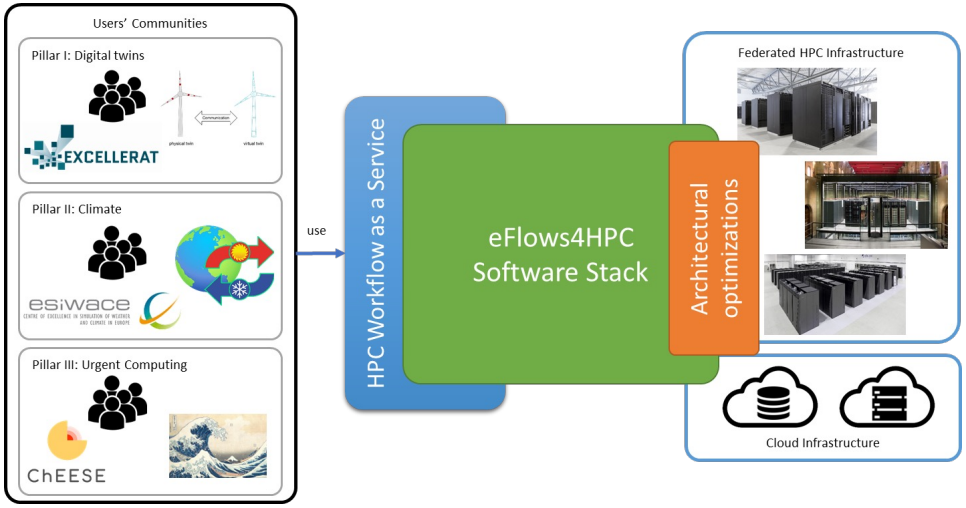

Pillar I: Manufacturing

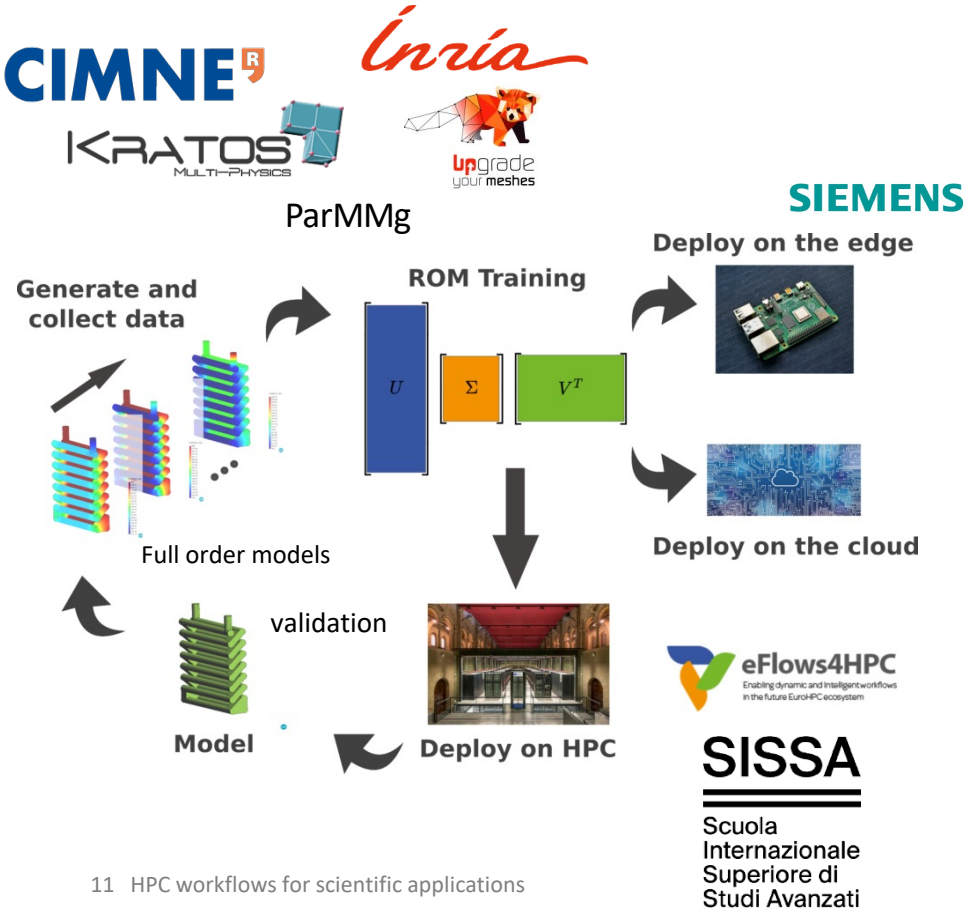

Pillar I focuses on the construction of DigitalTwins for the prototyping of complex manufactured objects:

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

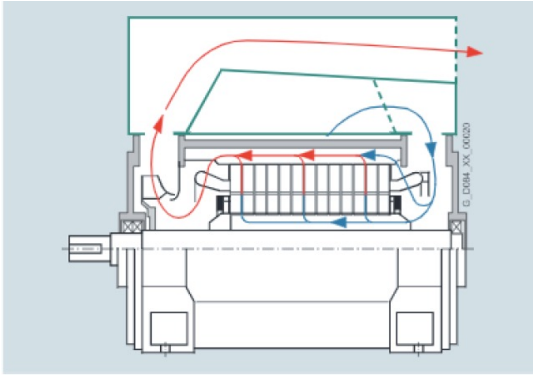

Pillar II: Climate

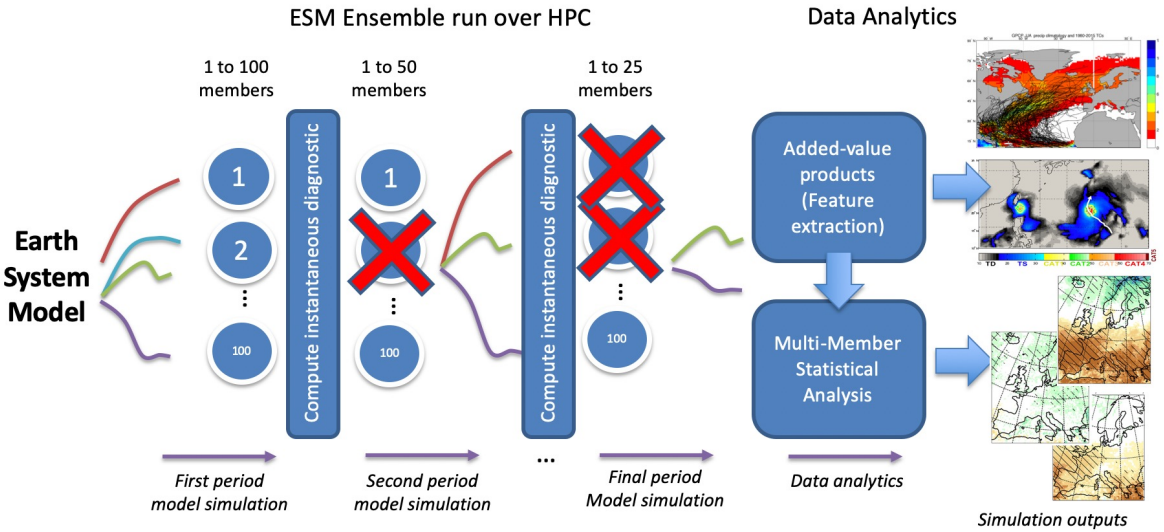

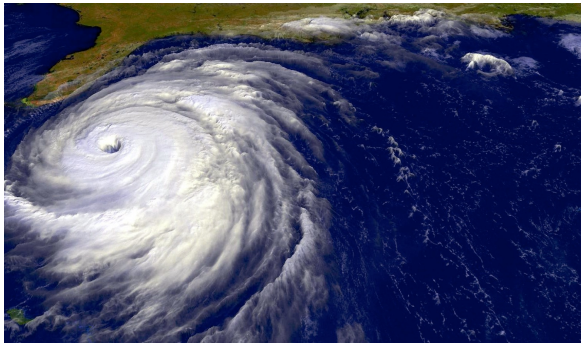

Dynamic (AI-assisted) workflow

HPDA & ML/DL

- Perform climate predictions: temperature, precipitation or wind speed
- AI-assisted pruning of the ESM workflow
- Study of Tropical Cyclones (TC) in the North Pacific, with in-situ analytics

Pillar III: Urgent computing for natural hazards

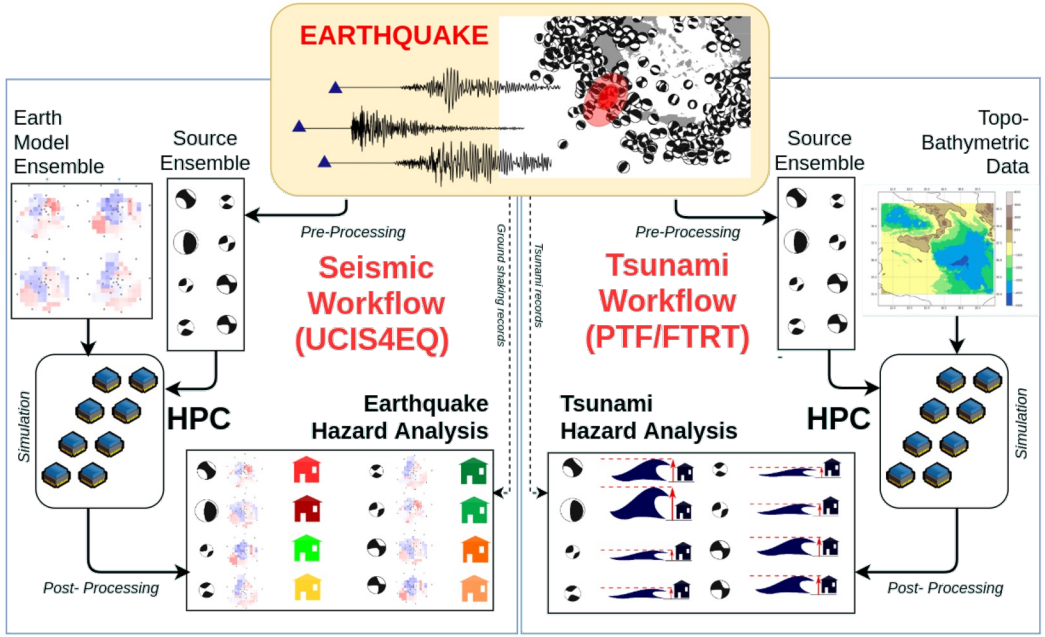

Pillar III explores the modelling of natural catastrophes:

- Earthquakes and their associated tsunamis shortly after such an event is recorded
- Use of AI to estimate intensity maps
- Use of DA and AI tools to enhance event diagnostics
- Areas: Mediterranean basin, Mexico, Iceland and Chile

Tsunami-HySEA GPU-based code

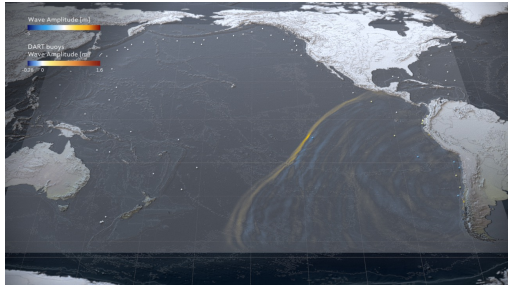

10/01/2024

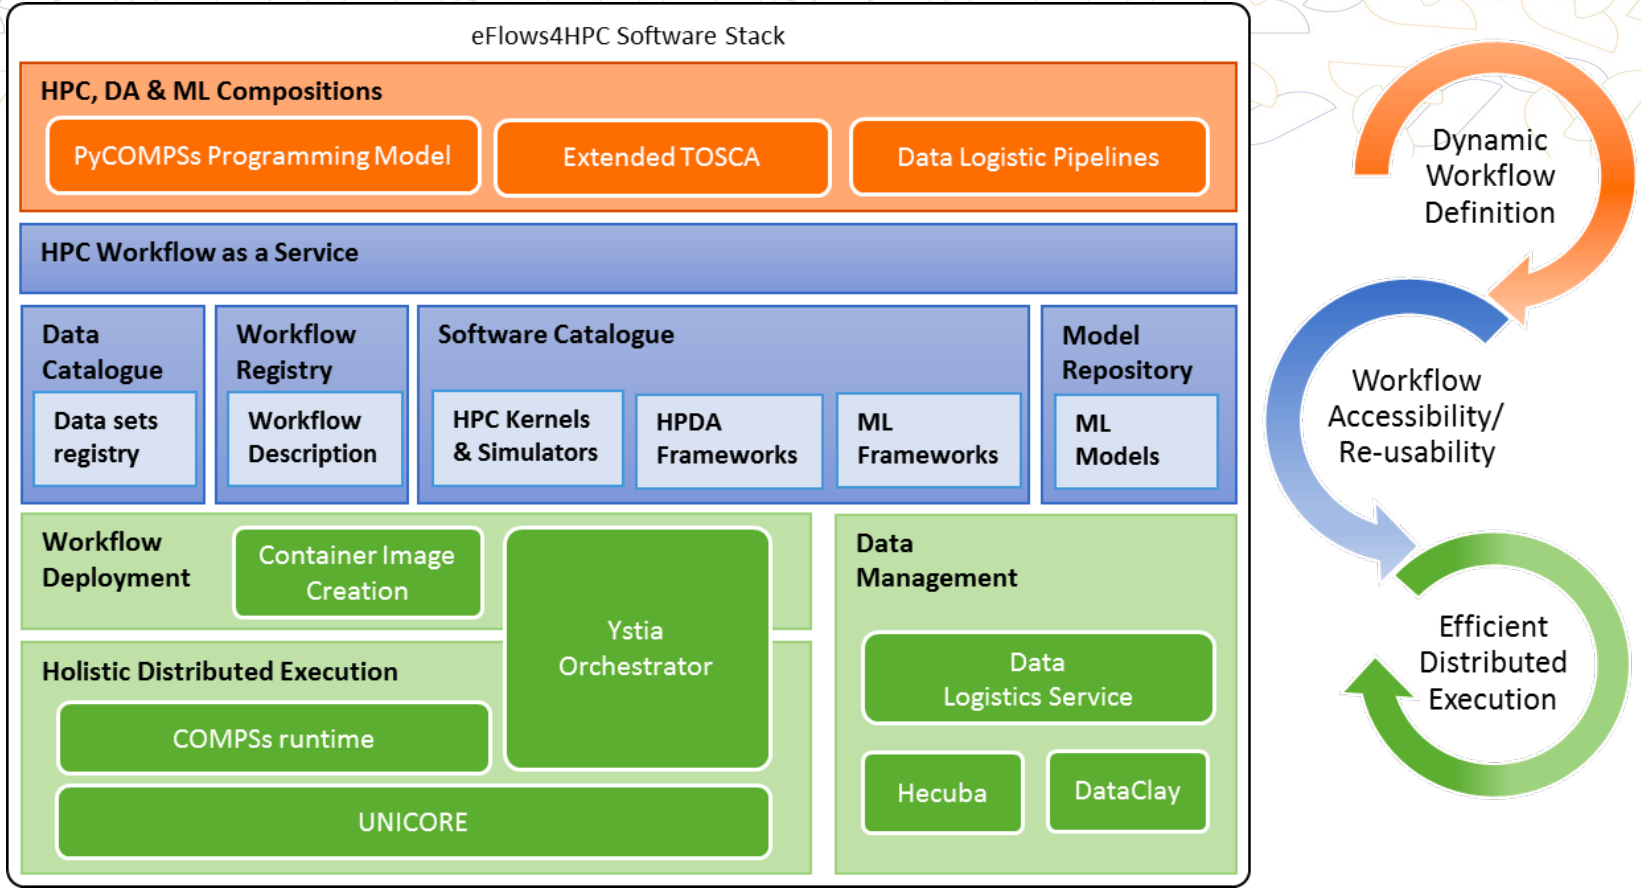
eFlows4HPC Software Stack
HPC, DA & ML Compositions
PyCOMPSs Programming Model
Extended TOSCA
Data Logistic Pipelines
HPC Workflow as a Service
Data Catalogue
Data sets registry
Workflow Registry
Workflow Description
Software Catalogue
HPC Kernels & Simulators
HPDA Frameworks
ML Frameworks
Model Repository
ML Models
Workflow Deployment
Container Image Creation
Ystia Orchestrator
Holistic Distributed Execution
COMPSs runtime
UNICORE
Data Management
Data Logistics Service
Hecuba
DataClay
Dynamic Workflow Definition
Workflow Accessibility/ Re-usability
Efficient Distributed Execution

Software stack deployment

Gateway services

- Components deployed outside the HPC infrastructure.
- Managing external interactions and workflow lifecycle

HPC and runtime Components

- Deployed inside the HPC infrastructure to manage the workflow execution

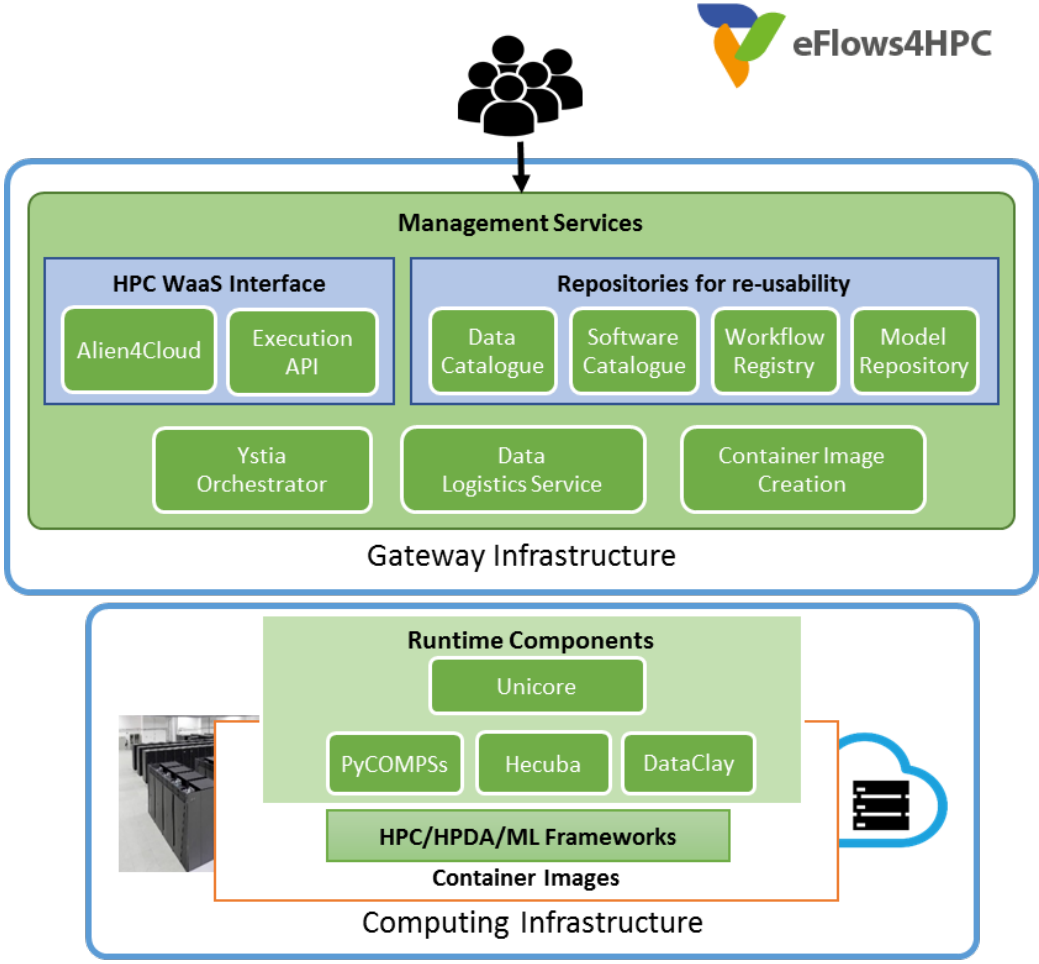

Workflow development overview

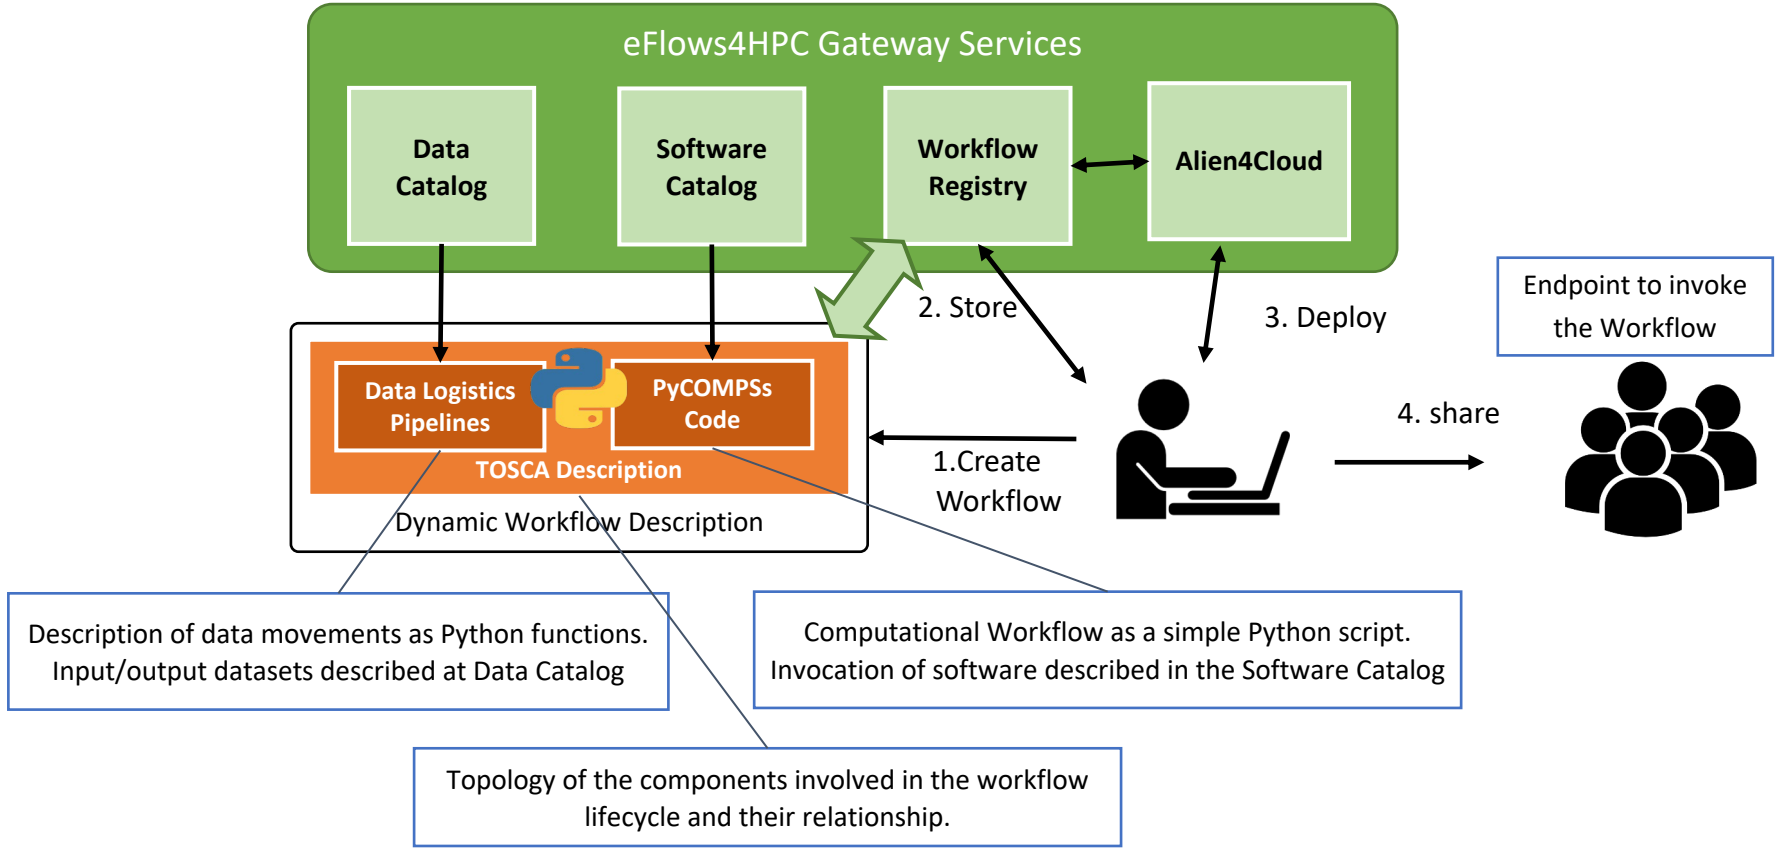

Main element: Workflows in PyCOMPSs

- 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 and parallel (threaded or MPI)
- Offers to applications the illusion of a shared memory in a distributed system
  - The application can address larger data than storage space: support for Big Data apps
- Agnostic of computing platform
  - Enabled by the runtime for clusters, clouds and container managed clusters

```
@task(c=INOUT)
def multiply(a, b, c):
    c += a*b
```

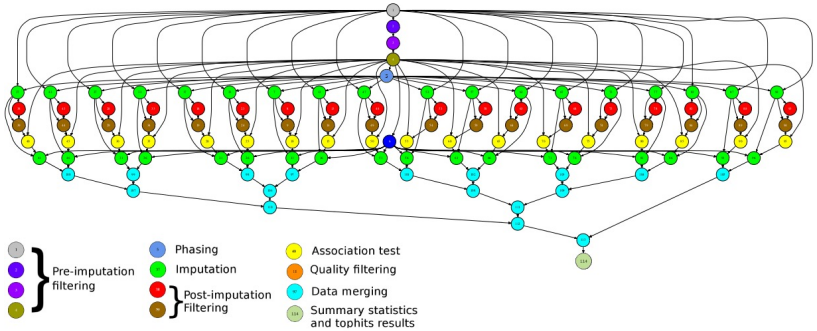

PyCOMPSs features and runtime

- Support for tasks' constraints – support for heterogeneous infrastructure
- Support for tasks' faults and tasks' exceptions
  - Enlarges the dynamicity of the type of workflows that we support
- Streamed data
  - ... and many others
- Runtime deployed as a distributed master-worker
- All data scheduling decisions and data transfers are performed by the runtime
- Support for elasticity

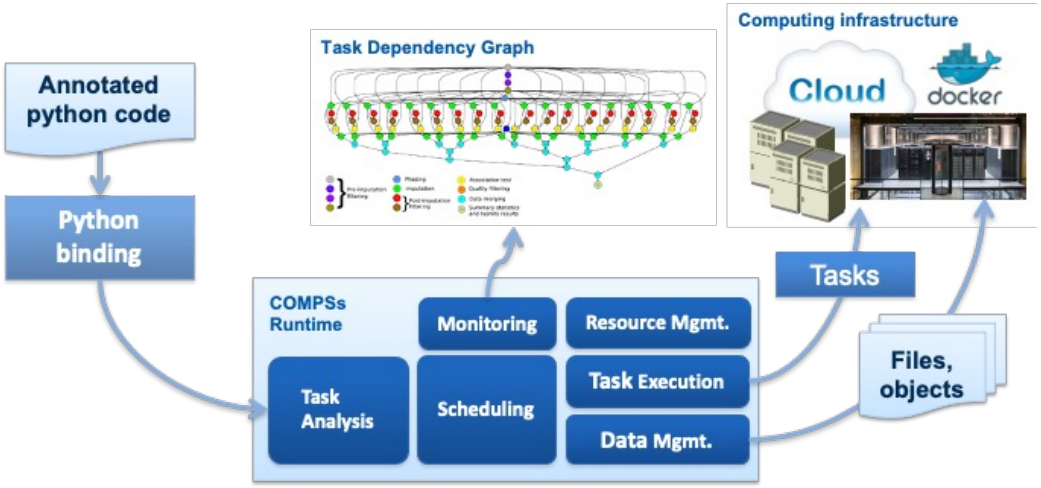

Addressing challenge 2 and 3

Interfaces to integrate HPC/DA/ML

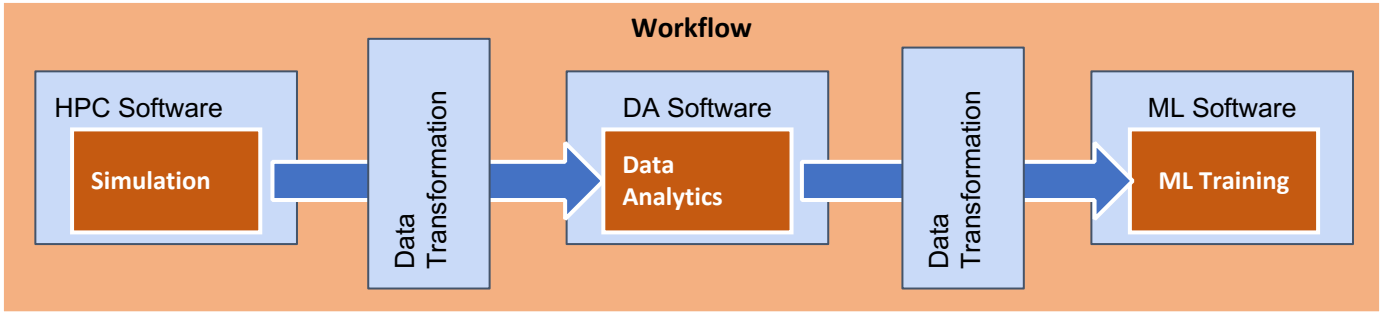

- Goal:
  - Reduce the required glue code to invoke multiple complex software steps
  - Developer can focus in the functionality, not in the integration
  - Enables reusability
- Two paradigms:
  - Software invocation
  - Data transformations

```
#workflow steps defined as tasks
@data_transformation (input_data, transformation description)
@software (invocation description)
def data_analytics (input_data, result):
    pass

#workflow body
simulation (input_cfg, sim_out)
data_analytics (sim_out, analysis_result)
ml_training (analysis_result, ml_model)
```


Addressing challenge 1 and 3

Data Catalogue and Data Logistics Service

Data Catalogue:

- Lists datasets used and created by the workflow according to FAIR principles
- Provides metadata to make data movement pipelines more generic

Data Pipelines:

- Formalization of data movements for transparency and reusability
- Stage-in/out, image transfer

Data Logistics Services (DLS):

- Performs the execution of data pipelines at deployment and execution time

Production Ready Services:

- https://datacatalogue.eflows4hpc.eu
- https://datalogistics.eflows4hpc.eu/

Data pipeline

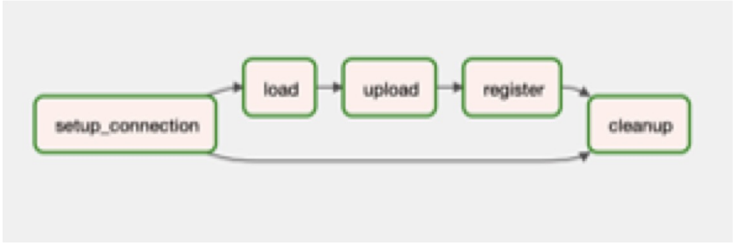

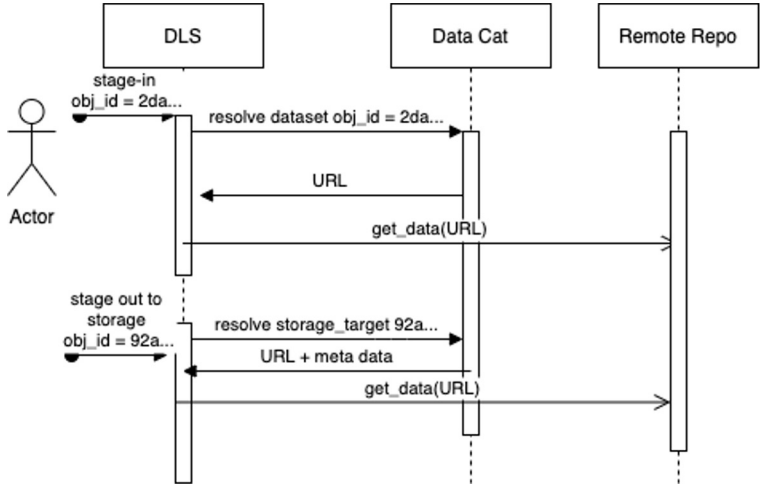

Top-level workflows approach

- Requires a description for workflow lifecycle management
  - TOSCA:
    - Model to describe cloud application topologies and its lifecycle orchestration
- Interface for deploying and running the workflows
  - HPC Workflows as a Service (HPCWaaS):
    - Deployment based on containers
    - Execution: HPCWaaS API

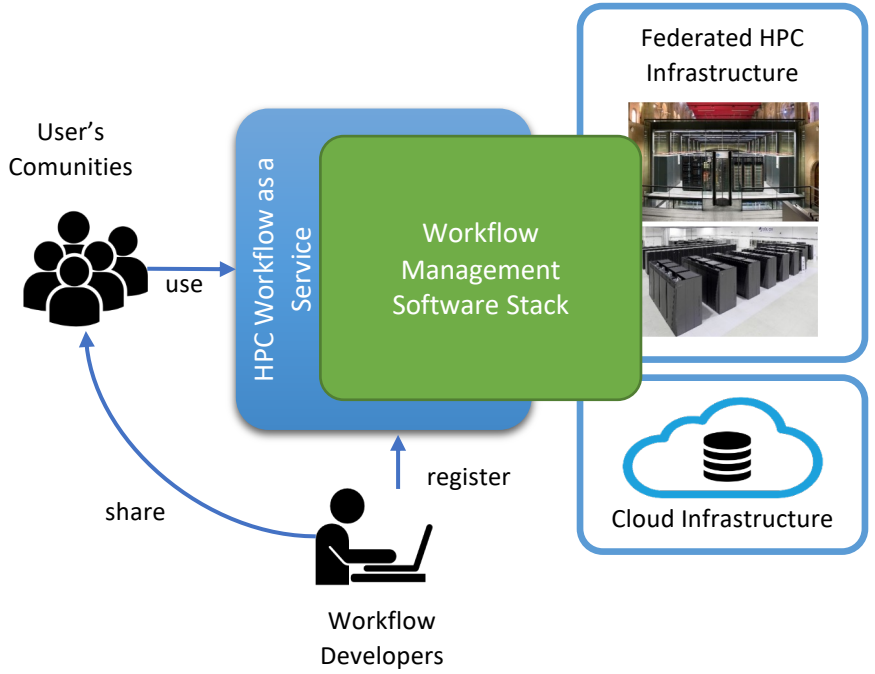

TOSCA Modelization

Topology of the different components involved in the Workflow lifecycle management

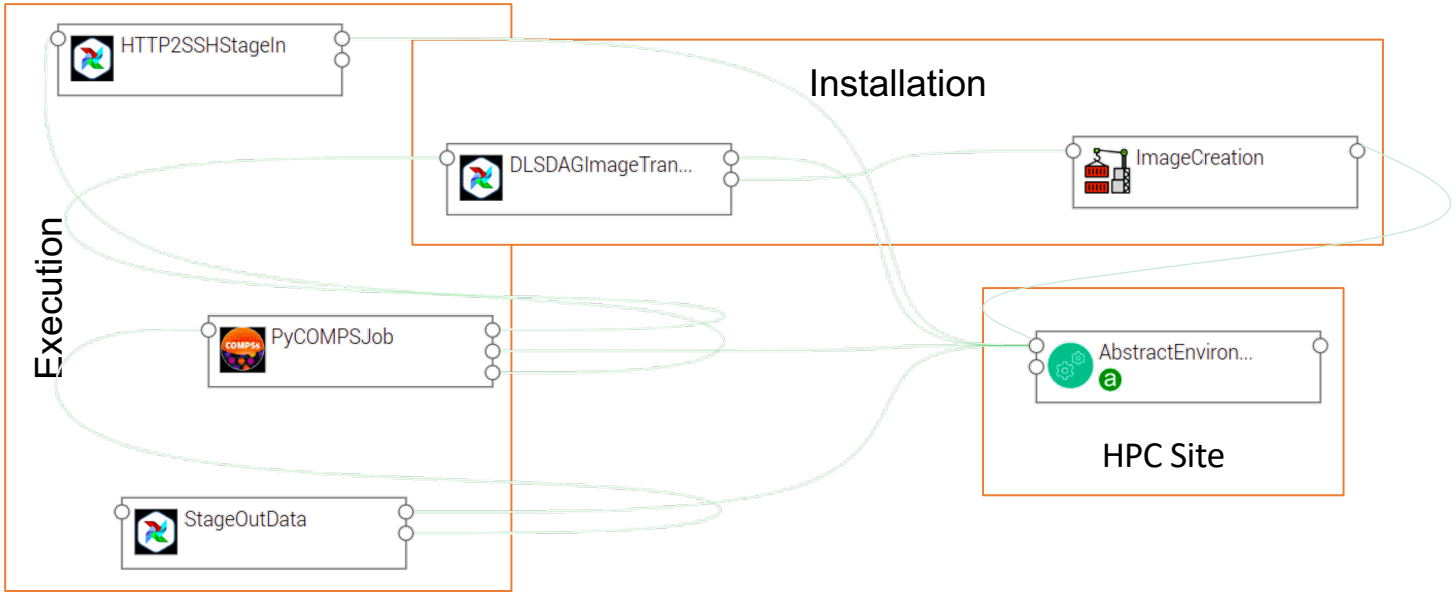

HPC Ready Containers

- Methodology to allow the creation containers for specific HPC system

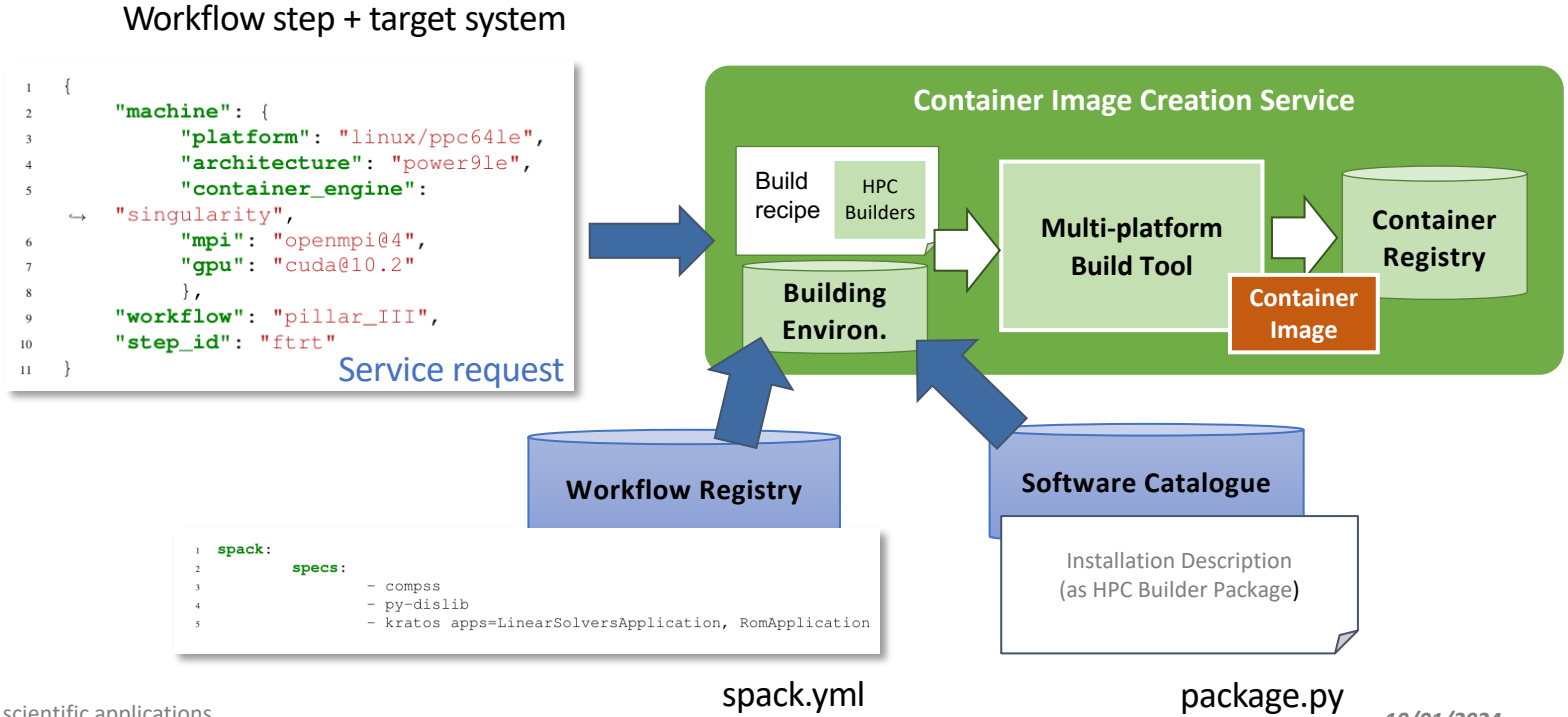

Project main achievements

https://eflows4hpc.eu/software/

- Requirements and software architecture. Reviewed at beginning of second iteration
- Definition and implementation of abstractions to support the integration of different stack components
- Development of a step-by-step example.
- Design and implementation of the HPCWaaS API
- Design and implementation of project services: Data Catalogue, Workflow Registry, Software catalogue
- Design and implementation of two versions of Pillars' workflows.
- Two releases of project software and documentation available
- Set of internal trainings about software stack components and HPCWaaS, including an ICS-HPC tutorial
- Good visibility: articles, keynote presentations, media

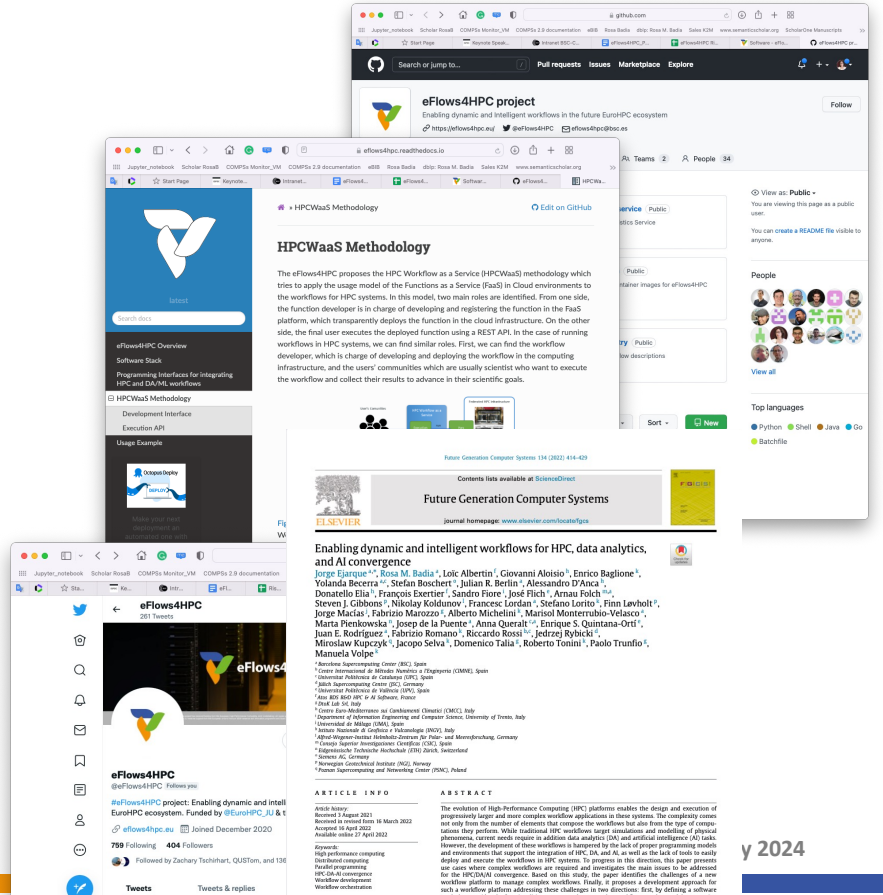

https://eflows4hpc.eu/software/

Adoption of eFlows4HPC tools in other projects

DT-GEO: prototype for a digital twin on geophysical extremes

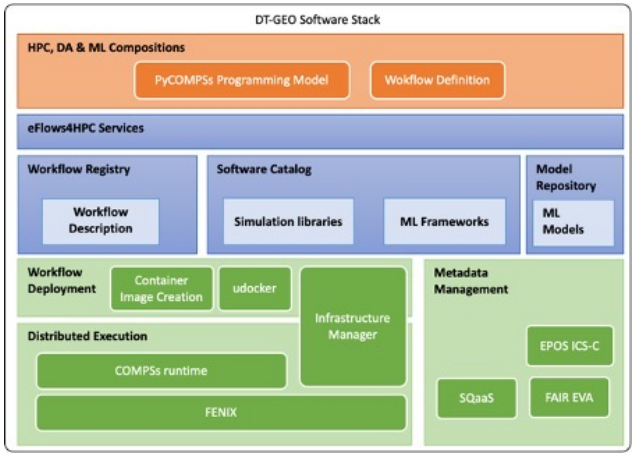

DT-GEO

CAELESTIS: simulation ecosystem of next generation aircraft structures

CAELESTIS

CEEC: Center of Excellence for Exascale CFD

CEEC

Project partners

eFlows4HPC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CIMNE

JÜLICH
Forschungszentrum

UNIVERSITAT POLITÈCNICA DE VALÈNCIA

Bull
atos technologies

DtoK LAB

cmcc
Centro Euro-Mediterraneo sui Cambiamenti Climatici

Inria

SISSA
Scuola Internazionale Superiore di Studi Avanzati

INSTYTUT CHEMII BIOORGANICZNEJ
Polskiej Akademii Nauk

UNIVERSIDAD DE MÁLAGA

AWI
ALFRED-WEGENER-INSTITUT
HELMHOLTZ-ZENTRUM FÜR POLAR- UND MEERESFORSCHUNG

ISTITUTO NAZIONALE DI GEOFISICA E VULCANOLOGIA

ETH zürich

SIEMENS

NGI

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.