Destination Earth -DestinE

HPC-IODC: HPC I/O in the Data Center Workshop

July 2, 2021

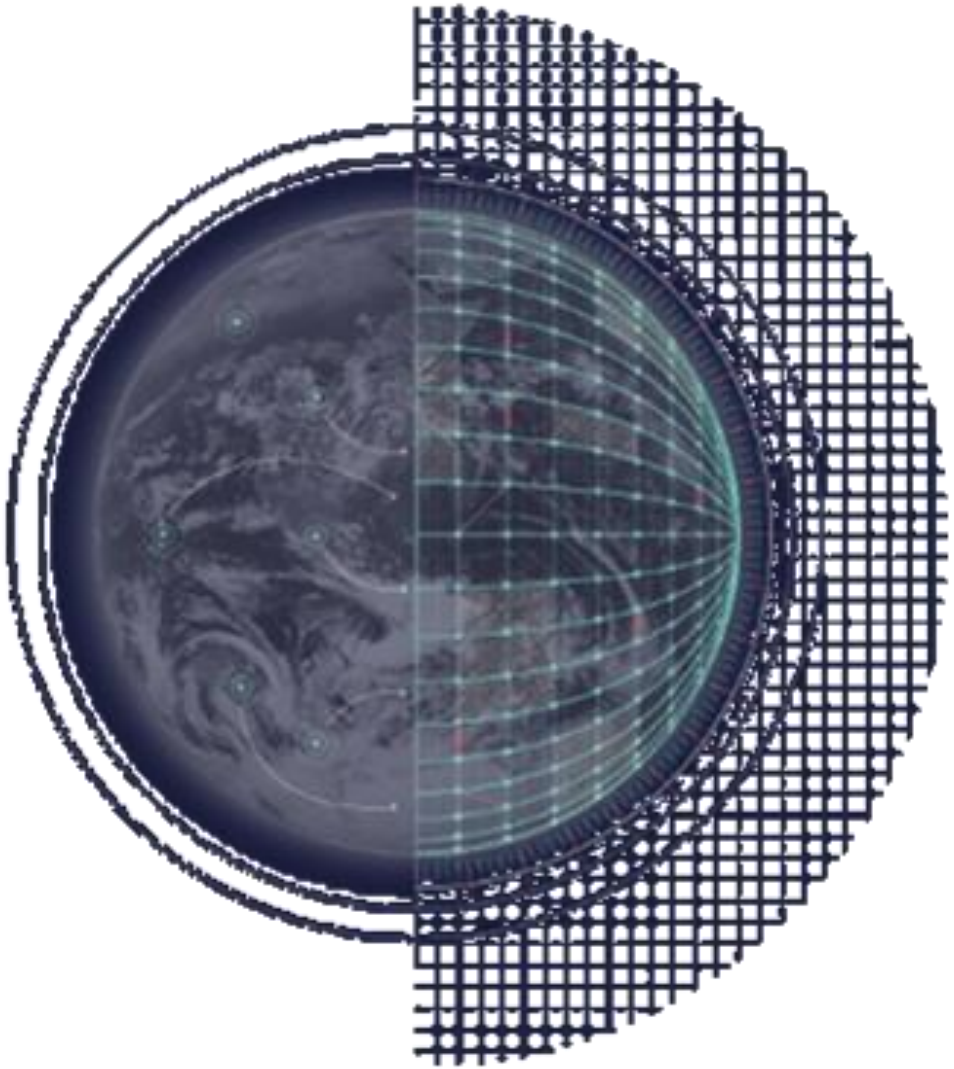

Vasileios A. Baousis (PhD)

ECMWF

Agenda

- Destination Earth (DestinE) initiative objectives and partnership
- About ECMWF : Organization, HPC, Cloud computing
- DestinE components, data and technological challenges.

*The objective of the **Destination Earth initiative** is to develop a very high precision digital model of the Earth to monitor and simulate natural and human activity, and to develop and test scenarios that would enable more sustainable development and support European environmental policies*.*

Users of DestinE will be able to access vast amounts of natural and socio-economic information in order to:

- continuously monitor the health of the planet
- support EU policy-making and implementation
- perform high precision, dynamic simulations of the Earth's natural systems
- improve modelling and predictive capacities
- reinforce Europe's industrial and technological capabilities in simulation, modelling, predictive data analytics and artificial intelligence (AI), as well as high performance computing.

* https://digital-strategy.ec.europa.eu/en/policies/destination-earth

Overall timeline of DestinE

- **Lead and business owner: European Commission (DG CONNECT)**
- **Strategic partnership** with ESA-ECMWF-EUMETSAT
- Responsibilities:
  - **ESA : key role of system integrator and implementer of the core platform**
  - ECMWF: Digital Twin implementer
  - EUMETSAT: responsible for the big data lakes and data integration
- Formal organization: "**contribution agreements**" by summer 2021

2021-2023

- **Operational cloud-based platform**
- **First two digital twins**

2023-2025

Platform integrates the next operational digital twins and offers services to public sector users

2025-2027+

Towards a full "digital twin of the Earth" through a convergence of multiple digital twins on the platform

DestinE partnership

European Commission

Governance incl. external advice

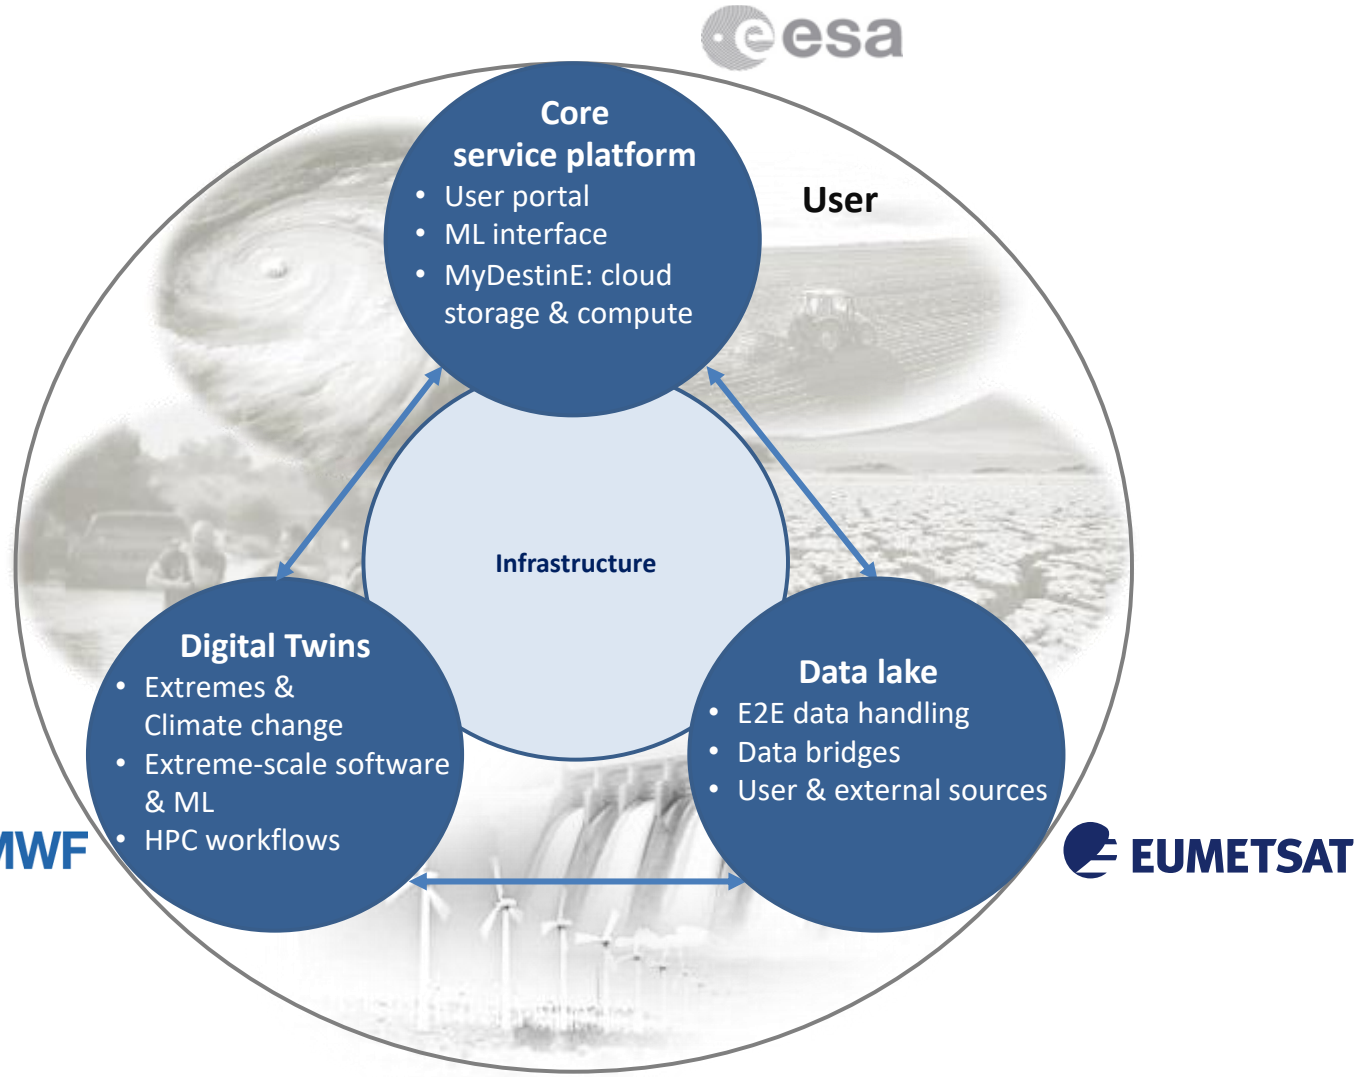

Partnership

DestinE will engage in continuous partnerships to co-evolve its components

- NMHS
- Science
- Technology
- Services
- Infrastructures

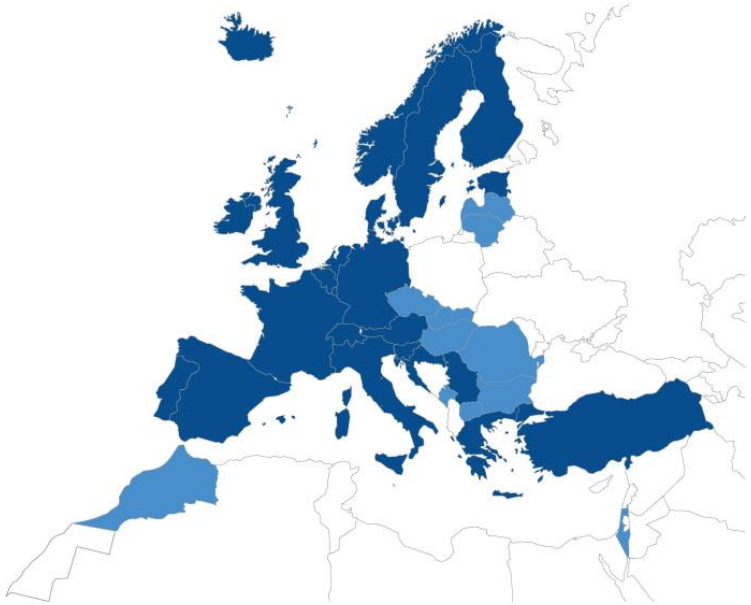

ECMWF

European Centre for Medium-Range Weather Forecasts (ECMWF)

Inter-governmental Organisation, Established in 1975

- 34 States (23 Member & 11 Co-operating States)
- Reading (UK), Bologna (Italy), Bonn (Germany)

Operational Numerical Weather Predication(NWP) centre

- Research institute and 24/7 operational service
- Twice daily generation of operational weather forecasts
- Assimilation of 60-80 million observations/day
- Archival of hundreds of Petabytes of observations and forecast data

IT infrastructure

- HPC facility is one of the largest global weather sites
- Cloud infrastructure for C3S, CAMS and WEkEO (DIAS) and the European Weather Cloud
- Climatological data : ~250 PB (daily growth of 250TB)

ECMWF

ECMWF's production workflow and European Weather Cloud

800 million observations/day

Data acquisition

Pre-processing

Forecast model run

Product generation

HPC

Data dissemination

Web Services

Internet
Leased lines
RMDCN

Public CSP

Cloud infrastructure

Storage

European Weather Cloud

Archive

Data Handing system

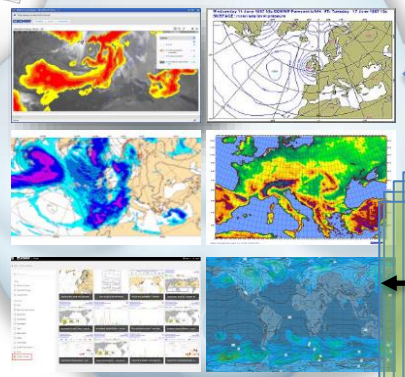

Frontend Backend

Applications

High level architecture

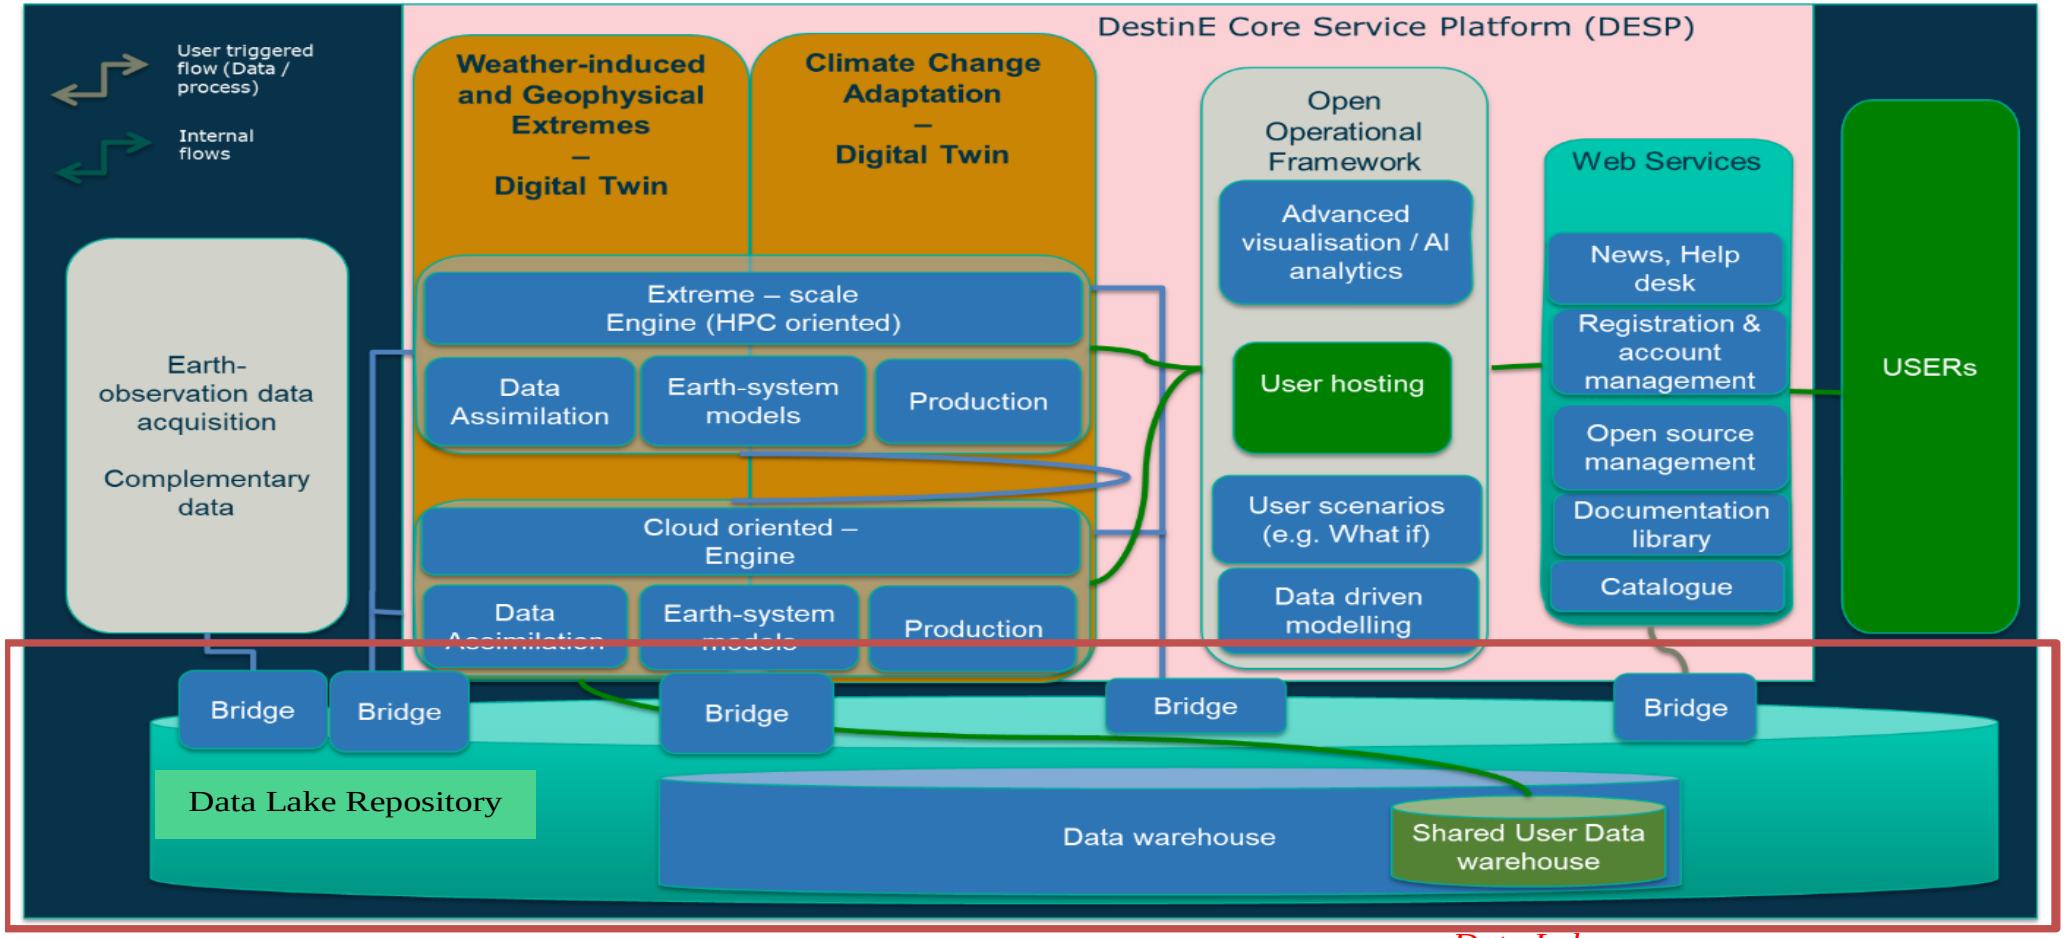

Digital Twin Engine (DTE)

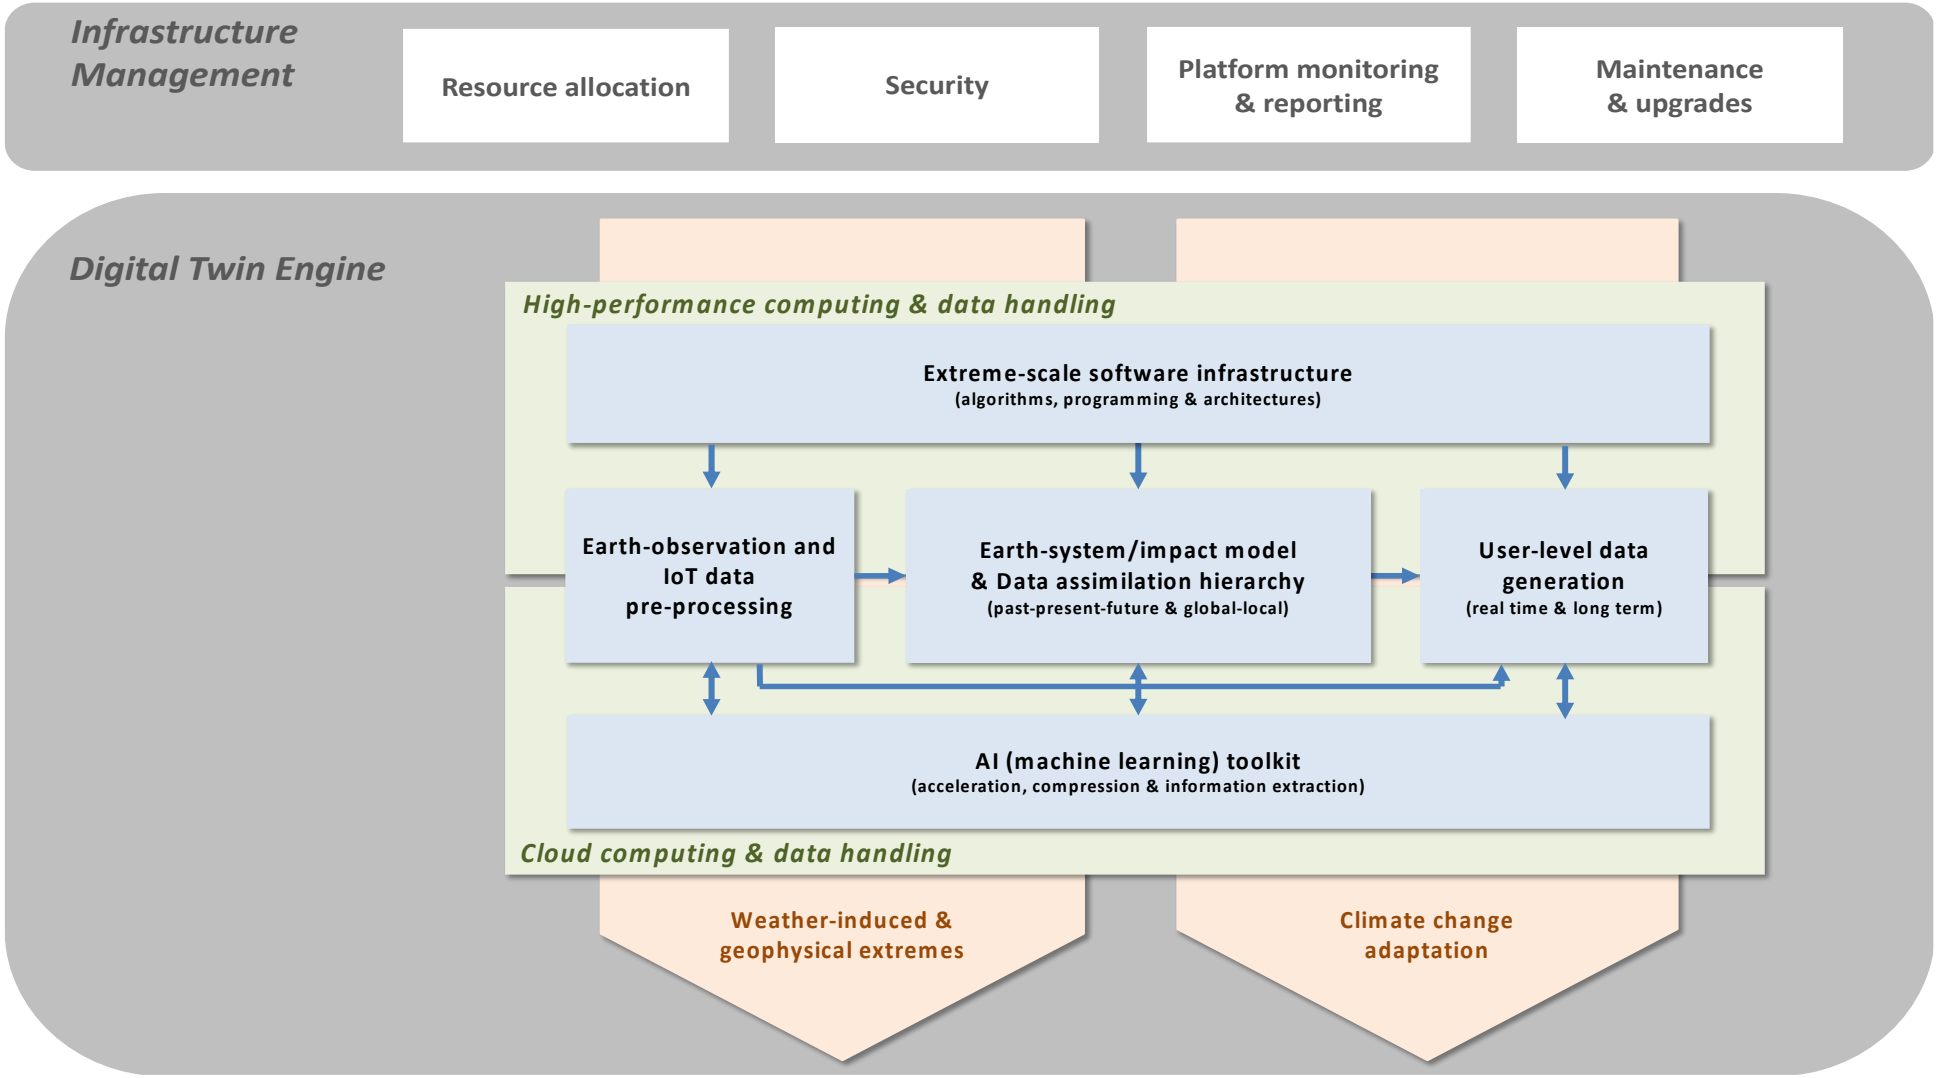

DestinE high level network organisation

- DestinE Core Service Platform (DESP)
- DestinE Data Lake (DEDL)
- Digital Twin Engine (DTE) & Digital Twins (DT)
- National Research/Education Network (NREN)

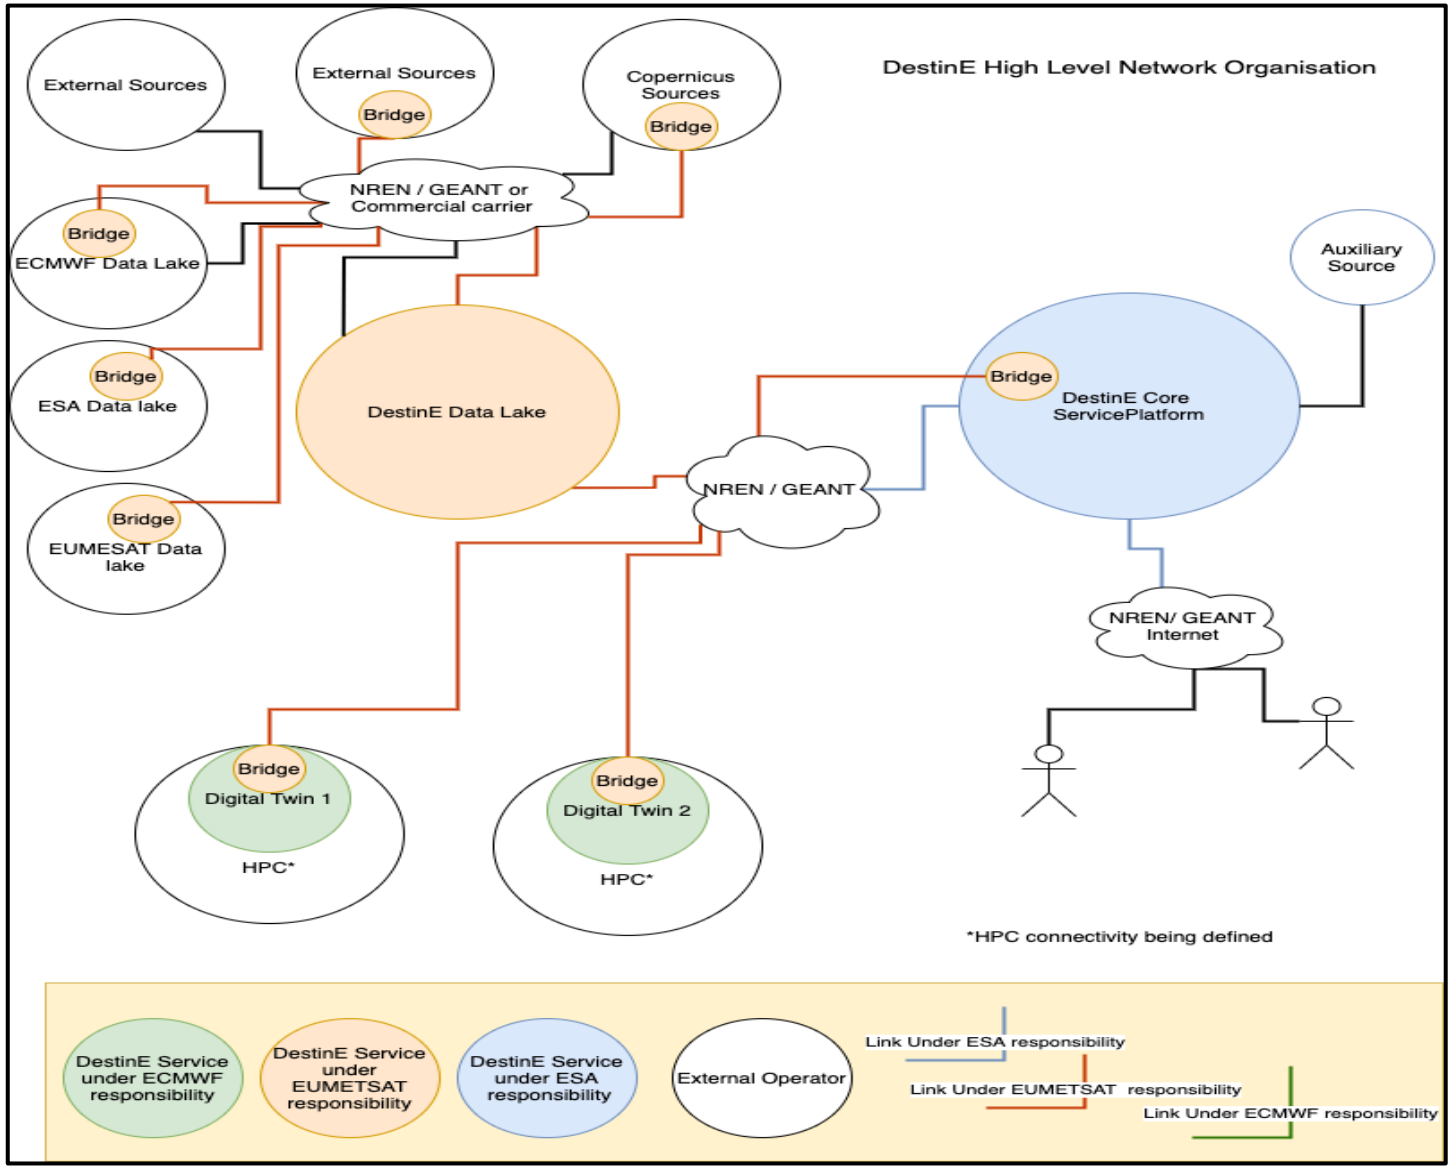

Digital Twins and Digital Twin Engine/Capabilities & Technologies

- Big data lifecycle management and governance (pertain mainly to DEDL but also important for DTE/DTs)
- Data access & data analytics tools and visualization
- Interactive, single/multi-site workflows (workflow scheduling/orchestrator with AI/ML; efficient resource allocation and load balancing)
- Earth observation analysis & Earth-system models
- Quality and uncertainty tracing
- Impact models

- HPC, Cloud, distributed storage
- Security essential to all functionalities
- Tools & systems maturity and operational readiness

Key challenges (1/2)

- **Input data:**
  - Volume : Millions of observations/day: data (Atmosphere, satellite, meteorological, user data, IoT, marine, climate change, environmental etc.); o**nly 5%* are currently used with a trend to be ~1%**
  - Data fusion (global-regional-local, past-present-future)
  - Quality (data selection, validation, uncertainty estimation, bias correction, error correlations),
  - Heterogeneity/Standardization, data interoperability improvement
  - Metadata management & automatic ways of labeling and describing data => Semantic annotation-common ontologies, vocabularies (DEDL but also necessary for the data fusion)
  - Security, privacy and integrity, and reliability (FAIR principles),

- **Output data (most of the above plus):**
  - Volume: can't be stored as before, needs recompute options/new approaches for Cloud/distributed environment
  - Should be managed flexibly in federated data infrastructures; mainly cloud environments – links to HPC
  - Data handling software & AI/ML
  - Data access methods -**Sovereign** data exchange (source identity & access rights)

* 5% of volume but much more of contained information

Key challenges (2/2)

- **Compute**
  - HPC and data handling resource allocation: reliable (continuous and peaks)
  - Observation handling, simulation, post-processing
  - Heterogeneous processor/memory/interconnect hardware => highly configurable/portable software
- **Data access, data analytics tools and visualization**
  - How to transfer (DTE ➔ DEDL)? From DEDL➔DESP) and visualize vast amount of data ?

	Weather-induced and geophysical extremes /day	Climate change adaptation (total for CNT+TRN+ENS)
Computing footprint	**8,500 node-hours**	**30,700,000 node-hours (~4 times more than WI DT)**
Data footprint	3-d atmospheric fields: 100 TBytes 2-d surface fields: 12 TBytes 3-d ocean/sea-ice fields: 1 Tbytes **Total : 113 TB/day**	3-d atmospheric fields: 24 PBytes 2-d surface fields: 4.8 PBytes 3-d ocean/sea-ice fields: 4.8 PBytes Total: 33.6 PBytes

DT Weather Data (GBytes)	113000				
DT Weather Data (GBits)	904000				
Tranfer Line capacity (Gbit/s)	100	**80**	**70**	50	10
Time (Secs)	9040	**11300**	**12914.29**	18080	90400
Time (Min)	150.6666667	**188.33333**	**215.2381**	301.3333	1506.667
Time (Hours)	2.511111111	**3.1388889**	**3.587302**	5.022222	25.11111
Days	0.10462963	**0.130787**	**0.149471**	0.209259	1.046296

DT Climate Data (GB)	33600000				
DT Climate Data (Gbits)	268800000				
Tranfer Line capacity (Gbit/s)	100	**80**	**70**	50	10
Time (Secs)	2688000	**3360000**	**3840000**	5376000	26880000
Time (Min)	44800	**56000**	**64000**	89600	448000
Time (Hours)	746.6666667	**933.33333**	**1066.667**	1493.333	7466.667
Days	31.11111111	**38.888889**	**44.44444**	62.22222	311.1111

- **Operationalisation**
  - Transition from Research to Operations (documentation, training, monitoring, process/task definitions)

DestinE at a glance (1/2)

- Core technologies : HPC, Cloud, Distributed storage.
- **Coupling of Cloud and HPC (HPC to the Cloud or Cloud to HPC?)**
- **HPC**
- More than one HPC Centre (e.g. EuroHPC)
  - Portable execution environment of the DTs and some parts of DTE; considerable workload envisaged for the DTs
  - Pre-/post-processing, data assimilation, predictions/projections, DT user data generation
- **Cloud**
  - Host DTE functionalities (technologies like SDN, SDS, federation, containerization & microservices, K8 etc)
  - Management of the DTE (pre-post processing, monitoring, orchestration, workflow management, interfaces/interconnections with DEDL & DESP
  - Data staging between DT and DestinE components (data fusion, compression etc)
  - Mainly IaaS and partially PaaS, and container & K8 execution environment.
  - Execution environment for less-demanding applications (more appropriate for the DESP Cloud environment)
    - Model as-a-Resource , Model As a Service, Model as Multi-functional system)

DestinE at a glance (2/2)

- **Distributed storage**
  - Software Defined Storage (SDS) like Ceph
  - DEDL Bridges and DEDL core will be based on SDS

- **Data workflows** need to process, transfer, store and retrieve vast amount of diverse data
  - post-processing and compression. (Earth-observation and IoT data pre-processing, Earth-system/impact model and data assimilation, User-level data)
    - DTE from/to Data Lake (cloud-based implementation)
    - DTE to/from Core Component (cloud-based implementation)

- **DT Continuous and on-demand production modes:**
  - Monitoring and utilization are important for workflow scheduling (who, what, when, where, why, how)
  - ML for workflow/load management (completely different scale of Load Balancing tasks)

Destination Earth -DestinE

HPC-IODC: HPC I/O in the Data Center Workshop

July 2, 2021

Q/A

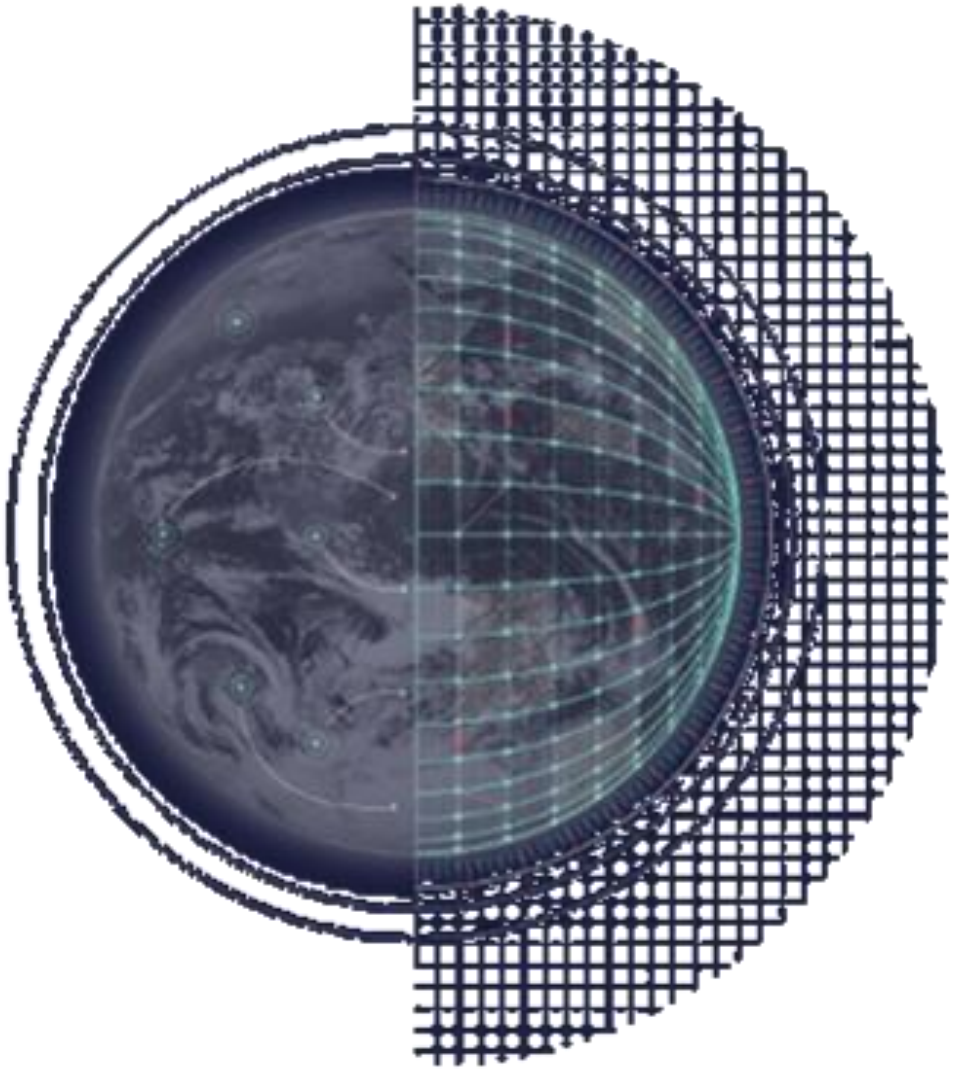

Vasileios A. Baousis (PhD)

ECMWF