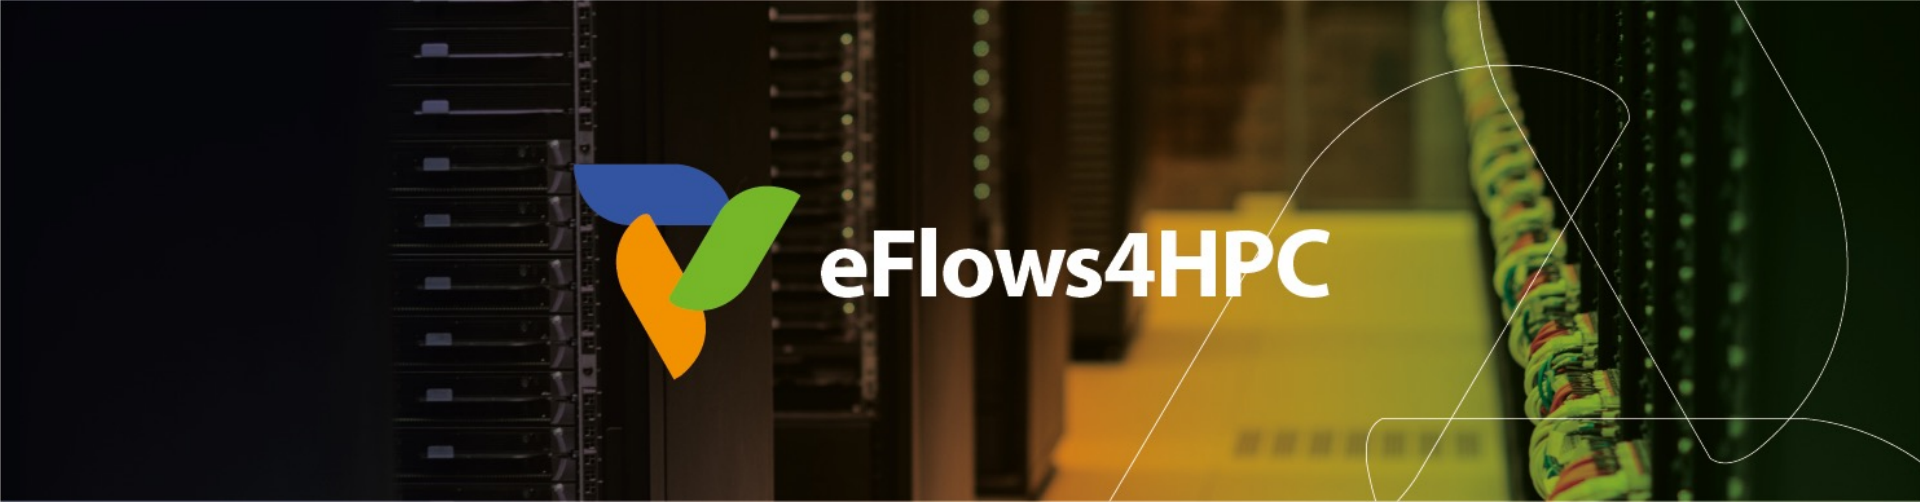

Easy Development and Execution of Workflows with eFlows4HPC

Rosa M Badia – BSC

21ST INTERNATIONAL SYMPOSIUM ON PARALLEL & DISTRIBUTED COMPUTING,
JULY 11-13, 2022 BASEL, SWITZERLAND

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

- 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

EuroHPC systems

eFlows4HPC

		Status	Country	Peak performance	Architecture
Pre-Exascale	LUMI	Operational	Finland	552 petaflops	64-core AMD EPYC™ CPUs + AMD Instinct™ GPU
	Leonardo	Under construction	Italy	322.6 petaflops	Intel Ice-Lake, Intel Sapphire Rapids + NVIDIA Ampere
	MareNostrum 5	Contract signed	Spain	314 petaflops	Bull Sequana XH3000 and Lenovo ThinkSystem
Petascale	Meluxina	Operational	Luxembourg	15 petaflops	AMD [illegible] NVIDIA Ampere A100
	Vega	Operational	Sl[illegible]		[illegible]12 + Nvidia A100
	Karolina	[illegible]	[illegible]		[illegible]Nvidia A100
	Discoverer	C[illegible]	Bulgaria	6 petaflops	AMD EPYC
	Deucalion	Under construction	Portugal	10 petaflops	A64FX, AMD EPYC +Nvidia Ampere

First European Exascale Supercomputer announced to be installed in Jülich

https://eurohpc-ju.europa.eu/about/our-supercomputers_en

Main objectives

- Software tools stack that make it easier the development of workflows
  - HPC, AI + data analytics
  - Reactive and dynamic workflows
  - Efficient resource management
- HPC Workflows as a Service:
  - Mechanisms to make it easier the use and reuse of HPC by wider communities

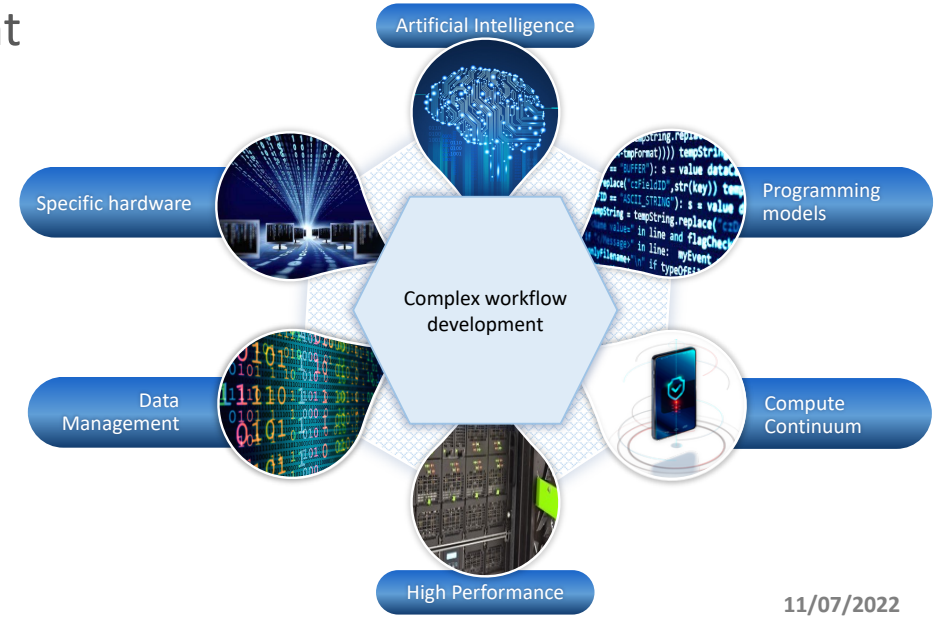

Outline

- Project architecture
  - Software stack
  - HPC Workflows as a Service
- Pillar applications

PROJECT ARCHITECTURE

Users' Communities

Pillar I: Digital twins

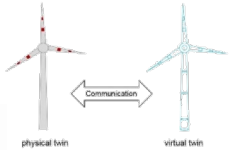

EXCELLERAT

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

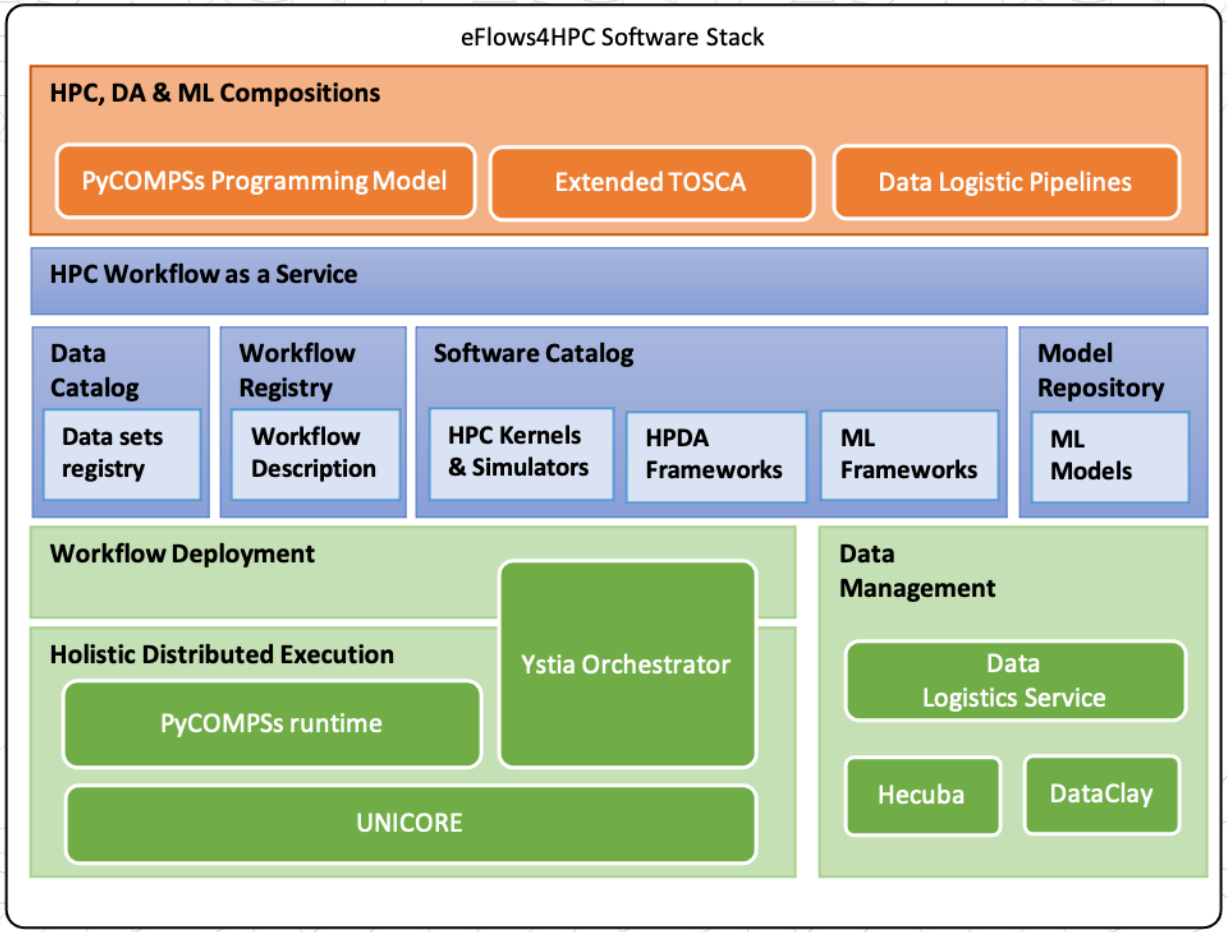
eFlows4HPC Software Stack
HPC, DA & ML Compositions
PyCOMPSs Programming Model
Extended TOSCA
Data Logistic Pipelines
HPC Workflow as a Service
Data Catalog
Data sets registry
Workflow Registry
Workflow Description
Software Catalog
HPC Kernels & Simulators
HPDA Frameworks
ML Frameworks
Model Repository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

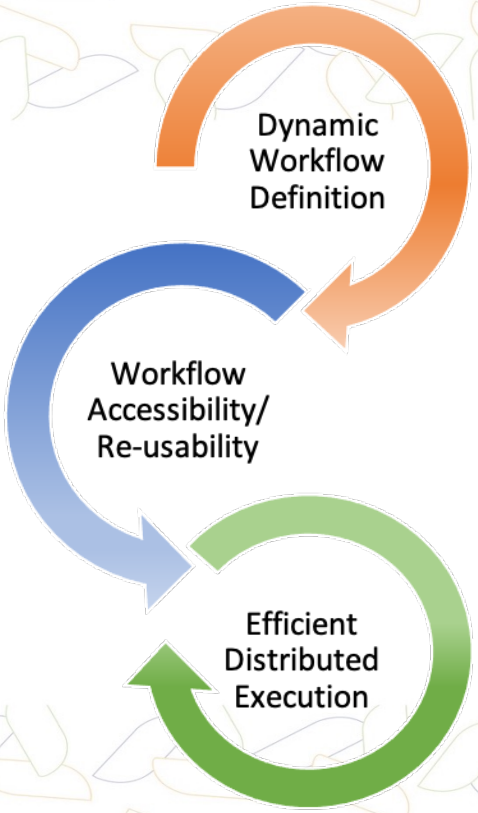
Dynamic Workflow Definition
Workflow Accessibility/ Re-usability
Efficient Distributed Execution

eFlows4HPC software stack and HPCWaaS

- Gateway services
  - Components deployed outside the computing infrastructure.
  - Managing external interactions and workflow lifecycle
- Runtime Components
  - Deployed inside the computing infrastructure to manage the workflow execution

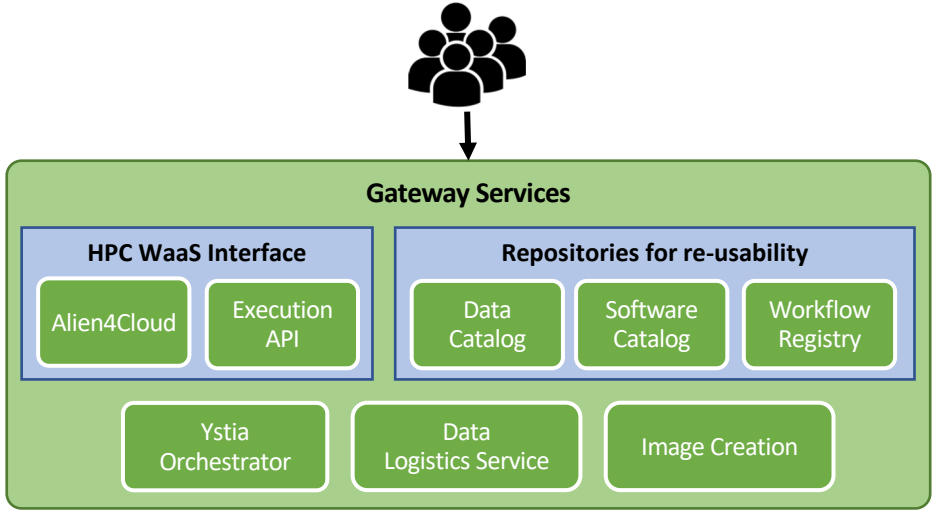

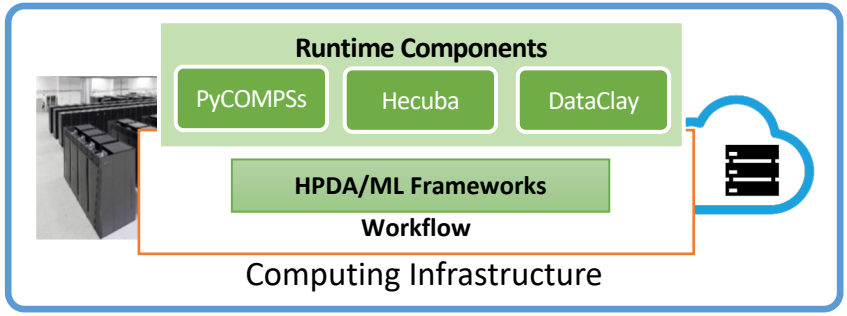

Workflow development overview

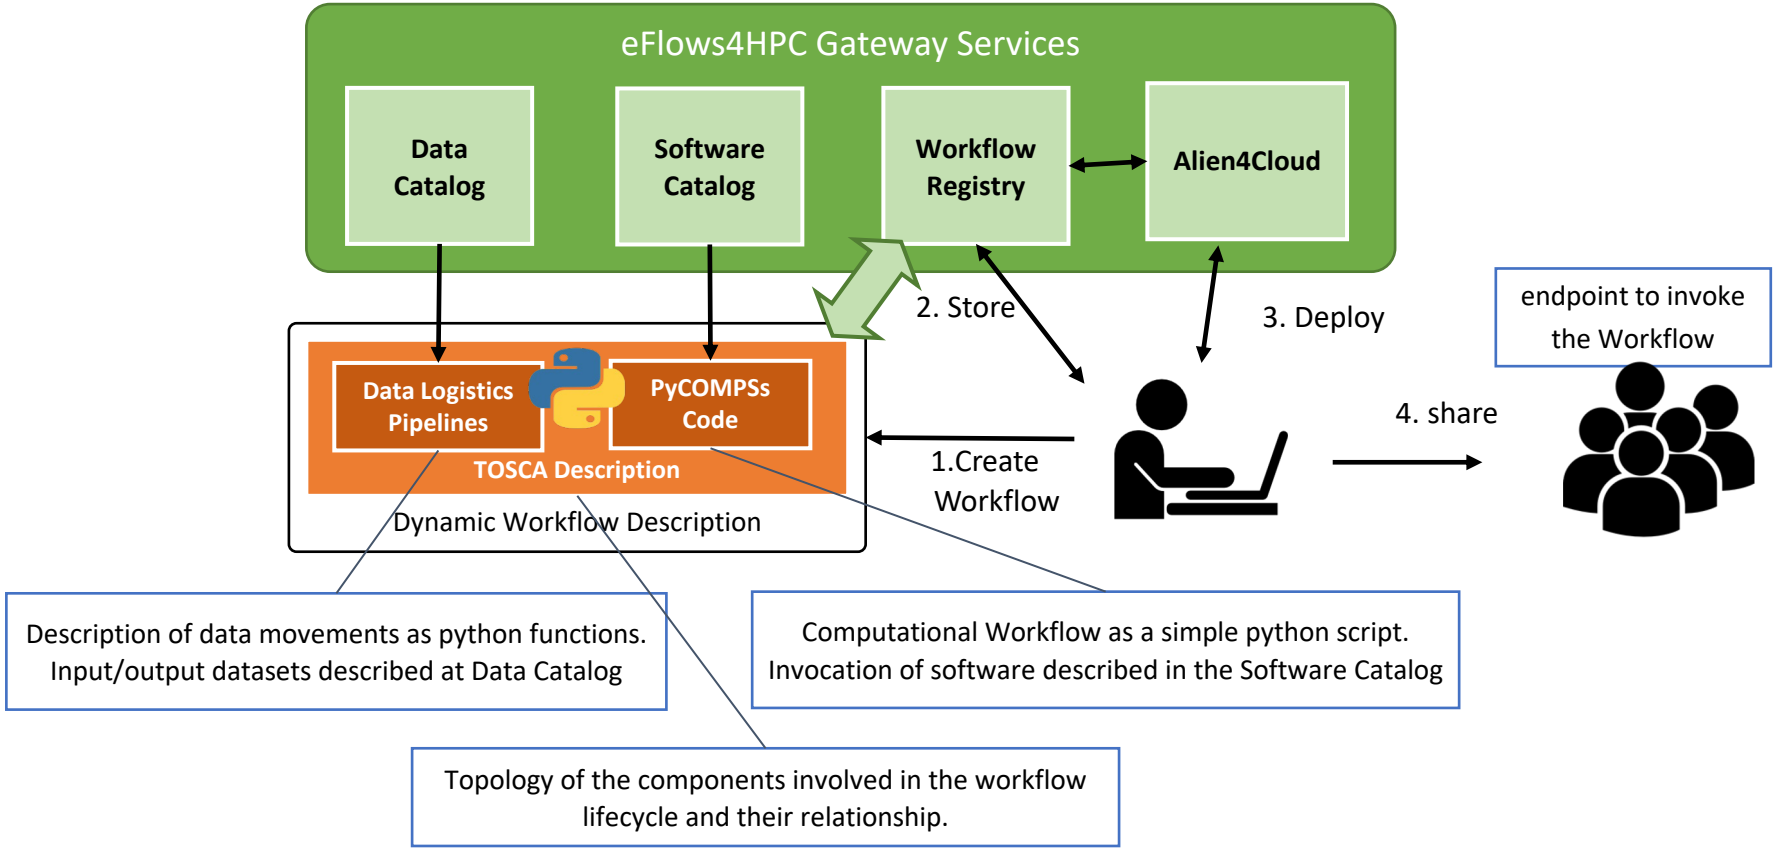

Main element: Workflows in PyCOMPSs

- 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, parallel (threaded or MPI)
- Offers to applications a shared memory illusion in a distributed system
  - The application can address larger data than 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

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

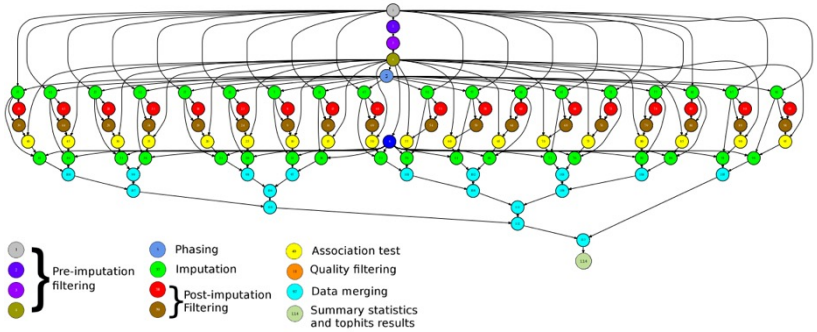

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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

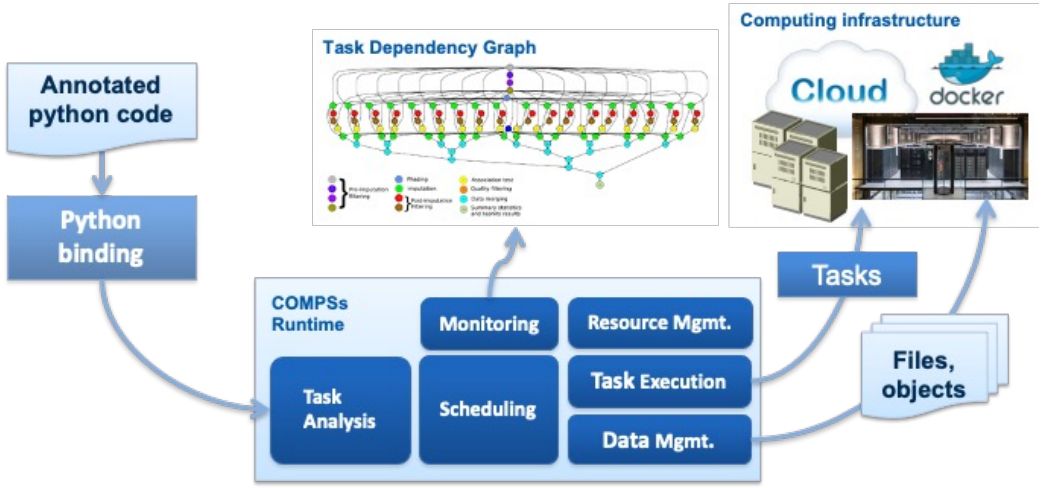

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Interfaces to integrate HPC/DA/ML

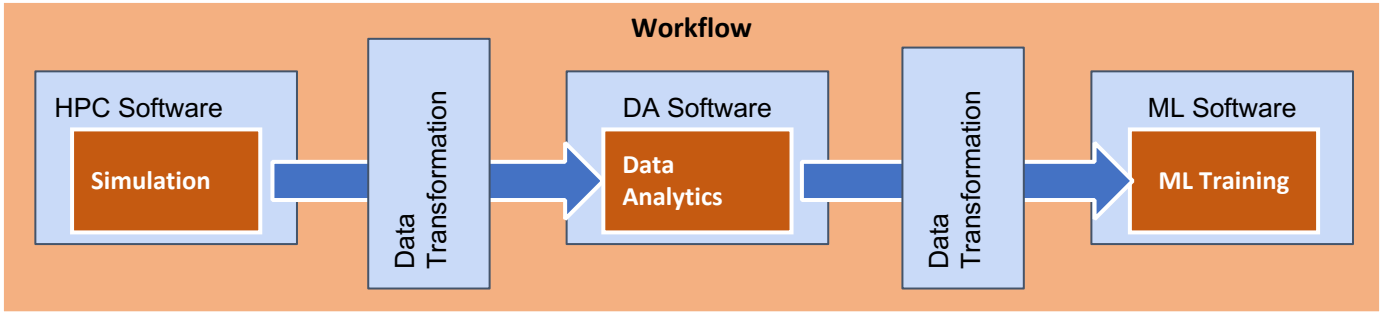

- Goal:
  - Reduce glue code
  - Developer focuses in the functionality, not in the integration
  - Reusability
- Two paradigms:
  - Software integration
  - Data transformations

```
@data_tranformation(input_data, transformation description)
@software(invocation description)
def data_analytics (input_data, result):
    pass

#Worflow

simulation(input_cfg, sim_out)
data_analytics(sim_out, analysis_result)
ml_training(analysis_result, ml_model)
```

Software Invocation description

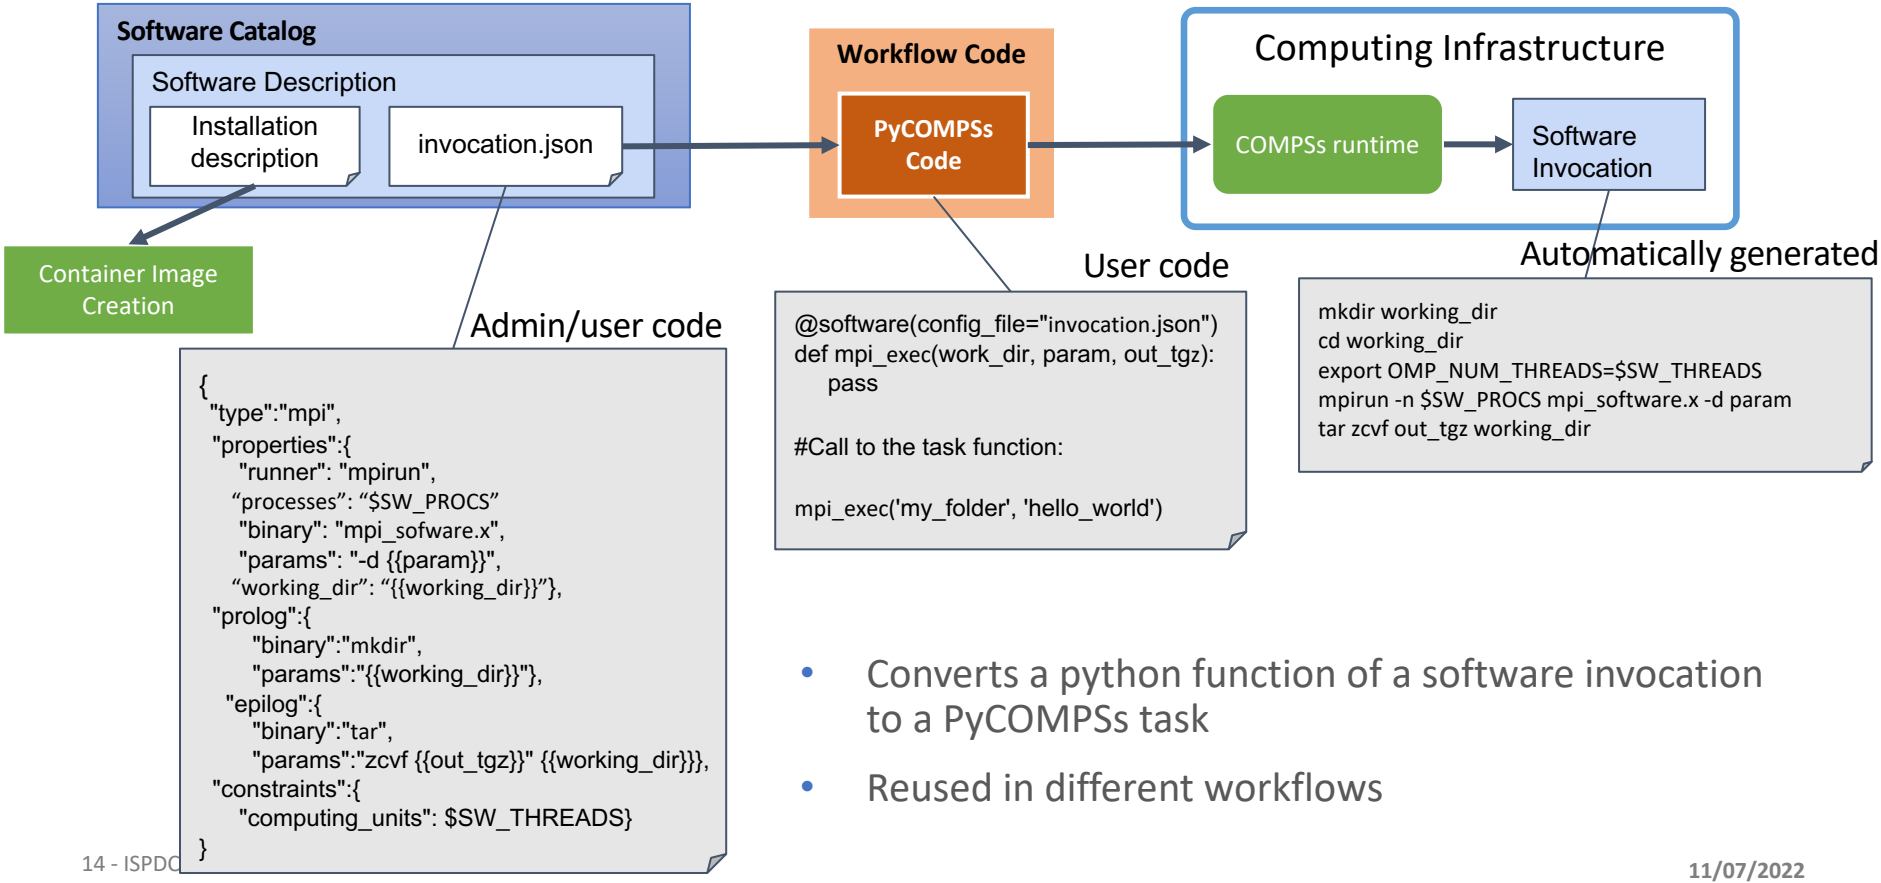

- Converts a python function of a software invocation to a PyCOMPSs task
- Reused in different workflows

Data Catalogue and Data Logistics Service

eFlows4HPC

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 to fuel Pillars' computations

Data pipeline

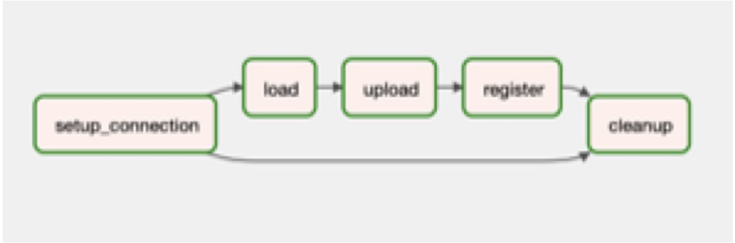

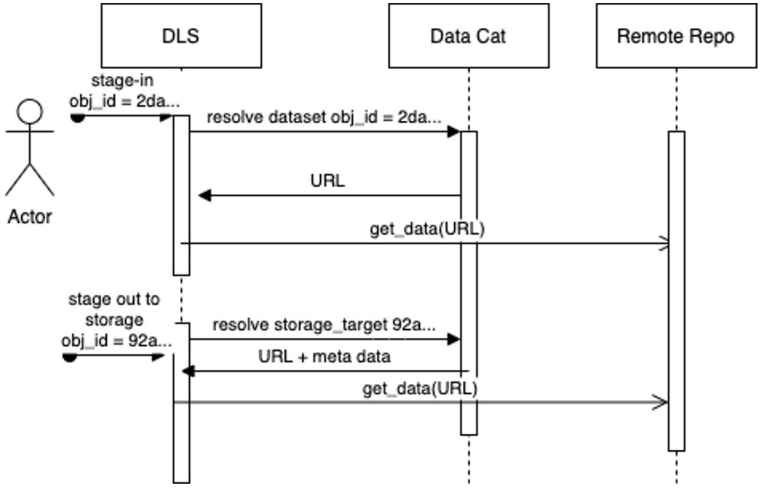

TOSCA Modelization

Topology of the different components involved in the Workflow lifecycle

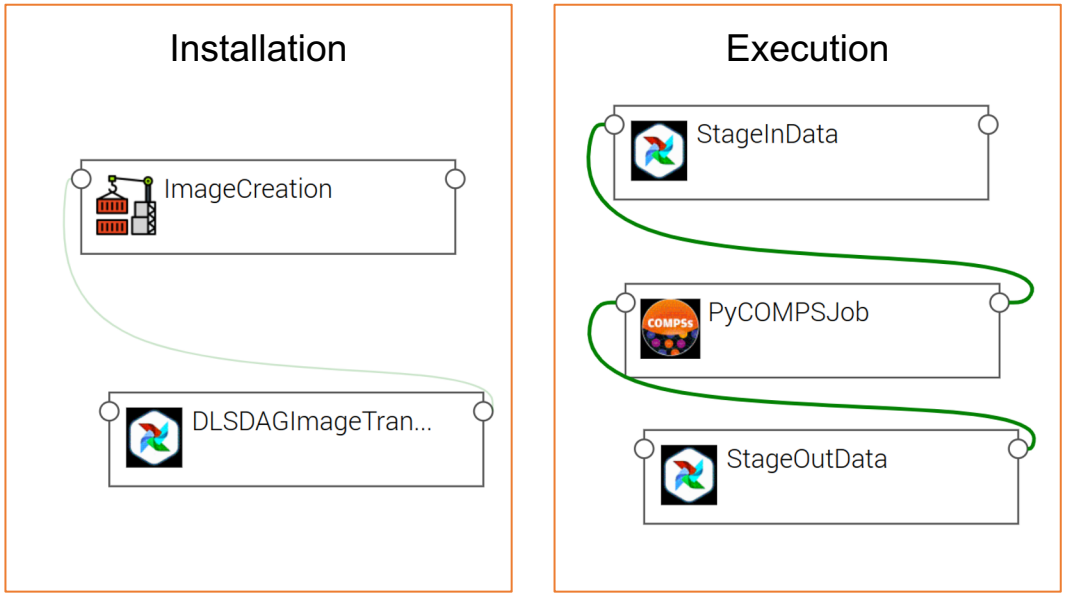

Deployment

- Invocation of a workflow
- Deployment orchestrated by Ystia Orchestrator (Yorc)
- Workflow retrieved from registry
- Data Logistic Service – data stage-in and stage-out
  - Periodical transfers of data outside HPC systems
- Deployment of workflow components in the computing infrastructures
  - HPC containers built with easybuild/Spack

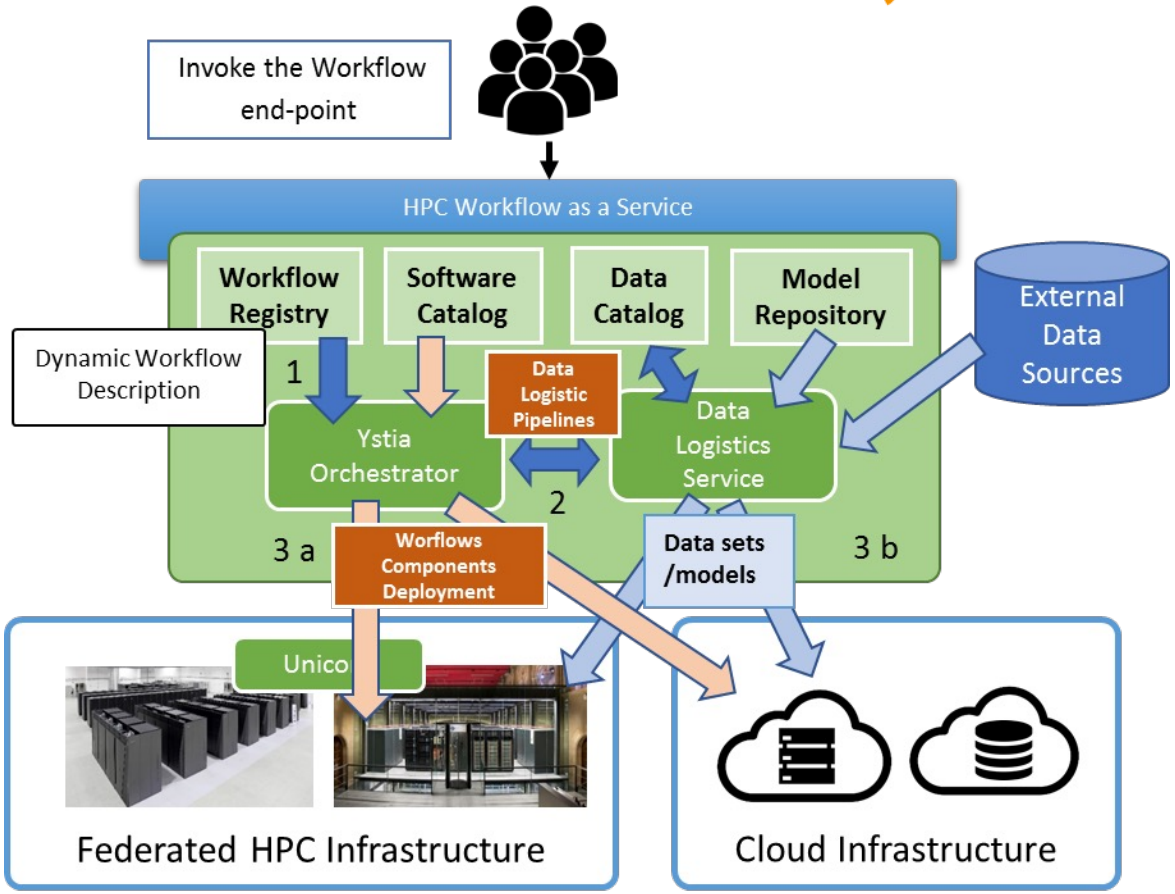

HPC Ready Containers

eFlows4HPC

Standard container image creation

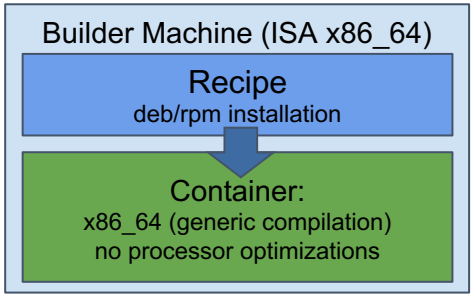

eFlows4HPC approach

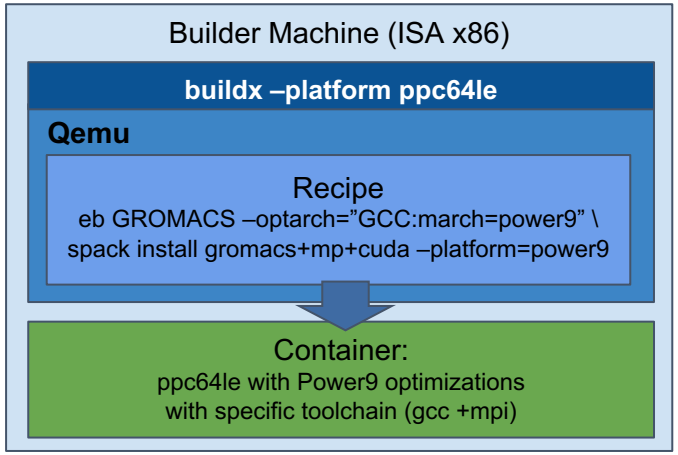

Service to automate the Container Image Creation

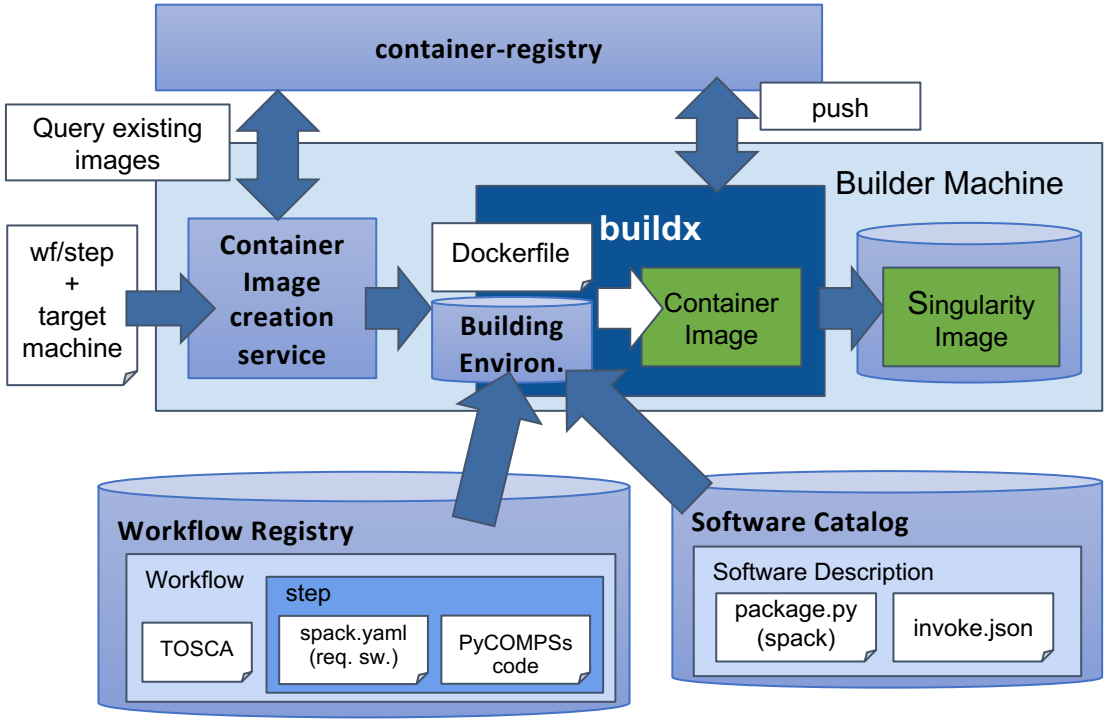

Credential management

- Prior to executing the registered workflows, the users have to configure the infrastructure access credentials
- Usernames, public-key certificates, passwords
- Users' certificates managed by an Execution API
  - Provides a few methods to register and access credentials or generate a new secret
  - HashiCorp Vault for secret (SSH keys) management
- User authorizes adding credentials in the HPC cluster
- Credentials identified by a token attached to the user's workflow invocation.

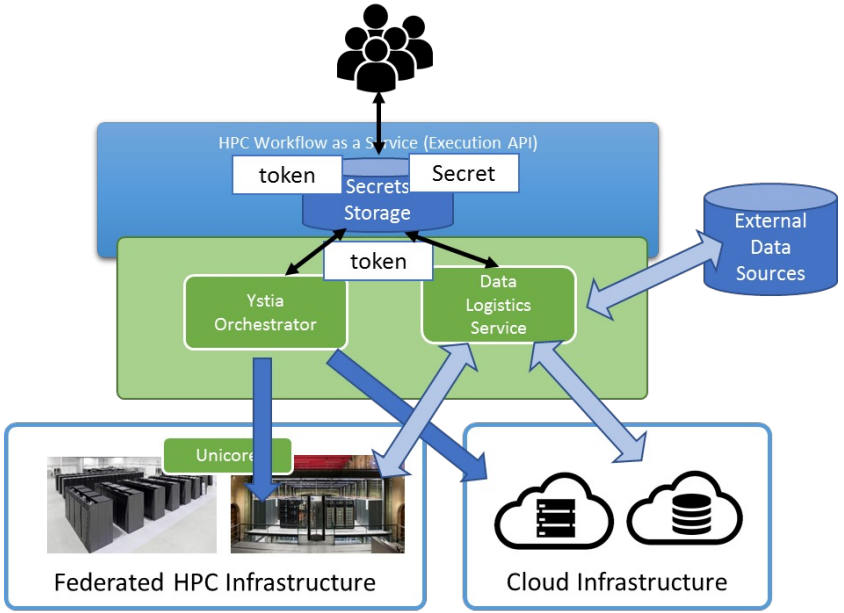

Operation- Workflow Execution

- Submission of the execution of the workflow processes to the HPC infrastructure
- PyCOMPSs orchestrates different task types
  - HPC (MPI), ML, DA
- Dynamic execution
  - Runtime task-graph
  - Task-level FT
  - Exceptions
- Data management
  - Persistent storage
- Optimized kernels
  - EPI, GPU, FPGA

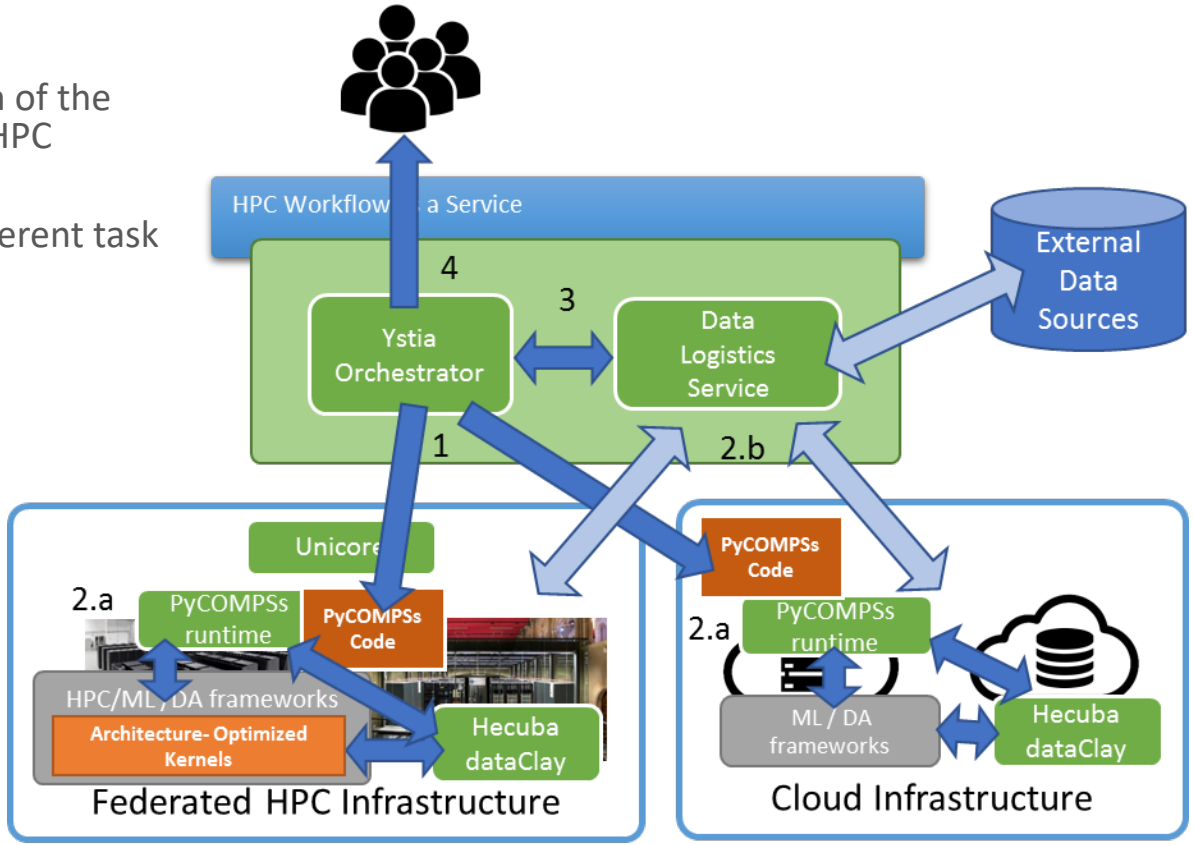

PILLAR APPLICATIONS

Pillar I: Manufacturing

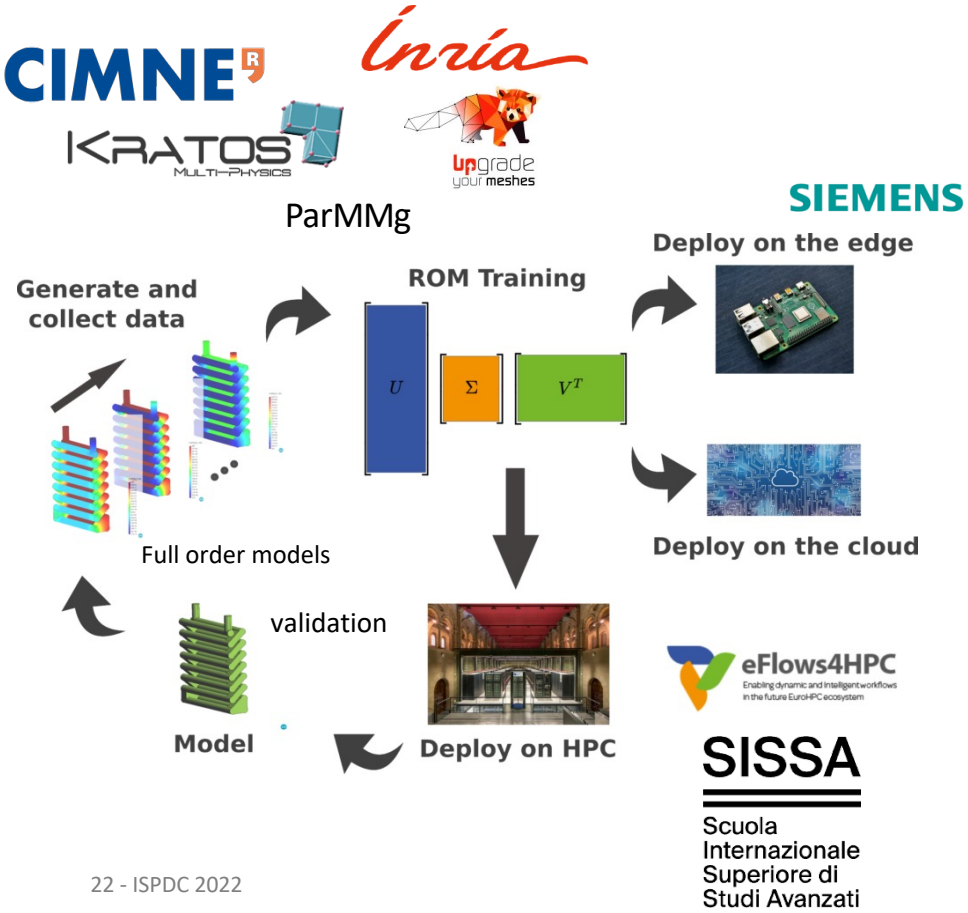

Pillar I focuses on the construction of DigitalTwins for the prototyping of complex manufactured objects:

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

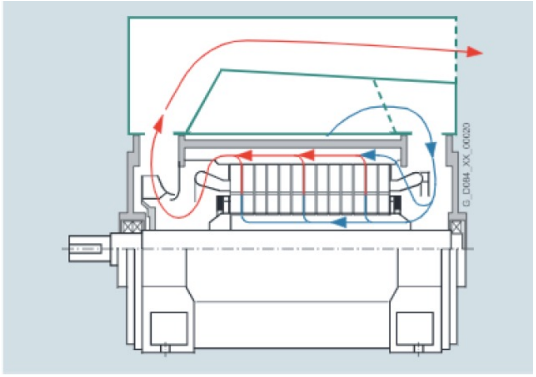

Integration of HPC and data analytics in a single workflow

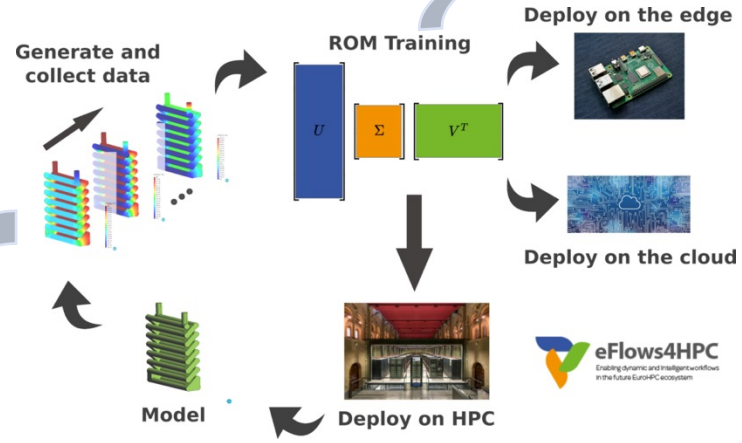


```
import dislib as ds

@constraint(computingUnits=8)
@task(Y_blocks={Type: COLLECTION_IN, Depth: 2},
def my_qr(Y_blocks):
    Y = np.block(Y_blocks)
    Q,R = np.linalg.qr(Y, mode='reduced')
    return Q,R

def rsvd(A, desired_rank):
    k = desired_rank
    ...
    Y = A @ Omega
    Q,R = my_qr(Y._blocks)
    Q=load_blocks_rechunk([Q], ...)
    ...
```

```
@constraint(computing_units="8")
@mpi(runner="mpirun", processes="16")
@task(returns = np.array)
def ExecuteInstance_Task(
    ...
    current_parameters =
    ...
    simulation = GetTrain
    simulation.Run()
    return simulation.GetSnapshotsMatrix()
```

```
....
for instance in range (0,TotalNumberOFCases):
    blocks.append(ExecuteInstance_Task(model, parameters, pars, inst
...
U, s = rsvd(A, desired_rank)
...
```

Pillar I sample workflow – initial version

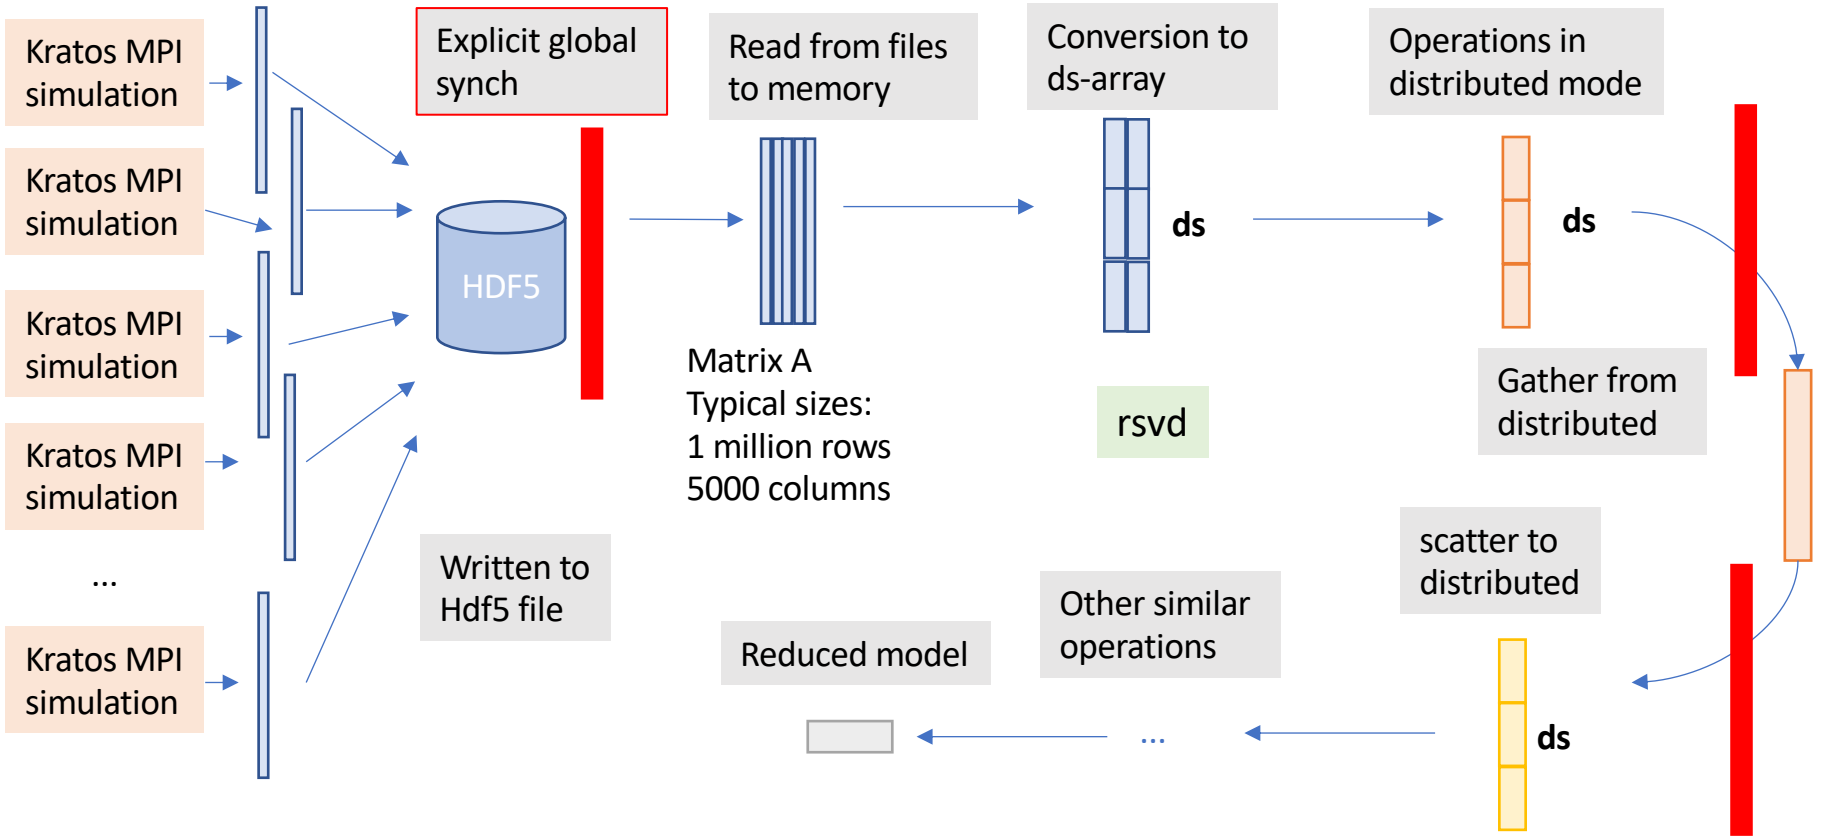

Pillar I sample workflow – optimized version

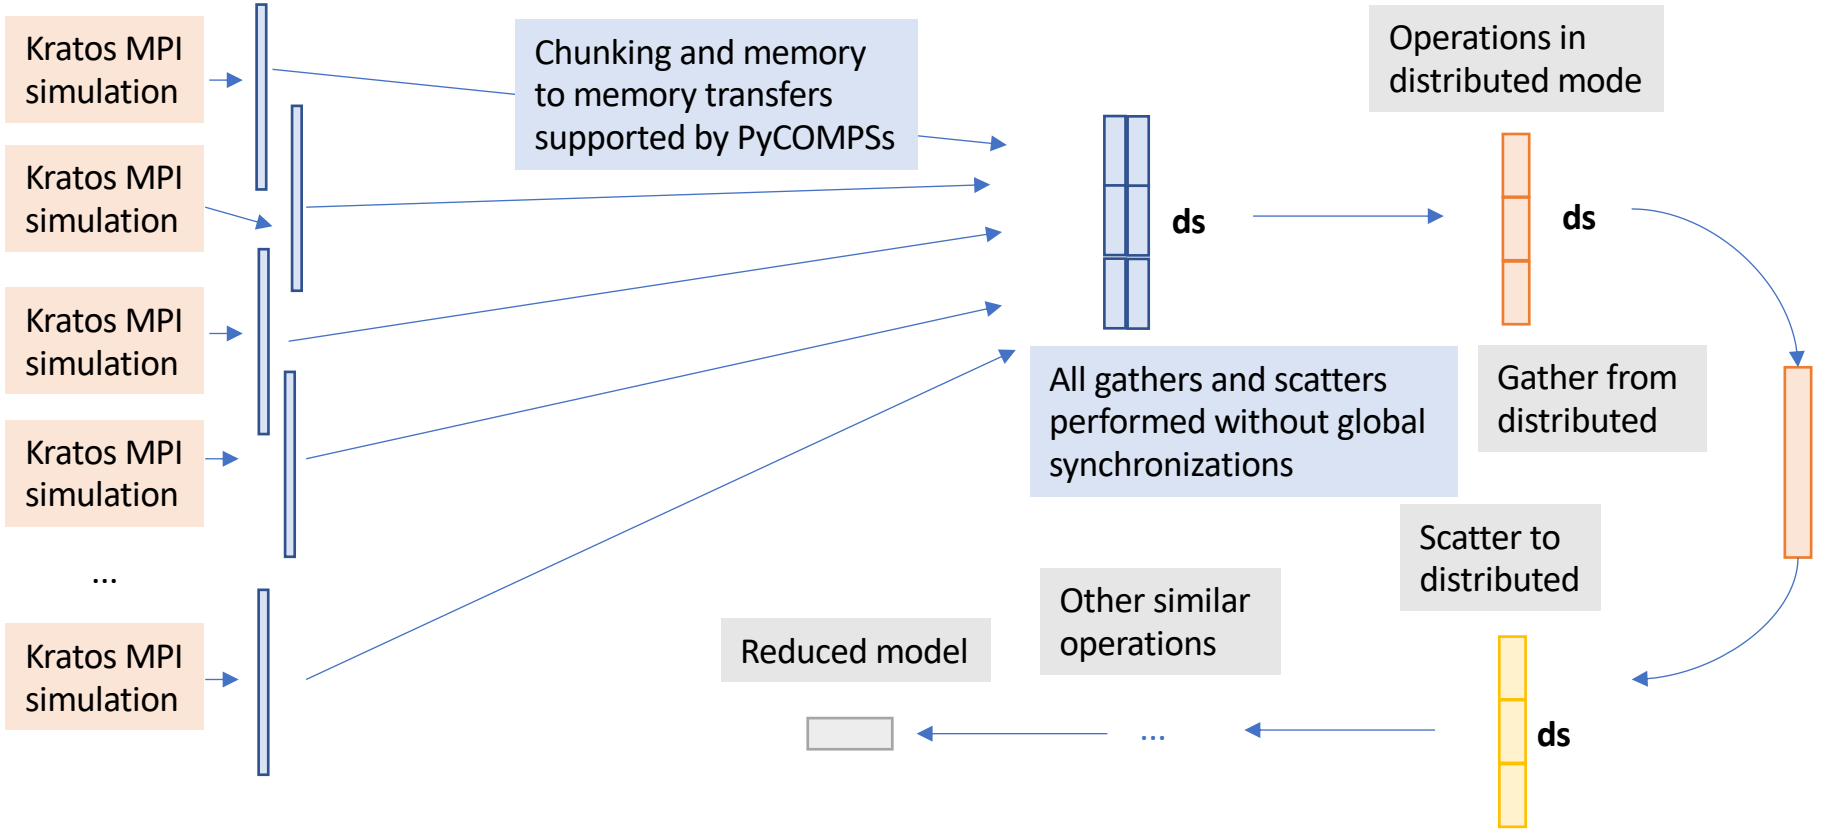

Pillar II: Climate

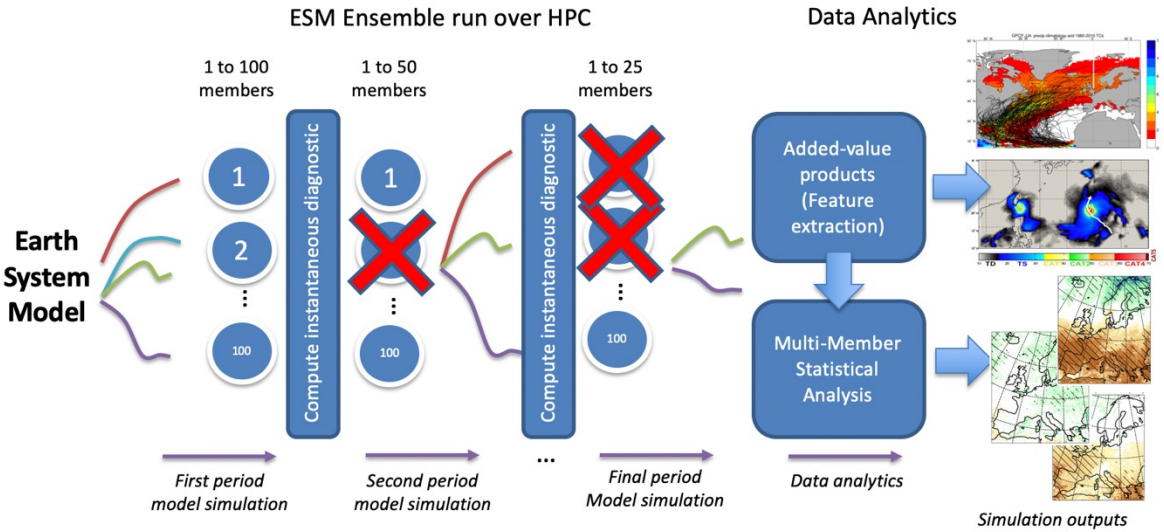

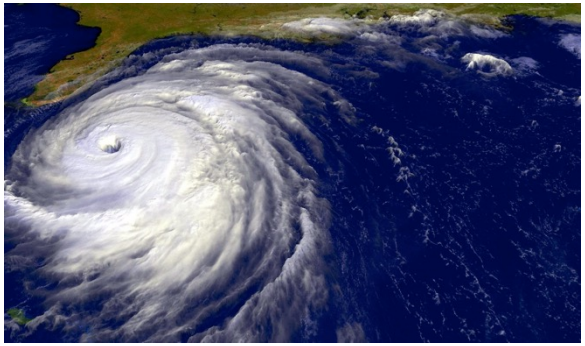

Dynamic (AI-assisted) workflow

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

ESM Dynamic Workflow

- Ensemble of FESOM simulations
- Intermediate simulation results stored persistently to Hecuba
- Pruning mechanism cancels given simulations based on dynamic anaysis of data stored in Hecuba

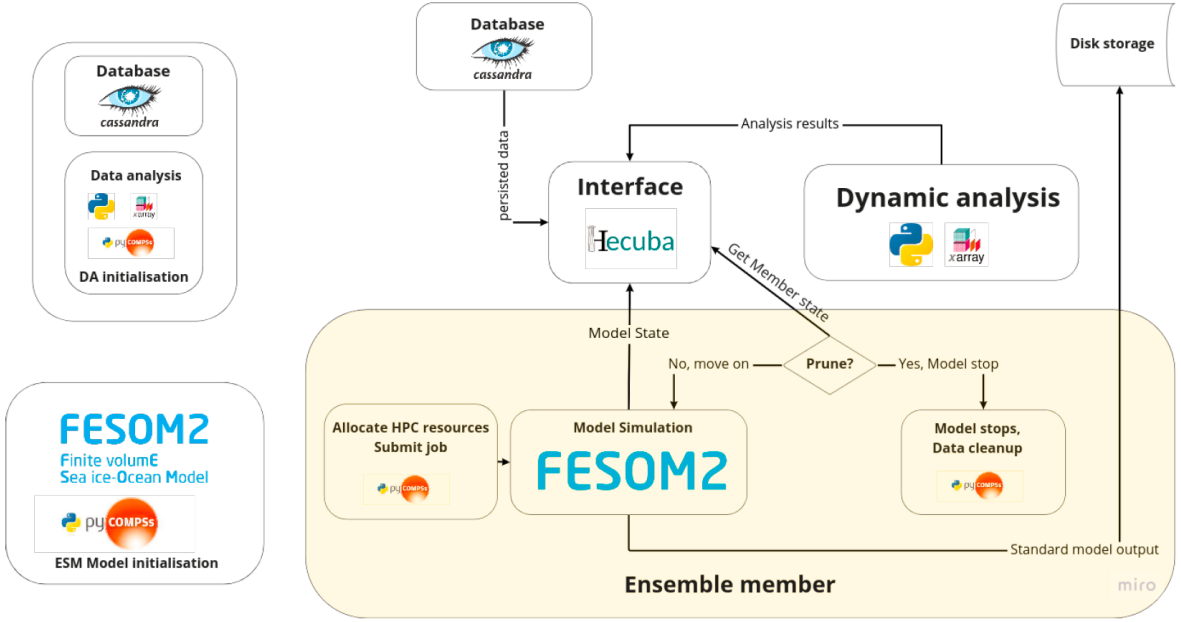

Use case 1 – CMCC-CM3 based

- Selection of input data for NN training ERA5 NetCDF (North Pacific region): ~240GB
- Patches generation (40x40): ~20GB

- NN defined and trained
- Inference on CMCC-CM3 datasets under testing

- Based on deterministic algorithm implemented in the TSTORMS tool

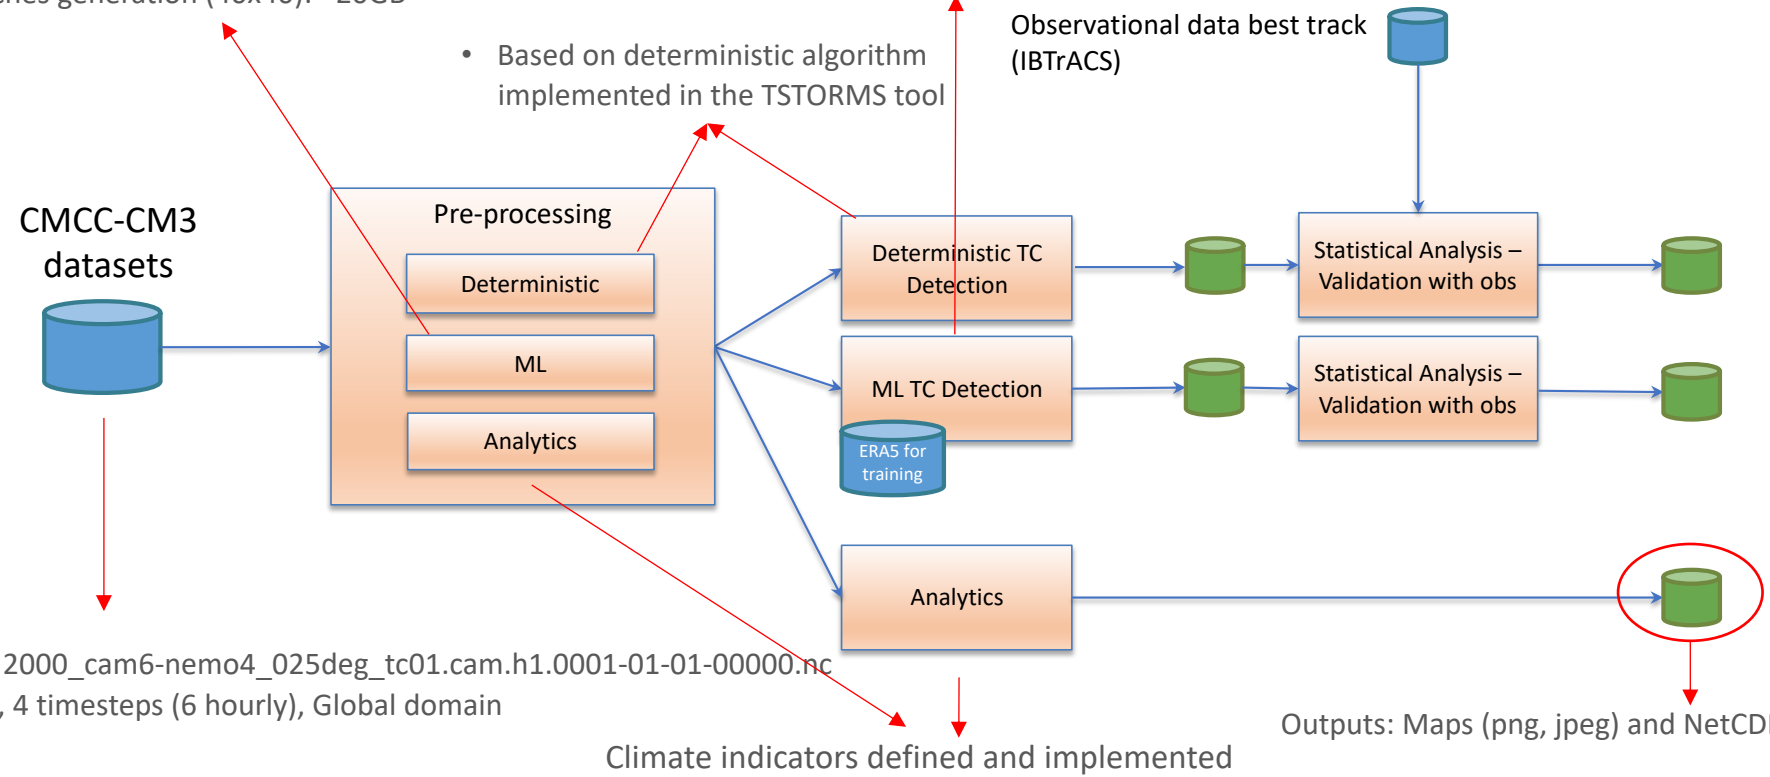

NetCDF: 2000_cam6-nemo4_025deg_tc01.cam.h1.0001-01-01-00000.nc
~200MB, 4 timesteps (6 hourly), Global domain

Climate indicators defined and implemented

Outputs: Maps (png, jpeg) and NetCDF

Tropical Cyclones Detection ML Workflow: Training Phase

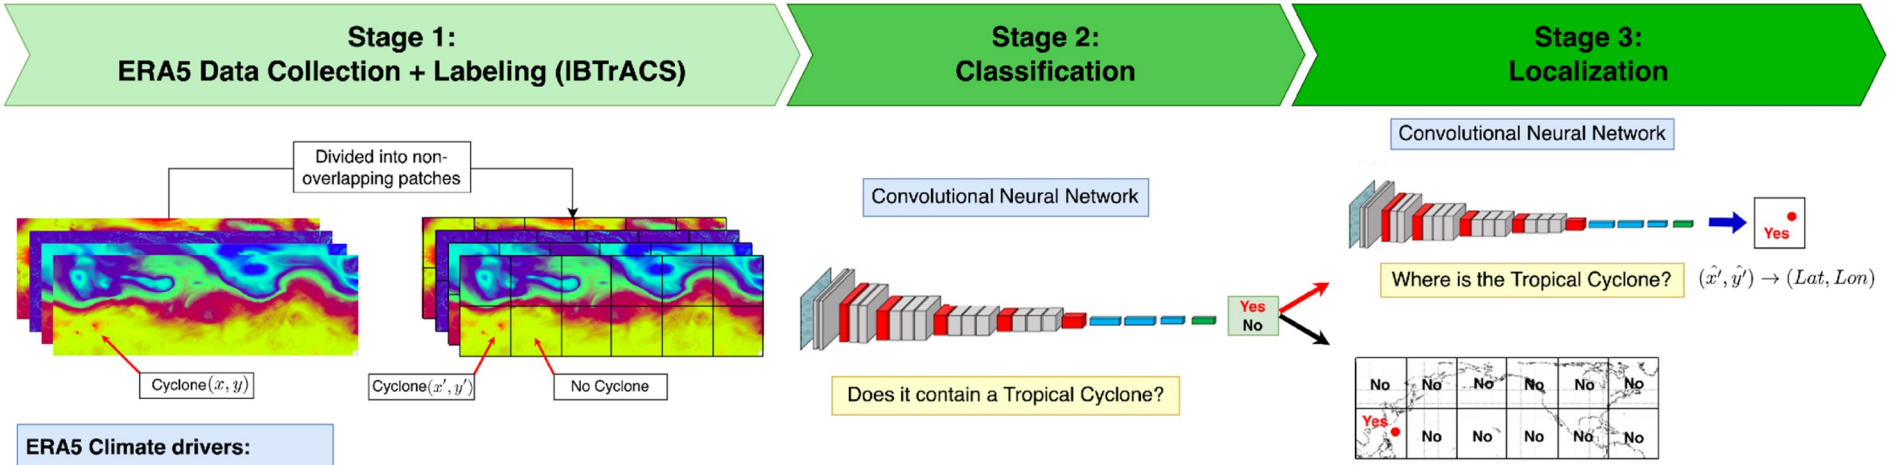

- Three main stages:
  - STAGE 1: Gathering of ERA5 climatic maps and generation of patches containing at most 1 tropical Cyclone (TC) each
  - STAGE 2: Classification of TC presence/absence inside the patch
  - STAGE 3: Localization of TC center coordinates in the patches in which the TC was previously classified as present

Pillar III: Urgent computing for natural hazards

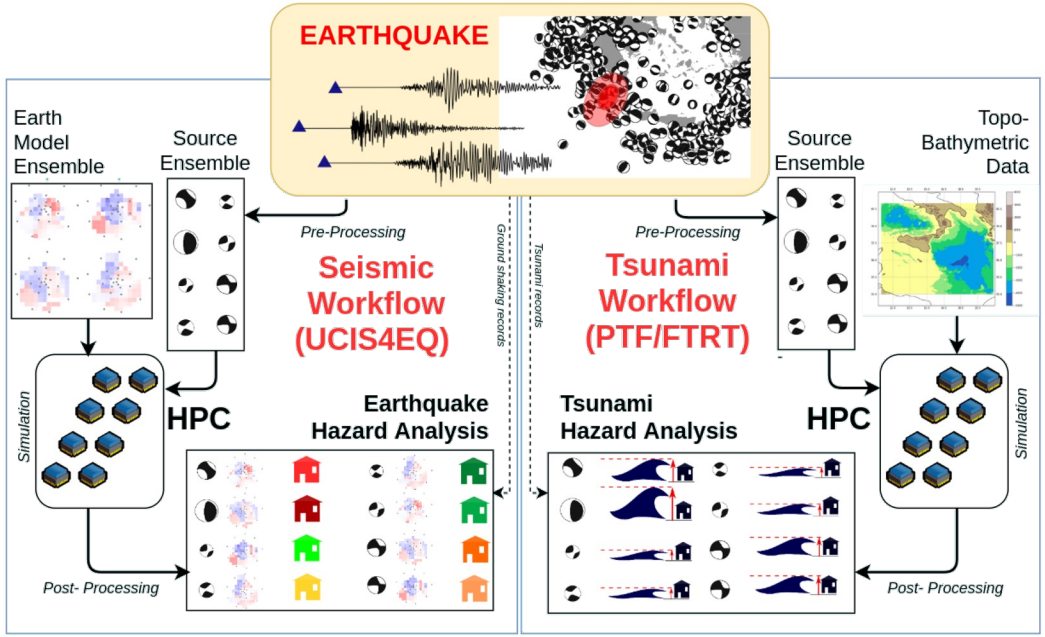

Pillar III explores the modelling of natural catastrophes:

- 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

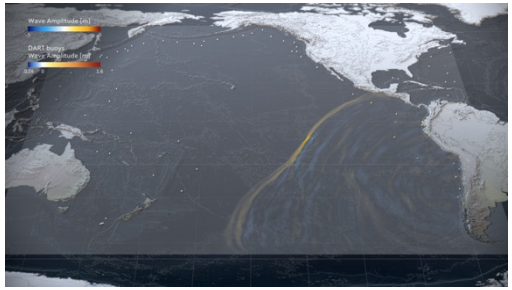

UCIS4EQ workflow: http services as tasks

Urgent Computing Integrated Services for Earthquakes: UCIS4EQ

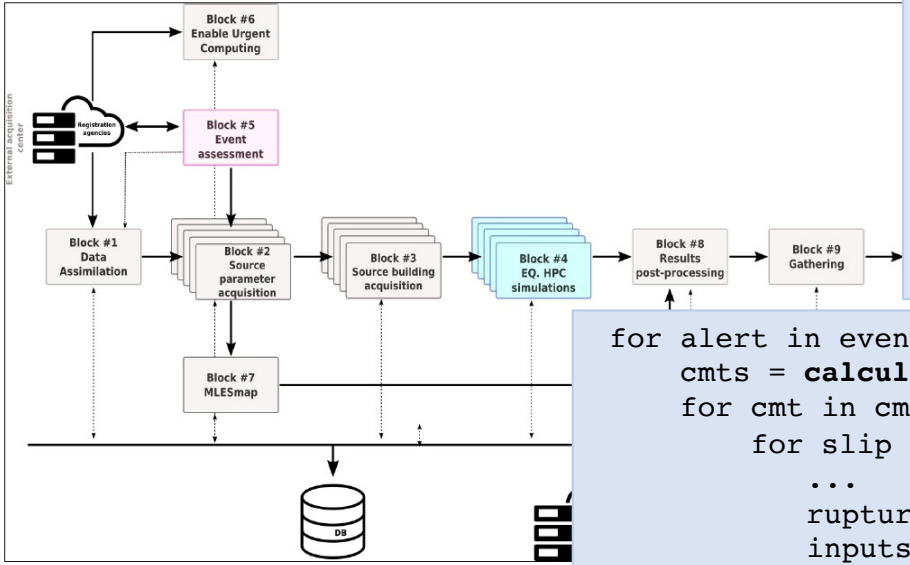


```
@http(request="POST", resource="SalvusRun", ...')
@task(returns=1)
def run_salvus(event_id, trial, input, resources):
    """
    """
    pass
@http(request="POST", resource="cmt",...)
@task(returns=1)
def calculate_cmt(alert, event_id, domain, precmt):
    """
    """
    pass
...
```

```
for alert in event['alerts']:
    cmts = calculate_cmt(alert, eid, domain, precmt)
    for cmt in cmts.keys():
        for slip in range(1, region['GPSetup']['trials']+1)
            ...
            rupture = compute_graves_pitarka(eid, alert, ...)
            inputs = build_input_parameters( eid, alert, ...)
            salvus_inputs = build_salvus_parameters( eid, path, ...)
            result = run_salvus( eid, path, ...)
            all_results.append(run_salvus_post(eid, result, ...))
result = run_salvus_plots(eid, basename, domain, resources)
```

Project main achievements

https://eflows4hpc.eu/software/

- Requirements and software architecture
- Definition and implementation of abstractions to support the integration of different stack components
- Design and development of a minimal workflow
- Design and first version of the HPCWaaS methodology
- Design and implementation of the Data Catalogue
- Definition of first version of Pillars' workflows. Implementation in progress
- First release of project software and documentation available
- Set of internal trainings about software stack components and HPCWaaS
- Good visibility: articles, keynote presentations, media

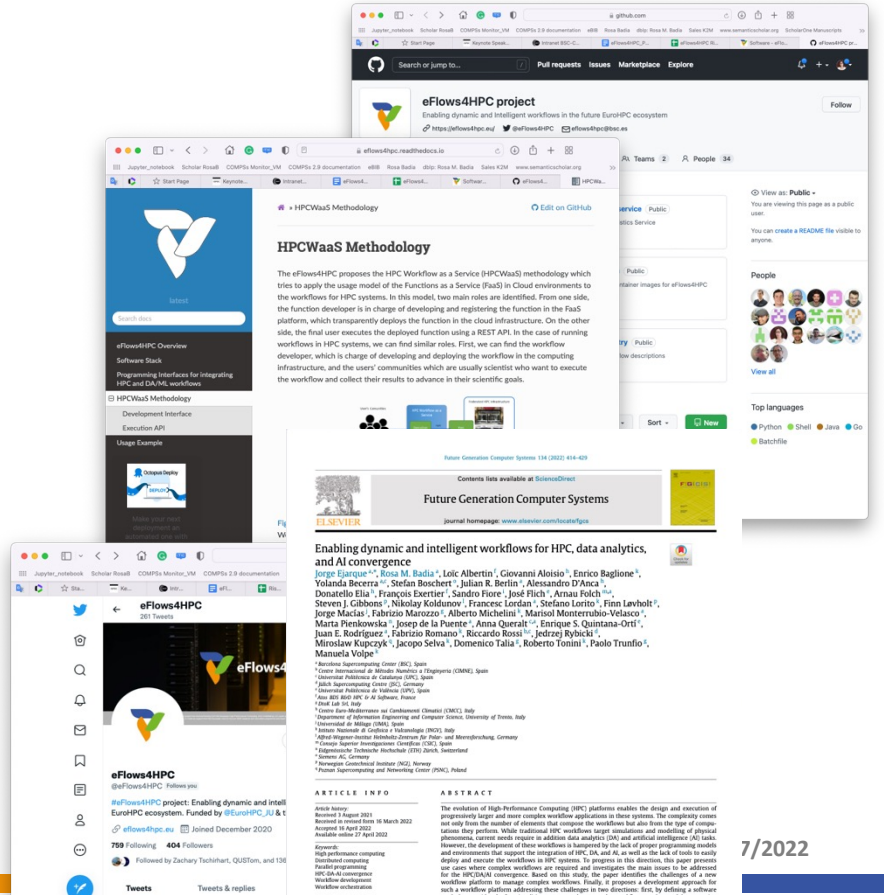

Conclusions

- 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 to make it easier the adoption of HPC technologies

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.

Project WP Structure

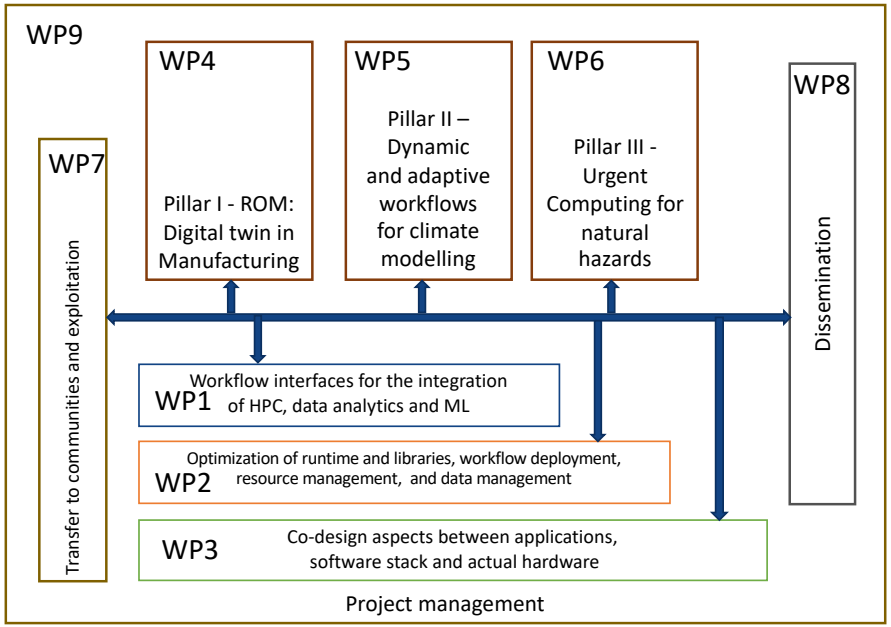

Project RoadMap