


ECMWF Feature article

from Newsletter Number 171 – Spring 2022

EARTH SYSTEM SCIENCE

From the Scalability Programme to Destination Earth



www.ecmwf.int/en/about/media-centre/media-resources

doi: 10.21957/pb2vnp59ks

This article appeared in the Earth system science section of ECMWF Newsletter No. 171 – Spring 2022, pp. 15–22.

From the Scalability Programme to Destination Earth

Peter Bauer, Tiago Quintino, Nils Wedi

This article recalls the origins of ECMWF's Scalability Programme eight years ago and how it helped to prepare ECMWF for the EU-funded Destination Earth initiative, which emerged through a considerable effort from some of the leading scientists in the field of weather and climate prediction and computational science. The Scalability Programme was designed to adapt weather prediction codes to emerging computing paradigms, and Destination Earth is an initiative to develop a highly accurate digital twin of our planet. Both initiatives illustrate how fruitful international collaboration can be, but they also highlight the challenges for devising and implementing large and ambitious programmes supporting the digital revolution of Earth system science.

The scalability challenge

ECMWF launched its Scalability Programme in 2014 as a result of the realisation that the efficiency of our substantial codes and complex workflows would not be future-proof in view of new, emerging computing paradigms. In particular, CPU processor performance was being capped to reduce power consumption. At the same time, substantial investments in software were needed so that emerging, alternative processor technologies like GPUs could be explored. GPUs have been around in high-performance computing (HPC) since 2006, but they presented serious code adaptation challenges that limited their use as a general-purpose computing device and delayed their wider uptake in operational weather and climate prediction. A notable exception was MeteoSwiss, which invested in a substantial, manual GPU code rewrite early and operationalised its limited-area prediction system on CPUs and GPUs in 2016.

It also became clear that enhanced computing requires enhanced data handling capabilities, as big computing creates big data. The classical approach of storing all model output on tapes and performing frequent read and write operations was going to change radically in the future. The data handling challenge appeared to be more critical at the model output than at the observational data input stage as model output increases at least by an order of magnitude when resolution is doubled.

Advances in research and operational prediction lead to a doubling of horizontal resolution about every eight years. Upgrades in vertical resolution and the use of ensembles in assimilation and forecasts add to the multiplication of the computing and data burden. With no sign of such demands reaching saturation, there was increased pressure to make a dedicated investment in scalability.

In the past, continuous efforts were made to exploit shared and distributed processor memory in ECMWF's Integrated Forecasting System (IFS). However, a rethink of algorithmic choices and computational efficiency became necessary because of a number of developments: the complex setup of semi-implicit spectral and grid-point calculations; the semi-Lagrangian advective transport and its necessary data communication; and the coupling to increasingly complex Earth-system-component models for waves, ocean, sea ice and land, based on different grids, solvers and computational patterns. Any improvement also needed to be sustainable and thus portable to emerging and future computing architectures without repeatedly redesigning models, algorithms and work- and dataflows.

Scalability Programme – the research phase

The Scalability Programme was launched as a concerted effort across ECMWF and its Member States. Contrary to most projects run at the same time at most weather services and climate prediction centres, the ECMWF programme targeted both efficiency and sustainability of the entire work- and dataflow.

Initially, the programme was formulated around several projects focusing on observational data ingestion and pre-processing, data assimilation and the forecast model, model output data handling and product generation.

Finding resources

Once initiated, the programme needed resources. Our Member States were very supportive as performance and technology concerns were widely shared across the weather and climate community. The ECMWF Council agreed to add five staff members in 2014. The first generation of scalability staff came mostly from diverse, non-meteorological backgrounds. This delivered multi-disciplinary innovation from other fields that also rely on strong computational science expertise, such as acoustics, aerospace and fluid dynamics. It became a template that has worked very successfully ever since.

Another pillar of collaboration was created through externally funded projects supported by the European Commission's Future and Emerging Technologies for HPC (FET-HPC) programme. This allowed us to draw in the expertise of likeminded weather prediction centres, academia like the nearby University of Oxford, HPC centres, and computer technology providers like NVIDIA and Atos. It also helped to create a new ecosystem where this expertise could accelerate progress in a concerted way and create community solutions that benefited more than one centre. ECMWF's membership in the European Technology Platform for HPC (ETP4HPC) allowed suitable partners to be found and requirements to be defined for the ETP4HPC strategic research agenda, which stimulated new ideas for the European Commission's funding programmes.

Following the broad scope of the Scalability Programme, the first generation of these European projects, called ESCAPE and NextGenIO, aimed to revisit basic numerical methods through so-called weather and climate dwarfs and adapting them to GPUs. They also explored different levels of the memory hierarchy, where new technologies offered alternative programming solutions and speed-ups to perform costly calculations, such as model output post-processing on the fly without the need for accessing spinning disks. Exploring processor-focused parallelism and memory-stack-focused task management together produced numerous options for efficiency gains. ECMWF partnered with several Member States, universities and leading vendors, such as NVIDIA, Atos, Cray, Fujitsu and Intel, to facilitate co-design with the latest available technology.

Since then, the dwarf concept has reached wide acceptance in the community. This is because it breaks down the overall complexity of full models but still represents realistic model components with intensive and diverse computational patterns that can be tested at scale. Always considering compute tasks and dataflow together created much more realistic solutions for real-life prediction systems.

ECMWF also participated in formulating the programme for the ESIWACE Centre of Excellence, which created a new hub for developments serving both weather and climate prediction, thus joining these communities for seeking common solutions to identical challenges. These external activities were paired with internal projects to align all efforts towards fulfilling ECMWF's operational needs.

Research focus

The next generation of projects continued and intensified the work on these topics, and they further extended partnerships. Many projects were formulated around new and emerging, general-purpose technology where ECMWF contributed one of several applications (MAESTRO, LEXIS, ACROSS, DEEP-SEA, IO-SEA, EUPEX, HiDALGO). A few projects had weather and climate prediction itself as the focus. A variety of new methods and technologies were trialed and assessed together (ESCAPE-2, ESIWACE-2, MAELSTROM). So far, ECMWF has been involved in 15 such projects drawing in over 10 million euros for developments at ECMWF since 2015 (Figure 1).

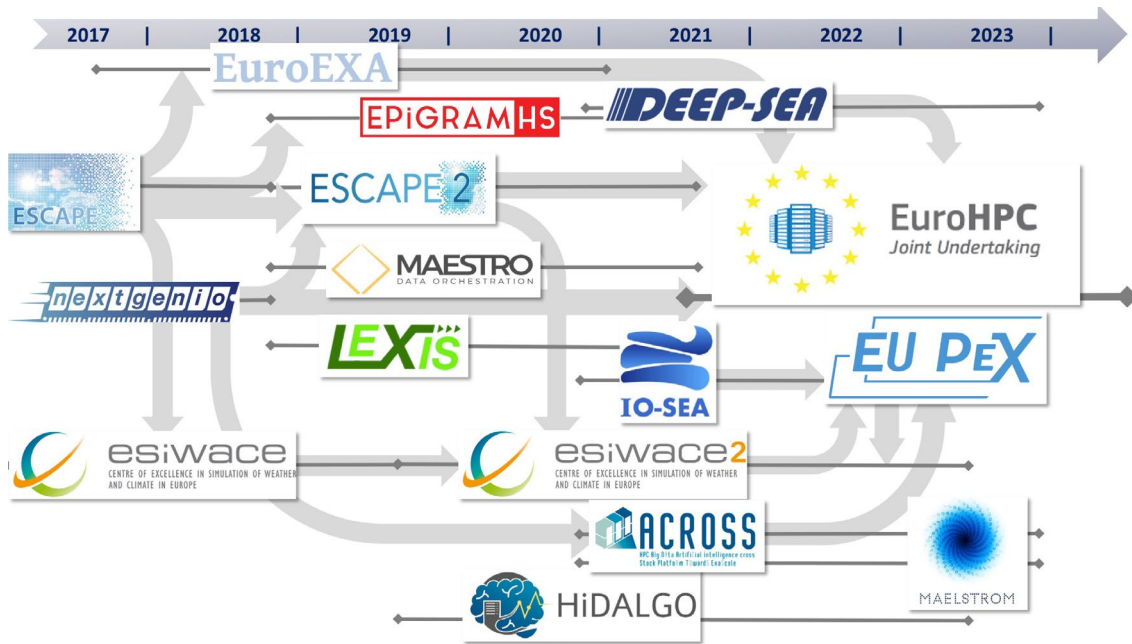


Figure 1 Many Scalability Programme projects with ECMWF involvement have been funded by the European Commission. The projects support basic developments for numerical methods and algorithms, adaptation to different hardware architectures, extreme-scale data handling, and community-serving centres of excellence. This created numerous collaboration opportunities with ECMWF Member States, academia, HPC centres and digital technology vendors. For more details, consult <https://www.ecmwf.int/en/research/projects>.

Another cornerstone has been the testing of the evolving IFS at scale on some of the largest HPC infrastructures in the world. While these tests were started before the Scalability Programme existed, they were carried out more frequently and at larger scale through the programme and also took in our scalability developments as they evolved. For example, we used what were then Europe’s fastest supercomputer, Piz Daint at the Swiss National Supercomputing Centre (CSCS, June 2017, top500), and the world’s fastest supercomputer, Summit (June 2018 to November 2019, top500). More recently, we carried out work on Europe’s latest, fastest machine, the JUWELS Booster (November 2021, top500), and on the world’s fastest machine, Fugaku (November 2021, top500). Access was granted through several competitive US INCITE awards and Europe’s PRACE programme, but also through bilateral agreements with e.g. Japan’s research institution RIKEN and the Swiss CSCS.

Over time, the complexity of these tests increased from shorter model integrations to seasonal and ensemble simulations at 1 km spatial resolution, while solving associated data challenges at these scales at the same time. The trials have been crucial for understanding new architectures and the remote deployment of workflows. They also helped to understand the limits of scalability, and they created unique datasets for gauging what science challenges and opportunities lie ahead, also in support of machine learning.

Scalability Programme – the implementation phase

With the completion of the first, research-oriented phase, ECMWF took stock of the results obtained so far, assessed the options for further development, and identified a roadmap for implementing research into operations. Some of the highlights include:

- establishing the quantitative limits for business-as-usual performance speed-ups, i.e. the baseline against which intrusive changes need to be compared
- developing the concept of dwarfs, tested on different processor technologies and optimised to the extreme
- preparing alternative formulations of model components, including the IFS-Finite Volume Model (FVM) dynamical core, which would overcome fundamental limitations of the spectral core
- developing and testing the new domain-specific language (DSL) concept for separating legible science from highly optimised code layers targeting different processor types, and automated source-to-source code transformation tools (Loki; GridTools for Python, GT4PY-FVM)

- first successful adaptation of dwarfs to more exotic dataflow processors, such as field-programmable gate arrays (FPGAs)
- testing and implementing mixed-precision developments (atmosphere and ocean models), which become increasingly effective due to machine-learning-targeted processor developments by industry
- developing ECMWF's new data structure library Atlas, which greatly facilitates operations on diverse grids and meshes, also for coupling different model components
- implementing the object-based datastore for workflow acceleration and product generation, ECMWF's Fields Database (FDB5), into operations
- full demonstration of model output post-processing for a 1 km ensemble based on FDB5, using the latest high-bandwidth, non-volatile memory extension technology
- developing the new IO-server supporting all Earth-system model data (MultIO)
- implementing a new, data-driven simulator for realistic workflows (Kronos), successfully used in ECMWF's recent HPC procurement.

The Scalability Programme also spun off a new ECMWF activity on machine learning that has established a roadmap for applications across the entire prediction workflow. This activity has created a new level of partnerships with Member States, academia and vendors and served as a template for similar activities run by other centres and agencies.

The rather wide scope of projects, the demonstration of successful developments, and the new levels of collaboration also planted the seeds for the ECMWF-Atos Center of Excellence in HPC, AI and Quantum computing for Weather & Climate. The centre of excellence has selected performance-portable codes and machine learning as the main topics for future co-development.

Road to operations

These outcomes were used to focus the second phase of the programme on the two main challenges for ECMWF's forecasting system: performance-portable codes, and data-centric workflows including distributed computing and data access models.

To turn the first task into operations, the new Hybrid2024 project was created. It aims to produce a hybrid CPU-GPU version of the IFS that is suitable for the post-Atos HPC procurement with a readiness date in 2024. Reduced precision is already operational at ECMWF today and has facilitated an increase in vertical resolution for the ensemble without requiring a new supercomputer. The OOPS data assimilation framework is entering operations in 2023. It provides a high-level control structure for the IFS, separates concerns of data assimilation and model, and makes the IFS more modular in preparation for Hybrid2024. Modularity and Member State support are also enshrined through Atlas. It is an open-source community library that serves as a common foundation to develop new algorithms, serves as possible data structures for DSLs, enhances Member State collaboration and interoperability between Earth-system model components, and is already used by a growing number of Member States and external partners. The Atlas developments helped to inspire ECMWF's operational triangular-cubic-octahedral (TCO) grid, and to support the rapid development of the IFS-FVM dynamical core.

The second focus is being established without creating a new project but by defining individual implementation targets for the key data-centric workflow components. These are FDB5 (already operational); the parallel I/O-server MultIO; the transition of the observational database to an FDB-type object store; and the connectivity to external, cloud-based infrastructures pioneered by the European Weather Cloud (EWC).

Increasingly, distributed computing and data handling are becoming necessary. This would ensure that even large model ensembles could be run externally (e.g. through containerised workflows) with dedicated cloud computing infrastructures. Data analysis would be large-scale too, therefore requiring dedicated computing resources close to the data and providing seamless access for users. This is part of ECMWF's future cloud computing strategy and includes the further development of the EWC and its convergence with the HPC infrastructure.

In addition, the second phase of the programme continues to invest in basic research on new numerical methods (finite volume and discontinuous Galerkin methods). New processor generations including non-x86, or non-GPU types like Fujitsu's A64FX, and emerging low-power processors of the European processor initiative (EPI) will likely require reduced precision by default and may be more suited to

execution models based on machine learning. Similar to the emergence of GPUs, these additional technology options also create challenges, and early access to prototypes combined with innovative, computer-science-driven thinking continues to be one of the leading driving forces for the Scalability Programme's next phase.

The Scalability Programme involved many individuals at ECMWF and greatly benefited from our collaboration with Member States, academia, HPC centres and vendors, which also created the foundation for the partnerships envisioned for Destination Earth.

Destination Earth and its origins

Destination Earth (DestinE) is funded by the European Commission's new Digital Europe programme. It was signed by DG-CNECT on behalf of the Commission and the European Space Agency (ESA), ECMWF and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) as entrusted entities on 15 December 2021. The programme supports the European Green Deal by delivering a new type of information system based on so-called digital twins of Earth. ECMWF produces the digital twin infrastructure and digital twins on weather and geophysical hazard-induced extremes and climate change adaptation, ESA contributes the core service platform, and EUMETSAT the data repository, called data lake. The first phase of the programme is 2.5 years long. It focuses on developing and demonstrating the functionality and novelty of different infrastructure and software components contributing to the wider DestinE digital environment. It will be followed by future phases in which the scope and operational level will be increased and used in a wide range of services.

Consolidating science and computing

DestinE is rooted in several initiatives, all of which ECMWF was involved in, sometimes in a leading way. The ground was laid in 2008 at the World Modelling Summit. T. Palmer (ECMWF at the time) and J. Shukla (George Mason University) were the driving forces behind the Summit, which was co-organised by Martin Miller (ECMWF). The event focused on climate prediction and concluded that a dedicated effort was needed to close apparent gaps in climate models, merge efforts across weather and climate communities and invest in dedicated computing resources to accelerate progress. The Summit was held at ECMWF, co-organised by the World Climate Research Programme (WCRP), the World Weather Research Programme (WWRP), and the International Geosphere-Biosphere Programme (IGBP). DestinE-type targets, such as 'kilometer-scale modelling' and 'regional adaptation and decision making', were already included in the Summit's report in 2009.

In the years following this event, the recommendations were turned into a proposal to create a new, centralised but international institution. This would create the critical mass for producing a much-required breakthrough for operationalising climate science production around a central piece of technological infrastructure, including dedicated HPC resources. It took more than five years and an opportunity created by the European Commission to provide this concept with a more solid underpinning.

The next step was an international workshop at Chicheley Hall in the UK in 2016, which gathered 21 of Europe's leading climate scientists to revise the original Summit's concept and incorporate the latest developments in Earth system science, numerical modelling and computing. The workshop was conceived by T. Palmer (then University of Oxford) and B. Stevens (Max Planck Society). The outcome was called the European Programme on Extreme Computing and Climate (EPECC). EPECC gave our efforts a sharper scientific focus on the topic of extremes and their origins, the recognition that resolving small scales would be essential for obtaining large-scale predictive skill, and the realisation that refining parametrizations would not overcome key model deficiencies.

EPECC was led by T. Palmer and was submitted in response to the European Commission's consultation on new ideas for so-called Flagship initiatives: large and ambitious research programmes aiming for moon-shot type targets and equipped with 1 billion euros over ten years. The consultation phase was followed by a call for preparatory projects for such Flagships. If successful, they would receive 1 million euros for one year to prepare the fully-fledged Flagship concept.

The preparatory project was coordinated by ECMWF because of our leading role in numerical weather prediction and our substantial investments in scalability, further emphasizing the synergy between Earth system and computational science. The fact that ECMWF is a Member State organisation that represents the successful process of pooling financial and intellectual resources in a single effort also helped, as EPECC was aiming to repeat this partnership concept in a different form. The proposal included a partnership with 18 institutions from weather services in our Member States, HPC centres, national research and public-serving institutions, and university labs.

Flagships received substantial attention, partly due to their high budget and visibility, and Earth-system and environmental change topics were likely to appear in several, parallel submissions. This is why EPECC was eventually merged with two other efforts, one led by UK Research and Innovation (UKRI) and the Helmholtz Association in Germany, and the other led by the University of Helsinki in Finland. The resulting project was called ExtremeEarth (<https://extremearth.eu/>) and had expanded from weather and climate to geophysical extremes. It also contained a stronger inclusion of observational infrastructures through these partnerships. However, it retained its focus on very high-resolution simulations of the fluid and solid Earth; advanced simulation–observation data fusion, enabling digital technologies like extreme-scale computing and data handling; and the creation of an ExtremeEarth Science Cloud, which was an important product of the merger process. The latter laid the foundation for the digital twin concept and aimed at a similar functionality as DestinE today. ExtremeEarth was supported by many organisations and scientists in Europe and received over 120 endorsements from organisations worldwide. Outside ECMWF, the core writing team comprised B. Stevens, W. Hazeleger (escience Centre, now Utrecht University), T. Jung (Helmholtz Association), T. Schulthess (CSCS), D. Giardini (ETH Zurich), P. Papale (INGV), J. Trampert (Utrecht University) and T. Palmer, and many others contributed.

ExtremeEarth passed the first round of reviews, coming second out of 33 proposals. One year later, however, the entire Flagship programme was terminated as Flagships were deemed ineffective tools for the new Horizon Europe programme. This meant a new programmatic home was needed.

Shaping Destination Earth

It took more than two years to disassemble ExtremeEarth and reassemble DestinE so that the original spirit was retained, but the changing boundary conditions were taken into account as well. An example of a changing boundary condition was that Brexit happened and the UK was no longer an EU member state and therefore not associated with the Digital Europe Programme, from which DestinE was funded. The European Commission also proposed that ESA and EUMETSAT partner with ECMWF so that their combined experience in running large programmes could be brought together. Another boundary condition was that Flagships as a sole instrument were discontinued and DestinE needed to be funded by a new programme that did not yet exist.

The thematic hook was given by the new European Green Deal announced in December 2019, for which DestinE would provide a unique information system in support of decision-making. Soon after, the European Commission also published a new digital strategy. This became the basis for DestinE's ambition to implement a highly flexible digital infrastructure hosting advanced simulations that are informed by observations. This infrastructure became labelled as the 'digital twin of Earth'. It took some time to exactly capture what the definition of such a digital twin is, and how it would be different from existing Earth system simulation and data assimilation systems that many meteorological services and the EU's Copernicus Earth observation programme already operate. However, the components of ExtremeEarth, in particular the ExtremeEarth Science Cloud, had paved the way for this definition.

Digital twins have a history in engineering, where digital models of a factory or aircraft help to design assets, but also help to optimise their production and enhance performance and reliability during operation. There is a clear distinction between a digital simulation, a digital shadow and a digital twin. The simulation only casts the real asset in a digital model, while the digital shadow uses real-world observations to inform the model. The twin goes one step further as the model informs actions in the real world. Thus the asset is modified as a result of decisions made based on the model, and the modifications become part of the model.

In our domain, weather prediction systems are already digital shadows as we use numerical models supported by hundreds of millions of observations to constrain the simulations and create analyses serving as initial conditions for forecasts. Making the step to a digital twin requires three major upgrades:

1. Enhancing the simulation and observation capabilities, such that the model becomes much more realistic and therefore creates a better global system providing reliable information at local scales where impacts are felt
2. Including and ultimately interacting with models and data from sectors that matter for society, namely food, water, energy, health and risk management
3. Creating all digital elements for an information system that allows intervention, such that users can create information that is optimized for their specific purpose, add their own data, build new digital components, and deploy them for decision-making.

(1) was the core message of EPECC while (2) and (3) emerged most strongly through ExtremeEarth. All predecessor projects made a strong point about extreme-scale computing and data handling being critical for success. Going further, ExtremeEarth proposed so-called extreme-scale laboratories as research melting pots where the computing and data challenges would be solved across disciplines.

DestinE replaced this by the notion of the digital twin engine, which is the extreme-scale technology heart of any Earth system digital twin. In the digital twin engine, different models and twins share software and infrastructure (e.g. a platform or software as a service), providing portable solutions between architectures and permitting scaling from small to very large applications. The engine supports interactive access to the full twin data stream in real time, where access window, modes of use and applications can be configured as required.

As can be easily seen, all this closely links to the development goals of the Scalability Programme. This link brings a well-established engineering approach into Earth system science to support how society adapts to change and mitigates the consequences of change based on the best available science empowered by the best available technology. Figure 2 provides an overview of the timeline of the Scalability Project and Destination Earth.



Figure 2 Timeline of events leading to DestinE. It moves from first concepts of a dedicated programme for climate prediction and extreme-scale computing and the foundation of the Scalability Programme to efforts to create a European-level funding framework for joining weather, climate and computational science within DestinE.

The first phase

The three main components of DestinE are the digital twin framework, the core service platform, and the data lake. They will be implemented by ECMWF, ESA and EUMETSAT, respectively. The first phase, December 2021 to June 2024, is dedicated to developing the individual digital technology components and demonstrating their individual functionality, but also how they work together as a system of systems.

The digital twin framework comprises the digital twin engine, which is the simulation and data fusion system. It makes it possible to run very advanced models informed by more and better observations, and it increasingly drives more applications across more impact sectors through the synergy of science-technology capabilities (see Figure 3). The initial focus on weather-induced extremes and climate change adaptation will be extended towards other themes later. The core service platform will be the main hub that users interface with and through which they gain access to all data and tools that DestinE will facilitate. The data lake will be a repository providing access to digital twin data. It will also give access to a large product diversity that DestinE will generate and accumulate from other sources, such as Copernicus.

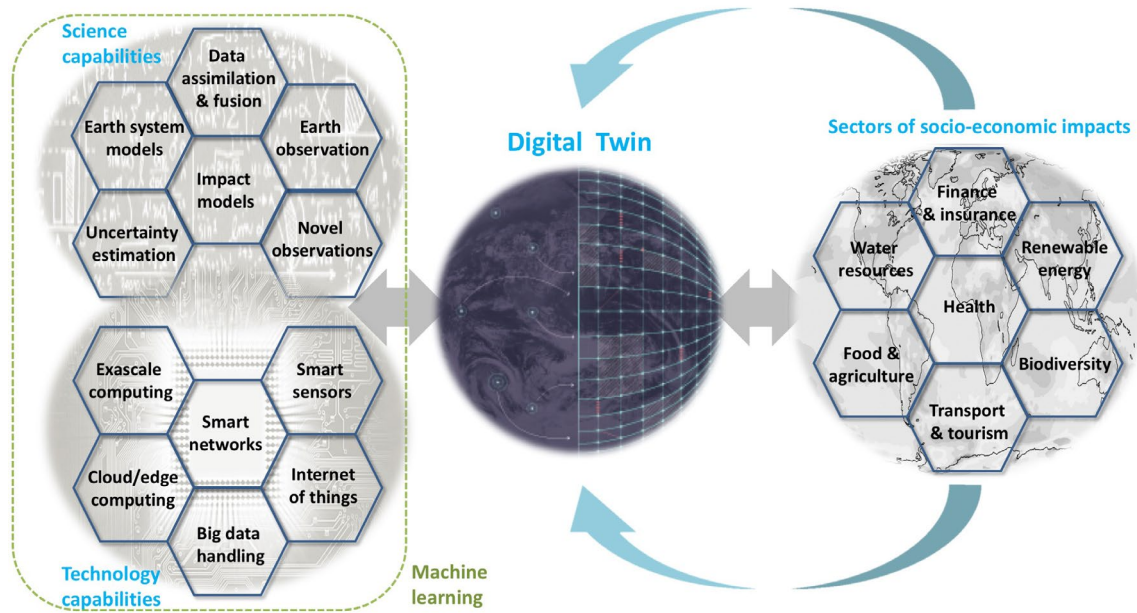


Figure 3 Digital twins are a tool to enable sectors of socio-economic importance to gain access to the best quality of data and to create tools to extract application-specific information. A digital twin builds on the convergence of leading science and technology capabilities based on simulations and observations through exploiting the entire range of digital technology. It is supported by machine-learning methodologies across all components.

The European Commission has allocated a total of 150 million euros to the first phase, with individual shares of 60, 55 and 35 million euros for ECMWF, ESA and EUMETSAT, respectively. DestinE will therefore benefit the Member States of the three organisations involved in DestinE in several ways. First, all three entities are intergovernmental organisations and have been cooperating with their Member States for decades. This will ensure that DestinE will produce direct benefits that many countries can share in terms of data, software and science by default. Second, the procurements offer a wide range of funded collaborations for creating the DestinE digital technology capabilities, which Member States can use themselves for their own benefit. Third, DestinE will create a substantial partnership programme that ensures that its aims and objectives will co-evolve with existing service providers, such as Copernicus and national meteorological and hydrological services. This is further strengthened by new funding opportunities at European (Horizon Europe) and national levels, which will support DestinE with basic research and infrastructures so that a continuous innovation cycle is created.

DestinE aims to establish its role as a unique accelerator and capability provider that sits at the interface between science, applications and technology. It feeds innovation into services, and it ingests specialist research, development and application partnerships. ECMWF’s expertise in Earth system modelling, the assimilation of vast amounts of diverse observational data into models, and its leading role in Copernicus service delivery, combined with its decade-long investments in extreme-scale computing, data handling and machine learning through the Scalability Programme, have therefore created a unique foundation to provide a new level of monitoring and prediction capabilities.

The Scalability Programme has positioned ECMWF to be a leader in the digital revolution sweeping through European climate and weather science, and in so doing it gives ECMWF and its Member States the opportunity to shape this revolution.

Further reading

- Bauer, P., T. Quintino, N. Wedi, A. Bonanni, M. Chrust, W. Deconinck et al.**, 2020: The ECMWF Scalability Programme: progress and plans, *ECMWF Technical Memorandum No. 857*.
- Bauer, P., B. Stevens & W. Hazeleger**, 2021: A digital twin of Earth for the green transition, *Nature Clim. Change*, **11**, 80–83.
- Bauer, P., P.D. Dueben, T. Hoefler, T. Quintino, T.C. Schulthess & N.P. Wedi**, 2021: The digital revolution of Earth-system science, *Nature Comp. Sci.*, **1**, 104–113.
- Deconinck, W., P. Bauer, M. Diamantakis, M. Hamrud, C. Kühnlein, P. Maciel et al.**, 2017: Atlas: A library for numerical weather prediction and climate modelling, *Comp. Phys. Comm.*, **220**, 188–204.
- Kühnlein, C., W. Deconinck, R. Klein, S. Malardel, Z.P. Piotrowski, P.K. Smolarkiewicz et al.**, 2019: FVM 1.0: A nonhydrostatic finite-volume dynamical core for the IFS, *Geosci. Model Dev.*, **12**, 651–676, <https://doi.org/10.5194/gmd-12-651-2019>.
- Müller, A., W. Deconinck, C. Kühnlein, G. Mengaldo, M. Lange, N. Wedi, P. Bauer et al.**, 2019: The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale. *Geosci. Model Dev.*, **12**, 4425–4441, <https://doi.org/10.5194/gmd-12-4425-2019>.
- Palmer, T.N.**, 2014: Climate forecasting: Build high-resolution global climate models, *Nature*, **515**, 338–339.
- Palmer, T.N. & B. Stevens**, 2019: The scientific challenge of understanding and estimating climate change. *Proc. Nat. Acad. Sci.*, **116** (49), <https://doi.org/10.1073/pnas.1906691116>.
- Schulthess, T.C., P. Bauer, N. Wedi, O. Fuhrer, T. Hoefler & C. Schär**, 2019: Reflecting on the goal and baseline for exascale computing: A roadmap based on weather and climate simulations, *Comp. Sci. Eng.*, **21**, 30–41, doi:10.1109/MCSE.2018.2888788.
- Shukla, J., T.N. Palmer, R. Hagedorn, B. Hoskins, J. Kinter, J. Marotzke et al.**, 2010: Toward a New Generation of World Climate Research and Computing Facilities, *Bull. Amer. Meteor. Soc.*, **91**, 1407–1412, <https://doi.org/10.1175/2010BAMS2900.1>.
- Smart, S.D., T. Quintino & B. Raoult**, 2019: A high-performance distributed object-store for exascale numerical weather prediction and climate. PASC '19: *Proceedings of the Platform for Advanced Scientific Computing Conference*, **No. 16**, 1–11, <https://doi.org/10.1145/3324989.3325726>.
- Wedi, N.P., I. Polichtchouk, P. Dueben, V.G. Anantharaj, P. Bauer, S. Boussetta et al.**, 2020: A baseline for global weather and climate simulations at 1 km resolution, *J. Adv. Model. Earth Sys.*, **12**, <https://doi.org/10.1029/2020MS002192>.

© Copyright 2022

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, UK

The content of this document, excluding images representing individuals, is available for use under a Creative Commons Attribution 4.0 International Public License. See the terms at <https://creativecommons.org/licenses/by/4.0/>. To request permission to use images representing individuals, please contact pressoffice@ecmwf.int.

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.