

Newsletter

No. 180 | Summer 2024

Destination Earth's digital twins

New ocean and sea-ice model

New ensemble reanalysis system for
ocean and sea ice

European Weather Cloud is
operational



© Copyright 2024

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.

Publication policy

The ECMWF Newsletter is published quarterly. Its purpose is to make users of ECMWF products, collaborators with ECMWF and the wider meteorological community aware of new developments at ECMWF and the use that can be made of ECMWF products. Most articles are prepared by staff at ECMWF, but articles are also welcome from people working elsewhere, especially those from Member States and Co-operating States.

The ECMWF Newsletter is not peer-reviewed.

Any queries about the content or distribution of the ECMWF Newsletter should be sent to Georg.Lentze@ecmwf.int

Guidance about submitting an article and the option to subscribe to email alerts for new Newsletters are available at www.ecmwf.int/en/about/media-centre/media-resources

The Destination Earth initiative



The EU's Destination Earth (DestinE) initiative, which aims to build a digital replica of our planet, has entered a new phase. ECMWF is one of three key players in DestinE, together with the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) and the European Space Agency (ESA). We are providing the first two high-priority digital twins of Earth and the Digital Twin Engine, which are at the heart of DestinE. The aim is to achieve highly accurate representations of climate developments and extreme weather-related events that can be used in decision-making, with unprecedented levels of interactivity and the possibility to run bespoke simulations to test 'what if' scenarios. To this end, we have involved 90 institutions across Europe, with a leading role played by our Member and Co-operating States in DestinE. Last month, the European Commission activated the initial DestinE system. In this second phase, our first two digital twins, ESA's core service platform, and EUMETSAT's data lake are becoming operational. DestinE is expected to continuously evolve by extending its operations and developing further components through co-design of applications with a wide range of users. By 2030, it should be able to make a full interactive digital replica of Earth available.

It is clear that DestinE is closely connected to our core activities: developing the best possible global numerical weather predictions and making them available to a wide range of users. This will enable DestinE to benefit from advances in weather prediction we develop in-house, including the increasing use of machine-learning techniques. But DestinE goes beyond that by focusing in on small regional scales, in the Weather-Induced Extremes Digital

Twin, and on long timescales, in the Climate Change Adaptation Digital Twin. A feature article in this Newsletter presents the current state of play of our contribution to DestinE.

The high-resolution global modelling used in DestinE relies on progress in our physical modelling systems. An example of this is our work on an updated ocean and sea-ice model. This will first be used to produce a new generation of ocean reanalysis, ORAS6, which will be used in our next reanalysis ERA6, planned to start production early next year. It will also be introduced into our weather forecasting operations in next year's upgrade to Cycle 50r1 of the Integrated Forecasting System (IFS). Details on this work can be found in two feature articles in this Newsletter. Another article sets out the operation of the European Weather Cloud (EWC): a community cloud computing platform jointly operated by ECMWF and EUMETSAT. The EWC is, for example, used by ECMWF for a joint machine learning project with Member and Co-operating States. One news article is a review of the ECMWF 'Weather Discussion', which celebrated its tenth anniversary last month. This weekly internal event focuses on improving forecasts for the benefit of all users. It covers a wide range of our activities, ranging from our operational and DestinE weather forecasting outputs to machine learning and Copernicus results.

Florence Rabier
Director-General

Contents

Editorial

The Destination Earth initiative 1

News

Unprecedented rainfall in the United Arab Emirates 2
 A new view of Earth's land surface and clouds with the IFS . 4
 ECMWF moves to final stages of open data transition 5
 Introducing JupyterHub at ECMWF 7
 A year in review of CESOC–ECMWF collaboration and initiatives 8
 A new way of handling observation errors for infrared measurements 9
 Tenth anniversary of the ECMWF 'Weather Discussion' . . . 11
 UEF2024 online attracts a global audience 12

Earth system science

Destination Earth's digital twins and Digital Twin Engine – state of play 14
 Introduction of a new ocean and sea-ice model based on NEMO4-SI³ 24
 ECMWF's next ensemble reanalysis system for ocean and sea ice: ORAS6 30

Computing

The European Weather Cloud is now operational 37

General

ECMWF publications 44
 ECMWF Calendar 2024 44
 Contact information 44

Unprecedented rainfall in the United Arab Emirates

Linus Magnusson, Tim Hewson, Matthieu Chevallier

On 16 April 2024, the United Arab Emirates (UAE) experienced unprecedented rainfall. The highest 24-hour rainfall recorded was 259.5 mm, reportedly setting a new all-time UAE record. Many stations exceeded 100 mm, e.g. Dubai Airport with 144 mm, corresponding to 1.5 years' worth of rain. The media reported chaotic scenes at the airport and across the region. In the days before, Oman was also hit by severe rainfall, with fatalities reported there.

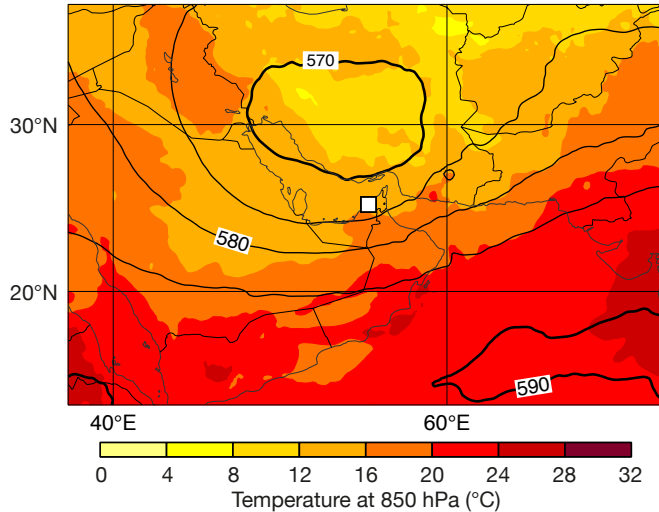
Weather conditions

On 14 April, a trough was centred over Iran, also with extreme rainfall over Oman on its southern flank. At the same time, a new trough developed over Türkiye and moved southeast. By the time of the peak of the storms in the UAE and nearby countries, it had developed into a cut-off low that provided the dynamical uplift that helped drive the main rainfall event (see the analysis chart). At lower levels, the flow into the Persian Gulf was from the southeast, providing moist air from the Indian Ocean.

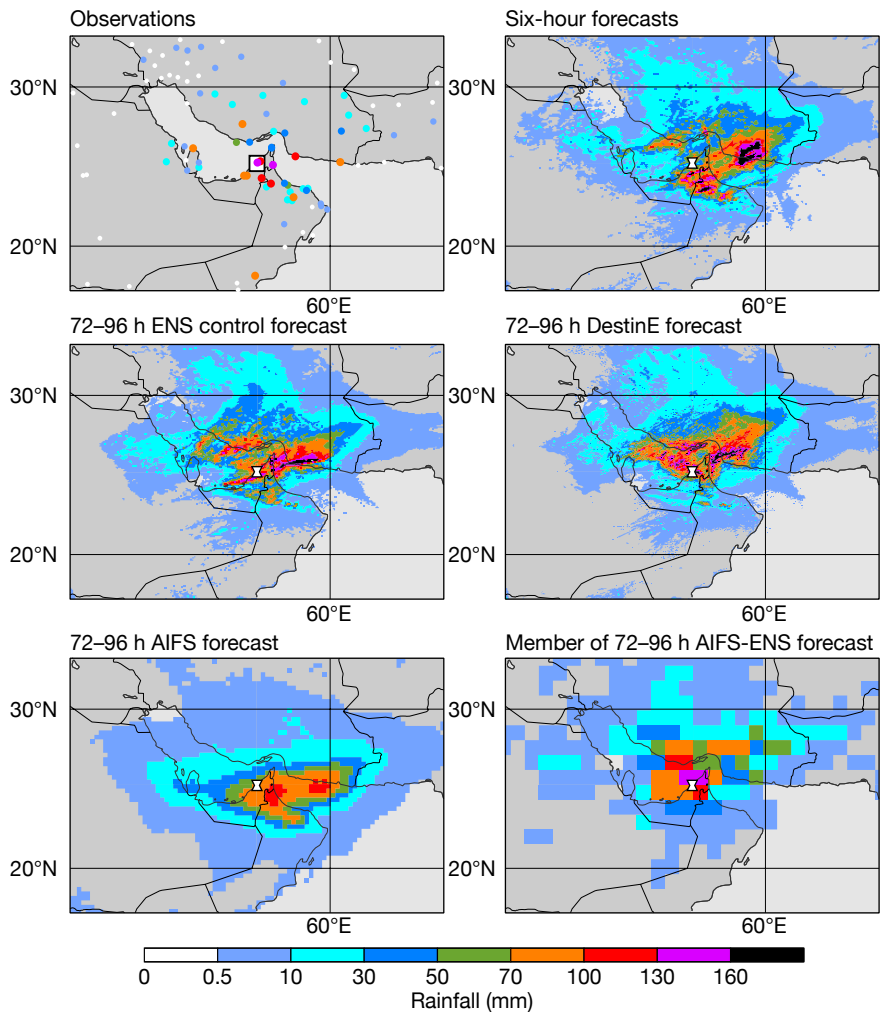
Here we focus on precipitation from 16 April 00 UTC to 17 April 00 UTC over Dubai. As a proxy for a precipitation analysis, we calculated 24-hour precipitation based on four concatenated six-hour forecasts from the ensemble (ENS) control forecast. We find that a wide region was affected by the rainfall, but with significant local variations in the magnitude, which is also hinted at in the sparse observations (see the top panels in the second figure). The analysis proxy for Dubai was 48 mm compared to 144 mm in the SYNOP weather station report from Dubai Airport, but there are values above 160 mm in the proxy analysis to the south. This local variability was probably due to the convective nature of the rainfall, with strong convective cells apparent on imagery.

Our forecasts

Looking at the forecasts issued on 13 April 00 UTC from the ENS 9 km control, 4.4 km forecasts produced for



Analysis chart. Analysis of 500 hPa geopotential height (contours, in decametres) and temperature at 850 hPa (shading) on 17 April 00 UTC. The location of Dubai is marked with a square.




Twenty-four-hour rainfall from 16 April 00 UTC to 17 April 00 UTC. The charts show observations (top left); concatenated 6-hour forecasts (top right); 72–96 h forecasts from the ENS control at 9 km (middle left), DestinE at 4.4 km (middle right), and the AIFS at 28 km (bottom left); and the first perturbed member from the AIFS-ENS at 111 km (bottom right).

the EU's Destination Earth (DestinE) initiative, and 28 km forecasts of the Artificial Intelligence Forecasting System (AIFS), we find that all forecasts predicted the rainfall region well. ENS control and DestinE had large variability around the UAE, causing large local deviations from the SYNOP observations. Since January, we have also been producing rainfall from the ECMWF machine-learning model, the AIFS. While producing a much smoother field, this model also predicted the region of the rainfall well, although it missed the extension to the northwest. The local maximum values are lower in the AIFS compared to the two versions of the Integrated Forecasting System (IFS), but the maximum in the region was still above 100 mm.

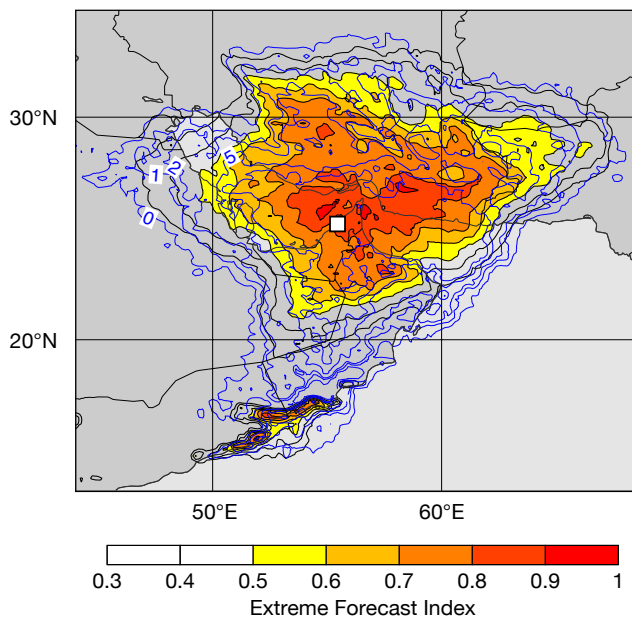
The forecast evolution plot for 24-hour rainfall over the gridbox including Dubai shows that even down to the shortest range (00–24 h) the uncertainty remained large, with IFS ensemble members ranging from 20 to almost 200 mm. It means that the observation from Dubai Airport was inside the ensemble range. The ensemble mean was around 80 mm, which is about four times the maximum in ECMWF re-forecasts (valid day 5) for this time of the year. While the ensemble highlighted large uncertainty, the ensemble mean was stable at around 70–80 mm/24 h from 9–10 April. This early signal can also be seen in the Extreme Forecast Index (EFI) from 10 April (6–7 days before the event), which already at this stage highlighted the risk of a rainfall extreme in the region (see the last figure). Going further backward in time, we see that the ensemble mean was consistently at or above the re-forecast maximum from 6 April onwards, i.e. from 10 days before the event.

All AIFS deterministic forecasts from 9 April 12 UTC onwards fell into the 25–75 percentile range of the IFS ensemble (except one that was slightly above 75th). For longer forecasts, the AIFS precipitation tended to be lower than most of the ensemble members, meaning that this model picked up the signal of the extreme 1–2 days later than the ensemble.

Since June 2024, ECMWF has also been running a real-time ensemble version of the AIFS. This first version has the output on a 1x1 degree grid, is using the same initial perturbations as



Forecast evolution plot for Dubai. The plot shows the evolution of 24-hour rainfall forecasts on 16 April over the grid box including Dubai. The model climate is about zero precipitation, with a maximum of less than 25 mm based on 1,800 forecasts (marked by the black triangle).



EFI and SOT. The Extreme Forecast Index (EFI – shading) and Shift of Tails (SOT – contours) for 24-hour rainfall on 16 April in the forecast from 10 April 00 UTC.

the IFS ensemble, and is using a diffusion model approach to simulate model uncertainties (see <https://www.ecmwf.int/en/about/media-centre/aifs-blog/2024/enter-ensembles>).

While this ensemble will be described in the next ECMWF Newsletter, we provide here a first glimpse of some results of retrospective forecasts. Looking at the spatial pattern of the first perturbed member from the AIFS-ENS issued on 13 April (bottom-right of the second figure), we see extreme rainfall (> 70 mm/24 h) in the region, but the low resolution (111 km grid spacing) is very apparent. Despite the low resolution, we see that the ensemble mean for precipitation over Dubai is of a similar magnitude as the IFS ensemble mean from 12 April onwards (see the forecast evolution

plot). Before that date, the AIFS-ENS predicted less rainfall, and it captured the first signal about 1.5 days later than the IFS ensemble, in line with the deterministic AIFS.

Conclusion

In summary, this case showed extraordinary predictability on the synoptic scale, with a signal for extreme precipitation ten days before the event. This predictability was due to the large-scale forcing of the event. However, even in the shortest forecasts large uncertainties remained about the mesoscale distribution of the rainfall. This case also showed that the machine-learning-based AIFS forecast was able to predict the large-scale features of the event while lacking local details.

A new view of Earth’s land surface and clouds with the IFS

Josef Schrötle, Cristina Lupu, Philippe Lopez, Patrick Gillies, Cihan Sahin (all ECMWF), Alessio Canessa (Centro Nazionale per la Meteorologia e Climatologia Aerospaziale – CNMCA)

From Cycle 48r1 of ECMWF’s Integrated Forecasting System (IFS), implemented in June 2023, we have released simulated images for top-of-the-atmosphere visible reflectance operationally (top panel of the first figure). From this autumn, we will provide a version of these images based on visible reflectance from a range of wavelengths of reflected sunlight,

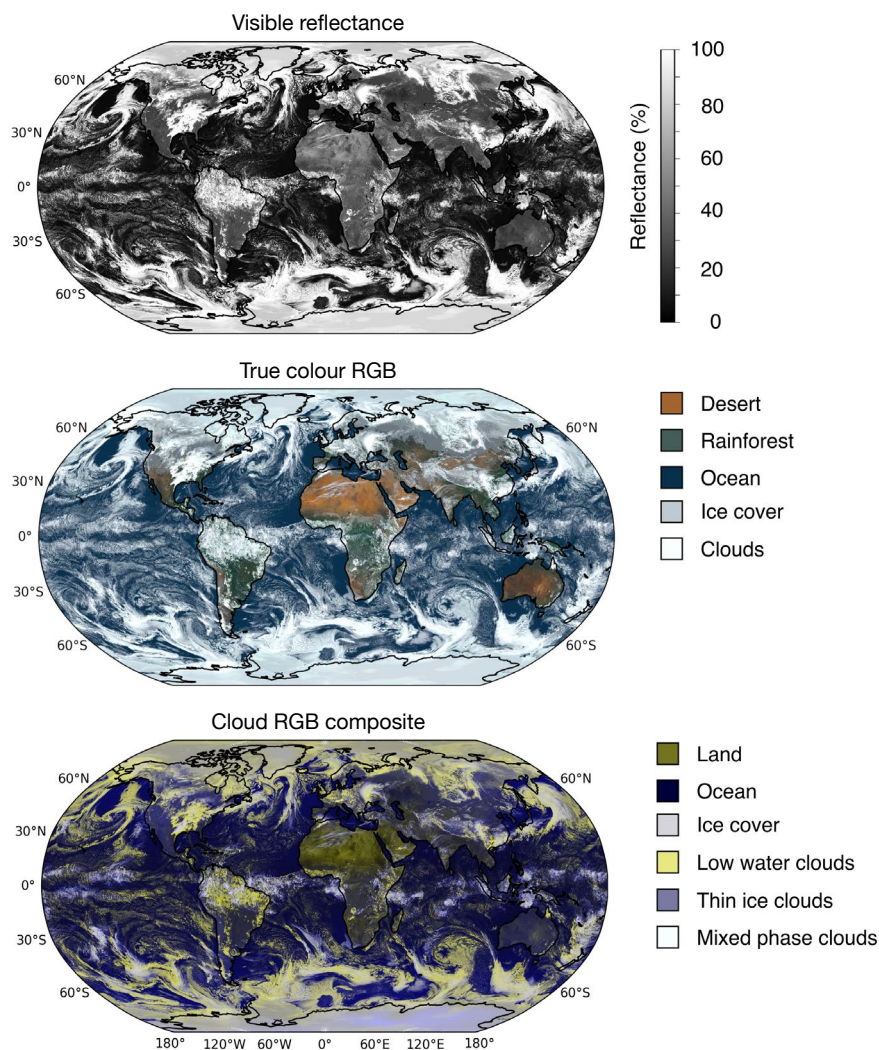
representative of the images to be provided by the new Flexible Combined Imager (FCI) on the MTG-I1 satellite. The simulated images in colour can provide a better contrast between the Earth’s surface and the clouds, and they can better reveal the multiscale structure of the clouds above the ocean surface. By including a broader range of visible wavelengths of reflected light, the

Saharan desert and deserts in Australia can be clearly distinguished from tropical rain forest coverage (middle panel). These images can also provide information on the cloud structure at different altitudes when including the infrared window channel (bottom panel). Forecasting clouds at the correct altitude and location is crucial for many applications, including aviation.

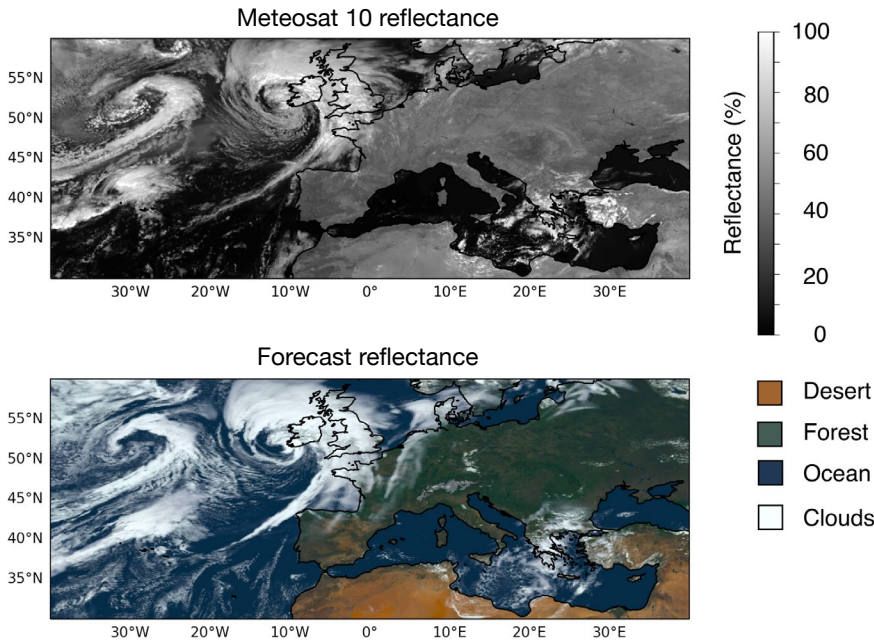
Visualising forecasts of vertical cloud structures

Clouds cover many altitudes, from the Earth’s boundary layer and stratiform clouds near coastal regions over the sea to thin cirrus clouds in the upper troposphere, where aircraft reach their cruising altitude. In the visible spectrum, water clouds, such as low stratus clouds near the west coast of South Africa, Chile, or California, can be easily discerned from the underlying surface. These clouds stand out from the underlying surface even though their respective temperature is comparable in the infrared spectrum. Further up in the troposphere, thin cirrus clouds can be almost transparent in the visible spectrum, but they are much colder than the underlying sea surface and can thus be clearly distinguished from the warmer ocean in the infrared window channel (10.5 micrometres). This means that visible and infrared wavelengths are very complementary in nature.

The information conveyed by these different wavelengths is used in a novel reflectance composite (bottom panel of the first figure). The yellow colour, which is a combination of the 640 and 865 nm bands, represents low-level clouds which strongly reflect solar light, while the intensity of the blue colour varies with the inverse of the normalised infrared brightness temperature in the 10.5 micrometre band. As a result, blue colours appear more intense for colder and thin cirrus clouds higher up in the troposphere, at the edge of very deep convective clouds. In regions with thick clouds (for example inside a tropical



ECMWF simulated satellite imagery. Global simulated imagery for a 10-day forecast initialised at 12 UTC on 14 April 2024 from the IFS Cycle 49r1 e-suite forecast shows visible reflectance in the 865 nm spectral band (top), true colour RGB created from reflectance factors in FCI channels 1 (444 nm, blue), 2 (510 nm, green), 3 (640 nm, red), and 4 (865 nm) (middle), and a cloud RGB composite created from reflectance factors in FCI channels 3 (640 nm), 4 (865 nm), and infrared 10.5 micrometres (bottom). All panels depict the same atmospheric state as simulated by the IFS.



Satellite reflectance and forecast reflectance. Two cyclones with strong precipitation, enhanced wind gusts and thunderstorms formed upstream of the United Kingdom on 26 September 2023. To capture this scenario, we show the observed reflectance by Meteosat 10 on 27 September at 12 UTC (top) and the corresponding 36-hour operational forecast as a natural colour image (bottom).

convective systems evolve at different locations, even though they all occur over the central Mediterranean.

Usefulness and availability

These newly enhanced products will become available as part of IFS Cycle 49r1, to be released this autumn. In this way, we will be able to visualise our forecasts of severe weather events across all simulated scales to support operational forecasters when estimating the probability of occurrence and location of clouds. In addition to other forecasting applications, these enhanced products are very useful in aviation, where thunderstorms inside mesoscale convective systems can have a severe impact on aircraft at cruising altitude, as well as in scenarios where very accurate forecasts of visibility are needed during start and landing.

storm), the clouds appear white as a combination of high reflectance in the visible spectrum and cold temperatures in the infrared.

An illustration of a severe weather event over Europe

In September, a series of storms impacted the west coast of Great Britain (see the second figure).

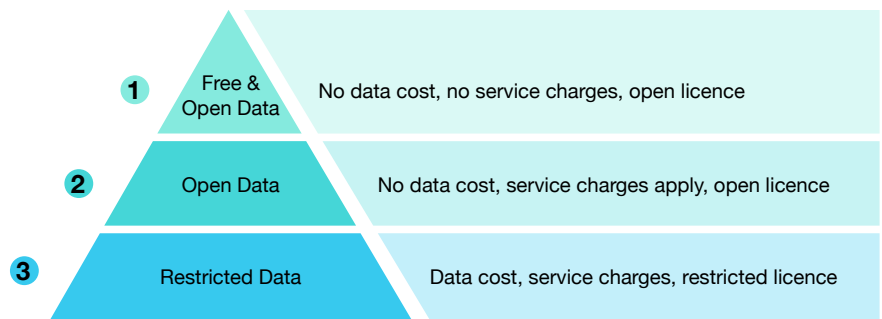
Two cyclones followed each other within a timespan of one week. These cyclones carried intense precipitation, strong thunderstorms, and enhanced wind gusts. In a 36-hour forecast, the predicted synoptic scale structures, such as frontal systems spanning across the North Atlantic or the location of the cyclone cores, match well with observations. Over the course of the 36-hour forecast, local

ECMWF moves to final stages of open data transition

Emma Pidduck, Umberto Modigliani, Victoria Bennett, Florian Pappenberger, Fabio Venuti, Xiaobo Yang, Ruth Coughlan, Maartje Kuilman, Ilaria Parodi

ECMWF is moving towards the final stages of its open data transition, with 2024 marking a significant year in achieving this goal. Recognising the economic and societal benefits demonstrated by existing open data initiatives and EU directives, the decision to move to open data was approved by the ECMWF Council in December 2020, with an ambitious target to achieve a fully open status by 2026.

During this transition, a tiered approach has been designed to provide balanced services for all users (see the first figure). Starting with data on a free and open tier, this is a subset of the full ECMWF Catalogue and is available with no charges under the CC-BY-4 open



Tiered access. Access to ECMWF data and products is currently governed by a three-tier access structure, which will be reduced to the top two tiers by 2026.

licence. This popular open-source licence allows the use of the data for commercial purposes. It permits redistribution, and the only requirement is that users must attribute the data

owner and source. Due to the size of the ECMWF Real-time Catalogue, only a popular subset is freely available on the data portal. It is provided with limited support. Data is also published by

popular third-party providers, such as Open-meteo.com, Amazon, Google and Microsoft. It is available with a one-hour delay after the operational dissemination. Also on this access tier are the Open Charts. These are available under the same CC-BY-4 licence, supported by tools such as Jupyter Notebooks to enable replication and adaptation as required.

In the next tier, open data is available with no data costs, but delivery costs may apply. The conditions of use for this data are consistent with the first tier, but the parameter options are wider, and the delivery time meets the operational schedule.


The final tier is the remaining restricted data at higher temporal and spatial resolutions. The conditions of use for the restricted dataset are governed by the ECMWF Standard Licence Agreement. Once the open data goal is fully achieved in 2026, the data accessible in this tier will move to the second level as 'open data' but will not necessarily be free (see the first figure). Users opting for the second and third tiers can pay for one of four service packs with a wide range of features, depending on their specific requirements.

Updates for 2024

In March 2024, the free and open data resolution was increased to 0.25 degrees (28 km), and the parameter subset was expanded to enable users to run limited-area models. In addition, data from the new Artificial Intelligence Forecasting System (AIFS) was added to the open data catalogue, allowing direct comparison with outputs from our Integrated Forecasting System (IFS).

From July 2024, ECMWF has made available all real-time data coarser than or equal to 0.4 degrees (~36 km) under the CC-BY-4 open licence. This means both extended-range ensemble forecast (ENS extended) and seasonal forecast datasets are now open and free of data cost at their native resolution, although they may be subject to service charges.

Additionally, data costs for the remaining restricted data have been reduced by up to 40%, enabling more users to take larger volumes of data to meet their requirements. The unit cost for data has been reduced by 40%, while the Maximum Charge cost has been decreased by 28%.



Monthly data usage in 2024. Since January, the data volume downloaded by users has increased to more than 400 TB per month, and the number of retrievals has gone up to about 47 million.

In addition to applying an open licence to much of the data, making data affordable while ensuring service quality has been part of the overall transition. Service charges, which comprise a volume charge and a service pack cost, have been reduced by an average of 60%, with the upper ceiling of volume limits increasing to 4 TiB per day.

Open data usage

Since releasing a much wider open dataset in March, ECMWF has seen a 30% increase in downloads retrieved via the ECMWF Data Portal, which serves over 400 TB per month and reached ~47 million retrievals in May 2024 (see the second figure). Distribution via third parties further enhances the popularity of these datasets.

Future plans – 2025 and 2026

The goal is to achieve a fully open status in 2026 by removing the Information Cost (data cost) for all parameters in the ECMWF Real-time Catalogue and providing a subset of the catalogue at no cost to users. The ECMWF Council recently approved proposals to further reduce data costs and to further open up data availability in 2025 as part of the next steps to meet the final goal in 2026.

This includes an increase in the spatial resolution of the current free and open data available at 0.25 degrees to a

resolution to match the new ERA6 dataset, which is planned to be 0.125 degrees (~14 km). Supporting this increase in resolution, ECMWF will further decrease the information cost for the remaining restricted data in the catalogue. The unit cost of the data will decrease by 33%, and the Maximum Charge fee will reduce by almost 50% in 2025. These changes will likely be introduced in the second half of 2025 and will be the final price decrease before removing the Information Cost in 2026.

Expanding support to WMO Members

As the open data offering expands, ECMWF will also remove barriers for World Meteorological Organization (WMO) Members by removing all data and service charges for full-resolution data in support of UN initiatives such as Early Warnings for All. Between now and 2027, the quality and accessibility of data will be improved for all WMO Members, first focusing on less developed nations in the WMO's Regional Association for Europe (WMO-RA6) and the countries supported by the WMO SOFF (Systematic Observations Financing Facility) initiative. SOFF aims to support some countries to generate and exchange GBON (Global Basic Observing Network) data, which are critical for improved weather forecasts and climate services.

Introducing JupyterHub at ECMWF

Xavier Abellan, Manuel Martins, Milana Vučković

This spring has seen the opening of the new JupyterHub service at ECMWF, a new way of accessing ECMWF resources in an interactive and modern way. With JupyterHub, authorised users can now launch Jupyter sessions on multiple backends, including ECMWF's high-performance computing facility (HPCF) and the ECGATE Service (ECS). They can thus leverage the computational capacity and data resources available at ECMWF without leaving their browser.

Jupyter is an open-source platform that provides an interactive computing environment where users can manipulate data, run code, and visualise results in real time, fostering a seamless and efficient workflow. It enables users to create and share documents – known as notebooks – that contain live code, equations, visualisations, and narrative text. It supports various programming languages, including Python, R, and Julia, making it a versatile tool for data analysis, machine learning, scientific computing, and more. Jupyter also offers rich visualisation capabilities, allowing users to create interactive plots, charts, and graphs directly within their notebooks.

Use of JupyterHub

The JupyterHub service had been running in beta mode for over six months. Early adopters amongst ECMWF staff identified issues for the new service and refined the final offering before it was made available more generally. The creation of this service addresses a requirement highlighted in a recent internal survey at ECMWF: a great number of scientists expressed their desire for improved and more modern interactive computing, data analysis and plotting capabilities.

In its current configuration, Jupyter users can enjoy one session at a time on either the HPCF or ECS, depending on their permissions, with a fixed duration of up to seven days. Sessions can be tailored to fit any particular requirements for the task at hand.

Backend Options

Select an Environment

ECMWF ATOS HPCF

Run a JupyterLab session on the ECMWF Atos HPCF under your user account. You will have access to the standard software stack, environment and files in the HPCF computing service.

Configurations selected below will be persisted and will be re-used the next time you use this profile.

<p>Memory</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">8 GB RAM</div> <p>Session Duration</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">12 Hours</div> <p>Version</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">New</div> <p>Debug</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">No</div>	<p>CPU</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">2 CPU</div> <p>Temporary Storage</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">3 GB</div> <p>Project Account</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 10px;">ecus</div>
--	---

Public SSH Key

The following key will be used to spawn Jupyter on your HPC environment.

If you did not configure it before, please append the following Public SSH Key onto your `~/.ssh/authorized_keys` file on your HPC home:

```
ssh-ed25519 [redacted]
```

Start

JupyterHub options. This screenshot shows JupyterHub options to configure a Jupyter session on the HPCF backend.

Resources, such as the number of CPUs, memory and temporary storage requirements, can be customised. Once they are active, users can take advantage of the rich software stack readily available for their activities, as well as using their own Conda or Python virtual environments.

Sessions are based on JupyterLab, an evolution of the classic Jupyter Notebook web-based application. It comes with additional features and customisation options to offer a truly interactive remote development environment. Users can leverage JupyterLab's multiple panes to use text editors with syntax highlighting and auto-completion, terminal consoles for direct system interaction and command execution, and image/document displayers for quick inspection of files. Seamlessly integrated with version control

systems like Git, JupyterLab promotes collaborative coding endeavours while facilitating reproducible research practices.

The JupyterHub system is designed with flexibility in mind, with a clear separation between the frontend service, responsible for managing user sessions, and the different backends, where those sessions are run. This modular architecture ensures a future-proof solution, with the possibility of expanding our backend catalogue with additional resources, such as those based on cloud technologies.

Overall, JupyterHub offers an intuitive and efficient environment for exploratory data analysis, visualisation, and sharing reproducible research. Its accessibility via a web browser not only enables remote collaboration but also ensures

seamless access to computational resources across various devices, underpinning its utility as a cornerstone tool for modern data-driven workflows.

Access to JupyterHub

JupyterHub is currently available to all registered users with access to

ECMWF’s HPCF and ECS, and more than 200 of them have already tried it. In order to use those services, non-registered users need to apply for access. People who work for a national meteorological service or scientists/researchers who work for a research organisation within one of our Member States or Co-operating States may follow the standard

guidelines to access our computing facilities: <https://www.ecmwf.int/en/computing/access-computing-facilities>.

You can find more information about the service and how to use it on our JupyterHub documentation page: <https://confluence.ecmwf.int/display/UDOC/JupyterHub>.

A year in review of CESOC–ECMWF collaboration and initiatives

Florentine Weber

Nearly two years ago, in the autumn of 2022, ECMWF and the Center for Earth System Observation and Computational Analysis (CESOC – <https://cesoc.net/>) signed a memorandum of understanding after ECMWF established a branch in Bonn, Germany. This marked a significant milestone in our collaboration with universities and research institutes around our new site. Since then, there has been substantial activity between the two Centres, showcasing the close alliances formed in Earth system science.

CESOC is a research partnership between the University of Bonn, the University of Cologne, and Forschungszentrum Jülich. Established in October 2020, its primary aim was to support Bonn's bid to host the new premises of ECMWF. The trio of partners brings to the table a wealth of collective experience, with years of scientific cooperation and involvement in various pioneering research endeavours. The scientific alliance between ECMWF and CESOC serves as a component within the broader

framework of academic partnerships currently taking shape in Germany.

Significant milestones

During the summer of 2023, CESOC held its inaugural Members' Assembly in Bonn and invited several ECMWF staff members to join the occasion. In addition to the 49 CESOC members present, 14 ECMWF guests participated. Jean-Noël Thépaut, Florian Pappenberger, and Irina Sandu from ECMWF shared research and development highlights, igniting discussions on potential collaborations and exchanges with CESOC members. The official part was kept short to enable further networking opportunities at the CESOC summer party later that day.

Following two previous visits to the University of Bonn and Forschungszentrum Jülich in 2022, in November 2023 the University of Cologne hosted ECMWF staff. These visits aimed to showcase the diverse spectrum of research undertaken by the CESOC institutions.

At the half-time event at the end of January, our inaugural group of STEP UP! Fellows – Florentine Weber, Luise Schulte, Katerina Anesiadou, and Paolo Androzzi – wrapped up their debut year at ECMWF in Bonn. STEP UP! is a two-year fellowship programme organised by the German Meteorological Service (DWD) with close cooperation between ECMWF and CESOC. ECMWF executives, the Fellows’ supervisors, and speakers and guests from CESOC, DWD, and the Federal Ministry for Digital Affairs



Half-time event. The STEP UP! Fellowship half-time event reflected the collective efforts of representatives from ECMWF, CESOC, the German Meteorological Service (DWD), and the Federal Ministry for Digital Affairs and Transport (BMDV), who have been supporting the first cohort of Fellows for over a year. *Credit: DWD*

and Transport (BMDV) supported the Fellows, showcasing their ECMWF journey. They also delved deeper into discussions on the role of machine learning in the digital transformation of Earth system modelling and on the DWD Hans Ertel Centre for Weather Research. A special guest was Insa Thiele-Eich, the first German female astronaut, who addressed diversity in space and sustainability. With a duration of two to three years, preparations are under way for the next Fellowship cohort, scheduled for the first quarter of 2025.

In the same month, an intergovernmental agreement was signed, deepening meteorological and climatological collaboration between Germany and Italy. This pact, endorsed in Berlin by officials from both nations, aims to strengthen research and bilateral cooperation in weather and Earth observation. Stefan Schnorr, State Secretary at the BMDV, which oversees the DWD, emphasised the necessity of transcending national borders to address weather phenomena and climate change challenges. The agreement facilitates various initiatives, including the establishment of a joint master's programme between CESOC and the

University of Bologna. The master's programme is part of the 'Italia – Deutschland science-4-services network in weather and climate (IDEA-S4S)', a joint research and education network with four-year funding periods, which currently focuses on high-impact events such as droughts and flooding. This programme will cooperate with and strengthen the scientific environment for the ECMWF sites in Bonn and Bologna.

In February 2024, the ECMWF–CESOC Reanalyses Working Group embarked on a major collaborative initiative, uniting 70 experts from both organisations on current atmospheric reanalysis and data fusion techniques' strengths and limitations. As ECMWF develops the reanalysis product ERA6 to serve diverse sectors, it has integrated user feedback from CESOC as well as many other partners across our Member and Co-operating States.

Over the last few years, several colloquia, workshops, and lectures have been presented by researchers from both ECMWF and CESOC. During the university summer term 2024, the successfully established CESOC seminar series 'My Research' was held fortnightly and was being hosted by ECMWF in Bonn.

To ensure that the next generation benefits from the collaboration between ECMWF and CESOC, 16 students from various master's programmes within CESOC visited ECMWF in Bonn in June 2024 and listened to talks on ECMWF's broad project portfolio.

After the 2022 Machine Learning Crash Course held by Peter Düben, the head of ECMWF's Earth System Modelling section, a jointly organised workshop on 'Large-Scale Deep Learning for the Earth System' is scheduled for 29 and 30 August 2024. For more information, please visit <https://cesoc.net/workshop-on-large-scale-deep-learning-for-the-earth-system/>.

Honorary professorship

The recent appointment of Peter Düben as Honorary Professor at the University of Cologne in May 2024 serves as a remarkable testament to the enduring and fruitful collaboration between ECMWF and CESOC. This prestigious recognition not only honours Peter's contributions but also reinforces the strong bonds that drive innovation and excellence in our shared scientific endeavours.

A new way of handling observation errors for infrared measurements

Chris Burrows, Cristina Lupu

The forthcoming upgrade of ECMWF's Integrated Forecasting System (IFS) to Cycle 49r1 includes observation errors diagnosed using a simpler method than previously for the Cross-track Infrared Sounder (CrIS), which is on board several American satellites. Also, a more physically consistent method of prescribing the observation error covariances for CrIS, to describe correlations between errors in the different channels, has been implemented. This is achieved by prescribing the observation errors in units of radiance, and then dynamically mapping these errors into brightness temperature units for every scene.

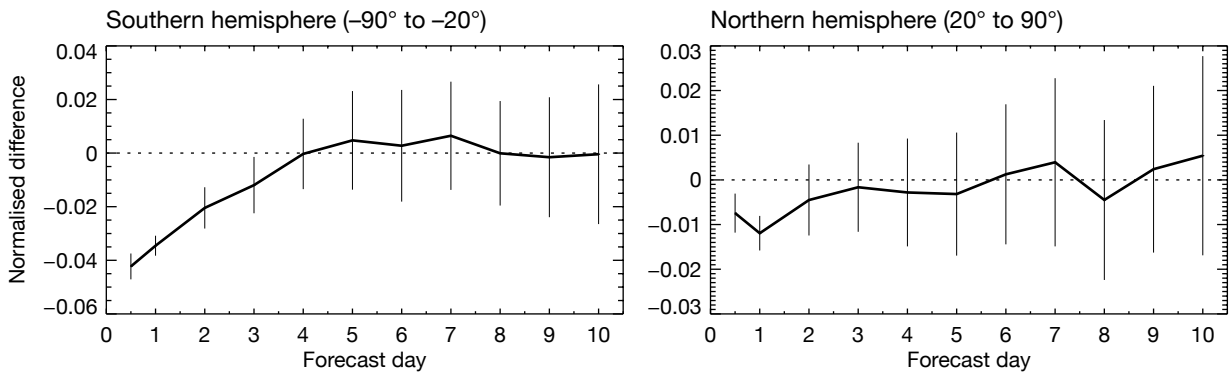
Diagnosing observation error covariances

For some time at ECMWF, observation error covariance matrices have been derived using the difference between observations and both a short-range forecast (the 'background') and the best estimate (the 'analysis') of the current state of the Earth system, using the method proposed by Desroziers et al., 2005 (<https://doi.org/10.1256/qj.05.108>). As with any method of diagnosing observation errors, some assumptions are required. These include unbiased errors in observations, short-range forecasts, and analyses. They also include the assumption that observation and

background errors are uncorrelated. The Desroziers diagnostic will result in suboptimal estimation if these criteria are not met.

In practice, these conditions are *not* met. For a new observation type, a time-consuming set of assimilation experiments with sub-optimal error covariances is used to generate the set of background and analysis departures from observations used in the method. Even then, the resulting matrix is not symmetrical, can be badly conditioned, and is likely to need arbitrary scaling to give good forecast impact.

Our proposal is to simply use the covariance of analysis departures,



Reduced errors. Percentage change in 500 hPa geopotential height root-mean-square error (RMSE) for different forecast lead times when the new observation error formulation is applied for CrIS (i.e. having used the new error diagnosis method, and the scene-dependent errors in the assimilation). The period of 3 December 2020 to 27 February 2021 was considered. Negative values show improvements.

which can come from an experiment with all observations actively assimilated, except, in this case, CrIS observations. The CrIS observations will be passive (i.e. monitored only and not affecting the analysis). This is effectively assuming the analysis to be a proxy for the *truth*. Benefits of this method include the fact that the resulting matrix is symmetrical and well-conditioned (i.e. it is a covariance matrix by definition), and that it can be applied without running an initial assimilation experiment to generate the analysis departures. No method for estimating observation errors is free from assumptions, but this method has been shown to be effective in terms of forecast impact, and it has the advantage that it can be applied quickly to observations from new instruments.

Specifying scene-dependent observation error covariances

Often, we refer to the assimilation of *radiance* observations, when in fact it is quantities of brightness temperature that are assimilated in the ECMWF system, having been converted using the Planck function (see the box). An attractive aspect of working in brightness temperature (units of Kelvin) is that the values are physically intuitive and closely related to geophysical temperature. In contrast, radiances have non-intuitive values, varying over several orders of magnitude depending on channel, making them an unattractive option for assimilation. Assimilating these quantities directly would involve a great deal of technical effort.

When a radiance measurement is

made by a satellite, the noise characteristics are approximately Gaussian in radiance units. However, when these measurements are fed into the non-linear Planck function, the resulting noise distribution in brightness temperature units becomes very non-Gaussian. This is particularly true at the high-frequency end of the infrared spectrum, and the non-Gaussian distribution of errors compromises one of the key assumptions in variational data assimilation.

Our suggested solution is to diagnose the observation error covariance in units of radiance, and then, at the time of the assimilation, *dynamically* convert these to brightness temperature units. This can be done using the derivative of the Planck function, in a way that is consistent with each scene. This compromise preserves inter-channel error correlations and allows a more appropriate specification of observation errors, while retaining the intuitive units of Kelvin for the monitoring of the observations.

Error covariances diagnosed by the new method radiance units, and dynamically mapped to brightness temperature units at the time of assimilation, have been tested for CrIS. The results are positive in the short range in terms of predictions of 500 hPa geopotential height (see the figure). This will become operational in Cycle 49r1, and it is intended that the method will be extended to other instruments in Cycle 50r1.

The Planck function gives the radiance spectrum emitted by an ideal black body at a particular temperature.

$$B = \frac{c_1 \nu^3}{e^{c_2 \nu/T} - 1}$$

Where B is the radiance in $\text{Wm}^{-2}\text{sr}^{-1}\text{cm}^{-1}$, ν is the wavenumber in cm^{-1} , T is the temperature in Kelvin, and $c_1 = 2hc^2$ and $c_2 = hc/k_B$ (h , c and k_B are the Planck constant, the speed of light and the Boltzmann constant, respectively).

The derivative of this equation with respect to temperature is:

$$\frac{dB}{dT} = \frac{c_1 c_2 \nu^4 e^{c_2 \nu/T}}{T^2 (e^{c_2 \nu/T} - 1)^2}$$

So, to convert small uncertainties in radiance units to those in units of Kelvin, for a given observed brightness temperature, we can simply calculate:

$$\delta T = \left(\frac{dB}{dT} \right)^{-1} \delta B$$

Tenth anniversary of the ECMWF ‘Weather Discussion’

Tim Hewson

The very first ‘Weather Discussion’ was delivered about ten years ago, on 20 June 2014, by Linus Magnusson at ECMWF’s Reading (UK) site. Ever since, this has been a very popular and productive Friday afternoon event in the internal ECMWF calendar. Accordingly, on 14 June 2024, ECMWF marked age ten with multi-site celebrations in Reading, Bonn (Germany) and Bologna (Italy).

How is this relevant to readers?

The key point is that the Weather Discussion focuses specifically, in a real-time-weather framework, on improving forecasts for the benefit of all users, and indeed issues discussed are often those raised directly by users. Discussions often evolve into entries in the severe event catalogue (<https://confluence.ecmwf.int/display/FCST/Severe+Event+Catalogue>) and/or Newsletter articles. Below we describe this activity in a historical context.

Origins

The inaugural Weather Discussion in 2014 linked to an ECMWF project instigated about two years earlier, under the leadership of the then Director-General Alan Thorpe. This project involved converting the old ECMWF library into a bright new open space to be called the ‘Weather Room’. It was to be a multi-function space, used for informal meetings, journal browsing, and breaks and poster displays during workshops. Staff and visitors alike could also check out and discuss some of the latest forecast information on a large, interactive, multi-screen ‘weather wall’ (front of room on the picture). Within the space, there is also a glass-walled ‘pod’ (towards the left on the picture) equipped with desk and workstation, where, for a week at a time, the rostered on-duty ‘Daily Report analyst’ is typically based. This enables easy interactions with other staff and visitors on current meteorological topics of interest, such as real-time forecast and analysis issues. Each week

culminates with the analyst leading the Weather Discussion on the weather wall.

Previously, ECMWF had had rostered analysts, but they were based in a remote corner of ECMWF, surrounded by a plethora of paper charts that had to be renewed daily. And whilst they did write a report every weekday, in-person interactions were disappointingly rare, and there were no Weather Discussions.

Since its inception, the Weather Discussion has been a meeting point – physical, then virtual, and nowadays hybrid – at which current meteorological topics of interest can be presented and then discussed with the audience. The goal is to improve ECMWF forecast performance, in many different ways, by working collaboratively. The initiative has been an unequivocal success, and the event has remained very popular, typically attracting about 100 attendees.

Issues

Many issues have been identified, and many corrective actions have resulted. Examples include physics-related problems in new model cycles. One such issue related to wet (sub-grid) tiles used in ECMWF models, which led to unrealistic 2 m temperatures at isolated gridpoints; another was runaway build-up of very deep snow cover in a few ensemble members, in exceptionally cold conditions, from cloud structures that could not support such snowfall. A specific problem class is observations. Radiosonde and buoy data have been found to be assigned to the wrong geographical locations, with worrying consequences for data assimilation. Meanwhile, observations themselves can sometimes go wrong in bizarre ways that only human intuition has been able to pin down. Other topics regularly investigated include extreme weather events, past and

forecast; product issues; tropical cyclone handling; ocean waves; forecast jumpiness/consistency; implausible forecasts; issues reported by our users (e.g. low humidity in fog); and checks of recently added variables. Verification scores also tend to be a hot topic, with ‘forecast busts’ that users sometimes highlight, and occasions when we perform notably worse than our international competitors, often triggering more detailed investigation!

Over the years, the range of work has grown substantially, and it continues to grow. Discussions now regularly encompass topics related to the EU’s Copernicus Atmosphere Monitoring Service (CAMS) implemented by ECMWF, such as dust storms; climate monitoring activities related to the EU’s Copernicus Climate Change Service (C3S) implemented by ECMWF; flood- and fire-related output run by ECMWF for the EU’s Copernicus Emergency Management Service (CEMS); 4.4 km resolution global forecast runs, which ECMWF operates for the EU’s Destination Earth initiative; and post-processing activities, such as ecPoint. For about a year now, we have also been examining the output of machine learning models, such as ECMWF’s Artificial Intelligence Forecasting System (AIFS).

We have also embraced technological advances, writing the daily reports wiki-style in an easy-to-maintain Confluence framework since 2016, and recording all Weather Discussions since the temporary move to an online format brought about by COVID. Occasionally, discussions are also now streamed from Bonn. Everything, including recordings, is archived, with search engines and labelling systems allowing ECMWF staff to connect with old topics without delay. This is all very conducive to productivity and efficiency.



Weather Room, Reading



Bonn

Reading

Bologna

Multi-site celebrations. The ‘Daily Report’ team, past and present, comprises Estibaliz Gascón (by the cake in Bonn) and, lined up in Reading (left to right), Fernando Prates, Ervin Zsótér, Linus Magnusson, Tim Hewson, Mohamed Dahoui, David Lavers, Ivan Tsonevsky and Jonny Day.

Different specialisms

All this has required a dedicated, diverse and talented workforce. Accordingly, the Daily Report team brings many different specialisms to the table, not to mention

wide-ranging meteorological experience from having lived in the four corners of Europe, and beyond. Many have something of a forecasting background, too. In fact, this job is about as close as it gets to being a forecaster at ECMWF,

but without the downside of night shifts! Perhaps it is a testament to the success of the initiative, and the high job satisfaction, that the team has been relatively static throughout the ten years.

UEF2024 online attracts a global audience

Becky Hemingway

The annual Using ECMWF’s Forecasts (UEF2024) event was held online on 5 and 6 June. It welcomed 294 attendees across the two days, with the livestream consistently reaching 110–160. Over 340 people from 66 countries registered to attend UEF2024, which was the largest ever number of registrations for a UEF event.

The event featured four sessions: Forecasting and Research updates, Copernicus and ECMWF Services updates, Machine Learning and the Artificial Intelligence Forecasting System (AIFS), and Speakers Corner – Updates on ECMWF forecast products.


Due to the online and compact format of UEF2024, it was decided that only ECMWF updates would be given, plus one invited speaker. In total, 18 presentations were delivered across the four sessions. All recordings and slides are available on the online UEF2024 event page (<https://events.ecmwf.int/event/397/>).

Attendee feedback showed 100% found UEF2024 ‘Very Useful’ or ‘Useful’. Many participants joined to gain insight into ECMWF plans, developments and product updates, with the AIFS and machine learning of notable interest.

Meeting highlights

With many exciting things going on across ECMWF, Andy Brown, Director of Research, and Florian Pappenberger, Director of Forecasts and Services, joined forces to present on ECMWF’s progress and plans. They provided a summary of ECMWF’s Strategy, model cycles of our Integrated Forecasting System (IFS), open data, machine learning, our role in the EU’s Destination Earth initiative, the two Copernicus services we run for the EU, the European Weather Cloud, and the innovation programme Code for Earth – setting the stage for more detailed talks on many of these topics later in the programme.

2m temperature week 2



New extended-range maps. A new visualisation and structure for extended-range anomaly maps was proposed for user feedback by Tim Hewson (ECMWF).

ECMWF presented recent and upcoming forecasting product changes, a topic always of interest to users. Matthieu Chevallier summarised updates and showed new visibility meteograms, the new precipitation type freezing drizzle, updated wave products, and the development of hail products. In the Speaker's Corner, Cihan Sahin gave updates on graphical products ecCharts and OpenCharts and how to access new forecast products in these tools. Nigel Roberts presented a new diagnostic called CURV (Curvature Using Radial Variation) to identify cyclonic and anticyclonic curvature from surface pressure fields or geopotential height using ensembles. One attendee commented that "CURV looks really interesting and promising!!!!". Cristina

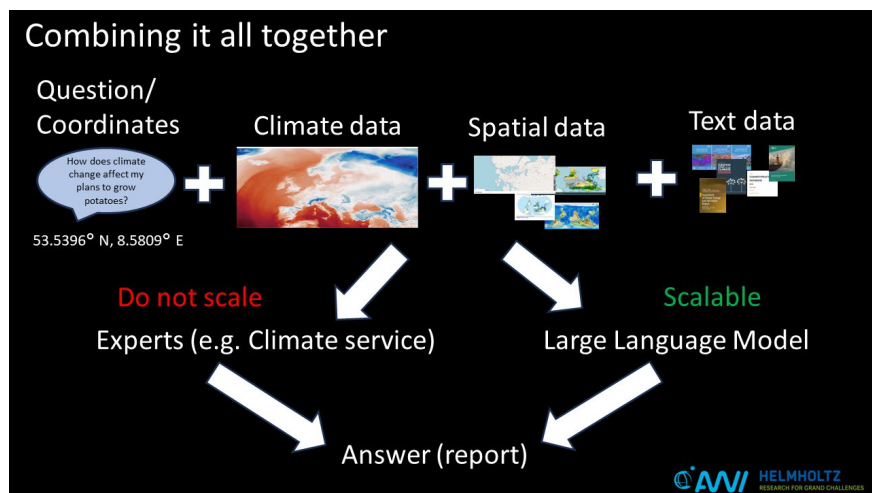
Lupu showed new real-time simulated imagery forecast products, which are available out to ten days. Tim Hewson proposed a new structure for extended-range anomaly maps and asked for user feedback on this (see the first figure). Tim also presented a collaboration with the UK Met Office on new extratropical cyclone products including front probabilities.

UEF2024's invited speaker addressed a topic of special interest to many in the meteorological field: the fast-evolving subject of machine learning. Nikolay Koldunov (Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung, AWI) talked about how to use Large Language Models (LLMs) for weather and climate information retrieval (see the second figure). Demonstrations were given on

how providing LLMs with relevant data, high-resolution modelling (such as provided by Destination Earth), and fine-tuning questions can result in useful weather and climate information for different sectors in a scalable way, using the open-source ClimSight prototype tool.

Mariana Clare, Simon Lang, and Linus Magnusson (all ECMWF) presented ECMWF's current machine learning and AIFS status, which consists of three elements: the role of machine learning in Destination Earth, the Machine Learning Project, and the Member State Pilot Project. AIFS scorecards, when compared to those of the IFS (2022), were remarkably good. The AIFS was shown to perform well in many cases with more consistent forecasts than the IFS. However, it verified less well in others, for example by predicting tropical cyclone intensities that were too weak. Simon Lang discussed an AIFS ensemble, which was of keen interest to attendees.

Updates on user-orientation advances in Copernicus were presented by Stijn Vermoote (ECMWF), showcasing the many collaborations and products and services available to users of the Copernicus Climate Change Service (C3S) and the Copernicus Atmosphere Monitoring Service (CAMS). The Copernicus Interactive Climate Atlas and Climate Pulse, online tools to visualise and analyse observed and projected climate information from C3S, were of particular interest to attendees. The section on National Collaboration Programmes, presented by Cristina Ananasso (ECMWF), highlighted strong collaborations with ECMWF Member and Co-operating States to enhance the uptake of C3S and CAMS products and services.



Large Language Models. Nikolay Koldunov (AWI) demonstrated how Large Language Models can answer weather and climate questions.

Future UEF events

During UEF2024, there were two exciting announcements. The first was that a side event will be held at the European Meteorological Society Annual Meeting (EMS2024) to engage in person with ECMWF forecast users. EMS2024 will take place from 2 to 6 September 2024 in Barcelona, Spain. More details on the side event can be found here: <https://events.ecmwf.int/event/423>. The second was that UEF2025 will be part of ECMWF's 50th anniversary celebrations in Bologna, Italy, in September 2025. More details on that event will follow in due course.

Destination Earth's digital twins and Digital Twin Engine – state of play

Irina Sandu, for the DestinE teams at ECMWF and at partner organisations

After Destination Earth (DestinE) transitioned into its second phase in mid-June 2024, this is a good time to reflect both on key achievements so far and on the next steps in this ambitious initiative funded by the European Union. DestinE aims to build a digital replica of our planet to enhance our capability to respond and adapt to the environmental challenges posed by extreme events and climate change. After two years of intense developments through a truly European collaborative effort implemented by ECMWF, the European Space Agency (ESA), and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) together with many partners across Europe, the novel information system DestinE is setting up will gradually open its gates to users. Its innovative elements will now continuously evolve and gradually transition towards operations. Here, we will mainly focus on the key elements of the DestinE system that ECMWF is delivering together with its 90 partners across Europe: the Digital Twin Engine and the first two high-priority digital twins, which are at the heart of the DestinE system. This article reflects the contributions of all the teams working on DestinE at ECMWF and at our partner organisations throughout Europe.

Building a digital replica of our planet

Recent years, and even more so the last 15 months, have shown that the climate emergency is amplifying and that we are entering uncharted territory. As shown by the EU-funded Copernicus Climate Change Service (C3S) implemented by ECMWF, the ERA5 reanalysis developed by ECMWF has shown that in 2023 the global mean temperature reached the highest average level since records are available. In January 2024, it became for the first time more than 1.5°C warmer on average over a period of 12 consecutive months than in the pre-industrial era. In other words, we are increasingly getting closer to the 1.5-degree global warming threshold set in the Paris Agreement. At the same time, certain extreme weather events are also becoming more frequent and intense, with often catastrophic impacts. There is an urgent need to enhance our capabilities to better respond and adapt to the environmental challenges posed by climate change and extreme events.

DestinE, an EU-funded initiative launched in 2022, responds to this need by aiming to build a digital replica of our planet by 2030. Digital twins of the Earth system can represent a game-changer as they provide an interactive capability to better understand how the Earth system will evolve following different scenarios, and to assess the implications of climate change and the impacts of extreme events, from the global scale to local scales. They can thus complement existing capabilities at national and European level and support adaptation policies and the implementation of the European Green Deal.

Under the lead of the European Commission's Directorate-General for Communications Networks, Content and Technology (DG CNECT), the initiative is being jointly implemented by three entrusted entities: ECMWF, which is responsible for delivering the first two high-priority 'digital twins' and the 'Digital Twin Engine'; ESA, which is responsible for the 'Core Service Platform'; and EUMETSAT, which is responsible for the 'Data Lake'.

DestinE has rapidly become a truly European endeavour as the three entrusted entities now work with more than 110 institutions from over 25 countries, including world-leading national meteorological and hydrological centres, climate centres, academia and private institutions, including small and medium-sized enterprises. At ECMWF, numerous staff members across the organisation work closely together with colleagues in our partner organisations in DestinE to develop the digital twins and the Digital Twin Engine.

DestinE crucially relies on a strategic partnership with the European High Performance Computing Joint Undertaking (EuroHPC JU). DestinE digital twins, as well as the artificial intelligence (AI) activities which are now ramping up, would not be possible without the world-leading facilities of the EuroHPC JU. DestinE is pioneering running the first complex Earth system applications for extreme weather events and climate change on Europe's first pre-exascale computers: LUMI, run by CSC – IT Center for Science in Finland (no. five in TOP 500); Leonardo, run by Cineca in Italy (no. seven in TOP 500); and MareNostrum 5, run by the Barcelona Supercomputing Center (BSC) in Spain (no. eight in TOP 500). It is also making use of medium-sized

EuroHPC facilities, such as Meluxina in Luxembourg.

Building DestinE's key components and bringing them together

Current Earth prediction systems rely on significant progress in weather and climate science, Earth system observations, and technology made over recent decades. DestinE builds on this progress, as well as on recent breakthroughs in digital technologies and supercomputing, and it brings in innovative features to address certain limitations of the current systems. In particular, it enhances:

- Earth system representation, through enhanced horizontal resolutions
- the flexibility to change simulation configurations as well as output products compared to operational systems today
- the ability of users to access, fuse and interact with data through standardised interfaces, and
- the coupling of impact-sector applications (e.g. for renewable energy, water management or health) and Earth system models.

DestinE also capitalises on recent breakthroughs made in artificial intelligence and machine learning (AI/ML) to boost two particularly important aspects: the

interactivity of the DestinE system and the quantification of uncertainties in the accuracy of simulations.

In partnership with EuroHPC, DestinE:

- sets up a unique, bespoke, cutting-edge high-resolution simulation capability allowing the testing of hypotheses and scenarios regarding the evolution of the Earth system from days to decades ahead
- provides globally consistent Earth system and impact-sector information at scales of a few km, where the impacts of many extreme events and climate change are observed
- establishes new end-to-end pathways of interacting with digital twin data
- fosters an AI-enabled digital ecosystem supporting decision-making.

The information system DestinE set up is composed of several innovative elements: the first two digital twins and the Digital Twin Engine; the Data Lake; and the Core Service Platform (Figure 1).

The digital twins of DestinE replicate different aspects of the Earth system. The initial high-priority digital twins on weather-induced extremes and climate change adaptation are highly accurate replicas of the Earth system, simulating the system behaviour at the scales

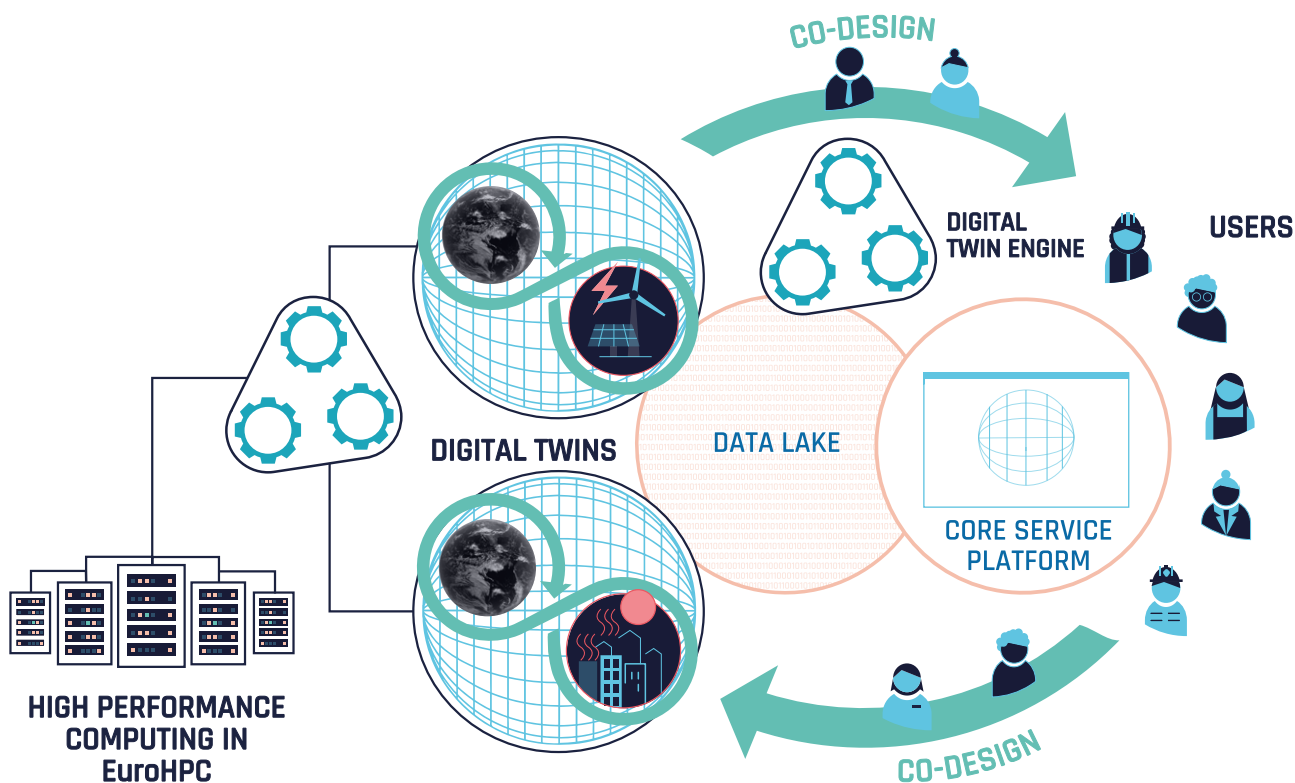


FIGURE 1 Key elements of the DestinE system, funded by the EU and implemented by ECMWF (digital twins and Digital Twin Engine), EUMETSAT (Data Lake) and ESA (Core Service Platform).

where extreme events happen, and many of the effects of climate change are felt. They combine several Earth system models and observations together with the most advanced digital technologies, machine learning and artificial intelligence. They are integrated and interoperable with applications for the impact sectors most affected by climate change and extreme events.

To efficiently deploy and embed the different Earth system digital twins within the overall DestinE system, the Digital Twin Engine brings together the software infrastructure necessary for extreme-scale simulations, data fusion, data handling, and machine learning. It provides a common system approach to a unified orchestration of Earth system digital twins and allows users to interact with them and their data in novel ways.

The Data Lake incorporates different data spaces and provides users with harmonised access to the datasets. Its three main pillars are: harmonisation of data access, facilitation of access to the digital twins' data outputs and other existing datasets, and big data and near-data processing.

The Destination Earth Service Platform is the entry point to the DestinE ecosystem, addressed to a diverse community of users, providing access to applications, tools and services supporting DestinE data exploitation.

During the first phase of DestinE, the focus was on building these elements and bringing them together, by defining novel interfaces and making the system work as a whole. The aim for the end of phase 1 was to demonstrate the system at scale in preparation for its opening to the first users, with an aim to transition it towards operations in phase 2.

DestinE digital twins – main achievements so far

DestinE's digital twins build on decades of advances in numerical weather prediction, climate sciences, digital technologies and AI/ML, and they exploit the world-leading supercomputers of the EuroHPC JU. They are characterised by three main features:

- **Quality:** they provide globally consistent information with local granularity, based on accurate km-scale simulations performed with cutting-edge Earth-system models and an enhanced observation–simulation fusion.
- **Impact:** they provide integrated Earth system and impact-sector information at the relevant spatial and temporal scales for key sectors impacted by climate change and extreme events.
- **Interactivity:** they provide interactive access to models, data and workflows, based on the Digital Twin Engine and innovative cloud-based solutions.

An important aspect is the co-design of the digital twins with users from impact sectors affected by climate change and extreme events. The involvement of selected users right from the start has been key to ensure that they obtain the information they require and that the digital twins are tailored to their needs.

Climate Change Adaptation Digital Twin

The Climate Change Adaptation Digital Twin (Climate DT) is a step change in the way information on climate change is provided. The Climate DT aims to produce operationally multi-decadal climate projections. These projections – from global to local scale, for several decades ahead – will enable decision-making in support of climate change adaptation and the implementation of the European Green Deal.

The Climate DT will bring added value to users compared to current climate modelling activities, by providing globally consistent data with higher temporal and spatial resolution (5 to 10 km globally instead of more than 100 km; hourly instead of 6-hourly output; Figure 2); by operationally producing updated quality-controlled climate information routinely (yearly or less compared to every seven to ten years today); by considering the needs of users in terms of output variables and simulation design; and by providing a framework to perform bespoke, on-demand climate simulations. This framework can be used to address 'what-if' questions related to the impact of certain scenarios or policy decisions on the evolution of our planet, or to create storylines of how extreme events already experienced by climate-vulnerable sectors may look in the future to support better adaptation measures. Figure 3 shows how a heatwave experienced by Europe in 2019 would change if the average global temperature were to rise by 2°C.

In practice, the Climate DT relies on global km-scale climate models and flexible workflows, including directly linked impact models. Key features include the development of an operational framework that includes monitoring and evaluation, and interactivity elements, allowing enhanced responsiveness to emerging user needs. The involvement of selected users right from the start has been key to ensure that they will obtain the climate information they require.

The Climate DT exploits and further evolves a new generation of global storm-resolving and eddy-rich models built through a cooperative model development approach. It is supported by the European Horizon projects nextGEMS and EERIE as well as national projects, such as WarmWorld in Germany and Gloria in Spain, involving dozens of leading climate and weather centres, supercomputing centres and academia throughout Europe. The three models used are ICON, IFS-NEMO and IFS-FESOM. ICON is the ICOSahedral




FIGURE 2 The image shows climate projections for temperature at different times in the summer of 2032 (a) at a grid spacing of 100 km from the CMIP models used for the sixth assessment report (AR6) of the Intergovernmental Panel on Climate Change (IPCC), and (b) at a grid spacing of 5 km from the Climate DT (using one of the prototype projections performed on LUMI with IFS-NEMO). The specific dates were chosen such that they represent a similar temperature distribution in Europe. The Climate DT provides consistent climate information with local granularity on a global scale, bridging the gap between large-scale climate projections and local climate impacts, and avoiding data gaps and inconsistencies that come with existing regional downscaling efforts.

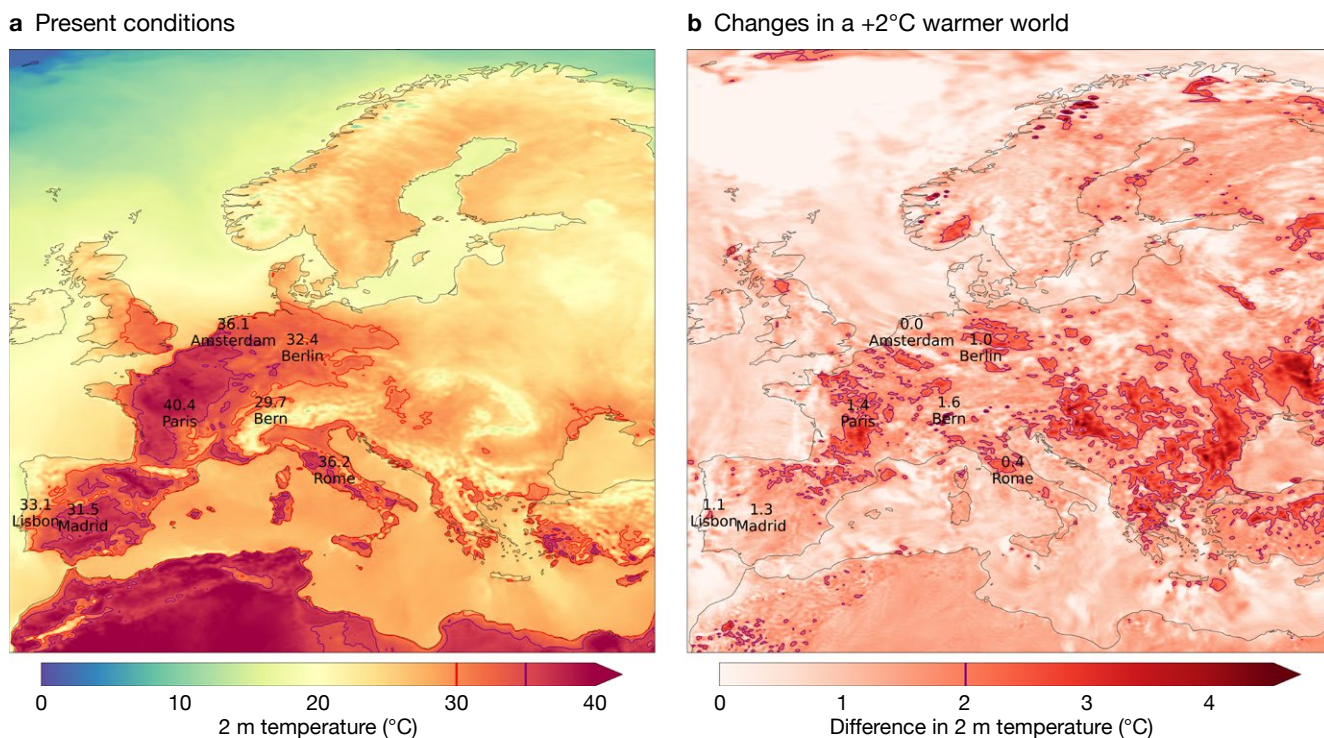


FIGURE 3 The figure shows (a) the European heatwave of 25 July 2019 and (b) the additional warming in a +2°C warmer world compared to present-day conditions. The results are based on kilometre-scale storyline simulations in which the large-scale atmospheric flow is nudged to ERA5, performed with IFS-FESOM (at 9 km in the atmosphere and approximately 5 km in the ocean) on LUMI. They suggest that the heatwave would be exacerbated in the interior of the continent.

Non-hydrostatic model developed by the German Meteorological Service (DWD), the Max Planck Institute for Meteorology (MPI-M), the German Climate Computing Centre (DKRZ), the Karlsruhe Institute of Technology (KIT), and the Center for Climate Systems Modeling (C2SM); IFS-NEMO is ECMWF's Integrated Forecasting System (IFS) atmospheric/land model coupled to the latest release of the NEMO ocean model implemented by BSC in collaboration with ECMWF; and IFS-FESOM is the IFS coupled to the Alfred Wegener Institute's Finite-VolumE Sea ice–Ocean Model. They are run at resolutions ranging between 5 and 10 km for the different components (atmosphere, land, ocean and sea ice).

The Climate DT is developed by a strong partnership led by CSC, currently involving 12 leading climate institutions, supercomputing centres, national meteorological services, academia and industrial partners, through a contract procured by ECMWF.

During phase 1 of DestinE, the Climate DT teams worked in close collaboration with ECMWF teams to build the key elements of the system, to deploy it on the EuroHPC LUMI supercomputer, to demonstrate its key innovative features at scale, and to perform the first ever global multi-decadal climate prototype projections at 5 km resolution.

The achievements of the teams implementing the Climate DT are remarkable given the short amount of time since the contract started (October 2022), and since access to LUMI became available (April 2023). These achievements can be summarized as follows.

Given the complexity of the task, developing and deploying the end-to-end Climate DT workflow on LUMI – the first pre-exascale supercomputer in Europe – in less than one year can be considered a real breakthrough. The Climate DT workflow covers the full chain from global km-scale climate simulations to impact-sector applications. It thus links climate models to applications for the sectors most impacted by climate change, via selected use cases.

To enable a step change in the usability of the model output for applications, the raw output of the different models is standardised by being transformed into a standardised Generic State Vector. This is an important novelty of the Climate DT, which ensures a unified model output format, in terms of parameters, units, and grid, facilitating consistency and interoperability across models and datasets.

The applications included in the workflow are an application for quality assessment and uncertainty quantification (AQUA), as well as impact-sector applications (e.g. a hydrological model or computations of relevant indicators). These transform the climate data

into actionable information on climate change impacts, through a co-design approach.

Another achievement of the Climate DT teams in phase 1 consisted in developing the basis for the future operationalisation of the system in phase 2, in close collaboration with ECMWF. This implied implementing the software infrastructure, workflow and data handling capabilities as part of the DestinE Digital Twin Engine. A first important and challenging task was to deploy the climate models on LUMI. They had to be adapted and optimised to run efficiently on this first pre-exascale EuroHPC platform. As a result, ICON is now fully running its atmosphere component on the AMD GPU partition, and the IFS also runs partially on the AMD GPUs. A second task was to develop and deploy the Autosubmit workflow manager, which allows the Climate DT components to be orchestrated on EuroHPC. A third task was to define and implement a common data portfolio, with data governance that follows international World Meteorological Organization (WMO) standards and WIS2 access patterns (<https://community.wmo.int/en/activity-areas/wis/wis2-implementation>), and a monitoring framework that contains a series of checks ensuring quality control.

Having the system ready for the production of the first-ever prototype projections at km scales in December 2023 was another major milestone during phase 1. The following simulations were carried out from February to April 2024, thanks to a special reservation on LUMI granted by the EuroHPC JU:

- The first ever multi-decadal prototype climate projections at ~5 km across Earth system components with the IFS-NEMO and ICON, streaming information to selected applications (e.g. wind energy and urban heat). ICON was run at 5 km across Earth system components and completed the period 2020–2032. The IFS was run at 4.4 km for the atmosphere and land, coupled with NEMO at 1/12 degree for the ocean and sea-ice, for the period 2020–2040.
- Historical simulations for the past period (starting in 1990) were initiated (at about 10 km resolution) for ICON and IFS-NEMO, and have completed so far 20 and 12 years, respectively.
- Control simulations (starting in 1950), which allow the quantification of model drift, will also be carried out in the next few months (at 10 km resolution).

Another notable innovation of the Climate DT is the production of storyline simulations using IFS-FESOM. This framework is designed to create global, kilometre-scale storylines of recent extreme weather events (2017 to today), illustrating the effects of climate change and supporting more effective adaptation and

mitigation strategies. The simulations reconstruct the unfolding of recent extreme events, such as heatwaves, floods, and droughts across various climate conditions (e.g. pre-industrial, present-day, a +2°C and +3°C world) by nudging the atmospheric circulation in IFS-FESOM for different climate conditions to ERA5 reanalysis data.

Weather-Induced Extremes Digital Twin

The Weather-Induced Extremes Digital Twin (Extremes DT) is a tool that enables the prediction of extreme weather events and their impacts two to four days in advance. It achieves this by using bespoke high-resolution numerical simulations and their fusion with Earth system observations. This capability will support decision-makers in their efforts to respond and adapt to meteorological, hydrological and air quality extremes. It relies on a global component developed by ECMWF and a regional component developed by a strong partnership led by Météo-France, involving 32 institutions in over 20 countries, including many national meteorological services throughout Europe.

The global component is based on IFS simulations at a 4.4 km resolution to produce four-day forecasts and is now initialised once a day on the EuroHPC supercomputer LUMI. In the meantime, Météo-France and its partners are concentrating on the regional component, which builds on the ACCORD system and will be configured and activated on demand to zoom in on selected extreme events over Europe at sub-kilometre resolutions (500 to 750 m). Several impact-sector elements for floods, air quality and renewable energy have already been integrated into the two components.

The Extremes DT will complement existing capabilities for extreme weather and impact prediction, supporting national authorities in making informed decisions to respond and adapt to the potentially devastating consequences of extreme events. The regional component of the Extremes DT can, for example, add value when some extreme events affect multiple countries simultaneously, such as windstorms or heatwaves, which can have a footprint of more than 1,000 km. This will be achieved by following the progression of the event in a more continuous way than is possible using existing regional prediction systems. In particular, the ability to configure any regional domain, in addition to the fixed domains used by national meteorological services, can provide local information where it is currently lacking. Likewise, the renewable energy market is vulnerable to weather variations. Enhanced management of production on a continental scale, encompassing all energy sources and potentially considering inter-country energy transfers, can offer an important advantage.

The integration of multiple sectoral models within the same digital twin workflow will facilitate progress in the forecasting of multi-hazard risk, a capability which has not yet been fully developed.

In addition to providing data close to real time, the Extremes DT can also be used to understand past and future extreme weather events, their change in a warmer climate, and the potential efficacy of adaptation measures in mitigating their impact. This can be achieved by exploiting the synergies of having the same underlying global model in the Extremes and Climate DTs to initialise, and drive, the regional component from either the global Extremes DT or from Climate DT simulations.

The achievements of the team implementing the two components of the Extremes DT are remarkable, especially given the short amount of time since the contract for the on-demand Extremes DT started (September 2022) and since access to LUMI became available (April 2023).

The first achievement was to design and deploy a global Extremes DT workflow that enables the delivery of information within a distributed environment. This involved overcoming challenges, such as using the operational ECMWF analysis produced on ECMWF's high-performance computing facility (HPCF) as initial conditions for the Extremes DT global simulations run on LUMI. Indeed, the two nested components of the Extremes DT had never before run on EuroHPC supercomputers. To a large extent, this is made possible by the capabilities of the Digital Twin Engine. Adapting and running the IFS on the mixed architecture of LUMI that includes the AMD GPUs is a major achievement, only made possible by building on the progress made in ECMWF's scalability programme (see Newsletter No. 171, Spring 2022).

Another achievement was to demonstrate the capacity to generate global km-scale and regional sub-km-scale weather forecasts on a regular basis within a new dedicated workflow. The global component has been run every day on ECMWF's HPCF since August 2023 and on LUMI since December 2023. Also, a mechanism to detect extreme weather in the ECMWF ensemble forecast (ENS) that can be used to trigger the regional component automatically is running every day on ECMWF's HPCF. The forecasts produced with the global component are pushed every day to ECMWF's web interface ecCharts, where they can be analysed by the ECMWF team in charge of the daily verification of ECMWF's operational forecasts.

About ten extreme events, mostly windstorms, were simulated with the regional component on ECMWF's HPCF, using initial conditions and lateral boundary conditions from the global component. An end-to-end

capability demonstration showed how the workflow could have predicted the drop in energy production from a wind farm off the coast of Belgium during Storm Eunice in February 2022 (Figure 4). This is invaluable information that would help ensure the continuity of the energy supply in northern Europe if a similar event happened. The first configurations of the regional

component will become available later in phase 2 of DestinE.

A key milestone was to perform the first comprehensive assessment of global forecasts at 4.4 km initialised every day over almost a year. It was found that large-scale skill scores of the global component of the

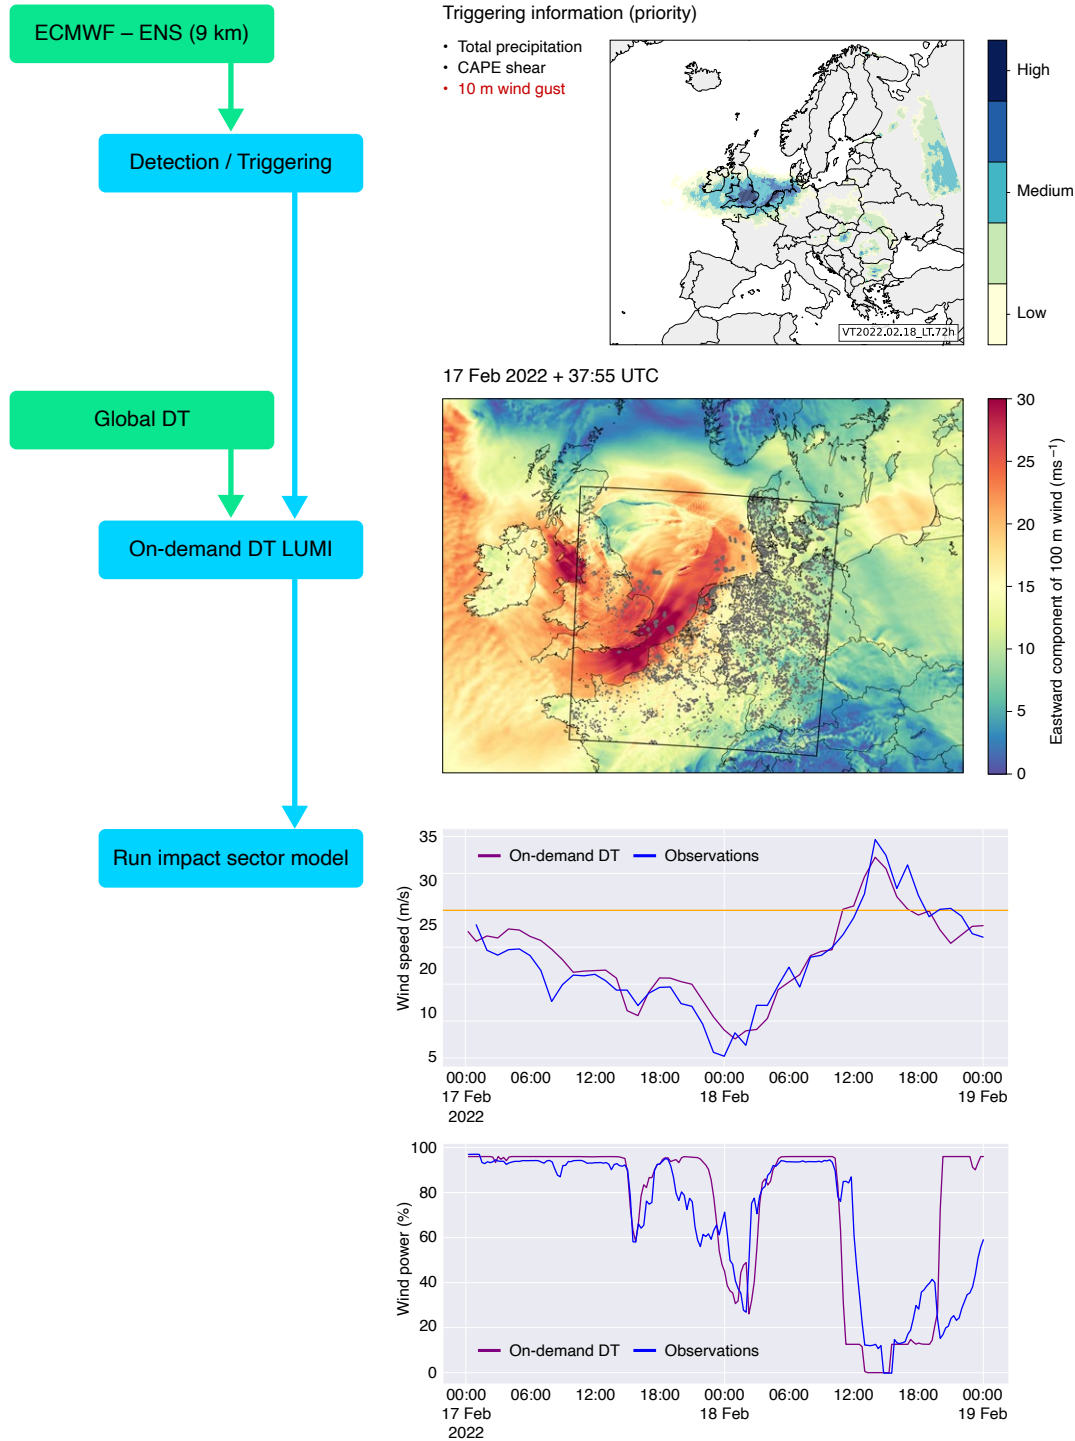


FIGURE 4 End-to-end demonstration of the Extremes DT. A detection mechanism for three hazards in the ECMWF ensemble forecast (ENS) has been running daily since October 2023. In the case of storm Eunice on 18 February 2022, the 10 m wind gust detection triggered the activation and configuration of a regional domain off the coast of Belgium with a resolution of 500 m, and the Global Extremes DT was used to provide the lateral boundary conditions. The high-frequency wind was used by a wind farm application that successfully predicted a drop in energy production due to the shutdown of wind farms two days in advance.

Extremes DT were not consistently better ‘out of the box’ than those of ECMWF operational forecasts, but significant progress is being made to reduce forecast errors. To achieve this, the model time-step and parameters in the convection scheme were tuned, and extra filtering on the small-scale orography was applied. Despite some remaining work to further improve large-scale skill, a consistent enhancement in the forecasting of tropical cyclones and medicanes was demonstrated compared to ECMWF operational forecasts. In addition, intense orographic precipitation was also found to be improved, benefiting from the steeper slopes of the resolved orography and a larger orographic uplift. Likewise, the representation of wind gusts and convective precipitation was often improved when simulating explicitly the sub-km scales in the regional component.

Finally, the added value of running a 4D-Var analysis at km-scale was explored. This involved increasing both the resolution of the background (short-range forecast) trajectory to 5 km and the resolution of the last minimisation to about 25 km. This was found to be beneficial in predicting the correct intensity of the tropical cyclone Otis, which had been missed by all global numerical prediction models, resulting in a major forecast bust. Initialising the tropical depression that led

to Otis's formation was challenging due to its small size, hence the improvement at higher resolution.

Digital Twin Engine – main achievements so far

The Destination Earth information system brings together a diverse range of complex components and interactions. Beyond the digital twins of the Earth system, this includes important elements like workflows, data handling, and supporting infrastructure that will allow selected users to become part of the processing chain while exploiting the most advanced accelerator hardware available on European supercomputers. The Digital Twin Engine enables digital twins to smoothly operate in a physically distributed computing landscape with data volumes and data production rates that go beyond numerical weather and climate prediction practices today.

DestinE developments have significantly accelerated the efforts of ECMWF and our Member and Co-operating States to future-proof their software in terms of both scalability and portability. Adaptation to a range of EuroHPC platforms with different accelerator and hardware technologies was certainly a challenge. The first outcomes included designing, developing, and providing software services that run the digital twins,

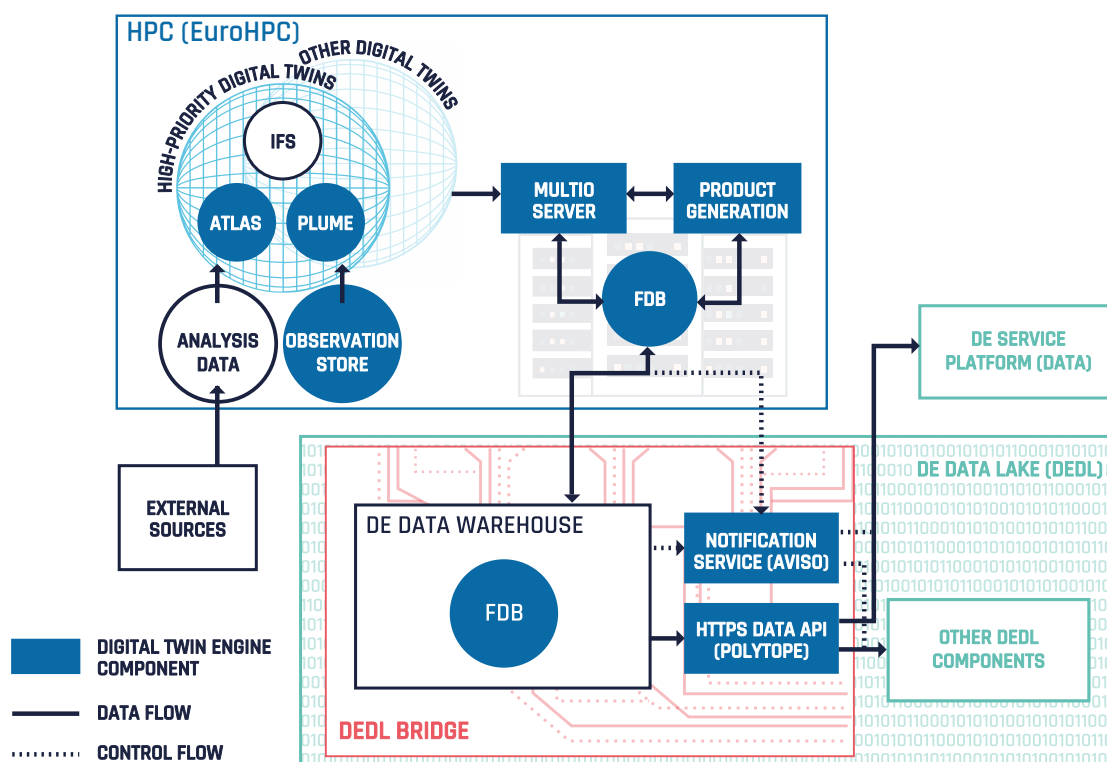


FIGURE 5 Components of the Digital Twin Engine (DTE), spanning multiple infrastructure components of Destination Earth: the digital twins operating on the high-performance computing facility (HPC), the data warehouse embedded within data bridges of the Data Lake (DEDL), and DTE software clients being deployed within the DEDL and the service platform (DESP) services. Atlas is ECMWF’s data structure library; Plume stands for model plugin architecture; MultiIO stands for multiplexing IO-server and on-the-fly post-processing; and FDB is key-value object storage with semantic data access.

provide the associated workflows, deliver and handle the data of the digital twins, and handle notifications in a decentralised, distributed computing/data environment. We created a blueprint for performing km- and sub-km-scale global and regional weather and climate simulations with an evolving process for the research-to-operations transition of weather-event and climate adaptation scenarios on distributed EuroHPC supercomputers. The data produced by the digital twins is being made accessible through standard access procedures. EUMETSAT and ECMWF are working together to ensure the availability and performance of the DestinE digital twin data handling services, e.g. with specific ECMWF software solutions being deployed on the EuroHPC and on the EUMETSAT-provided data lake infrastructure, as well as with client software available on the ESA core service platform.

ECMWF is also working closely with partners in the GLORI DT project (<https://leonardo-supercomputer.cineca.eu/glori-destine-national-dt/>) and a recent pilot project with MeteoSwiss and partners to explore the Digital Twin Engine software stack for connecting to national digital twin and data sharing initiatives. The recent ECMWF developments on novel data access paradigms for extracting features from hypercubes of data (Polytope) and notification services (Aviso) are offered on the DestinE platform (<https://platform.destine.eu>) and make it possible to explore the digital twin data.

Future developments will consolidate the different software components and enrich their feature set. Offering and sharing the developments at an early stage of development has its challenges, but it allows users in our Member and Co-operating States and partners to participate early in the development and feedback loop, using the latest in open software development practices and adhering to standard interfaces for interoperability. The latter is ensured through a rigorous data governance process, which accelerates ECMWF's GRIB2 transition across atmosphere/ocean/sea-ice/wave/land components; adds new features, such as a nested, hierarchical HealPix grid support (and the same for output grids in the ocean and atmosphere across different DestinE Climate DT models); and feeds into the development of ERA6.

Demonstrating the added value of the digital twins through use cases

DestinE digital twins have targeted support to decision-making in a variety of impact sectors from the start. To ensure that, even during early development phases, impact sector needs and user stories are understood and can shape technical and scientific developments, ECMWF has been piloting several 'use cases'. Their aim is:

- to demonstrate compelling applications of the advanced capabilities provided by the two high-priority digital twins of DestinE and the functionalities of the Digital Twin Engine
- to engage users and technology stakeholders from different sectors to help co-design different aspects of DestinE.

The early demonstrators in phase 1 have been implemented by partners contracted by ECMWF in air quality management, flood risk management, climate change adaptation (urban heat, flood risk, forestry), renewable energy management (on- and offshore wind energy, energy systems planning), water management and wildfire risk management (<https://destine.ecmwf.int>).

In each of these cases, one or more models – hydrologic models, energy system models, wildfire risk, etc. – were employed, driven by DestinE digital twin data. For example, under a contract developing a use case for energy systems, the German Aerospace Center (DLR) implemented a representative energy system simulation workflow around their REMix model, which can be driven by DestinE data or other sources. The team demonstrated that different climate datasets in use today strongly impact the results of resource adequacy assessments (Figure 6). During the contract, they are investigating the added value from using DestinE Climate DT data. By the start of phase 2 of DestinE, these demonstrators established workflows that couple to the data streams of the digital twin simulations and produce products for their users. In many cases, users are looking for improvements from the higher resolution of climate projections offered by DestinE, others are interested in the ability to easily configure (impact) model simulations and run them on up-to-date data.

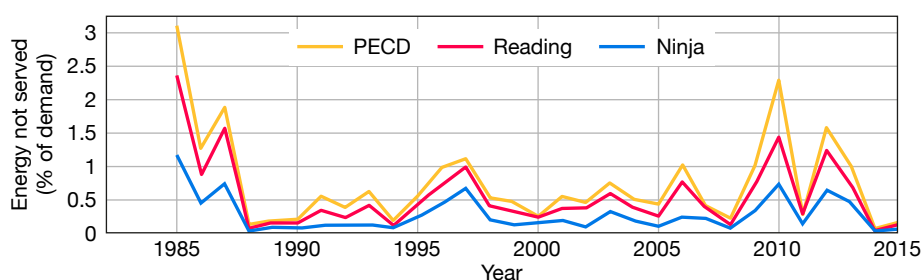


FIGURE 6 Energy not served from 1985–2015 in per cent of the annual demand for Germany under the National Estimates 2030 scenario and for three different meteorological datasets.



FIGURE 7 Director-General of ECMWF Florence Rabier, Director for Destination Earth at ECMWF Irina Sandu, and the European Commission Executive Vice-President for a Europe Fit for the Digital Age Margrethe Vestager (fourth to sixth in the front row from the left) were in Kajaani to share the latest developments and to explain how the DestinE system works.

Fostering innovation through European collaboration

In a short span of time, DestinE has become a European endeavour. ECMWF is working with 90 institutions in the framework of DestinE, and cooperating with many others on the topics underpinning the digital twins and the Digital Twin Engine.

The close cooperation with the involved national meteorological services and climate and supercomputing centres, particularly those involved in the regional component of the Extremes DT and the Climate DT, drove us to interact much more closely, learn from each other, and co-design systems together. It also fostered collaborations with our Member and Co-operating States who are not directly involved in DestinE on related topics through ECMWF pilot projects.

Achieving the objectives of the first phase of DestinE on such a tight timeline would have been impossible without bringing together expertise from a variety of fields – numerical weather prediction, climate science, Earth observations, Earth system and impact modelling, supercomputing, software development, and impact assessment.

DestinE is exploiting this common expertise and knowledge in the European Union to translate investments in science, digital technology and supercomputing, in particular in EuroHPC, into value for society. At the same time, it is enhancing expertise and retaining talent in Europe, by training a new generation of weather and climate scientists at ECMWF and directly within EU member state services, working together in a distributed operational environment.

Looking ahead

In June 2024, DestinE entered its second phase, which will evolve the DestinE system and ramp-up operations, with a focus on users, consolidation, maintenance, and

continuous evolution of the components of the DestinE system. A high-level event held on 10 June at the LUMI supercomputer centre in Kajaani, Finland, marked the start of phase 2. The European Commission Executive Vice-President for a Europe Fit for the Digital Age, Margrethe Vestager, led a ceremony that gathered representatives of the European Commission, key stakeholders of the initiative, and the three implementing entities. ECMWF was represented by a strong delegation to unveil the progress made in developing the first two high-priority digital twins and the Digital Twin Engine (Figure 7).

Efforts at ECMWF and its partner institutions will continue to focus on the two digital twins and the Digital Twin Engine. The overarching aim is to transition the Climate and Extremes DTs towards operational-level execution. This will be achieved through evolving the configuration developed in phase 1. The focus will be on increasing the efficiency of performing complex Earth system simulations on the EuroHPC supercomputers, refining the quality of the digital twins, and enhancing their interactivity elements. The operational-level execution depends of course on the availability, reliability, and support of EuroHPC supercomputers.

The Digital Twin Engine will also be further evolved with a focus on scalability, portability and interoperability. A key area will be the transition of end-to-end workflows of the digital twins on EuroHPC towards an operational-level execution.

An important aspect of phase 2 is the intensification of the use of recent game-changing breakthroughs in machine learning and artificial intelligence (ML/AI). In particular, ECMWF together with its partners will lead developments towards a machine-learning-based Earth system model that will support quantifying uncertainties of the Earth system digital twins and enhancing their interactivity capabilities.

For all the latest news and developments, visit ECMWF's new DestinE website: <https://destine.ecmwf.int/>.

Introduction of a new ocean and sea-ice model based on NEMO4-SI³

Sarah Keeley, Kristian Mogensen, Jean-Raymond Bidlot, Magdalena Alonso Balmaseda, Sam Hatfield

Since 2018, all operational forecast systems at ECMWF have modelled the evolution of the ocean and sea ice using a dynamical model, based on NEMO3.4.1 and a single-category ice model LIM2 (Keeley & Mogensen, 2018). For the last five years, we have been working to develop an updated ocean and sea-ice model and data assimilation system for use in all of ECMWF's physical modelling systems. The next upgrade of the ocean component of the model will be an updated ocean model based on the code base NEMO4.0.6 of the NEMO community model (<https://www.nemo-ocean.eu/>). This features, among other things, a new spatial grid, new options for the formulation of the surface fluxes, and a new community multi-category sea-ice model SI³ (Sea Ice modelling Integrated Initiative), which builds upon the models previously used within NEMO: LIM3, CICE and GELATO.

In addition to the scientific changes, several technical changes have been introduced, leading to significantly better computational performance in terms of scalability, which reduces the time to solution. The new version of NEMO has been integrated in the ECMWF infrastructure, and it has been further developed and configured. Thus, the upgrade of the ocean modelling component will mirror the fact that ECMWF's Integrated Forecasting System (IFS) makes use of a single-precision version of the model for forecasts while retaining double precision for the analysis. The ocean output will become available in ECMWF's Meteorological Archival and Retrieval System (MARS) in GRIB2 via the MultIO server. There will be no change in resolution; the challenge has been to improve the realism and variability of the ocean surface and to ensure consistency in the surface flux formulation between coupled and uncoupled experiments. The first operational application of NEMO4-SI³ will be the production of the new generation of ocean reanalysis (ORAS6, see article by Zuo et al. in this Newsletter), and its coupling with ERA6. It will also be used by the upcoming seasonal forecasting system SEAS6. The plan is for NEMO4-SI³ to be the ocean component of all ECMWF forecasting and monitoring systems from IFS Cycle 50r1, which is to be introduced in 2025. The details of the ocean data assimilation upgrades and the corresponding ocean reanalysis are discussed in more

detail in the article on ORAS6 in this Newsletter. This article focuses on the features of the updated elements of the model in terms of ocean forecasts and impacts on the atmosphere. For the rest of this article, we refer to the new model version as NEMO4 and the model it is replacing as NEMO3.4.

Main developments from NEMO3.4 to NEMO4

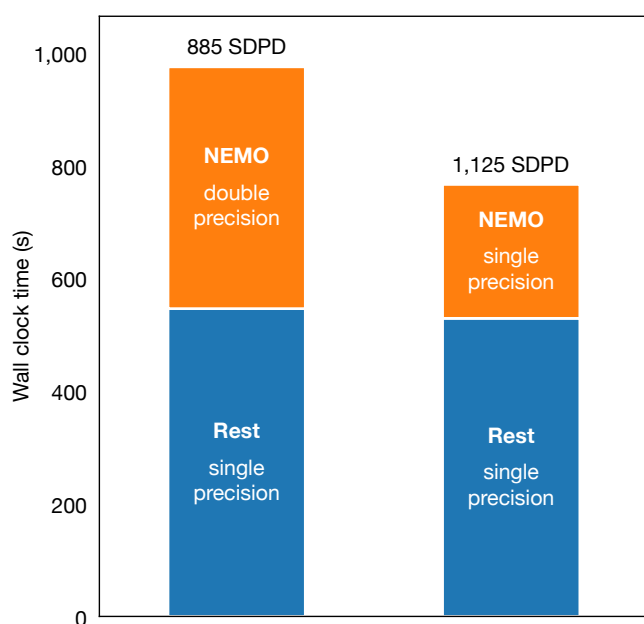
The NEMO community model has undergone much change since our last operational version was implemented. The code has many differences both technically and scientifically, including the replacement of LIM2 with SI³. When implementing the NEMO4 ocean model, we have tailored it for our use within our operational systems, alongside the reimplementing of changes we previously made in house to NEMO3.4 to improve computational efficiency. These changes are made based on the forecast time horizons we are interested in, robustness and computational efficiency of the code for operational implementation, and the ability to initialise the forecasts. This required many hundreds of model simulations to be run, assessing the impact of the choices that can be made to optimise the model setup. There are also a number of new options in NEMO4 that have been provisionally tested during the development phase. They include the representation of icebergs and new sea-ice rheologies, and they can be activated at a later date should they be found to improve performance. This article reviews the key features of NEMO4 and summarises the key performance metrics for the new system. The main change in physical representation in the upgrade from NEMO3.4 is the formulation of the way the ocean surface is represented. We now use a variable volume formulation for the upper surface layers of the ocean. The upper ocean turbulent mixing is modified when sea ice is present, in which case the impact of waves on the mixing is attenuated.

Collaboration with Member States

Implementation of the new model drew on ECMWF Member State expertise with ocean sea ice model development work. The NEMO team at LOCEAN in France provided support, especially with the new sea ice model code as we tested it for operational use. During much of the development phase, we have worked alongside the UK Met Office Joint Marine Modelling

Programme, which draws on oceanographic modelling expertise across the Met Office and UK research community. Our NEMO4 model grids and bathymetry are based on the UK Met Office configuration of its Global Ocean 8. The model grid and bathymetry use the extended ORCA (eORCA) grid (Madec & Imbard, 1996); the ocean model boundary is further south and now extends beyond 78°S to 86°S. The extended grid is required for any future development work to model the ocean interaction under Antarctic ice shelves, although this is not currently considered. The resolution remains the same as in the current operational model, with 75 levels in the vertical and a tri-polar grid with a nominal 0.25-degree horizontal resolution.

Building on earlier work by the Barcelona Supercomputing Center, we have also added the capability to run NEMO4 using single-precision arithmetic (see Rodwell et al., 2021). We carried out extensive testing to make sure that our version of NEMO4 can run stably with single precision, and that there is no impact on forecast scores. Nevertheless, as with the atmospheric component of the IFS, we will keep using double precision as a default for the ocean analysis. We find that the single-precision implementation can accelerate NEMO by up to 1.8 times, compared with the double-precision version. However, this depends strongly on the context in which NEMO is running. In cases where NEMO is running at the limits of its scalability (for example, when using many hundreds of nodes for eORCA12 simulations), the gains from using single precision diminish dramatically. For the seasonal and extended-range forecasting systems, where the ocean cost is a large part of the total run time, the effect is very important, as illustrated in Figure 1.



Atmospheric forcing of the ocean

When developing the model and generating the ocean analysis, we run the model with atmospheric forcing. This is referred to as bulk forcing of the ocean and represents the momentum, heat, and moisture exchange over the ocean. The exchange of fluxes between the atmosphere and ocean is based on roughness length scales for these three processes in the IFS. The quality, frequency and available variables of a driving dataset can influence the performance of the ocean-sea-ice model. Since the last ocean model was implemented, improvements have been made in the forcing data from the atmosphere, which are now from ECMWF's ERA5 reanalysis rather than ERA-Interim. Not only is the quality of the driving data improved on that of the previously used ERA-Interim reanalysis, but the higher frequency (hourly) forcing has contributed to a better representation of the diurnal cycle. Results of this can be seen in Figure 2 of the article on ORAS6 by Zuo et al. in this Newsletter. New bulk formulation options that have been added within NEMO4 allow increased consistency between the surface fluxes used in uncoupled ocean mode (which is used for the production of ocean reanalysis) and coupled mode (which is used for the coupled forecasts). Given that the wave model is an integral part of our forecasting system and mediates the surface exchange between the atmosphere and ocean, we run the NEMO4 system with the surface wave impacts. The aim is to improve the upper ocean representation and aid the development of the ocean model for coupled implementation within our forecasting systems. Wave data are also available from ERA5. The use of wave forcings was first introduced in our version of NEMO3.4 (Breivik et al., 2013), and many of these wave forcings are now available options in NEMO4. To be consistent with our earlier version of NEMO, we have re-introduced the impact of wave breaking on the turbulent kinetic energy (TKE) surface flux in the parametrization of the upper ocean mixing. Where data that is required is not available from ERA5, it has been recomputed from existing available data.

For momentum under windy conditions, the IFS uses a sea-state-dependent Charnock relation. We modified NEMO4 from using a constant Charnock coefficient to a varying Charnock field that is consistent with the other wave forcing fields. The surface fluxes in NEMO4 are adjusted iteratively to account for the presence of surface currents. Naively, when using atmospheric forcing that has


FIGURE 1 Impact of using single precision in NEMO instead of double precision on wall clock times in seconds and simulation speed in simulated days per day (SDPD) for a coupled IFS benchmark equivalent in resolution to the seasonal and extended-range forecasting systems (eORCA025 resolution in NEMO and TCo319 resolution in the IFS atmosphere). The boxes labelled 'Rest' show the wall clock times for all parts of the IFS excluding NEMO. These already use single precision.

a

Multi-category ice

The model represents the subgrid-scale heterogeneity in ice thickness through a multi-category formulation. Ice thickness is therefore another dimension in the prognostic variables, and the thickness distribution is discretised into five categories with fixed thickness ranges. Each category has an associated concentration, discretised into a number of evenly spaced layers, and the thermodynamics of the model is solved for each category. As ice forms or melts, there is a redistribution in the fraction of each thickness category that is present at each model point. It should be noted that the dynamics of the model is solved for each grid box of the model and not for individual categories, and it represents the ridging and rafting of ice when under compression.

Schematic of the five-category, 4 + 1 ice and snow layer model.



not used information on surface currents, such as ERA5, one would assume that the surface fluxes should be computed based on 10 m winds relative to the surface currents. However, we have shown that, if the atmospheric model had used surface current information, the whole near-surface wind profile would be modified in response to the modified surface condition. On average, in order to yield similar surface stress values using the data from the system without surface currents, one should compute the atmospheric forcing with about half the current effect. Therefore, we opted to only apply half of the surface current effect in the relative wind correction in NEMO4.

ERA5 and operational wave data are not available when sea-ice cover is above 30% due to the lack of a physically consistent representation of wave–sea-ice interactions. Therefore, when sea ice is present in the forcing field, the wave forcing data are determined using simple empirical formulations based on 10 m wind and sea-ice cover or simply using a constant default value.

New sea-ice model – SI³ (sea ice cubed)

Perhaps the largest change in the model setup is the way sea ice is represented. LIM2 is no longer available; NEMO4 contains the new community SI₃ model. Operationally LIM2 represents sea ice with a single category and two thermodynamic layers, with a single layer of snow on top. Sea ice has a fixed salt concentration and no possibility of melt pond formation.

The sea-ice rheology of LIM2 is based on continuum dynamics and uses a viscous plastic rheology. For SI³ we make use of a computationally more efficient formulation of this rheology, the elastic-viscous-plastic formulation (EVP).

SI³ brings many new features to the way the thermodynamics of sea ice is modelled at ECMWF that improve the representation of the seasonal evolution of the ice:

- multi-category ice (see Box A for more details)
- spatially and temporally varying sea-ice salinity
- melt ponds
- new prognostic variables to describe the thermodynamic sea-ice properties.

The model has been developed with five thickness categories, each with four thermodynamic layers and a single snow layer on top. This configuration is a compromise to represent the complexity of the ice without being too computationally expensive. This change in model formulation required a change in the way sea-ice increments are applied in the data assimilation system (see the article by Zuo et al. in this Newsletter).

The other improvement in the new model that is important for the surface exchange over sea ice is the

representation of various surface processes and properties. For example, a level pond scheme is activated, and each category has its own pond fraction and volume. The pond can have a frozen lid, which also affects the surface albedo of the pond. The albedo scheme has been slightly reformulated from that used in LIM2, and the effect of melt ponds is also included. Snow on sea ice is also considered in more detail: as snow falls onto the ice, some fraction gets blown into the ocean, and snow is also lost to the ocean as the ice ridges and rafts. The representation of snow loading has become more critical for the surface exchange as we wish to make use of the snow depth from the ocean analysis to initialise the snow on sea ice within the IFS surface analysis. Representing snow on sea ice helps to reduce warm temperature biases in winter that have been reported both operationally and for ERA5 (see Arduini et al., 2022, for more discussion).

Comparing standalone NEMO4 with NEMO3.4

To assess the performance of the ocean model, key model

fields are compared to observational datasets. In preparation for providing a model basis for the analysis system for ORAS6, we compared free-running experiments of the uncoupled ocean model, both the operational NEMO3.4 model and the NEMO4 model. Early development work was done using a 1° ocean model and then further refined and tested at 0.25° . We present results here for the latter. Each model is started from a climatological state and run from 1979 for 40 years with hourly ERA5 forcing without data assimilation. The focus of the development work was to provide a model setup that was as good as NEMO3.4, or better, on which to build the assimilation system. We present results with the final setup that was provided for the data assimilation system. To avoid any spin-up from the climatological state in the upper ocean and sea ice, our analysis focused on the second half of the integration (the last 20 years) and is compared to upper ocean observations. Figure 2 shows the differences in root-mean-square error (RMSE) between NEMO4 and NEMO3.4 simulations for sea-surface temperature (SST – verified against ESA-CCI SST v2 analyses), upper ocean 500-metre temperature and salinity (verified against in-situ observations), and sea-level

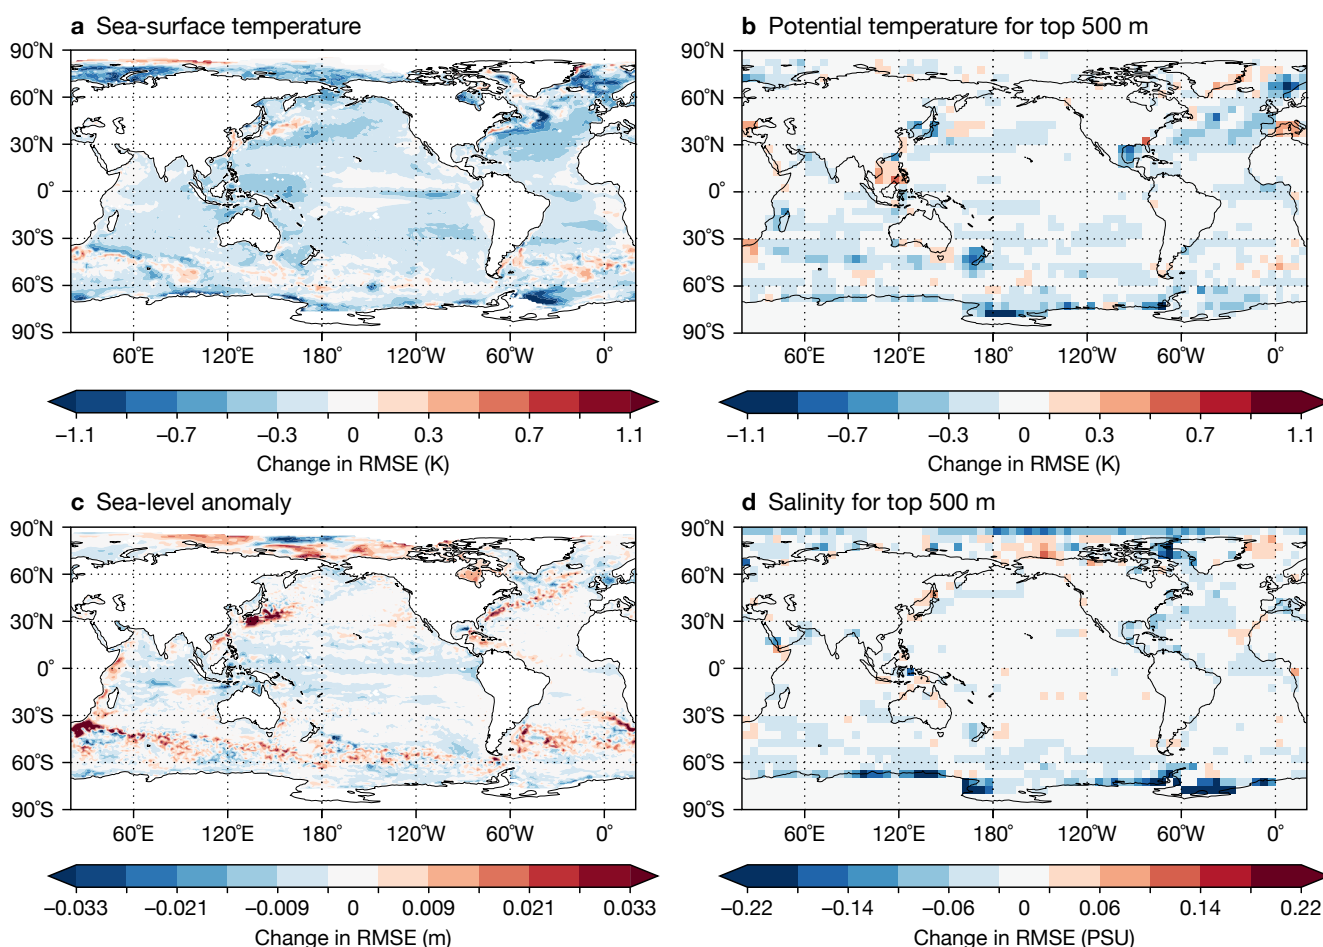


FIGURE 2 Annual mean changes in root-mean-square error (RMSE) for the period 1999–2019 for NEMO4 relative to NEMO3.4. Negative values show an improvement when using NEMO4. Results are shown for (a) satellite-derived sea-surface temperature (SST), (b) potential temperature for the top 500 m of the ocean from in-situ observations, (c) sea-level anomaly from in-situ observations, and (d) salinity for the top 500 m of the ocean from in-situ observations.

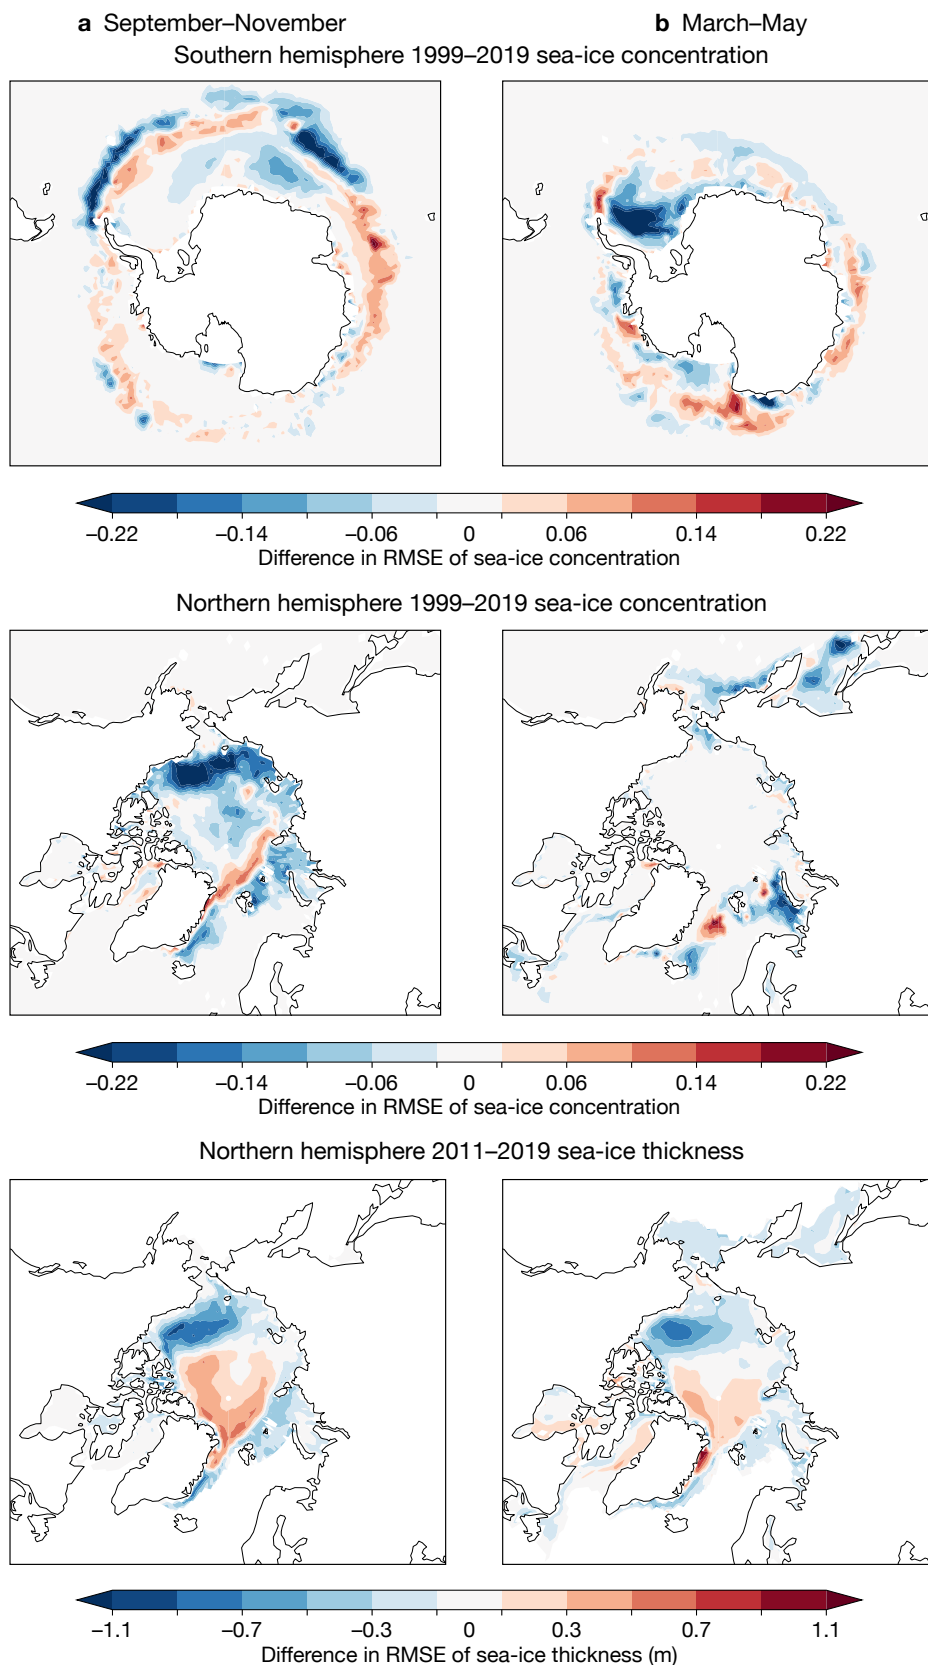


FIGURE 3 Mean changes in root-mean-square error (RMSE) relative to satellite estimates for NEMO4 relative to NEMO3.4 in (a) September–November, and (b) March–May. Negative values show an improvement when using NEMO4. Results are shown for sea-ice concentration for the southern hemisphere (top row) and northern hemisphere (middle row) for the period 1999–2019. Changes in RMSE sea-ice thickness are shown for the northern hemisphere for the period 2011–2019.

anomaly (verified against altimeter data). The free-running NEMO4 shows errors are broadly reduced compared to NEMO3.4 for all the observed quantities, except for an increase in sea-level anomaly RMSE in eddy-rich regions. Improvements can be seen in SST in areas with known

errors in the Southern Ocean and the North Atlantic Current region.

The multi-category sea-ice model freezes and melts sea ice more rapidly than a single category model. We would, therefore, expect the new model to help reduce errors in

the marginal ice zone (MIZ). Figure 3 shows the change in RMSE for NEMO4 compared to NEMO3.4 for autumn and spring, when the marginal ice zone is at its minimum and maximum, respectively, for concentration and thickness in the northern hemisphere, where satellite data products exist. In the northern hemisphere, there is a broad reduction in errors in sea-ice concentration, although there are some increased errors in the transition between the MIZ and the perennial ice in the region that extends from northeast Greenland towards northern Russia. In the southern hemisphere, the results are more mixed for the maximum ice extent, with increased errors in the Davis Sea. Large reductions in error are seen in autumn in the Weddell Sea, which is an area sensitive to the snow loading on the ice. The multicategory model is also found to have thinner ice on average in the Arctic (not shown), compared to satellite derived sea-ice thickness products from CryoSAT2 and SMOS. The model removed the positive thickness bias in the Beaufort gyre but has increased the negative thickness bias to the north of Greenland. This leads to increased RMSE in the central Arctic in the sea-ice thickness field, shown at the bottom of Figure 3.

Easier access to ocean data

The implementation of NEMO4 brings with it a change in the way we archive data from the ocean. In the past, the data was output in NetCDF format and could not be stored in the MARS archive. The technical work to deliver ocean output that is more accessible is described in more detail by Sarmany et al. (2024). It will allow users to access the data from across Earth system model components in the same way.

Impact on coupled system

One of the biggest improvements with the change of sea-ice model and the representation of surface processes is that we can now consider coupling more of the thermodynamic components of the ice. Early testing showed that the marginal ice zone (MIZ) being better represented by the multi-category ice was detrimental to coupled forecasts when using the fractional ice-cover coupling that is used operationally. The operational system assumes that the ice is 1.5 m thick and has no snow on top. These assumptions lead to a warm bias in the polar regions, where the snow insulation effects are not correctly represented, and a cold bias when the thinner ice in the marginal ice zone would allow heat fluxes from the ocean to the atmosphere. Without being able to represent the melt ponds, the summer albedo was not low enough in LIM2, so we made use of a uniform climatological albedo over sea-ice surfaces in the IFS. With this simplification, we found that the modelled sea-ice response is slower than observed. This was particularly true in summer, when there are strong albedo feedbacks. The new elements of SI³ mean that we can make improvements to the sea-ice/atmosphere interface and reduce model biases that are

seen in the reanalysis and forecast systems. We can now implement closer coupling between the sea-ice surface and the atmosphere, coupling the albedo, snow depth and temperatures. Ongoing testing is finalising the coupling framework for the use of NEMO4 within our forecast systems.

Conclusion

Considerable work was done to make sure that the NEMO4 code base that will be implemented in operations is robust and tuned, to provide a system with a baseline requirement to perform as well as its predecessor NEMO3.4. Alongside the scientific development, there have been technical developments to enable the output from the model to be archived in MARS using GRIB2.

NEMO4 is now being used in production of the next ocean reanalysis ORAS6 and will be used for the operational ocean analysis OCEAN6. More details about the specific development work for the assimilation system have been documented by Zuo et al. in this Newsletter. The model will also form the ocean component for all the coupled systems from IFS Cycle 49r2 onwards. The first time this will be implemented in the operational forecast will be in IFS Cycle 50r1.

Further reading

Arduini, G., S. Keeley, J. J. Day, I. Sandu, L. Zampieri & G. Balsamo, 2022: On the Importance of Representing Snow Over Sea-Ice for Simulating the Arctic Boundary Layer, *Journal of Advances in Modeling Earth Systems*, **14/7**. <https://doi.org/10.1029/2021MS002777>

Breivik, Ø., M. Alonso-Balmaseda, J.-R. Bidlot, P.A.E.M. Janssen, S. Keeley & K. Mogensen, 2013: Closer together: coupling the wave and ocean models. *ECMWF Newsletter No. 135*, 6–7. <https://doi.org/10.21957/ha77zhpn>

Keeley, S. & K. Mogensen, 2018: Dynamic sea ice in the IFS, *ECMWF Newsletter No. 156*, 23–29. <https://doi.org/10.21957/4ska25furb>

Storkey, D., A.T. Blaker, P. Mathiot, A. Megann, Y. Aksenov, E.W. Blockley et al., 2018: UK Global Ocean GO6 and GO7: a traceable hierarchy of model resolutions, *Geosci. Model Dev.*, **11**, 3187–3213. <https://doi.org/10.5194/gmd-11-3187-2018>

Rodwell, M., M. Diamantakis, P. Düben, M. Janoušek, S. Lang, I. Polichtchouk et al., 2021: IFS upgrade provides more skilful ensemble forecasts. *ECMWF Newsletter No. 168*, 18–23. <https://doi.org/10.21957/m830hnz27r>

Sármány, D., K. Mogensen, M. Valentini, R. Aguridan, P. Geier, J. Hawkes et al., 2024: New software libraries for ORAS6 ocean reanalysis, *ECMWF Newsletter No. 178*, 13–15. <https://doi.org/10.21957/1a8466ec2f>

ECMWF's next ensemble reanalysis system for ocean and sea ice: ORAS6

Hao Zuo, Magdalena Alonso Balmaseda, Eric de Boissésou, Philip Browne, Marcin Chrust, Sarah Keeley, Kristian Mogensen, Charles Pelletier, Patricia de Rosnay, Toshinari Takakura

ECMWF has a long history of development and production of ocean and sea-ice reanalysis. Initially conceived to provide initial conditions for seasonal re-forecasts, the user base of the ECMWF ocean and sea-ice reanalysis has increased over the years. This reanalysis supports climate monitoring, the initialisation of extended-range and decadal forecasts, and more recently the initialisation of medium-range forecasts. The 5th generation of ECMWF's ocean and sea-ice reanalysis system, ORAS5, went operational in 2016. It serves as the basis of OCEAN5, which has been providing ocean and sea-ice initial conditions for all ECMWF coupled forecasting systems ever since (Zuo et al., 2019). ORAS6 – the 6th generation of ocean and sea-ice ensemble reanalysis – will have an important addition: the provision of surface boundary conditions for the upcoming atmospheric reanalysis ERA6.

As a consequence, the faithful representation of sea-surface temperature (SST) and sea-ice mean state and variability, including fast processes such as the diurnal cycle, has been a priority in the development and design of ORAS6. For this to happen, variational assimilation of SST in the ocean model was implemented for the first time in the ECMWF system. This required the representation of flow-dependent error covariances via an Ensemble of Data Assimilations (EDA) in the ocean. The SST assimilation, together with the availability of ERA5 hourly surface forcing, has made it possible for ORAS6 to represent the diurnal cycle in SST, which is one of the main highlights of this new system. ORAS6 brings another important novelty: the availability of its output in the ECMWF Meteorological Archival and Retrieval System (MARS) in the generally used format for weather data, binary GRIB2 (see Sármany et al., 2024). Other innovations include: an update of the ocean and sea-ice model, the latter bringing significant physical upgrades; enhanced ocean and sea-ice observation usage; an updated bias correction scheme; and ensemble generation. All of these have been made possible by more than eight additional years of research and development at ECMWF and close collaborations with meteorological services and research institutions in

ECMWF Member States. ORAS6 has been developed as a dedicated cycle of ECMWF's Integrated Forecasting System (IFS), Cycle 48r2. The evolution of ORAS6 is strategically determined by the impact on coupled forecasting systems and performance assessments, including feedback from end user groups. A complete ORAS6 record, back to 1950, is expected to be available by the end of 2024.

System updates and performance

In ORAS6, the ocean and sea-ice model has been updated and is now driven by hourly ERA5 forcings. A novel ensemble-based variational data assimilation system has been developed that provides flow-dependent background error variances and vertical correlation scales. Other updates include the assimilation of level-4 (L4) SST and level-3 (L3) sea-ice concentration data, an improved ensemble generation method, and updates in the bias correction scheme and altimeter data assimilation. Figure 1 is a schematic plot of the ORAS6 system, highlighting major system improvements compared to ORAS5.

Summary of ocean and sea-ice model updates

The ocean and sea-ice model is now based on NEMOv4.0.6, which includes a new multicategory sea-ice model SI³, with prognostic salinity. This is a step change in modelling capability for sea ice. The assimilation system has had to be adapted to account for new prognostic variables in the sea ice, as well as for the concept of representing the ice with multiple thickness categories. The ocean model itself has also changed by extending the southerly limits of the grid around Antarctica, and by modifying the way atmospheric forcing is applied to the ocean in a standalone model (i.e. not coupled to an overlying atmosphere). The resolution of the ocean remains unchanged at ¼ degree and with 75 vertical levels. The forcing is applied hourly from ERA5 and has a new bulk formulation for the computation of fluxes that is much closer to the ECMWF operational forecasting system. Thanks to this hourly ERA5 forcing and switching to direct assimilation of SST observations, ORAS6 represents SST with a diurnal cycle closely matching that reconstructed from drifter data (see Figure 2). More details about the model developments

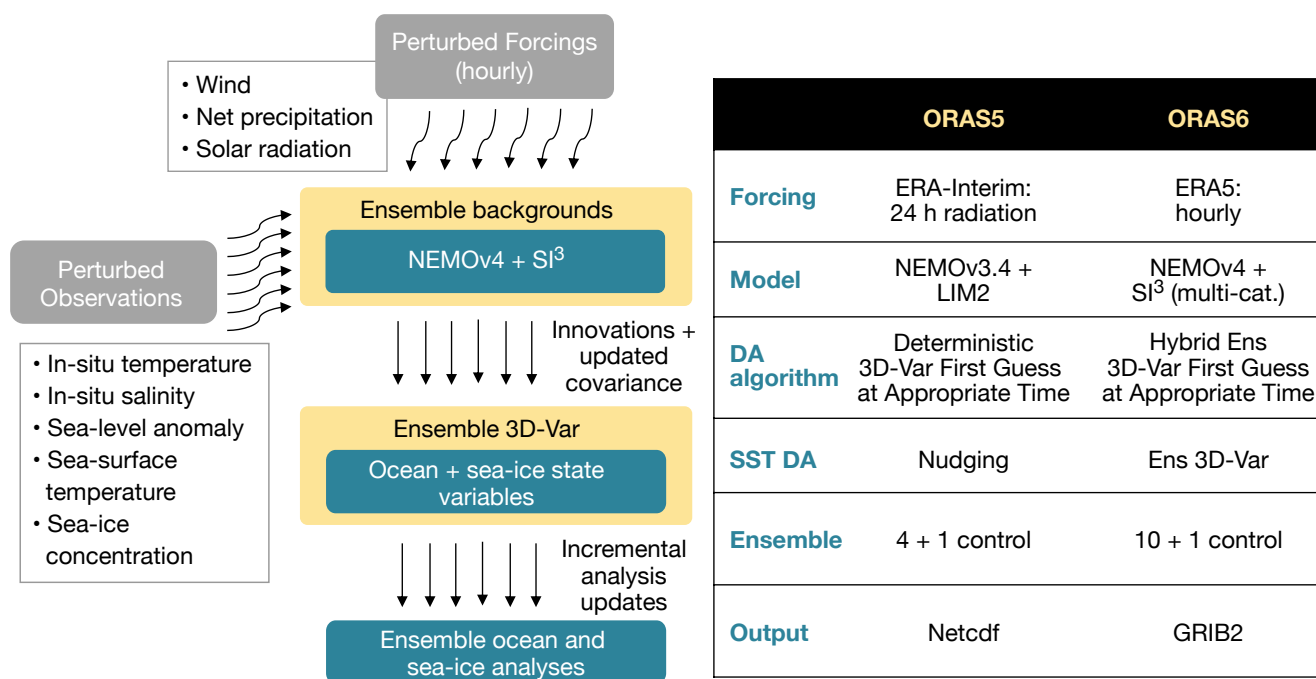


FIGURE 1 Schematic diagram of the ORAS6 system. Instead of being forced in 24-hour intervals by ECMWF's previous reanalysis, ERA-Interim, ORAS6 will be forced by hourly input of the ERA5 reanalysis. The ocean and sea-ice background are now based on NEMOv4 and the Sea Ice modelling Integrated Initiative (SI³), rather than NEMOv3.4 and the Louvain-la-Neuve Sea Ice Model (LIM2). Data assimilation (DA) of sea-surface temperature (SST) uses ensemble-based three-dimensional variational methods (Ens 3D-Var) rather than simple nudging. ORAS6 will include 11 ensemble members of ocean and sea-ice analyses in GRIB2 data format.

can be found in the NEMOv4 article in this Newsletter.

SST data assimilation with Ocean EDA system

Direct data assimilation of L4 SST data has been introduced in ORAS6. One major challenge in SST data assimilation is that the background error covariance formulation used in ORAS5 is not optimised to deal with very dense satellite SST observations together with a sparse ocean in-situ observing network. To solve that, ORAS6 features deployment of an Ensemble of Data Assimilations successfully developed by ECMWF in collaboration with Member State institutes. This was done in the framework of a contract of the EU-funded

Copernicus Climate Change Service (C3S) implemented by ECMWF (Chrust et al., 2021). The hybrid background error covariance matrix is a key element of the data assimilation system, determining the spatial structure of corrections to initial conditions arising from observations. It combines newly computed seasonally varying climatological variances and horizontal correlation length scales with flow-dependent errors of the day and, most importantly, vertical length scales estimated from an ensemble. The latter was enabled by an innovative normalisation algorithm and is crucial for effective assimilation of surface observations. A new version of a reprocessed OSTIA L4 SST product from the UK Met

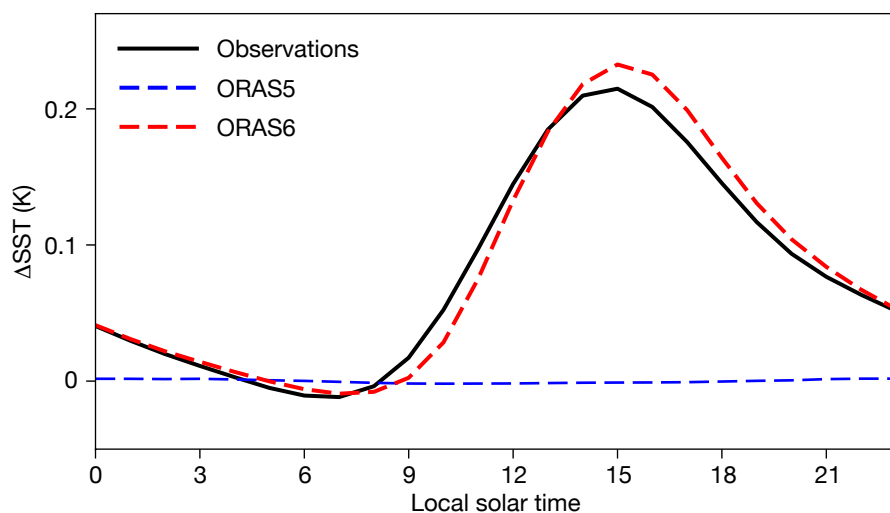


FIGURE 2 Global averaged sea-surface temperature (SST) diurnal cycle from ORAS5 (blue dashed line), the ORAS6 prototype (red dashed line) and drifter observations (black solid line) in 2019. The vertical axis indicates the SST difference (in K) from the mean value for the period around dawn.

Office is assimilated in ORAS6 to ensure temporal consistency to its homonymous real-time product. It has been found that, compared to SST nudging, direct assimilation of SST data with this new ocean EDA system has greatly reduced SST biases, especially in Gulf Stream regions to the north of Cape Hatteras, where the $\frac{1}{4}$ degree NEMO model has a persistent warm water over-shooting issue. Figure 3 shows SST biases in the Gulf Stream region, with clearly superior performance of ORAS6 compared to its predecessor.

Sea-ice data assimilation with SI³

Many developments were required to successfully use data assimilation with the new multcategory SI³ model. Care has been taken to ensure the application of single category concentration increments into the multcategory model does not unduly change the model ice thickness. A balance between sea-ice concentration and sea-water temperature was introduced to enable the model to retain concentration increments and in turn to provide a better fit to observations. ORAS6 assimilates the latest version of the OSI SAF global sea-ice concentration climate data record (L3) from EUMETSAT, based on observations from three satellite instruments: the Special Sensor Microwave/Imager (SSM/I), the Special Sensor Microwave Imager/Sounder (SSMIS), and the Scanning Multichannel Microwave Radiometer (SMMR). New methods on observation perturbations, and a way to account for observation

error variances using uncertainty information, have been developed in order to assimilate L3 sea-ice data. Compared to gridded OSTIA L4 sea-ice concentration data used by ORAS5, switching to L3 OSI SAF data could potentially increase our operational reliability by making our ocean and sea-ice analysis less susceptible to errors in external data production (e.g. data infilling method). The overall performance of sea-ice concentration in ORAS6 is improved when compared to ORAS5 (see Figure 4). Sea-ice thickness properties in ORAS6 have changed substantially due to differences in forcing and the sea-ice model. In some regions of the Arctic, the sea ice in ORAS6 can be biased thin. Attempts to constrain sea-ice thickness in ORAS6 have been made, including nudging towards a merged sea-ice thickness reconstructed dataset from a product based on observations of ESA's CryoSat2 and SMOS satellite instruments (CS2SMOS). This approach was not pursued mainly due to the product only being available in a fairly recent period (from 2011 on), and also due to discrepancies between this thickness product and the assimilated OSI SAF sea-ice concentration product.

Update on ensemble generation

Sampling the uncertainty in the ocean state is important to provide a reliable ensemble of i) ocean initial conditions for coupled forecasts from the medium range to the seasonal range, and ii) lower

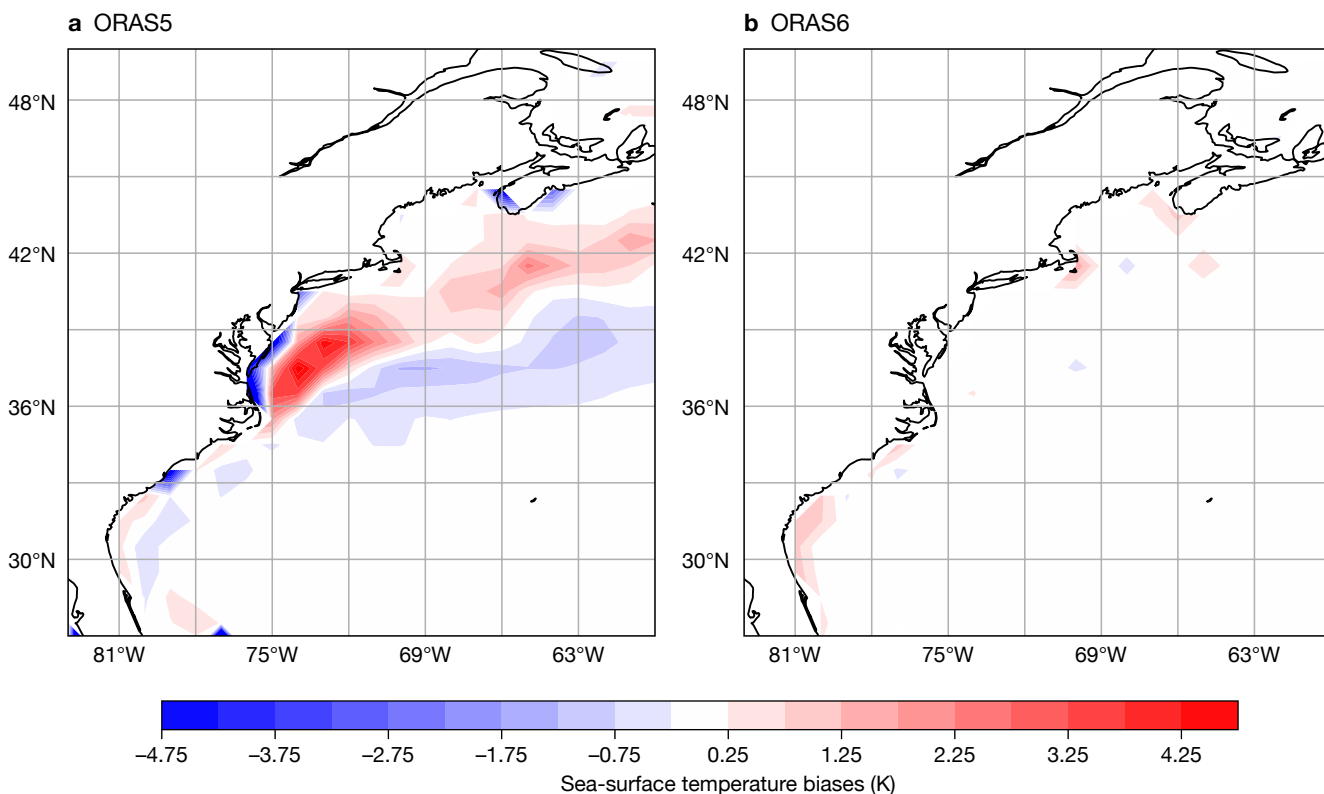


FIGURE 3 Sea-surface temperature (SST) biases (in K) in the Northern Atlantic Ocean in (a) ORAS5 and (b) ORAS6. SST is verified against the European Space Agency (ESA) CCI2 SST dataset in the period from 2005 to 2015.




FIGURE 4 Changes in sea-ice concentration root-mean-square error (RMSE) in the Arctic between ORAS6 and ORAS5. Sea-ice concentration (in percent) is verified against the ESA CCI2 dataset in the period from 2005 to 2015. Blue means a reduction in RMSE with the new system.

atmospheric boundary conditions over the ocean and sea-ice surface for the upcoming ERA6 reanalysis. Reliable model background ensemble variance is also critical in the ocean EDA system in estimating flow-dependent errors of the day. Many new developments have been implemented in the ORAS6 ensemble generation so perturbation can be applied to the newly introduced SST and L3 sea-ice concentration assimilation. ORAS6 uses surface flux forcing from the ERA5 reanalysis. Uncertainty of air–sea fluxes is known to be high and needs to be accounted for in an

ensemble of analysis. The perturbation of the sea-surface forcing is based on a pre-built repository that includes perturbations of wind stress, SST, sea-ice concentration (SIC), freshwater and heat fluxes. It accounts for differences between separate atmospheric reanalysis datasets (Structural Errors – SE), and for uncertainties by sampling all individual analysis members from the same atmospheric reanalysis product (Analysis Errors – AE). In ORAS6, SE are computed for the wind stress using monthly outputs from ERA5 and the Japanese 55-year reanalysis (JRA-55), and for SST and SIC using monthly fields from OSTIA reprocessed analyses and ESA-CCIv2 analyses. AE are computed for SST, evaporation-minus-precipitation (EMP), solar radiation, and wind stress, using all ERA5 ensemble members. The two types of error are considered independent and thus additive. The perturbation repository is stratified by calendar month and date range to capture uncertainties that correspond to the season and the level of sampling of the ocean by the observing system. In addition to surface forcing perturbations, the same perturbations on ocean in-situ observations as in ORAS5 are also activated in ORAS6, including additional stratified random sampling when assimilating SST observations. Figure 5 shows that the ORAS6 SST ensemble spread is generally larger than that estimated from its predecessor product, thanks to this updated ensemble generation scheme.

Other ocean data assimilation system updates

Sea-surface height (SSH) assimilation in ORAS6 has seen several major updates as well. Compared to ORAS5, ORAS6 assimilates a new reprocessed Sea Level Anomaly (SLA) dataset from the multi-satellite Data Unification and Altimeter Combination System (DUACS), made available through the Copernicus Marine Environment Monitoring Service (CMEMS). This provides altimeter data from more satellites with

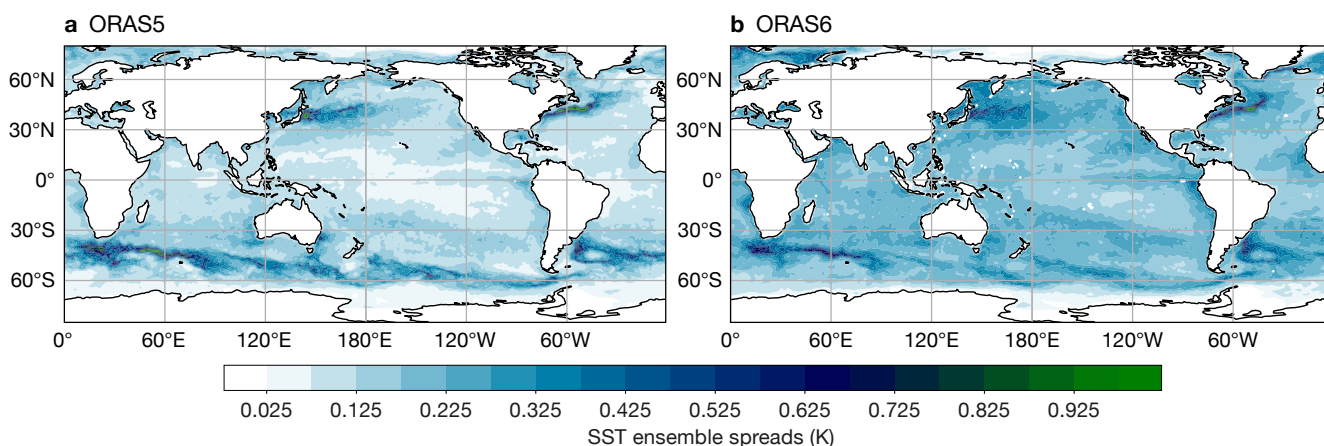


FIGURE 5 Sea-surface temperature ensemble spreads in (a) ORAS5 and (b) ORAS6. Ensemble spreads are computed with monthly mean SST from all ensemble members of ocean syntheses and averaged for 2010–2012.

improved temporal consistency. Unlike in ORAS5, SLA data in high latitudes (outside a latitude band from 50°S to 50°N) and near continental coasts are also assimilated in ORAS6, thanks to reduced regional biases in this new SLA product. As a result, SSH in ORAS6 is much improved compared to ORAS5, especially in regions with additional altimeter data assimilation (see Figure 6). We continue to use a model-based estimation of Mean Dynamic Topography (MDT), which is part of the forward model in altimeter assimilation. While errors in MDT are not explicitly accounted for in the ocean EDA cost function, we have introduced new methods to mitigate biases in MDT, namely a two-step approach and removal of SSH trend errors in pre-production runs used to estimate MDT.

The model bias correction scheme implemented in ORAS6 still comprises a priori bias and online bias to account for biases with different spatial and temporal scales. Compared to ORAS5, the a priori bias terms in ORAS6 are estimated with a multi-step approach, including an additional spatial filter. We have also introduced a new vertical element in addition to the original 2D partition coefficients matrix, which is used to specify partition of the assimilation increments into the different bias control variables. A new analytical function has also been introduced to compute partition coefficients for the online bias correction term. These new developments are required for the model bias correction scheme to work successfully together with other new ORAS6 elements, such as the ocean EDA and direct assimilation of SST observations.

We have introduced a new method to distribute freshwater budget (FWB) adjustments in ORAS6, in which global water mass and volume adjustments are

distributed according to the model background EMP value, by separating regions with positive EMP from those with negative EMP.

Forecast performance with ORAS6

To evaluate the medium-range forecast performance for predicting SSTs, we used a quality-controlled dataset from the Global Drifter Program of the US National Oceanic and Atmospheric Administration (NOAA) (<https://www.aoml.noaa.gov/phod/gdpl/>). Figure 7 shows global change in root-mean-square error (RMSE) with the prototype ORAS6 system compared to the operational ORAS5 system for a one-year period. The new system reduces the error in SST almost everywhere, with improvements being larger in boundary current regions and in the Southern Ocean. Results from testing with a copy of the operational forecasting system with an atmospheric resolution of 9 km (medium-range ensemble resolution) show that the improvement is maintained for longer lead times until day 9 (see Figure 8). Comparing the system performances at different lead times shows that errors in SST predictions at day three of the new system are comparable with those of the old ocean analysis system.

Conclusion

In summary, the design and development of ORAS6 has been guided by the following criteria:

- to maintain the climate quality required for its existing applications
- to improve the representation of the surface of the ocean to cater for the needs of medium-range forecasts and upcoming coupled atmospheric reanalyses

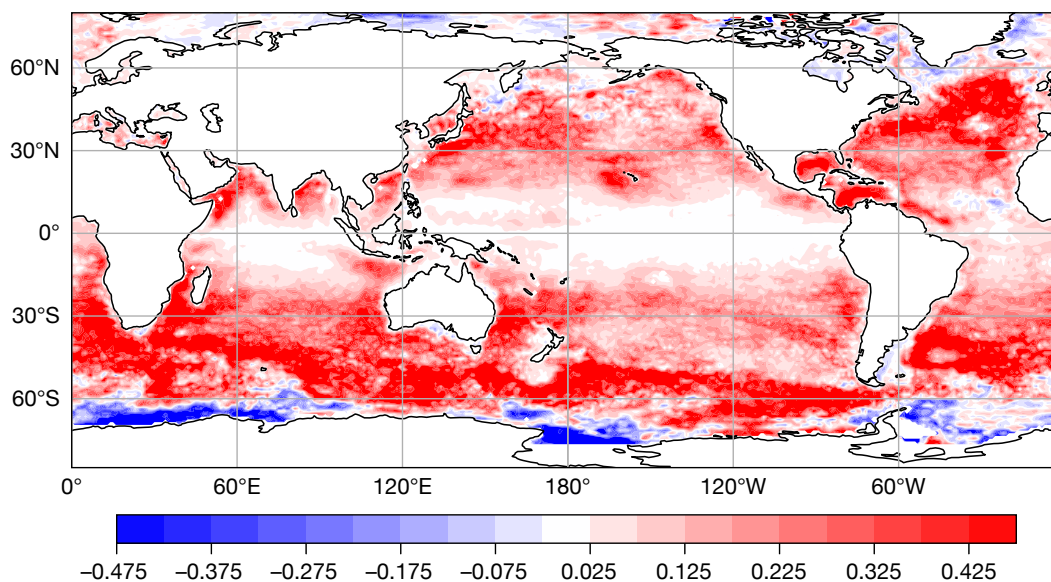


FIGURE 6 Changes in sea-surface height (SSH) temporal correlation between ORAS6 and ORAS5. Temporal correlations are computed with monthly-mean Sea Level Anomaly (SLA) data from ocean syntheses against the ESA CCI2 dataset between 2005 and 2015, after removal of the seasonal cycle. Red means improved temporal correlation with the new ORAS6 system.




FIGURE 7 Changes in sea-surface temperature (SST) root-mean-square error (RMSE) for the ORAS6 prototype compared to the operational ORAS5 system, according to a quality-controlled ocean drifter dataset from June 2020 to June 2021. Blue means a reduction in RMSE with the new system.

- to use the latest scientific developments and modelling tools, so the operational system and its users can benefit from years of research
- to facilitate data access to both external and internal users, as well as long-term data curation.

ORAS6 will span the period from 1950 to the present. To facilitate its timely production and to account for the evolving nature of the global ocean observing system, ORAS6 is being produced in separate parallel streams, which are connected via a continuous prior analysis to ensure there is climate continuity. At the time of writing, the operational production of the most recent ORAS6 stream (2005 to the present) is under way, and work is ongoing on preparing the settings for the early and pre-satellite streams. It is expected that the whole ORAS6 record will be completed by the end of 2024, and data should be available for external users shortly after. ORAS6 SST and sea-ice will be used as lower boundary conditions for ERA6. This has been possible thanks to the improvement in SSTs, especially in western boundary currents and the representation of the diurnal cycle. ORAS6 will also be used to initialise the re-forecasts of the upcoming seasonal forecasting system, SEAS6, and it will form the basis for the initialisation of the forecasts for Cycle 50r1, which will be implemented in 2025.

The priority given to surface fields in ORAS6 is in agreement with ECMWF’s strategy to exploit interface observations. It is a first step to enable further advances with coupled data assimilation (de Rosnay et al., 2022). ORAS6 will also provide a relevant dataset for emerging data-driven models of the ocean and sea ice.

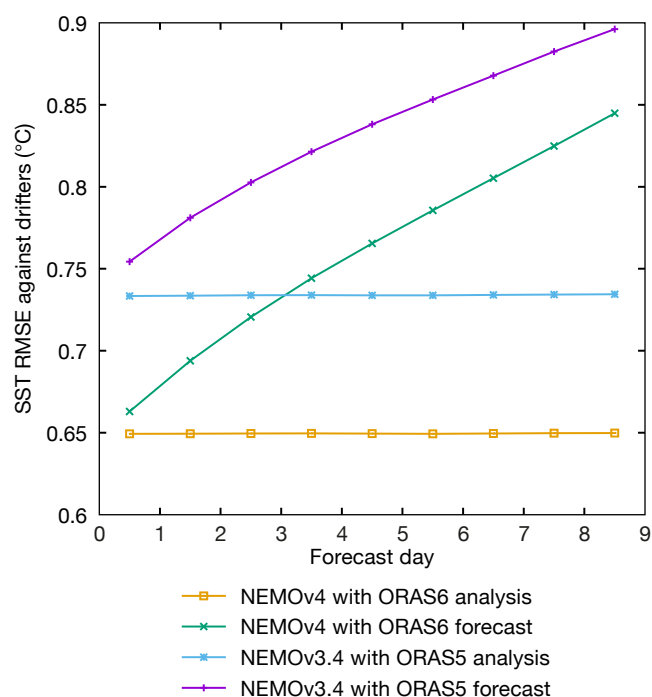


FIGURE 8 Root-mean-square error (RMSE) of sea-surface temperature (SST) forecasts and analyses against drifters. Medium-range forecasts have been carried out under TCo1279 setups based either on NEMOv4 with ORAS6 prototype initial conditions or NEMOv3.4 with ORAS5 operational initial conditions. SST RMSEs are calculated for different lead times using all available observations and start dates between 1 June 2020 and 1 June 2021. Analysis values are for the exact same sample as for the forecast values.

In the longer term, future generations of ocean reanalysis will be conducted at higher spatial resolution, with potential 4D-Var approaches to further improve the assimilation of SST and other surface observations. This will ideally make it possible to provide smoother and more consistent climate data records. We also intend to

devote efforts to the assimilation of more interface observations, such as sea-ice thickness and snow-over-sea-ice information, which have proven to be sources of predictability across timescales.

Acknowledgement

We acknowledge the contributions of a C3S tender (C3S_321b) for the development of the EDA in the NEMOVAR system. Contributions were made by the National Institute for Research in Digital Science and Technology (Inria, France), the Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS, France), the Centre for Maritime Research and Experimentation (CMRE, Italy), and the Met Office (UK).

The NEMOv4 and SI³ versions have been provided by the NEMO community model <https://www.nemo-ocean.eu/>. The specific configurations used in ORAS6 have benefited from interactions with the UK Joint Marine Modelling Programme.

The development of the ORAS6 system was partially supported by the CMEMS GLORAN project (third phase) under contract number 21003-COP-GLORAN-LOT3.

Further reading

Zuo, H., M.A. Balmaseda, S. Tietsche, K. Mogensen & M. Mayer, 2019: The ECMWF operational ensemble reanalysis–analysis system for ocean and sea ice: a description of the system and assessment. *Ocean Science*, **15**(3), 1–44. <https://doi.org/10.5194/os-15-779-2019>

Chrust, M., M.A. Balmaseda, P. Browne, M. Martin, A. Storto, A. Vidard et al., 2021: Ensemble of Data Assimilations in the ocean for better exploitation of surface observations. *ECMWF Newsletter No. 168*, 6–7. <https://www.ecmwf.int/en/newsletter/168/news/ensemble-data-assimilations-ocean-better-exploitation-surface-observations>

Sármány, D., K. Mogensen, M. Valentini, R. Aguridan, P. Geier, J. Hawkes et al., 2024: New software libraries for ORAS6 ocean reanalysis. *ECMWF Newsletter No. 178*, 13–15. <https://www.ecmwf.int/en/newsletter/178/news/new-software-libraries-oras6-ocean-reanalysis>

de Rosnay, P., P. Browne, E. de Boissésou, D. Fairbairn, Y. Hirahara, K. Ochi et al., 2022: Coupled data assimilation at ECMWF: current status, challenges and future developments. *Quarterly Journal of the Royal Meteorological Society*, **148**(747), 2672–2702. <https://doi.org/10.1002/qj.4330>

The European Weather Cloud is now operational

Xavier Abellan, Vasileios Baousis, Ricardo Correa, Roberto Cuccu, Cristina Duma, Charalampos Kominos, Samuel Langlois, Umberto Modigliani

The European Weather Cloud (EWC) is a community cloud computing platform jointly operated by ECMWF and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT). The concept was proposed by ECMWF, building on its extensive experience in providing computing access to its Member and Co-operating States, for example through the ecgate service. The EWC facilitates data access, increases data processing capability, and fosters new forms of collaboration among national meteorological and hydrological services and researchers in ECMWF and EUMETSAT Member and Co-operating States. The initiative started in 2019 with a pilot phase laying the foundation of the current service. The pilot phase had several objectives, including: i) establishing the overall governance and coordination processes among the federating partners; ii) setting up and testing a cloud infrastructure and gaining important operational experience; and iii) on-boarding and supporting a variety of use cases covering different thematic areas and applications in order to gather feedback and requirements to continue to develop and evolve the platform. In 2023, the operational ECMWF cloud infrastructure was deployed in the data centre in Bologna (Italy), co-located with the other ECMWF systems. Thereafter, the migration of existing users, from the pilot cloud infrastructure running in the original Reading (UK) data centre to the operational system in Bologna, was performed. The EWC was declared operational on 26 September 2023.

What is the EWC?

According to the original vision, the EWC is “the cloud-computing-based collaboration platform for meteorological applications development and operations in Europe enabling the digital transformation of the European Meteorological Infrastructure”.

The EWC brings together users from different countries and organisations to collaborate and share resources. It allows users to customise and deploy their applications and workflows and to build and expose services on the Internet. This flexibility in the deployment of user

applications and workflows is one of the main benefits of the EWC compared to a managed high-performance computing facility (HPCF) service. Users run applications and services next to where the data is produced, avoiding large data movements over the network.

Who can use the EWC?

The EWC service is mainly available to eligible users from the European meteorological community. In particular, it is accessible by:

- users from the national meteorological and hydrological services (NMHSs) and public institutions in ECMWF and EUMETSAT Member and Co-operating States
- research users in the context of ECMWF Special Projects and EUMETSAT R&D calls
- ECMWF and EUMETSAT for internal use in the scope of their mandates
- members of the European Meteorological Infrastructure (EMI), including organisations such as EUMETNET
- ECMWF Third Party Activities and Optional Programmes approved by ECMWF’s Council
- users from the NMHSs of the World Meteorological Organization (WMO) and research organisations for activities aligned with ECMWF’s mission.

Figure 1 illustrates the availability of the EWC.

What are the main features of the service?

The platform provides many benefits to the user community, including access to a cloud computing facility, flexibility in the provisioning and management of the deployed resources and, not least, all the advantages of a community environment, such as shared knowledge, experience, applications and data, as well as synergies and collaborations.

An important aspect of the platform is optimised access to the data repositories of the two organisations, which include services for the tailored retrieval and dissemination of petabytes of

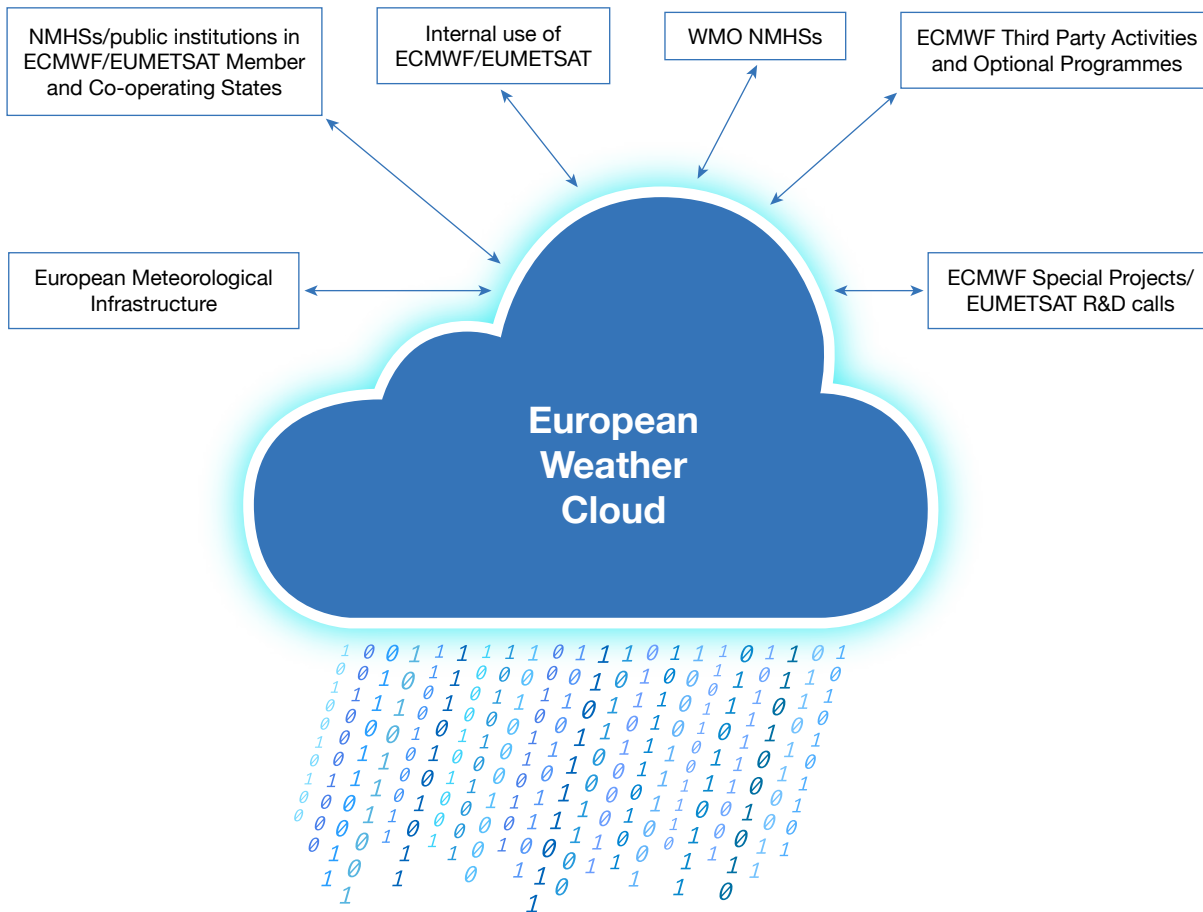


FIGURE 1 The European Weather Cloud service is available to different groups of users. For more details, see the text.

meteorological forecasts as well as climate, ocean, and satellite data.

The EWC provides users with the capability to access cloud computing infrastructure from ECMWF and EUMETSAT by deploying custom virtual environments. These range from virtual machines to more complex setups, including tailored storage and networking. Users can automate the management of resources with the help of blueprints, templates, and other automation tools to ensure the reproducibility of their workflows. Additional services are at the user's disposal for monitoring, reporting, and accounting of the provisioned resources.

Users can manage their assigned virtual environments by using a cloud orchestrator tool that enables the control and deployment of resources, workflows or services on the underlying cloud infrastructures operated at ECMWF and EUMETSAT.

Finally, a set of support and collaboration tools, such as a support portal, shared documentation, and a discussion platform, complements the solutions provided. All these services and features are constantly evolving following an established roadmap, taking into account users' feedback.

How is the EWC integrated with other services at ECMWF?

By running in the same ECMWF data centre, the EWC can leverage the co-location of its infrastructure together with the rest of the service portfolio operated by ECMWF.

In particular, a 'fast-track' data access to ECMWF data services brings the computation and services close to the data sources. Users can exploit a combined set of retrieval and dissemination services:

- The ECMWF Production Data Store (ECPDS) can be configured to disseminate the tailored data directly to the EWC storage resources to make them immediately available for processing within the cloud computing facility.
- The Meteorological Archival and Retrieval System (MARS) stores hundreds of petabytes of meteorological and climate data which can be retrieved by users leveraging the local network access.
- The Copernicus Data Stores enable access via APIs to data from the Copernicus Climate Change and Atmosphere Monitoring Services, harnessing local network access.

Additionally, the ECMWF Aviso Data Notification service can be configured on the EWC for receiving automatic notifications of the availability of model output data and the delivery of the products from the ECPDS dissemination system, to trigger further processing workflows.

The EWC is a complementary facility to ECMWF's HPCF. The HPCF remains the most suitable facility for highly intensive CPU, I/O or interconnect batch parallel workloads. Being a managed service, the HPCF does not enable the degree of customisation and flexibility that specific use cases may require. The EWC, on the other hand, can provide a higher degree of freedom and flexibility for user-configured environments.

For example, running a public web service available on the Internet would not be suitable on the HPCF. In the EWC, users can create their own virtual environment on demand (configuring the required CPU, memory, and disk resources; assigning a public IP and domain name), deploy their web service, and finally make it available to other end users.

While the EWC and HPCF may serve different use cases, they can also be integrated as components of a more comprehensive system or application, benefiting from the respective nature of the two services. Several use cases have shown, for example, how data could be produced by a model running on the HPCF, with the output delivered to the EWC, where it could be further post-processed and served by a web visualisation service.

How can the EWC foster collaboration across the meteorological community?

The EWC is, by design, a cloud-computing-based platform built with a strong focus on European meteorological community needs. It aims to enhance collaboration across its users. The EWC platform gathers in a single place computing resources, data access, and tools to facilitate and promote data exploitation as well as the generation of value-added products and services (Figure 2).

The platform enables its users to collaborate on a joint virtual environment which runs in close proximity to the data. This enables the co-development of new weather and climate applications and services, and the co-design and creation of new systems and their operations.

A notable example is the EUMETNET RODEO project, whose objective is the provision of open access to public meteorological data and the establishment of a shared federated data infrastructure to develop information products and services. It is a combined effort by 11 European NMHSs, ECMWF, and the overall EUMETNET network. In the context of RODEO, the EWC is used for the development, implementation and

hosting of some of the project's components. Moreover, available tools, such as the discussion platform (<https://confluence.ecmwf.int/display/EWCLOUDKB/EWC+Discussion+Platform>), allow users to work together, promoting information exchanges and the fundamental collaborative nature of the EWC.

The vision is to stimulate the creation of synergies and exchanges across users, to let them join forces for similar objectives, to increase reusability and reproducibility of systems and applications, and to share results and templates for common benefit.

What kind of training and outreach are offered to users?

The EWC teams at ECMWF and EUMETSAT have planned a training roadmap for the users of the platform consisting of a series of short webinars covering a wide range of topics. Their primary objective is to provide guidance to users on how to effectively use the different capabilities of the EWC. These webinars cover the basics to get started, progressively advancing to more focused thematic aspects and advanced examples. They are open to all existing and potential users of ECMWF and EUMETSAT Member and Co-operating States. All events are recorded and can be followed offline from the EWC Knowledge Base (<https://confluence.ecmwf.int/display/EWCLOUDKB/EWC+Training+and+Tutorials>).

Another important event in the calendar is the annual EWC User Workshop (<https://events.ecmwf.int/event/360/>). Jointly organised by ECMWF and EUMETSAT, these workshops are a perfect opportunity




FIGURE 2 The EWC provides computing resources, data access, and tools to facilitate data exploitation and the generation of value-added products and services.

to bring together the entire community of EWC users. The programme usually combines updates on the status of the platform and recent developments, discussions on key emerging topics, and a strong contribution from Member and Co-operating States. The latter can present their activities, inspire others, and share their experiences and feedback.

What are the main challenges and opportunities ahead?

The main challenge is to continue to ensure that the EWC offer is in line with evolving user needs. This can be achieved by providing a set of useful tools and features to keep the platform attractive and serve the meteorological community as required. To accomplish this goal, ECMWF and EUMETSAT have in place a continuous improvement process that helps to track users' requirements and to guide the development and implementation of the existing roadmap by prioritising the most critical items.

Regarding opportunities, the EWC plays a role in a wider strategy that foresees the possibility to combine the outcomes of different projects and initiatives carried out by ECMWF and EUMETSAT. The idea is to build upon these outcomes to maximise the benefits to the users of the platform. This is true especially with projects established in close cooperation with Member and Co-operating States where cloud technologies are being adopted. An example of these synergies is the EUMETNET RODEO project mentioned above and ECMWF machine learning pilot project activities.

How to get access to the service?

Eligible users may request access by contacting the ECMWF Computing Representative (<https://www.ecmwf.int/en/about/contact-us/computing-representatives>) in the corresponding Member or Co-operating State with a high-level description of their activity and an estimation of the resources required to fulfil their objectives. After assessment, if the usage is approved by the Computing Representative, the EWC support team will set up and provide access to the requested cloud resources and services.

Another way to access the EWC for users of ECMWF and EUMETSAT Member States is to apply to ECMWF Special Projects (<https://www.ecmwf.int/en/research/special-projects/special-project-application>) or the EUMETSAT Research and Development calls (<https://user.eumetsat.int/news-events/news/european-weather-cloud-research-and-development-call>) to perform research activities. This is also suitable for projects which are undertaken in cooperation between several institutions, nationally or internationally.

Any interested users, including those working internally at ECMWF and EUMETSAT, may contact the EWC

support team via the EWC Support Portal (<https://support.europeanweather.cloud>) to discuss their needs and the best approach to follow.

Success stories

In the past few years, there have been a variety of interesting use cases covering a wide spectrum of technical and scientific applications. At the time of writing, about 110 different projects have made use of the EWC, in one or both sides of the cloud at ECMWF and EUMETSAT, depending on the required data sources and target applications. The increasing number of use cases confirms the growing interest from Member and Co-operating States in using and adopting the cloud technologies provided by the EWC. At the same time, the EWC is actively used internally at ECMWF and EUMETSAT within the scope of their official mandates.

All these projects have provided invaluable feedback and insights, especially during the pilot phase. They enormously helped to improve the service and to shape its evolution, contributing to the community growth around the platform.

Data-driven machine learning applications

The increased volumes of numerical weather prediction and climate data, as well as of observations and satellite products, stimulate the research community to experiment and develop new efficient ways to exploit these data.

In this context, machine learning (ML) and artificial intelligence (AI) applications are becoming increasingly popular. This can be attributed in part to the availability of complementary computing capabilities including smaller-scale exploratory GPU environments, the abundance and accessibility of data, and finally, the availability of tools and frameworks that facilitate the development of some of these types of application. The EWC provides these elements within a coherent platform that complements the ECMWF HPCF and other resources needed for training large ML models.

Among the hosted projects, one example is the Météo-France LabIA team, which presented its activities and plans at the EWC User Workshop, held in September 2023. These plans included several research and development case studies exploiting different models and algorithms.

Another example is the transversal collaboration between Member and Co-operating States and ECMWF on a common ML pilot project to boost knowledge-sharing and target key activities, such as enhancing data-driven forecasting. As part of this work, the EWC complements the ECMWF service portfolio, particularly the HPCF, by offering collaborative environments for the participating community of experts in AI/ML.

Supporting training courses

The EWC can serve as a collaboration platform on which training courses can be run. It enables ECMWF and EUMETSAT Member and Co-operating States to conduct training sessions with a larger number of participants than would be possible in a physical classroom. For example, by using the EWC, a training course that would previously have been limited to 10 participants could have an increased capacity of 100 participants.

ECMWF has transitioned to virtual or hybrid training events, partly driven by COVID-related lockdowns. In this scenario, the EWC (including during the pilot phase) has provided support for practical and hands-on sessions through virtual classrooms. The system's flexibility means that Member and Co-operating States can also use the EWC as an infrastructure to run their own training.

EUMETSAT also runs virtual courses on the EWC platform, through which participants can easily practise using EUMETSAT data. During training courses for meteorologists, participants use virtual machines to run Jupyter notebooks, enabling them to display data from geostationary and polar-orbiting satellites quickly and easily.

The German Meteorological Service (Deutscher Wetterdienst) used the EWC during the 2020 edition of its training course on the ICON model – the German weather forecasting model. Using the EWC made the practical part of the training session go more smoothly, allowing participants to use the model to forecast the weather and put into practice what they had learned during the lectures.

International collaborations and system complementarity

The South-East European Multi-Hazard Early Warning Advisory System (SEE-MHEWS-A) was designed to improve forecasting in an area of Europe that has experienced a significant number of meteorological and hydrological hazards. Implemented by the WMO, and hosted at ECMWF during the demonstration phase, the system provides tools for forecasting such hazards and their potential impacts, on a single virtual platform.

The functionality of the EWC means that NMHSs of participating countries can deploy their own web-based applications for their users. Rather than being sent to the NMHS to be fed into their applications, the data produced on ECMWF's HPCF by Member and Co-operating States participating in SEE-MHEWS-A is sent directly into the EWC. Co-location of both services within the same data centre means that the data can be fed directly to applications efficiently and without saturating the external network. This is especially

important due to the increased resolution in ECMWF's current forecasting system, Integrated Forecasting System (IFS) Cycle 48r1.

Support to emergency response

In March 2020, an earthquake severely damaged the headquarters of the Croatian Meteorological and Hydrological Service (DHMZ). There were no casualties and the IT infrastructure remained intact, but there was severe damage to the building hosting it. Without any prior preparation, a backup system was established on the EWC within days. Copies of the services DHMZ ran were rapidly deployed, running in parallel, providing reassurance in case of further deterioration of the building and associated infrastructure.

"I am hugely impressed by the prompt, effective response and support by ECMWF and EUMETSAT to DHMZ after the earthquake," said Dr Branka Ivančan-Picek, former Director-General of DHMZ.

In 2021, DHMZ moved its offices and data centre to new premises. During this transition, the backup running on the EWC and ECMWF's HPCF became operational, ensuring that DHMZ was able to provide an uninterrupted service to its users.

System interoperability

Whilst use of the EWC is primarily intended for ECMWF's and EUMETSAT's Member and Co-operating States, it is also available for WMO activities. The WMO Information System 2.0 (WIS 2.0) provides a framework for WMO data sharing. Users can contribute to and download data from the system.

While WIS 2.0 components could be deployed on any cloud, the EWC was selected by Météo-France, one of the main contributors to WIS 2.0, to host one of the core components during the pre-operational phase. The EWC is designed to run applications that provide a service available to external users. This example and others, showing that the EWC can be interconnected with other clouds and services around the world, is invaluable and demonstrates the outward-facing aspects of the EWC.

What is the underlying cloud computing infrastructure at ECMWF?

The EWC at ECMWF is hosted within the ECMWF Common Cloud Infrastructure (CCI). This is a cloud-computing-based IT infrastructure hosting and serving multiple projects and services offered at ECMWF. The CCI has been set up in ECMWF's data centre in Bologna, Italy, co-located with other ECMWF facilities and services, such as the HPCF and the Data Handling System (DHS).

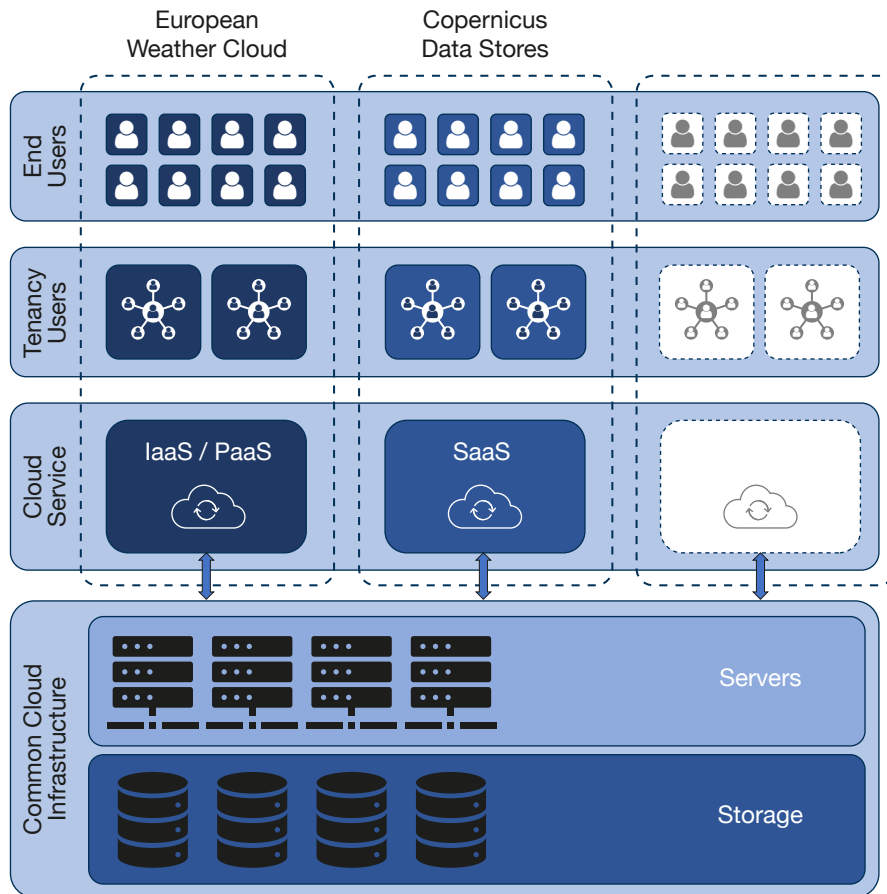


FIGURE 3 The CCI supports the European Weather Cloud, the Copernicus Data Stores and other applications and services. IaaS, PaaS and SaaS stand for 'infrastructure as a service', 'platform as a service', and 'software as a service'.

The CCI is also the home of the new ECMWF Copernicus Data Stores, the main data storage system and backend for the two Copernicus operational services implemented by ECMWF on behalf of the EU: the Climate Data Store (CDS) for the Copernicus Climate Change Service (C3S) and the Atmosphere Data Store (ADS) for the Copernicus Atmosphere Monitoring Service (CAMS). Figure 3 shows the relationship between the CCI, the EWC and the Copernicus Data Stores.

The cloud infrastructure is split into two production clouds, named CCI1 and CCI2. Each production cloud is hosted in a different computing hall at ECMWF's data centre for redundancy and resilience purposes. In the current configuration, the CCI features 68 compute nodes, with a total of 19,456 cores, and over 117 TiB of memory. This infrastructure offers access to 32 Nvidia Ampere A100 80 GB graphics processing units (GPUs), required by AI and ML users and applications. An overview of CCI resources is given in Table 1.

Openstack, a well-known open-source standard cloud computing software platform, is the technology used to manage those computing capabilities and expose them

TABLE 1 The European Weather Cloud is one of the services hosted on the CCI. It could use approximately 50% of the overall CCI resources displayed in the table.

CCI Resources	
Compute nodes	68
Compute node central processing units (CPUs)	2 x AMD 7713 2.0 GHz, 64C/128T
Graphics processing unit (GPU) nodes	16
GPU node CPUs	2 x AMD 7543 2.8 GHz, 32C/64T
Hyper-threading cores per GPU node	128
Total GPU cards	32 x NVIDIA Ampere A100 80 GB
Total hyper-threading cores	19,456
Total physical memory	117 TiB
Total storage	11 PiB
GPU card per node	2
GPU card type	NVIDIA Ampere A100 80 GB (partitioned as virtual GPUs)

to different applications, services, and users. Backing the computing capabilities, approximately 5.5 PiB of usable hard-disk-drive-backed (HDD) storage and around 300 TiB of solid-state-drive-backed (SSD) storage are available on each of the two CCI cloud clusters. Ceph is the storage solution behind them, ensuring scalable and robust access to the data. This storage capacity is used for both classic block storage by the virtual infrastructure deployed in the cloud, as well as object storage using popular APIs, such as Amazon S3 and Openstack Swift.

Conclusion

The EWC has served many use cases and a wide range of applications in its journey from a pilot project to an operational service. The community around the platform is growing, with strong participation from existing and prospective users during relevant events, such as training and user workshops, and in the discussion platform.

The EWC continues to evolve and enrich its offer with new tools and services, improving its users' experience and further enhancing collaboration within the community.

Further reading

Pappenberger, F. & M. Palkovič, 2020: Progress towards a European Weather Cloud, *ECMWF Newsletter* **No. 165**, 24–27. <https://www.ecmwf.int/en/newsletter/165/computing/progress-towards-european-weather-cloud>

Abellan, X., K. Horvath, I. Pelajić & A. Stanešić, 2020: Croatian met service backs up its production at ECMWF after earthquake, *ECMWF Newsletter* **No. 164**, 5–7. <https://www.ecmwf.int/en/newsletter/164/croatian-met-service-backs-up-its-production-at-ecmwf-after-earthquake>

[ecmwf.int/en/newsletter/164/news/croatian-met-service-backs-up-its-production-at-ecmwf-after-earthquake](https://www.ecmwf.int/en/newsletter/164/news/croatian-met-service-backs-up-its-production-at-ecmwf-after-earthquake)

Giraud, R., D. Podeur & T. Lacoste, 2023: Running a Global Broker as part of the new WMO data sharing solution. *ECMWF Newsletter* **No. 177**, 32–36. <https://www.ecmwf.int/en/newsletter/177/computing/running-global-broker-part-new-wmo-data-sharing-solution>

ECMWF publications

(see www.ecmwf.int/en/research/publications)

Technical Memoranda

- 917 **Pillosu, F.-M., A. Bucherie, A. Kruczkiwicz, T. Haiden, C. Baugh, C. Hultquist et al.:** Can global rainfall forecasts identify areas at flash flood risk? Proof of concept for Ecuador. *June 2024*
- 916 **Healy, S., N. Bormann, A. Geer, E. Holm, B. Ingleby, K. Lean et al.:** Methods for assessing the impact of current and future components of the global observing system. *April 2024*

ESA or EUMETSAT Contract Reports

Abdalla, S., S. Healy, G. De Chiara & M. Chantry: Revisiting the Direct Assimilation of Scatterometer ‘Sigma0’ over the Ocean. *June 2024*

Duncan, D., A. Geer, N. Bormann & M. Dahoui: Vicarious calibration monitoring for MWI and ICI using NWP fields. *May 2024*

Lean, K. & N. Bormann: Evaluation of the EPS-Sterna 325 GHz channels in the Ensemble of Data Assimilations. *April 2024*

Salonen, K., P. Weston & P. de Rosnay: Annual SMOS brightness temperature monitoring report 2022/23. *February 2024*

EUMETSAT/ECMWF Fellowship Programme Research Reports

63 **Warrick, F. & N. Bormann:** On the usefulness of additional cloud and tracking information for the assimilation of AMVs. *May 2024*

ECMWF Calendar 2024/25

2024

Sep 3	ECMWF’s UEF at EMS 2024: Side Event
Sep 9–12	Workshop on Diagnostics for Global Weather Prediction
Sep 10	European Weather Cloud webinar
Sep 30–Oct 3	Training course: Use and interpretation of ECMWF products
Oct 7–9	Scientific Advisory Committee
Oct 10–11	Technical Advisory Committee
Oct 21–22	Finance Committee
Oct 22	Policy Advisory Committee
Nov 4–8	Training course: Predictability and ensemble forecast systems

Nov 11–15	Training course: Numerical methods for weather prediction
Nov 19–22	NWP SAF Workshop on Satellite Observations of the Earth System
Dec 10–11	Council

2025

Apr 7–11	Annual Seminar 2025
Apr 30	Finance Committee – virtual
July 3–4	Council – virtual
Oct 20–21	Technical Advisory Committee – virtual
Oct 27–28	Finance Committee

Contact information

ECMWF, Shinfield Park, Reading, RG2 9AX, UK

Telephone National 0118 949 9000

Telephone International +44 118 949 9000

ECMWF’s public website www.ecmwf.int/

E-mail: The e-mail address of an individual at the Centre is firstinitial.lastname@ecmwf.int. For double-barrelled names use a hyphen (e.g. j-n.name-name@ecmwf.int).

For any query, issue or feedback, please contact ECMWF’s Service Desk at servicedesk@ecmwf.int. Please specify whether your query is related to forecast products, computing and archiving services, the installation of a software package, access to ECMWF data, or any other issue. The more precise you are, the more quickly we will be able to deal with your query.



Newsletter | No. 180 | Summer 2024

European Centre for Medium-Range Weather Forecasts

www.ecmwf.int