20 BSC YEARS

EuroHPC
Joint Undertaking

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as the Participating States Spain, Portugal and Türkiye

Hosting Consortium:

Spain Portugal Türkiye

Generalitat de Catalunya
Departament de Recerca i Universitats

GOBIERNO DE ESPAÑA MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES

UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH

Plan de Recuperación, Transformación y Resiliencia

UNIÓN EUROPEA
Fondo Europeo de Desarrollo Regional

From Supercomputers to Solutions: Application Workflows for Science and Society

Rosa M Badia

02/10/2025

BIG DATA FROM SPACE 2025

Presentation roadmap

- HPC* systems in Europe
- Workflows outlook
- Workflows lifecycle
- Geohazards' Digital Twins

HPC: High Performance Computing

EuroHPC JU systems

	TOP500	Country	Peak perf.	Architecture
JUPITER	**4**	Germany	930 PF	SiPearl Rhea1 processor, Nvidia GH200
LUMI	**9**	Finland	539.13 PF	AMD EPYC™ CPUs + AM[illegible]™ GPU
Leonardo	**10**	Italy	315.74 PF	Intel Ice-Lake [illegible]re Rapids + NVIDIA Ampere[illegible]
MareNostrum 5	**14**, 45	Spain	249.44 PF, 40.10 P[illegible]	[illegible]phire Rapids, NVIDIA H100
Meluxina	136	Luxembourg	[illegible]	AMD EPYC + NVIDIA Ampere A100
Vega	305	Slovenia	[illegible]10.05 PF	AMD Epyc 7H12 + Nvidia A100
Karolina	195	[illegible]public	12.91 PF	AMD + Nvidia A100
Discoverer	[illegible]	[illegible]ulgaria	5.94 PF	AMD EPYC
Deucalion	297	Portugal	9.76 PF	A64FX, AMD EPYC, Nvidia Ampere

One more exascale system to be deployed in France (Alice Recoque)

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

MN5

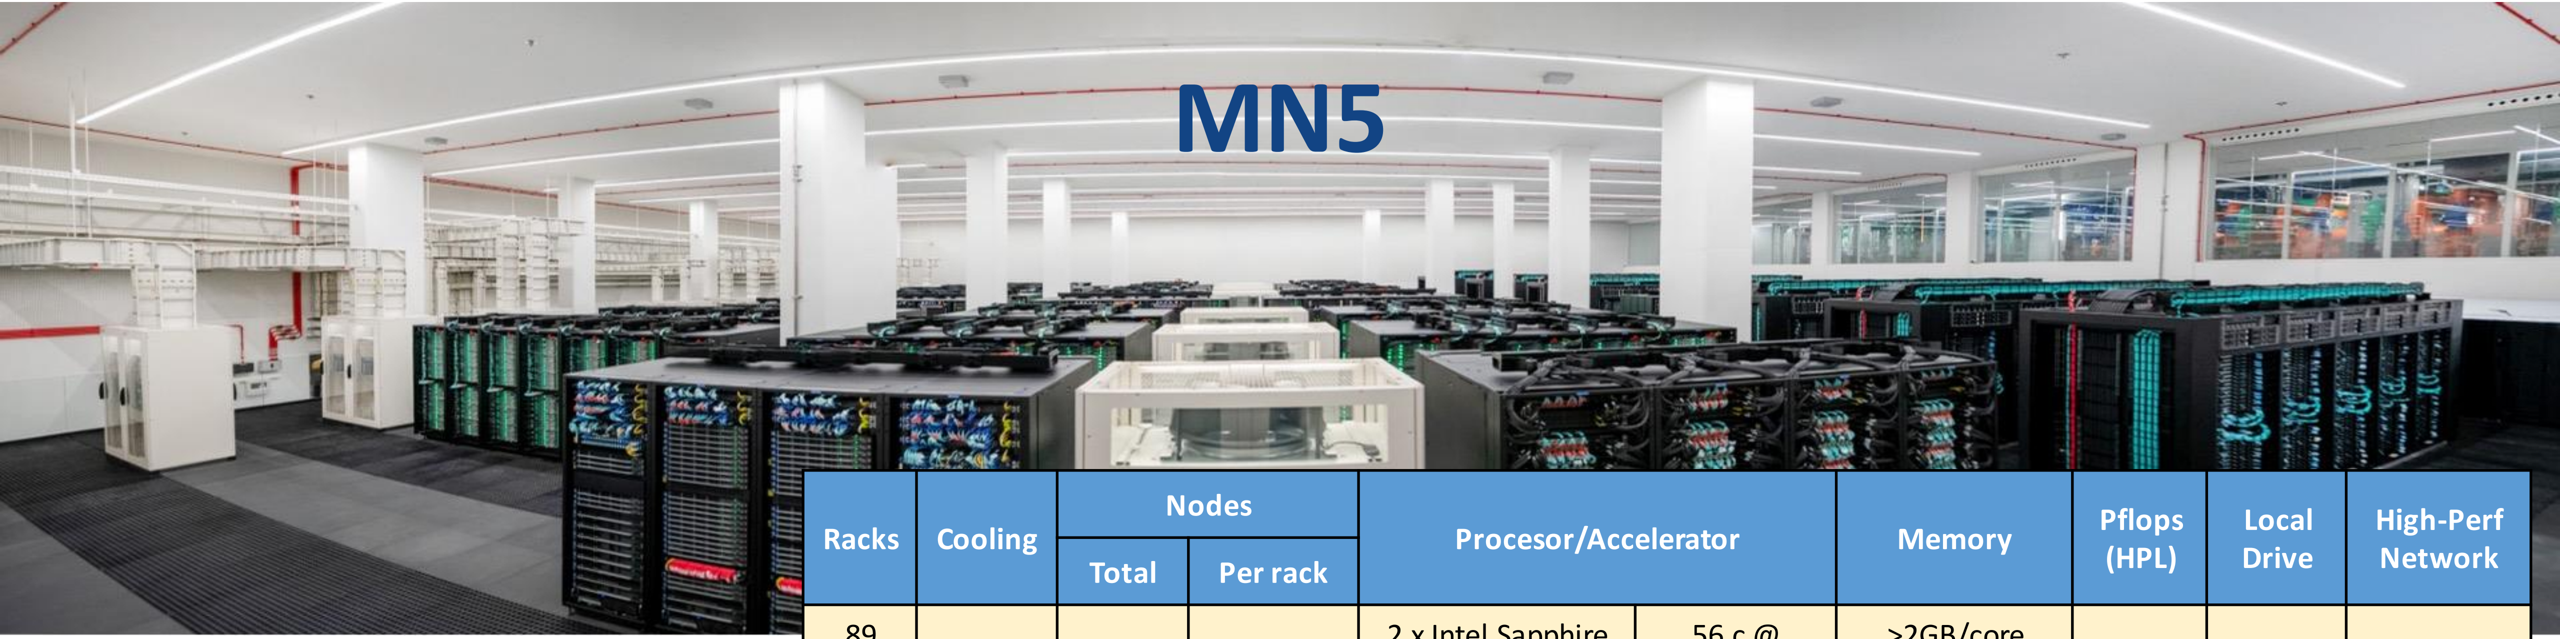

- Federated slurm:
    - Hybrid allocations enabled
- Shared General Parallel File System (GPFS)

GPP

Racks	Cooling	Nodes		Procesor/Accelerator		Memory	Pflops (HPL)	Local Drive	High-Perf Network
		Total	Per rack						
89	DCL + RDHX	6192	72 (6 x 6 x 2)	2 x Intel Sapphire Rapids 8480+	56 c @ 2GHz	>2GB/core 256 GB DDRS	40.10	960 GB NVMe	1 x NDR200 Shared between 2 nodes
		216				>8GB/core 1024 GB DDRS			
1		72		2 x Intel Sapphire Rapids 9480	56 c @ 1.9 GHz	> 0.5 GB HBM/core 128 GB HBM + 32 GB DDRS	0.34		

ACC

Racks	Cooling	Nodes		Procesor/Accelerator		Memory	Pflops (HPL)	Local Drive	High-Perf Network
		Total	Per rack						
89	DLC	1120	32	2 x Intel Sapphire Rapids 8460Y+	40c @ 2 GHZ	512 GB	175.3	480 GB NVMe	4 x NDR200
				4 Nvidia Hopper 64GB HBM					

Quantum SPAIN

MN5 - Ona

- 2 coolers for the various quantum chips installation
- Shared GPFS
- Multi-cluster SLURM

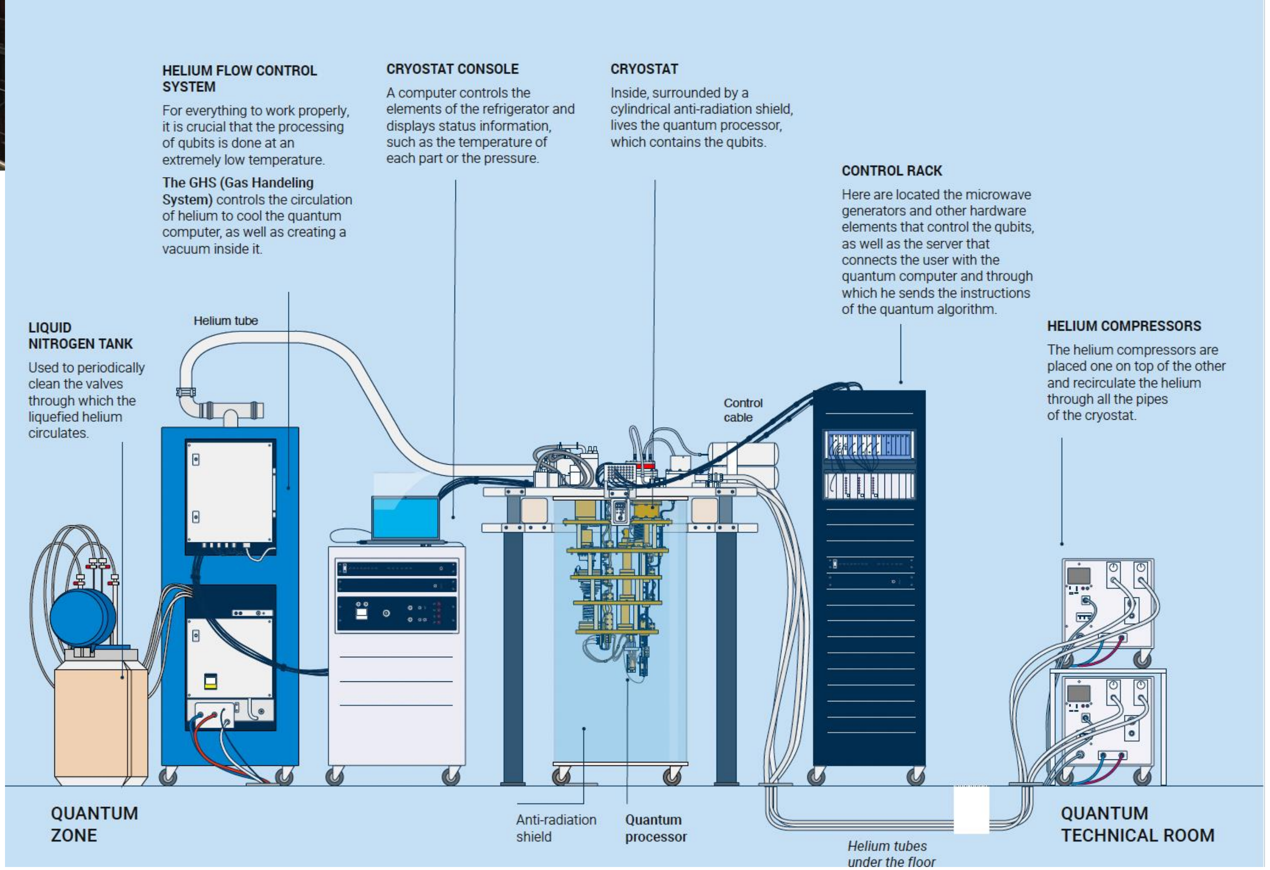

Presentation roadmap

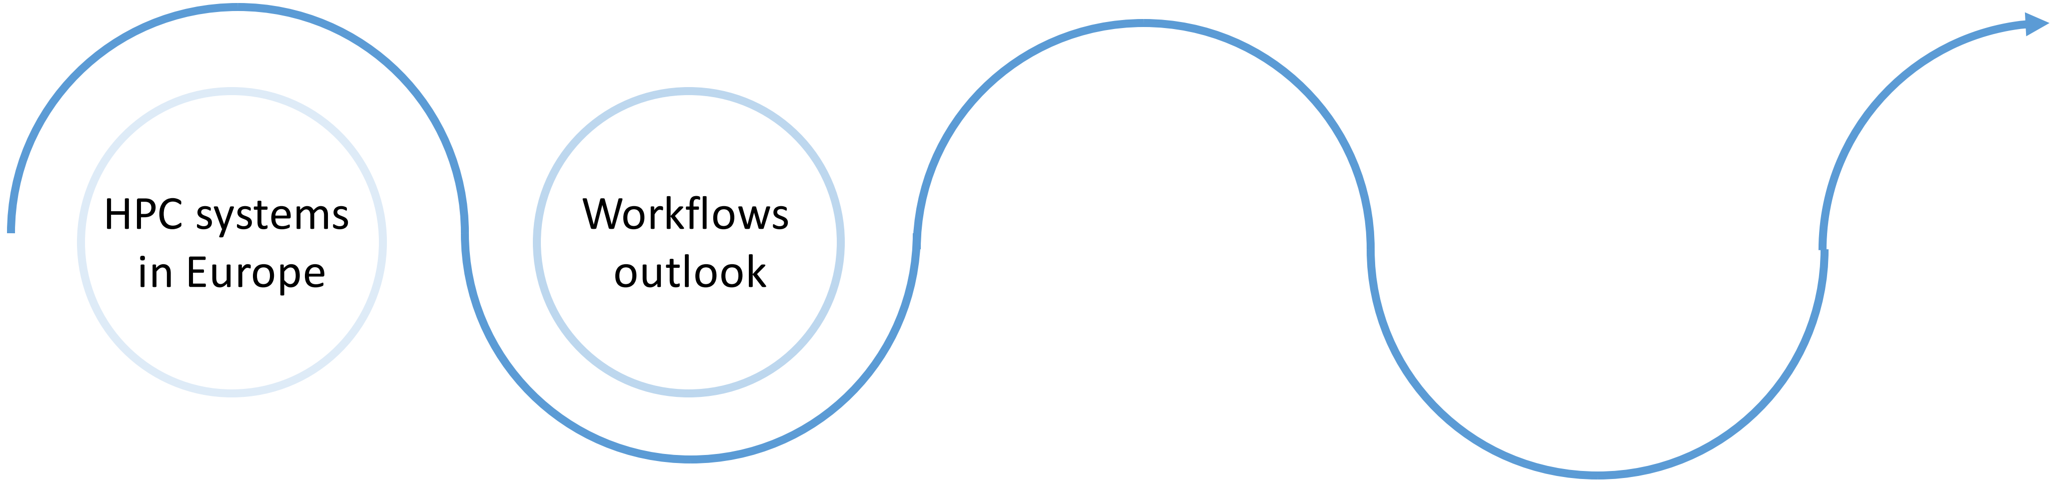

Complex problems for complex computing infrastructures

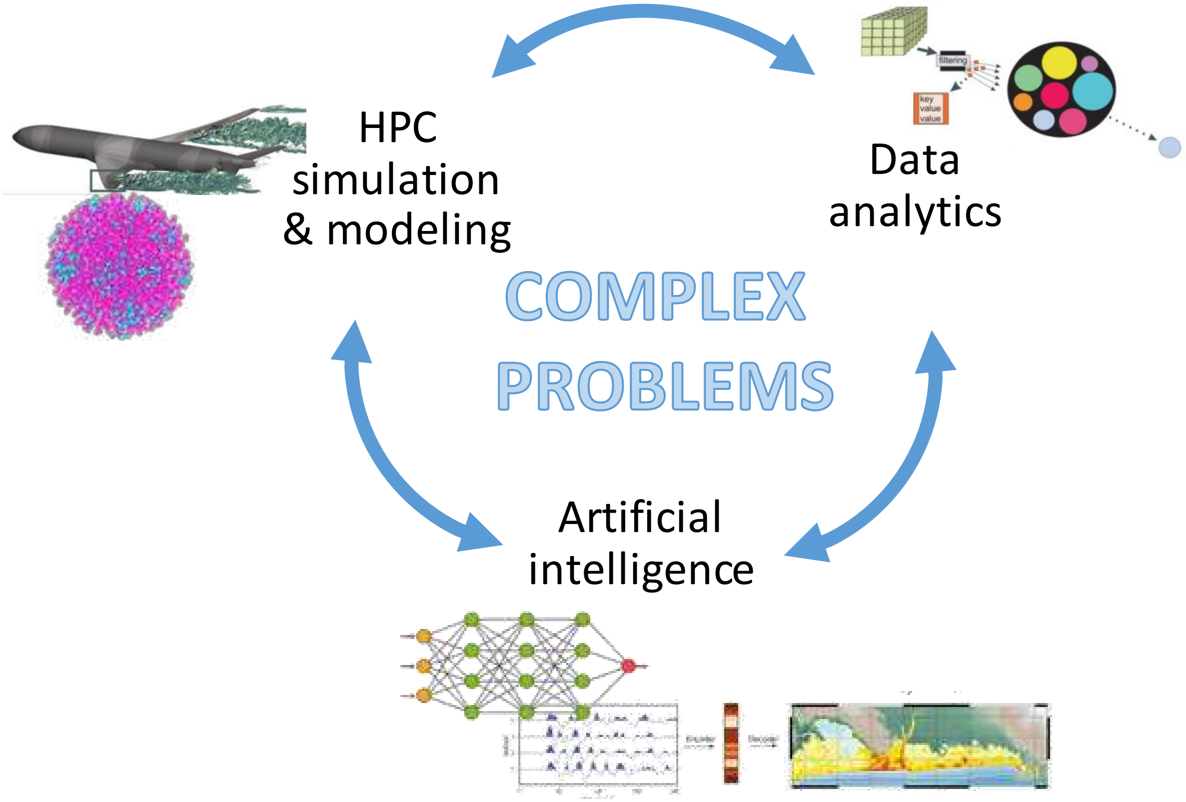

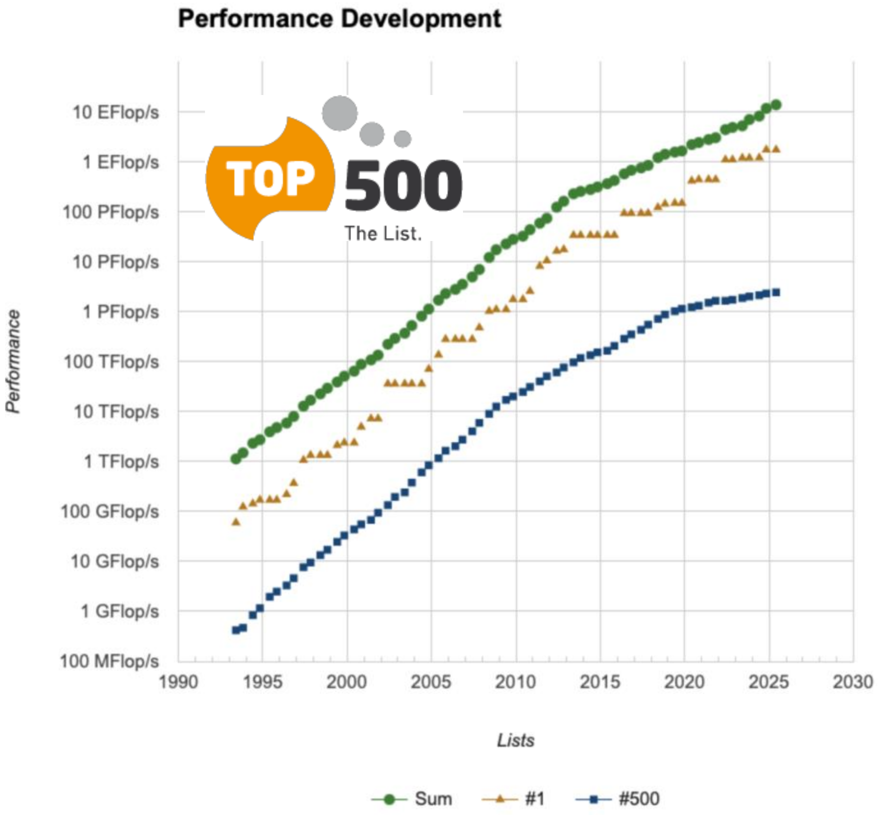

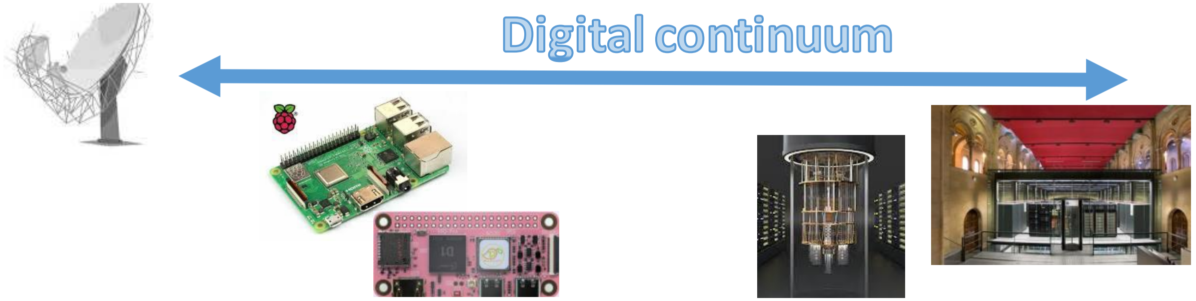

EuroHPC MASP* 2025 update

Pillar	Area	Component	Area	Component
Applications	Applications	Application codes, libraries		
Technology (Software)	Programming Environment	Programming Language and Frameworks	System Management	Runtime System
		Parallel Programming Models		Workflow manager, Resource management/ Job scheduling
	Tools	Compilers		Software deployment tools, S
		Debuggers		File
		Performance analysis tools		Operat
Technology (Hardware)	System Components	Processors (CPU, GPUs, other accelerators)		
		Memory and Storage		
		Network		
	Rack integration	Mechanics, Electric, Cooling		
	System deployment	Site preparation, Containers		
	Operations	System management, User access		
		ss,		

Software technologies priority domains:

"**Workflow managers** supporting scientific and AI pipelines **across modular HPC supercomputers**, federated HPC, cloud-HPC, and QC-HPC. The focus should be on workflow managers that support the **integration of HPC, AI and data analytics** in the same application workflow within a supercomputer. In addition, since many workflow solutions are available, **efforts towards their integration** should be sought."

Federation challenges:

"Orchestration and Workflow Scheduling: **Workflow management across HPC and cloud** platforms faces challenges in task prioritisation and resource utilisation. Hybrid schedulers like could provide seamless integration."

Recommendations for applications:

"**Workflows**: promote new frameworks for the **development of complex workflows** that combine different application components, possibly of different nature (**HPC, AI, HPDA**) to be executed inside large allocations in a single supercomputer. ... efficient use of the **new heterogeneous and complex HPC** systems, with **malleable, elastic execution and dynamic** allocation ... **performance as well as energy**....."

* Multi-Annual Strategic Programme

EuroHPC MASP* 2025 update

Hardware technologies:
With **growing volumes of digital data** and federation of their sources, **high performance interconnect hardware, storage and AI engines** are becoming increasingly valuable, in many cases eclipsing raw compute power

Data management services:
Data and storage have not been until now part of EuroHPC's scope, but it is considered **critical for the success of EuroHPC** that highly **efficient data management technologies for controlling and moving large volumes of data** within the HPC/AI/QC environments are established.

Software technologies:
Data management: Software needs to adapt to ever-increasing **variety and complexity of underlying memory and storage** subsystems, including accelerators/co-processors equipped with self-hosted memory and multi-tiered memory/storage systems. It shall also address challenges of **data distribution and processing across networks**.

Key challenges:
Data Movement and Bandwidth: Efficiently transferring **vast datasets across HPC systems**, AI Factories, clouds, and repositories is critical. Technologies like Globus and Nodeum, alongside high-bandwidth networks like the ones to be procured by EuroHPC, are potential solutions.

* Multi-Annual Strategic Programme

Presentation roadmap

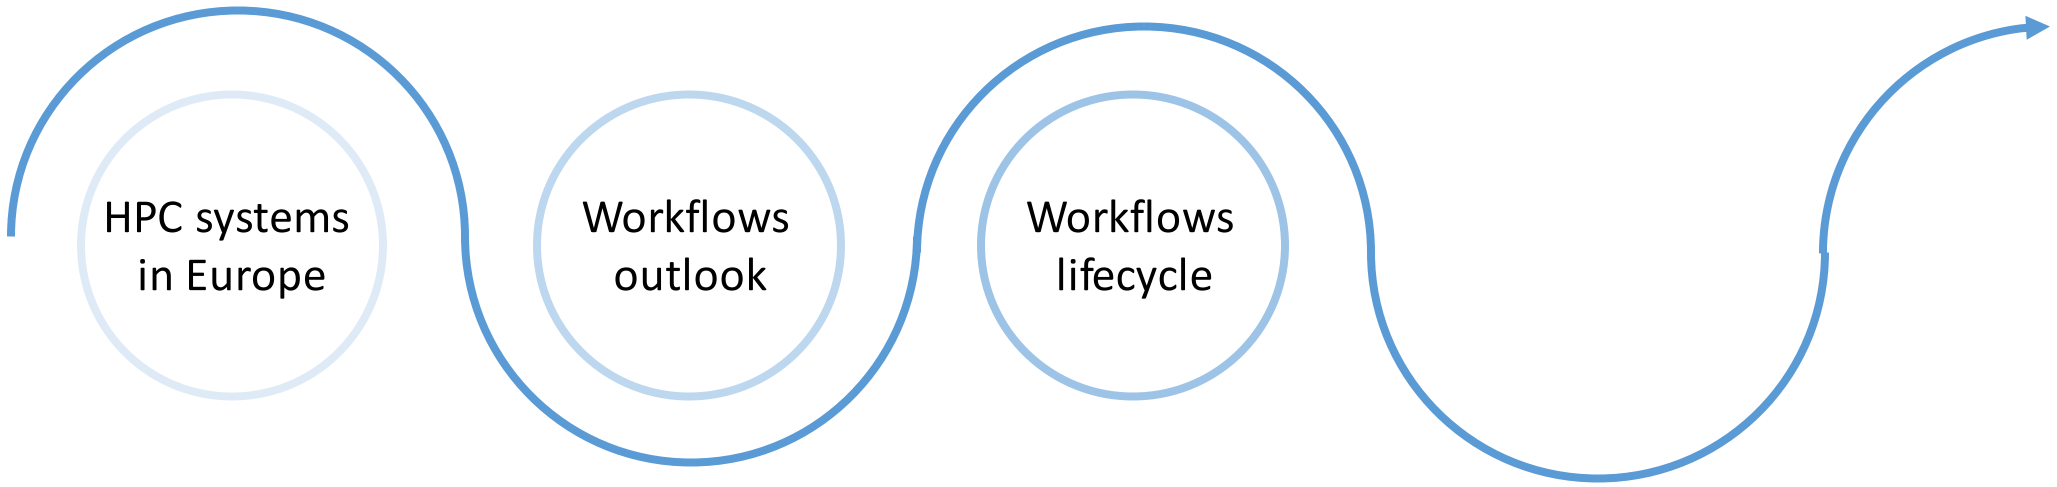

Workflow lifecyle challenges

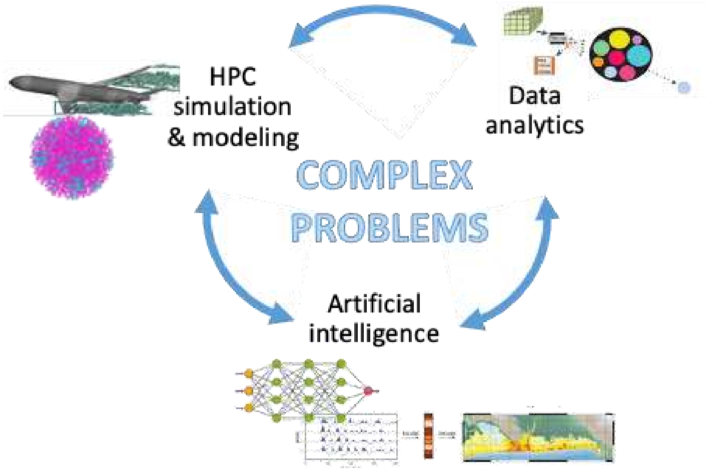

- Workflow development
  - Different programming models and environments
- Workflow deployment
  - Can we make it easier to new HPC users?
- Workflow operation
  - Go beyond static workflows
  - Not only computational aspects, data management as well
- Workflow reproducibility
  - Provenance

Specific challenges for EO data

- Earth Observation (EO) data processing challenges*
  - Large volumes
  - Multi-modality
  - Multiple sources
  - Variety of data types and formats
  - Temporal/spatial resolutions
  - Latency
- Aspects to be considered:
  - Data storage and management
  - Data integration and harmonization
  - Data processing for different applications
- Need for processing workflows based on parallel algorithms that can scale on heterogeneous HPC systems

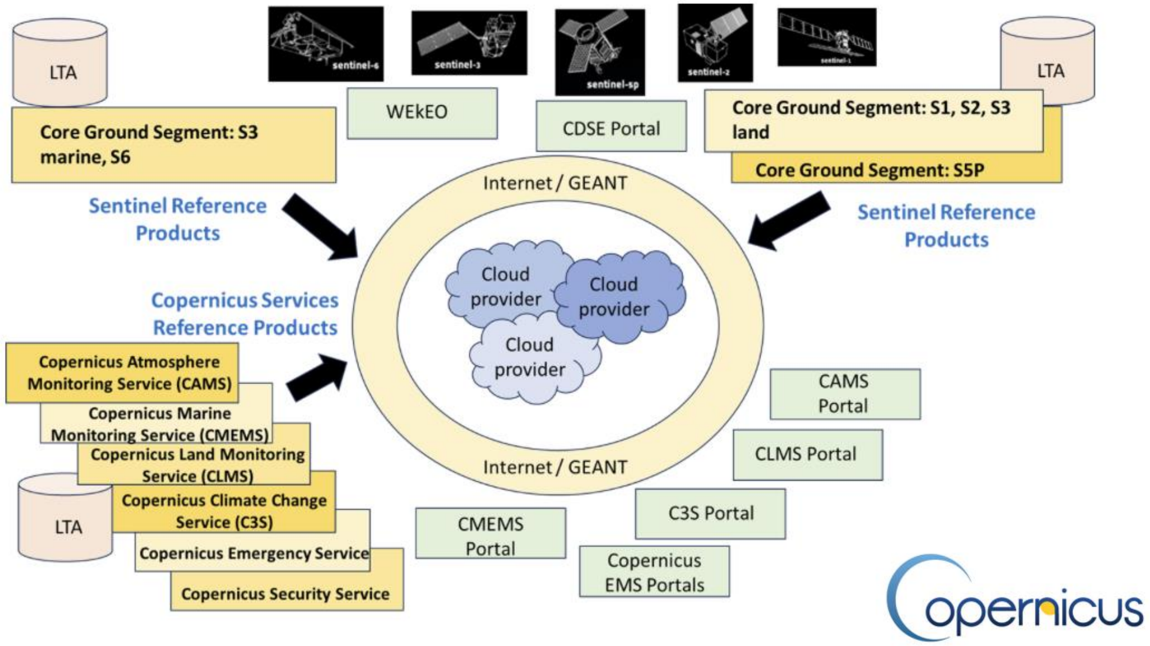

L- Tian, et al. "End-to-end process orchestration of earth observation data workflows with apache airflow on high performance computing." In IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium, 2023.

eFlows4HPC Workflow lifecycle: HPCWaaS

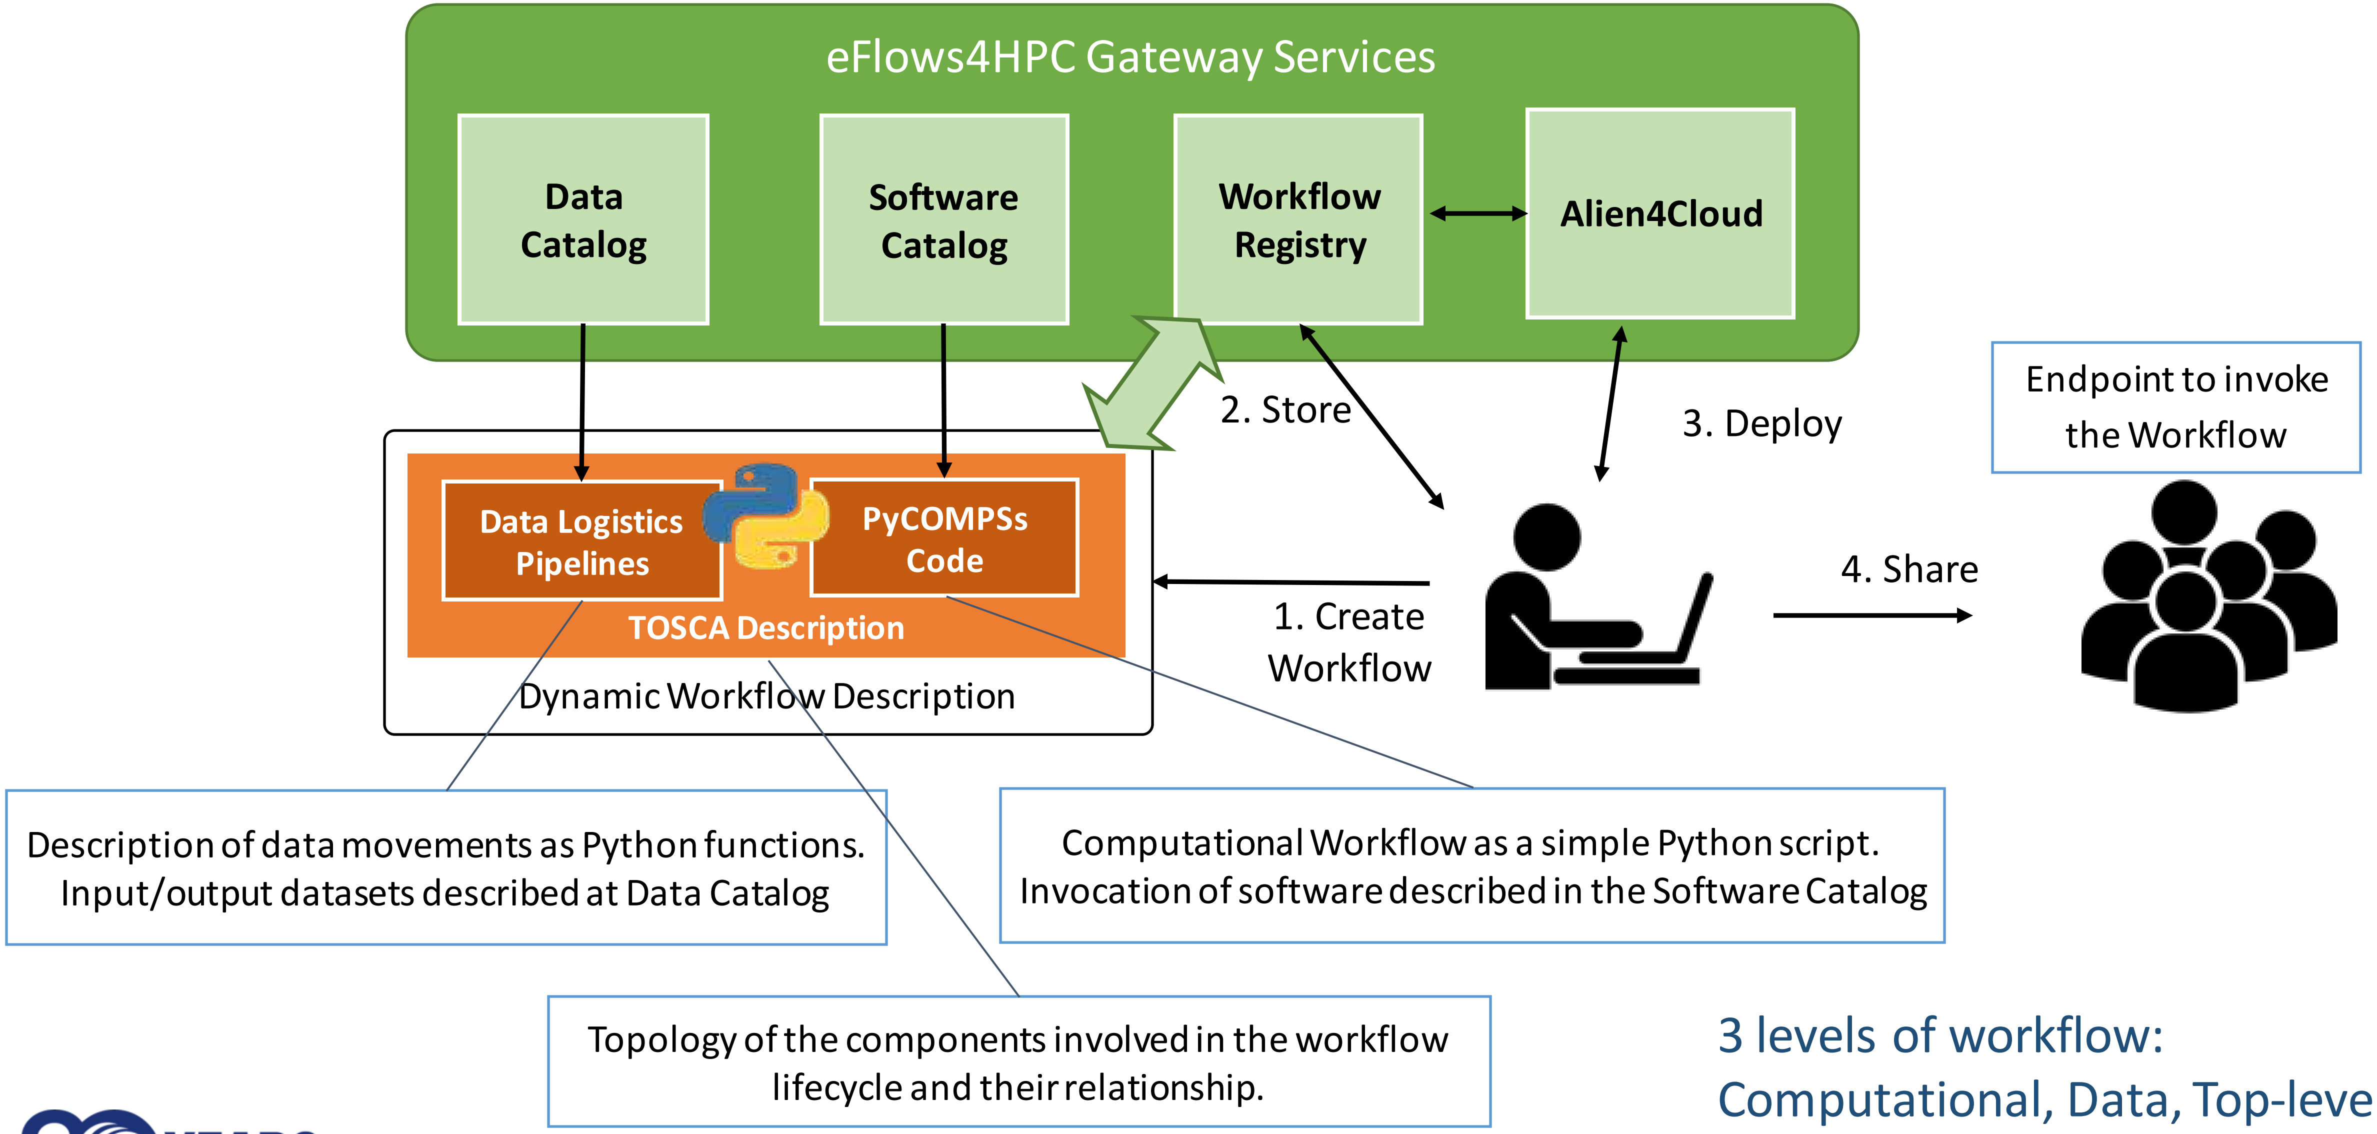

3 levels of workflow:
Computational, Data, Top-level

Workflow development: Coarse-grain task-based parallel programming

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

- Programming model
  - Sequential programming, parallel execution
  - General purpose programming language + annotations/hints
  - Task graph generated at runtime that express potential concurrency
  - Offers a shared memory illusion to applications in a distributed system
  - Agnostic of computing platform: clusters, clouds, containers
- Runtime deployed as a distributed master-worker
  - All data scheduling decisions and data transfers performed by the runtime
  - Elasticity in the cloud and in HPC
  - Support for heterogeneity through for tasks' constraints
  - Fault tolerance: tasks' faults, tasks' exceptions, checkpointing
  - Stream data
  - Workflow provenance recording

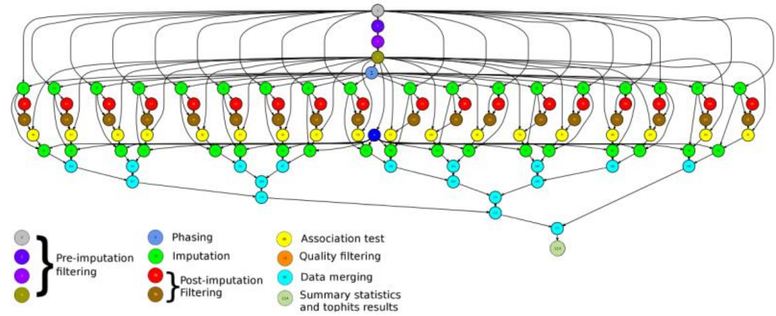

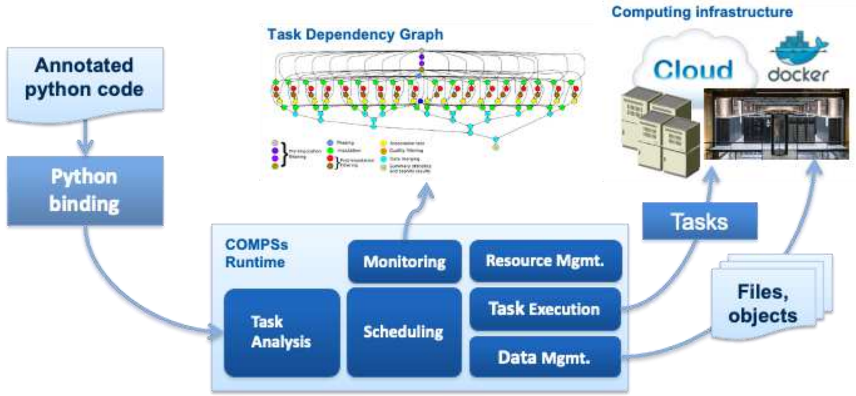

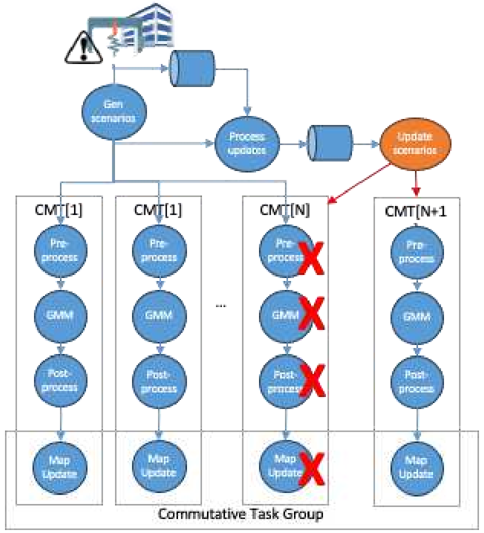

Data management

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:

- Python scripts (Apache Airflow)
- **Formalization** of data movements for transparency and reusability
- Data stage-in/out, **image transfer**, upload/download **AI models**
- From **services providers to HPC, from HPC** to model/data **repositories**, using multiple transfer protocols

Data Logistics Services (DLS):

- Based on Apache Airflow runtime
- Performs the execution of data pipelines at deployment and execution time

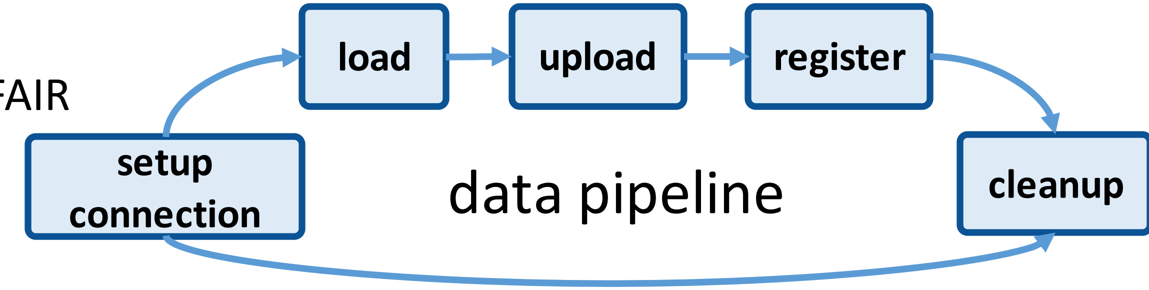

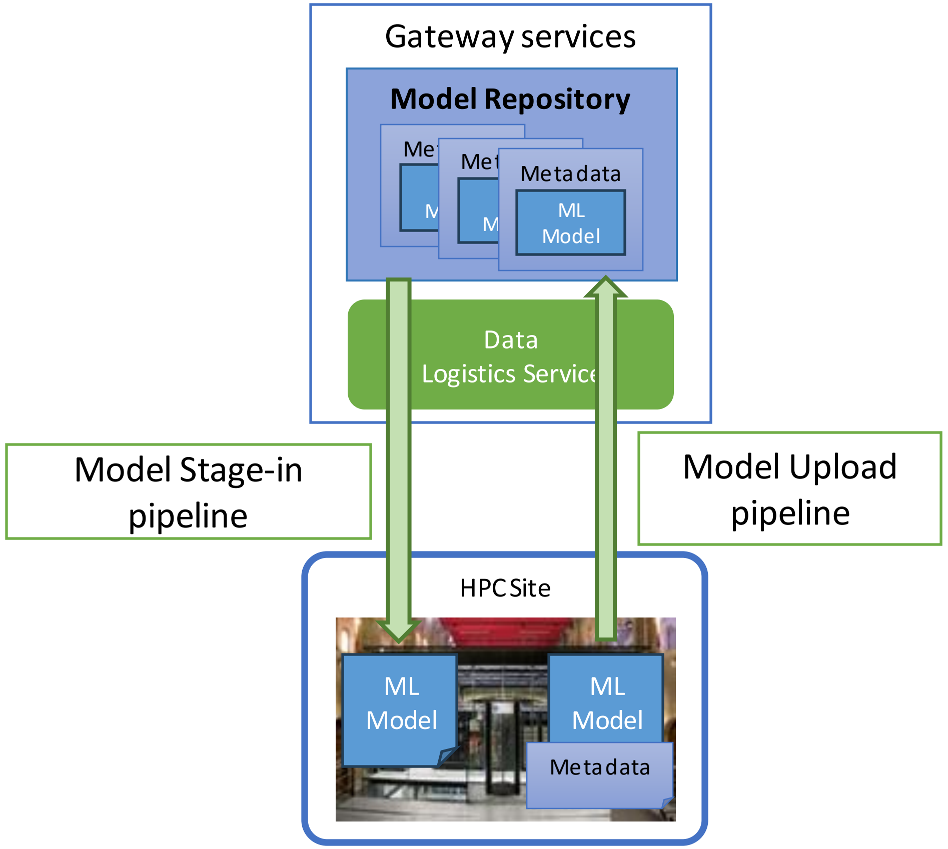

Data Provenance

- Support for automatic recording of Data Provenance from the execution of COMPSs applications:
  - To support **reproducibility** and **replicability** of COMPSs applications
- Approach: use RO-Crate 1.1 Specification
  - Designed a COMPSs RO-Crate structure from the standard, using as basis the Workflow RO-Crate profile
  - In the COMPSs runtime
    - A logger registers unique accesses to files and directories, to automatically identify inputs and outputs of the workflow, as well as all workflow activities
    - A post-process extracts the information to generate the RO-Crate (using ro-crate-py library)
  - Provenance data (ro-crate object) can be stored in WorkflowHub repository
- Reproducibility service: Reproduces a previously executed COMPSs experiment
- Knowledge graph visualizations, that enable users to explore, filter, and validate experiments interactively

Presentation roadmap

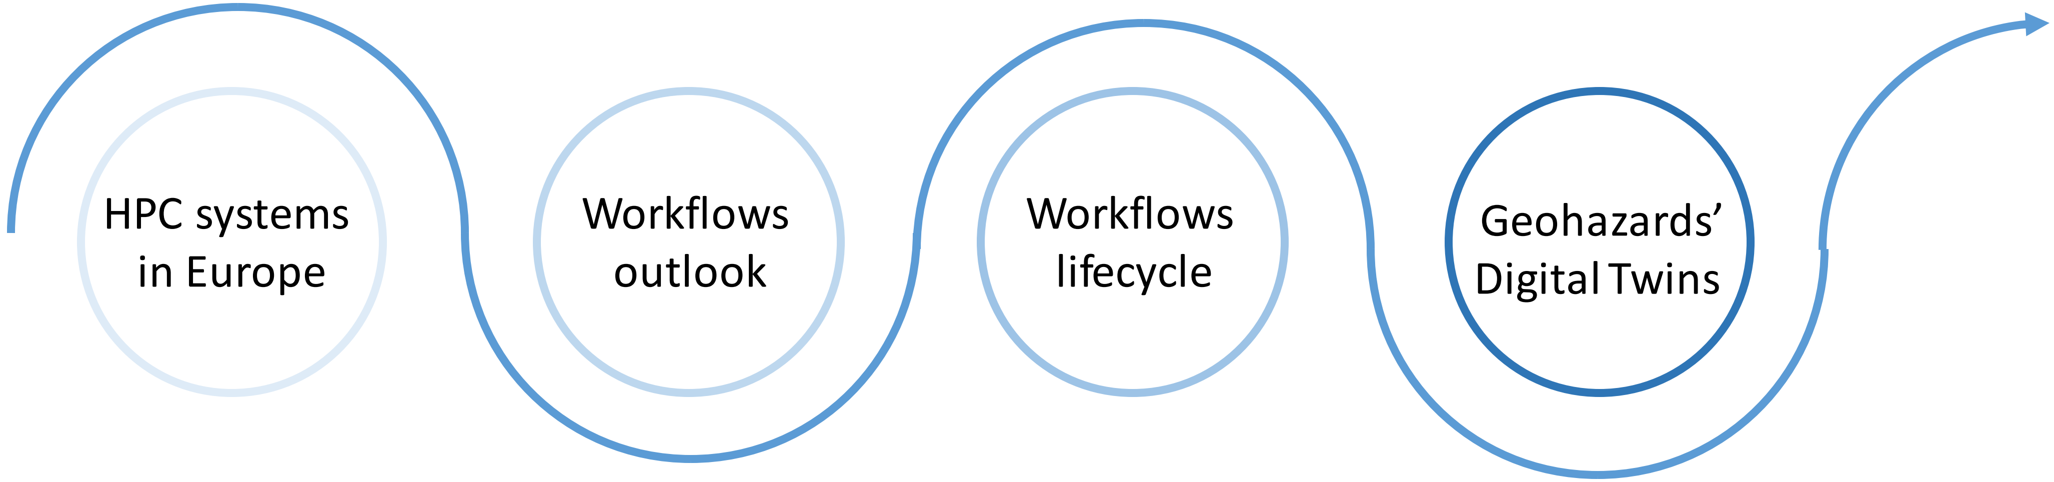

The geohazards community ecosystem

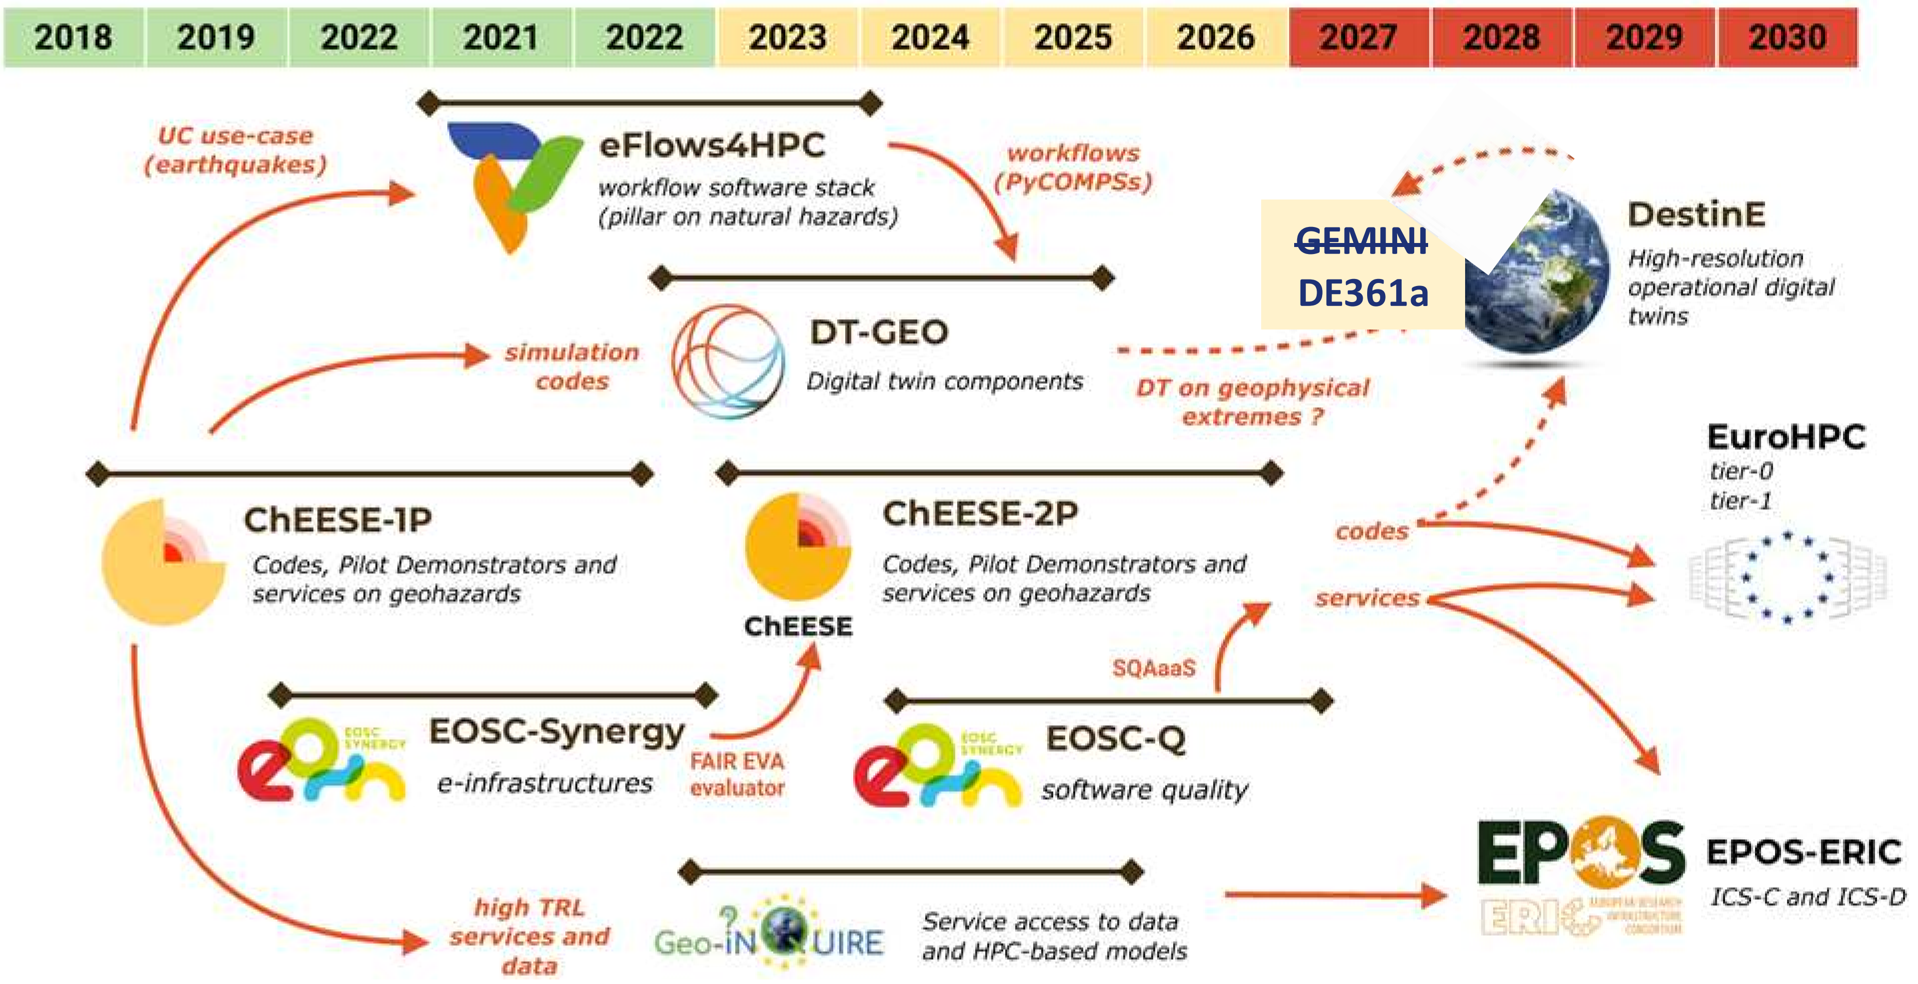

From simulations to workflows, from workflows to Digital Twins

DE361a: Geophysical Extremes Digital Twins Readiness Assessment for Destination Earth

- Assessing the readiness of the digital twins for geophysical hazards
  - Earthquakes, volcanic eruptions and tsunamis, geomagnetic storms
- Evaluate possible roadmap for integration in DestinE
- Assessment for each of the candidates w.r.t. existing mandates and associated operational services
- Explore the benefit for each candidate twin to run inside DestinE system.

Destination Earth

ECMWF

BSC   CSIC

R. M. Badia, J. de la Puente   A. Folch

ID	Subgrup ID	TRL (assessed)	Interest	Absence conflicts	TRL (revised)
DE_VA	DE_VA_1.1	7	yes	yes	7
	DE_VA_1.2	5	yes	yes	5
	DE_VA_1.3	4	yes	yes	4
	DE_VA_1.4	4	yes	yes	4
	DE_VA_1.5	4	yes	yes	4
	DE_VA_2.1	4-5	partial	yes	4-5
DE_VG	DE_VG_1.1	6	partial	yes	6
	DE_VG_1.2	5-6	yes	yes	5-6
	DE_VG_2.1	5-6	yes	yes	5-6
	DE_VG_2.2	4	yes	yes	4
	DE_VG_3.1	7	yes	yes	6 (*)
	DE_VG_3.2	7	yes	yes	6 (*)
DE_TS	DE_TS_1.1	7	yes	yes	7
	DE_TS_1.2	7	yes	yes	7
	DE_TS_1.3	7	yes	yes	7
DE_EP	DE_EP_1.1	7	yes	yes	7
	DE_EP_1.2	7	yes	yes	7
	DE_EP_1.3	6-7	yes	yes	6 (**)
	DE_EP_2.1				
	DE_EP_3.1	6-7	yes	yes	6 (**)

20 YEARS BSC

DTC-V2: volcanic ash clouds and ground accumulation

- Designed to automate the **forecasting of atmospheric dispersal and ground deposition** (fallout) of **volcanic ash**
- Configures and runs an ensemble of **FALL3D model realisations**, assimilating in the model different sources of data
  - Ground-based or satellite-based observations
- Allows delivering both **deterministic** (e.g. ensemble mean) and **probabilistic** (e.g. fraction of ensemble members exceeding a given condition) **forecasts**, reflecting the inherent uncertainties of Eruption Source Parameters (ESP)
- Coarse grain 4 main steps:

1. Get atmospheric data from Numerical Weather Prediction (NWP) models, delivered by a third party
2. Get the Eruption Source Parameters (ESP) from different sources
3. Setup FALL3D model and run an ensemble of simulations in the FENIX RI or HPC (leonardo@CINECA)
   - ✓ **PyCOMPSs** orchestrates the ensembles in HPC
   - ✓ **FALL3D** containers
4. Post-process the results

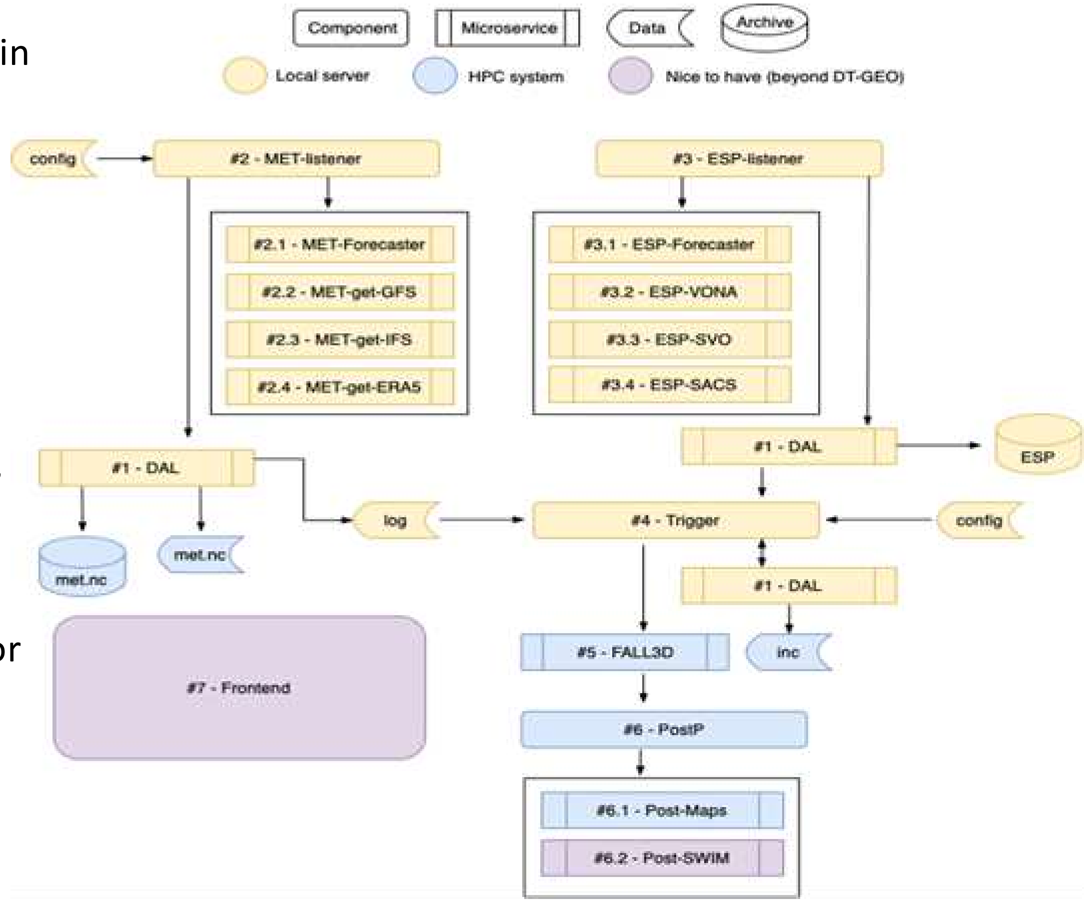

GET-IT project

GET—IT Geohazards Early Digital Twin Component

- DTC-V2
  - + assimilation of Copernicus data – Sentinel 1A and Sentinel 1B
- Result:
  - https://dea.destine.eu/web/stories/viewer/6809f951796e1d677c826e73

20 BSC YEARS

DEA User stories – GET IT

- https://dea.destine.eu/web/stories/viewer/6809f951796e1d677c826e73

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

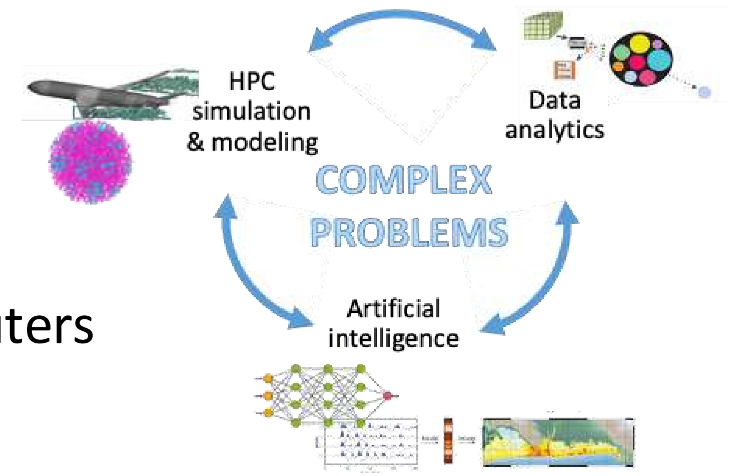

- Europe has managed to deploy a collection of first-class supercomputers through the EuroHPC JU initiative
- Computational workflows are mainstream and at the backbone of scientific discovery
- Workflows enable to tackle grand challenges, integrate diverse infrastructures, and accelerate innovation across domains
- Scientific applications require complex workflows that combine HPC, AI and High-performance data analytics
- Earth Observation data is an important component of domain applications such as digital twins for natural hazards
- Methodologies to standardize and simplify workflow lifecycle management are needed

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum 5

rosa.m.badia@bsc.es

ACKs

MINISTERIO
DE CIENCIA
E INNOVACIÓN

Financiado por
la Unión Europea
NextGenerationEU

Plan de Recuperación,
Transformación y
Resiliencia

AGENCIA
ESTATAL DE
INVESTIGACIÓN

Further Information

- Project page: http://www.bsc.es/compss
    - Documentation
    - Virtual Appliance for testing & sample applications
    - Tutorials
- Source Code

  https://github.com/bsc-wdc/compss
- Docker Image

  https://hub.docker.com/r/compss/compss
- Applications

  https://github.com/bsc-wdc/apps

  https://github.com/bsc-wdc/dislib
- Dislib
    - https://dislib.readthedocs.io/en/latest/

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Thanks!

rosa.m.badia@bsc.es