ECMWF
European Centre for Medium-Range
Weather Forecasts

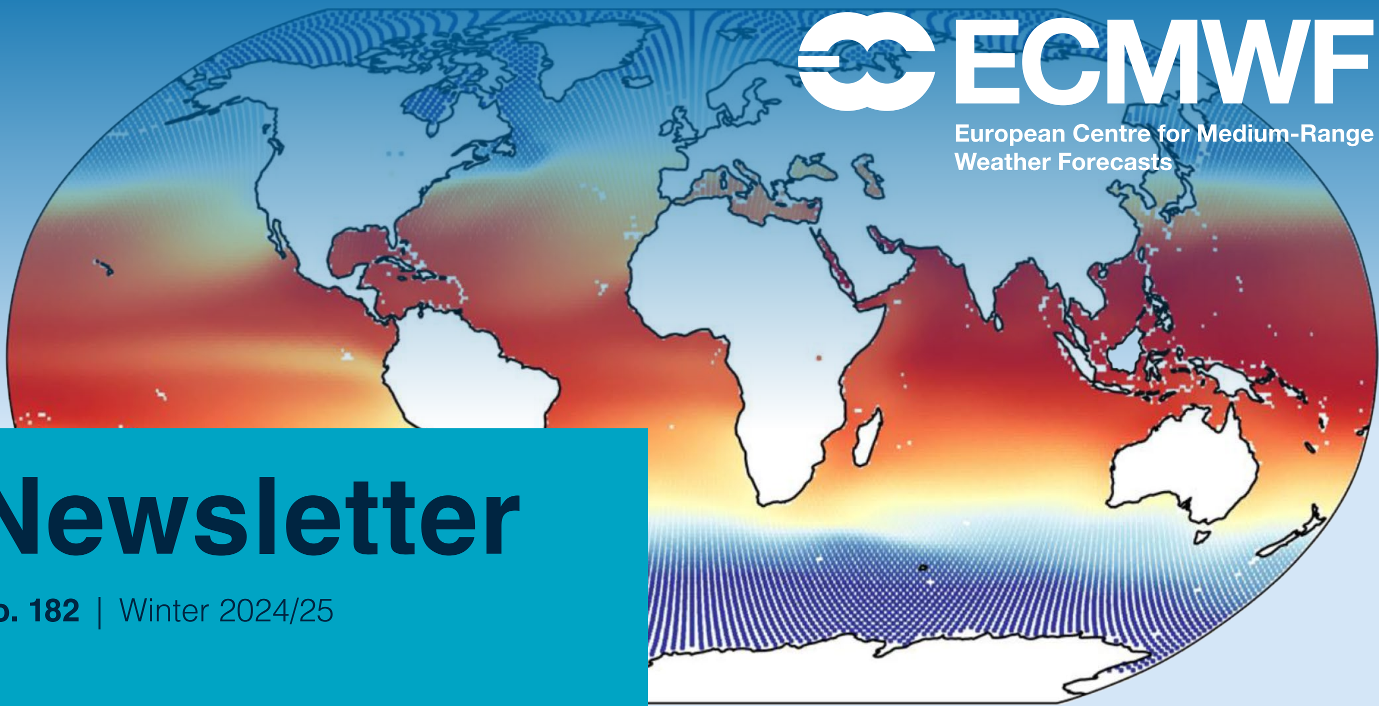

Newsletter

No. 182 | Winter 2024/25

Weather forecasts directly from observations

Modernisation of the IFS

Framework for Member State time-critical applications

ECMWF Sites: websites as a service

Forecast performance 2024

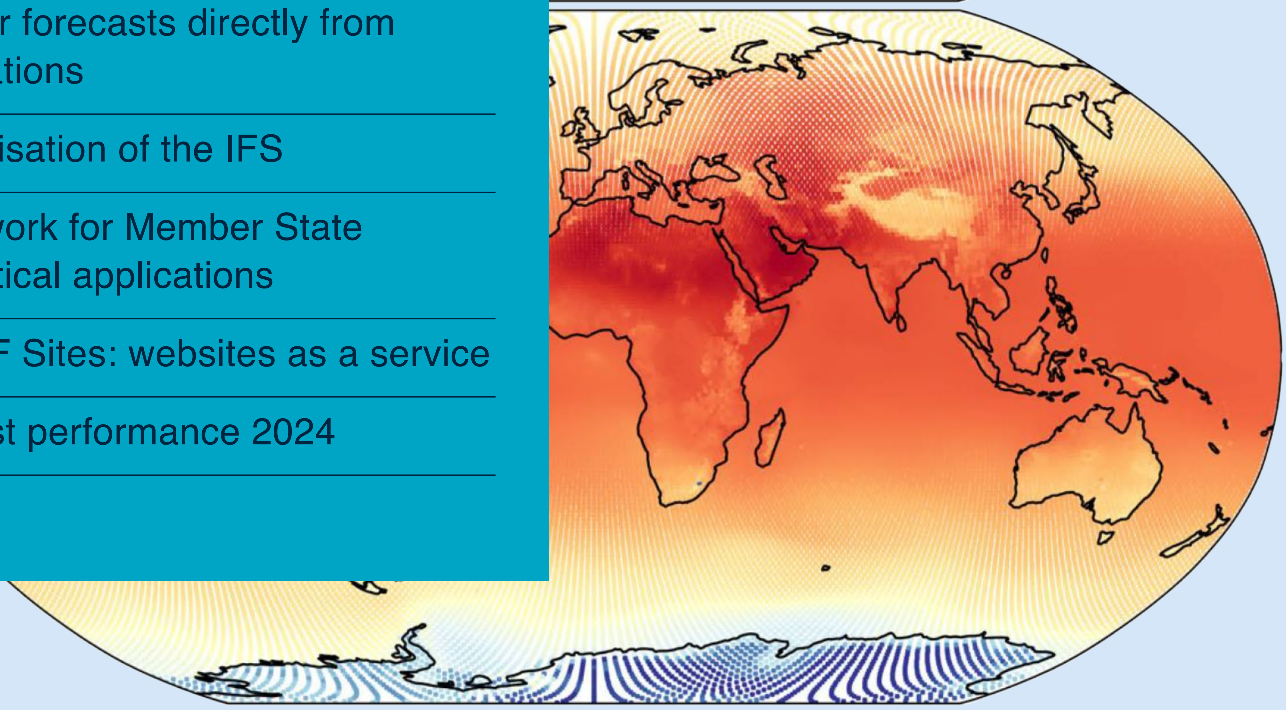

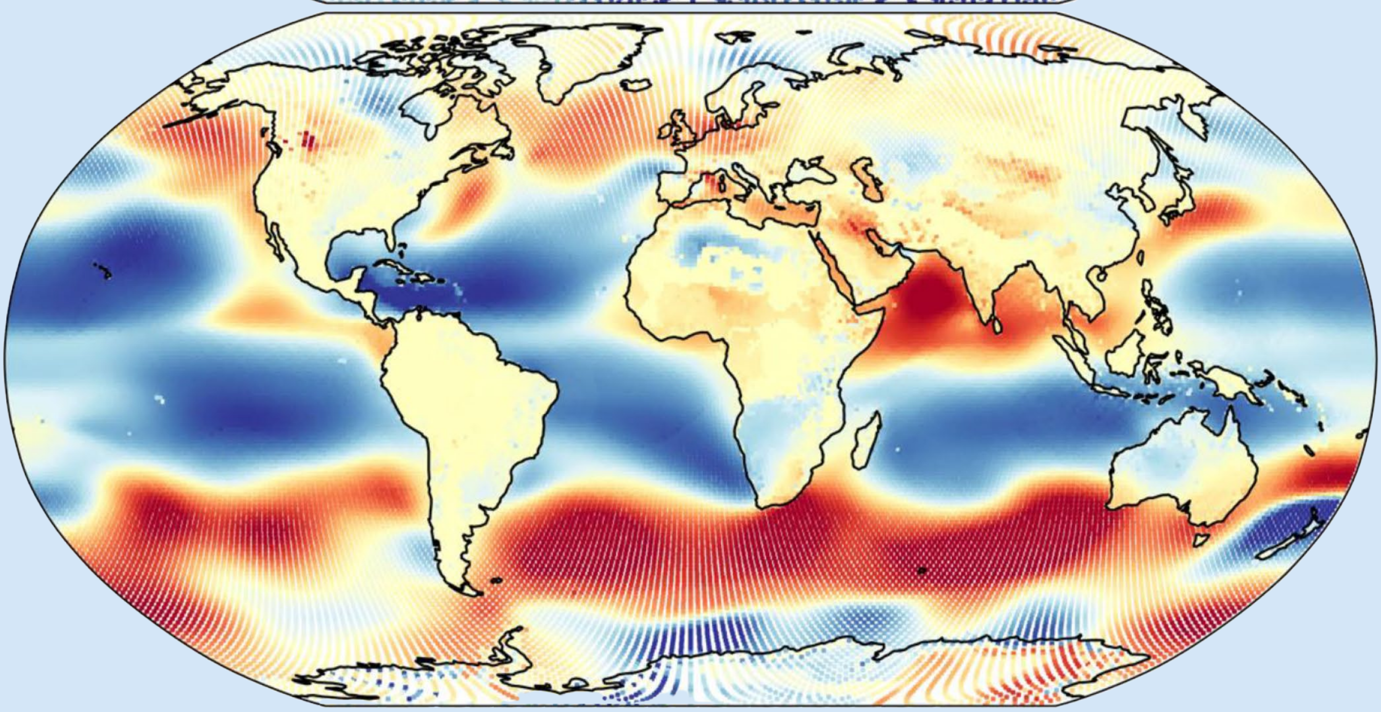

Publication policy

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The ECMWF Newsletter is not peer-reviewed.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*

Forecasts from observations

Weather observations have always been a key ingredient in weather prediction. They are used to help determine the starting point for any forecast. In traditional numerical weather prediction, they are combined with a previous short-term forecast to produce a complete initial state for the whole of the atmosphere and related Earth system components. This is done in a complex process called data assimilation, and much research into weather prediction is dedicated to improving this data assimilation step. Recently, there have been a number of developments performing weather forecasting through machine learning, for example as implemented in ECMWF's Artificial Intelligence Forecasting System (AIFS), but these methods currently do not do away with data assimilation: they still require an optimal starting point for their forecasts, and this is provided by data assimilation. However, some of our staff have been working hard on a system, based on machine learning, that does not use data assimilation at all: neither in training nor in inference mode. Instead, forecasts are based on observations alone, without a separate step of establishing a complete initial state and without relying on reanalyses for the training.

Such an approach avoids two issues associated with data assimilation: the first is the detailed knowledge that is required of uncertainty in the observations and in the prior forecast state; and the second is the need to have a very accurate mapping between the observations, which may be radiation measured by satellites, and the model states – this is often not straightforward or even impossible. The forecasting method under development, called Artificial Intelligence–Direct Observation Prediction (AI–DOP), avoids these issues. It does so by operating directly on the physical quantities that are actually measured by meteorological observing systems. First results of this method are presented in this Newsletter. They show impressive performance for the first two days, although after that performance is currently not as good as that of other forecasting methods. This innovative approach is the most radical AI departure from physics-based forecasting methods we are investigating, and it will be interesting to see where it takes us.

The other feature articles in this Newsletter indicate the importance of computing at ECMWF. They cover plans for the modernisation of our Integrated Forecasting System (IFS); a review of 20 years of the Framework for Member State time-critical applications, which enables our Member and Co-operating State users to run time-critical work on our high-performance computing facility with varying levels of monitoring and support; and an introduction to ECMWF Sites, a service that enables our staff and visiting scientists to create and publish websites.

This Newsletter also reviews our forecast performance in 2024. In addition to other highlights in our forecast skill, that article mentions the excellent results of the AIFS in experimental mode. The AIFS is due to be made operational later this year. This will be a fantastic operational milestone on our road to embracing machine learning in weather forecasting.

Florence Rabier
Director-General

Contents

Editorial
Forecasts from observations . . . 1

News
Severe rain in central Europe from Storm Boris . . . 2
Forecast performance 2024 . . . 4
Using ECMWF's observation processing (SAPP) in four countries in northwest Europe . . . 6
ECMWF forecasts supported Dutch Solar Challenge team . . . 7
Training in support of Member and Co-operating States . . . 8
ECMWF and ESA start project to better exploit Earth system satellite data . . . 10
High-altitude balloons: steering the future of navigation . . . 11
Five years of the CAMS Weather Room . . . 13
New observations October – December 2024 . . . 14

Earth system science
An update on AI–DOP: skilful weather forecasts produced directly from observations . . . 15

Computing
Modernisation of the Integrated Forecasting System . . . 19
Twenty years of the Framework for Member State time-critical applications . . . 24
ECMWF Sites: websites as a service . . . 33

General
ECMWF Council and its committees . . . 42
ECMWF publications . . . 43
ECMWF Calendar 2025 . . . 43
Contact information . . . 44

Editor Georg Lentze • **Typesetting & Graphics** Anabel Bowen • **Cover** See figure on page 17

Severe rain in central Europe from Storm Boris

Linus Magnusson, Matthieu Chevallier, Thomas Haiden, Tim Hewson, Cincia Mazzetti, Maliko Tanguy, Shaun Harrigan, Christel Prudhomme, Michael Mayer, Retish Senan, Tim Stockdale

During autumn 2024, different parts of Europe experienced severe rainfall, such as Spain, France, and Italy. While there are many important aspects to evaluate for all these events, we will here focus on the rainfall related to Storm Boris that affected central Europe.

Between 13 and 16 September 2024, the low-pressure system Boris brought record-breaking rainfall that led to severe flooding in parts of Austria, Czechia, Slovakia, Poland, Romania and Hungary. The flooding caused widespread damage and disruption with costs estimated in the billions of euros. Sadly, 27 lives were reported lost as a result.

Based on ECMWF forecasts and other sources of information, national meteorological and hydrological centres and other civil authorities within Member and Co-operating States issue weather, flood and other emergency warnings at the national and regional level.

About Storm Boris

During 11 and 12 September cold air was driven south across western Europe into the Mediterranean to meet warm, moist air to the east, creating a strong temperature contrast across central Europe. A low-pressure system formed along the interface over northern Italy on 11 September and was named Boris. It moved east and then north, and an upper-level cut-off low developed. This followed the classical 'Vb' ('five-b') pattern, which typically leads to large amounts of precipitation in the eastern Alpine region. High-pressure ridges to the northeast and the northwest helped to anchor the system and its very active rainfall zone for several days over central Europe. The largest three-day rainfall totals (13 September 00 UTC to 16 September 00 UTC) occurred over eastern Austria and along the border between Czechia and Poland (see the top panel of the first figure). A three-day total of 442 mm in the mountains in northern Czechia was the highest value found in observed station data available at ECMWF.

On top of the precipitation observations included in SYNOP weather station reports, ECMWF regularly receives additional observations from Member and Co-operating States. Triggered by this case, observations from Czechia started to be collected as well, and Slovakia also provided additional observations for evaluation purposes at ECMWF. The benefit of these additional observations for the evaluation of this case is clearly seen in the top panel of the first figure.

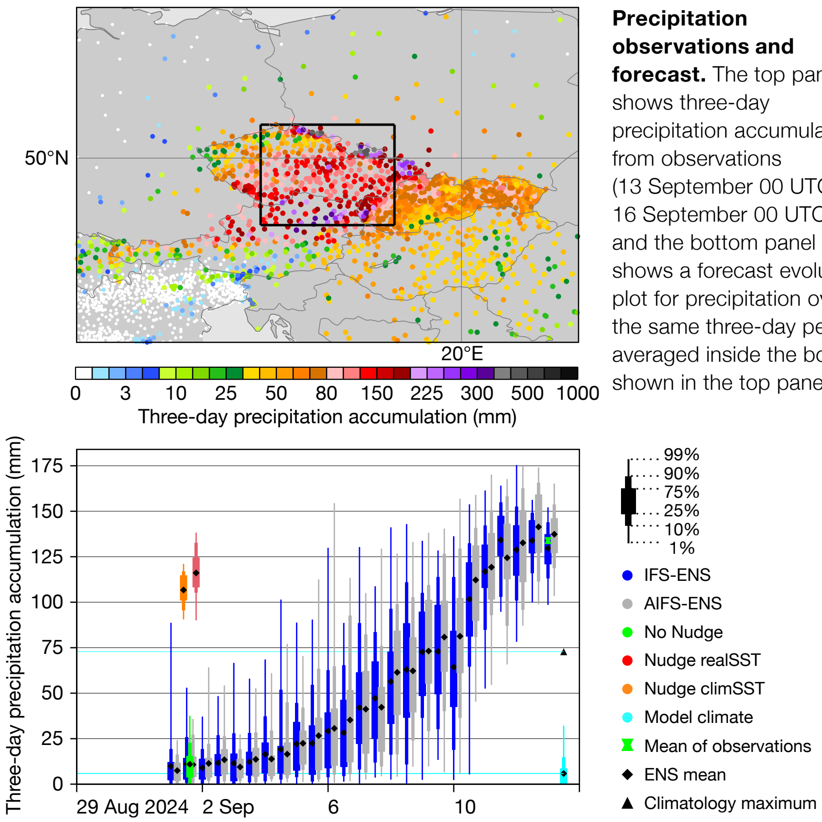

Precipitation observations and forecast. The top panel shows three-day precipitation accumulation from observations (13 September 00 UTC to 16 September 00 UTC), and the bottom panel shows a forecast evolution plot for precipitation over the same three-day period averaged inside the box shown in the top panel.

Prediction of the rainfall from Storm Boris

The bottom panel of the first figure shows the mean precipitation in the box outlined in the observation map, covering large parts of Czechia, northeastern Austria, and parts of Poland and Slovakia for 13 to 16 September. Approaching the event, the signal in the ensemble (ENS) became gradually stronger.
The ensemble mean for the Integrated Forecasting System (IFS) ensemble passed the 99th percentile of the model climate around 6 September, a week before the onset of the event. After a slight dip in the signal on 10 September 00 UTC, the ensemble mean was above the model climate maximum (based on 1,800 re-forecasts valid at the same time of the year as the event). One can note that the Artificial Intelligence Forecasting System (AIFS) ensemble (currently running with a 1° resolution) did not suffer from the drop in the signal on 10 September, giving a more consistent evolution of the signal compared to the IFS.

Role of sea-surface temperature

The very warm summer of 2024 over southeastern Europe (see the previous Newsletter: https://www.ecmwf.int/en/newsletter/181/news/extremely-warm-summer-southeast-europe) resulted in strong anomalies of sea-surface temperature (SST) in the Mediterranean and the Black Sea. It raises the question of whether these anomalies increased evaporation and contributed to the extreme rainfall during Storm Boris. One way to test this is to run forecast experiments where, to constrain the dynamics, the wind is nudged to ECMWF's ERA5 reanalysis globally, while the moisture-related variables are freely evolving. (Note that temperature and surface pressure will be affected by the nudging due to geostrophic adjustments.) In a first test, we ran three lower-resolution 11-member ensembles initialised on 1 September, where one did not use any nudging (No Nudge), one used nudging together with the observed SST (Nudge realSST), and one used nudging but with the climatological SST in the Mediterranean and the Black Sea (Nudge climSST). The resulting ensemble distributions are included in the forecast evolution plot, placed on 1 September. While the 'No Nudge' experiment resembles the longest (15-day) ensemble forecasts, both nudged experiments reproduce well the extreme seen in the short-range forecasts for the event. But by replacing the SST with the climatology, the ensemble mean precipitation in the box is reduced by around 10 percent. It indicates that the warm SST contributed to the extreme precipitation. However, one aspect not tested by this setup is the feedback on the atmospheric flow of the warm SST and increased latent heat release, which could have an additional effect.

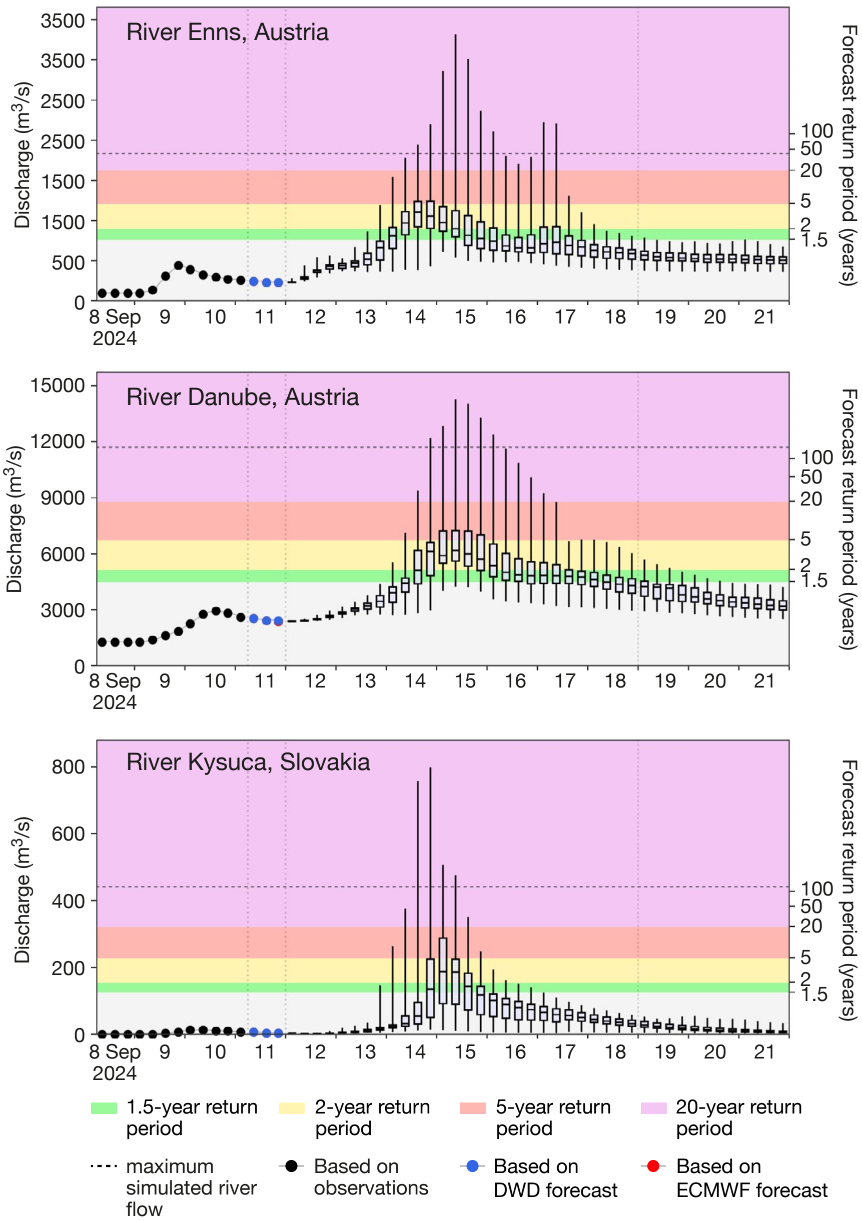

Discharge hydrographs. The plots show EFAS river discharge hydrographs from 12 September 00 UTC for three locations. Coloured areas show the 1.5-year (green), 2-year (yellow), 5-year (red) and 20-year (purple) return period thresholds derived from the model climatology. The horizontal dashed line is the maximum simulated river discharge in the model climatology. Black dots show the simulated river discharge forced with latest available meteorological observations. EFAS bridges the gap between the latest available meteorological observations and forecast initialisation by forcing the hydrological model with short-range weather forecasts, which are a DWD deterministic forecast (blue dots) and an ECMWF deterministic forecast (red dots – behind the blue dots).

Flood forecasts through EFAS

The European Flood Awareness System (EFAS) is managed by the European Commission's Joint Research Centre and is part of the EU's Copernicus Emergency Management Service (CEMS). As the computational centre, ECMWF generates flood forecasts using an open-source hydrological model, LISFLOOD, based on meteorological forecasts from ECMWF, the German National Meteorological Service (DWD) and the COSMO Limited-Area Ensemble Prediction System consortium, and meteorological and hydrological observations.

The first strong signal of a severe flood associated with Storm Boris was suggested by EFAS on 9 September with the 12 UTC forecast run. Large parts of the Oder River basin (in western Poland and the border with Czechia and Germany) were forecast to be at risk of river discharge exceeding a 20-year return period. The flood signal extended to the upper Danube (in Austria) with the 10 September 12 UTC forecast run. The hydrographs shown in the second figure, based on forecasts initialised on 12 September 00 UTC, show the risk of exceptional flood levels – with some ensemble members well beyond the maximum level within the model climatology (shown by the dashed line in the figure).

In this article we showed a few examples of evaluating forecasts of Storm Boris, mainly focusing on large-scale aspects of the precipitation and on flood forecasts. However, there are many more details, especially about local impacts, to be explored.

Forecast performance 2024

Thomas Haiden, Matthieu Chevallier

ECMWF maintains a comprehensive range of verification statistics to evaluate the accuracy of its forecasts. Each year, a summary of verification results is presented to ECMWF's Technical Advisory Committee (TAC). Their views about the performance of the operational forecasting system in 2024 are given in the box.

In 2024, upper-air ensemble forecast (ENS) skill reached a new high point due to the positive effect of Cycle 48r1 of the Integrated Forecasting System (IFS), which was implemented on 27 June 2023 and which included, among other changes, an increase in the horizontal resolution of the ENS from 18 to 9 km. In Europe, for example, the 12-month average lead time at which the Continuous Ranked Probability Skill Score for 850 hPa temperature drops below 25% exceeded 10 days for the first time (red curve in the Figure). More generally, ENS scores of upper-air variables improved by 1% to 3% in the extratropics, and surface scores improved by 2% to 6% after Cycle 48r1 was implemented. This has enabled ECMWF to maintain and consolidate its overall lead for upper-air parameters in the medium range.

For 2 m temperature and 10 m wind speed, the number of large errors in the ENS in the extratropics has decreased by 5–10% over the last 12 months. ENS precipitation forecasts have improved by about a quarter of a day in terms of lead time in the medium range. Despite these improvements, ECMWF is currently not leading in terms of surface forecast skill, especially at shorter ranges. However, Cycle 49r1, which was implemented on 12 November 2024 and includes, among other changes, 2 m temperature data assimilation in 4D-Var and the use of the Stochastically Perturbed Parametrizations (SPP) scheme, is expected to further reduce errors in 2 m temperature, 10 m wind speed and, to a smaller degree, precipitation.

The performance of the DestinE Global Continuous Extremes Digital Twin (DT) run at 4.4 km resolution is very similar to that of the ENS control forecast for upper-air variables, and 1–3% better for surface variables, except for precipitation, where the larger variability at small scales makes it more difficult to see gains in realism reflected in the overall point statistics.

Tropical cyclone (TC) forecasts by the IFS have improved as well, such that in 2024 track forecasts in the Atlantic as well as forecasts of TC genesis compared favourably to those of other global models.

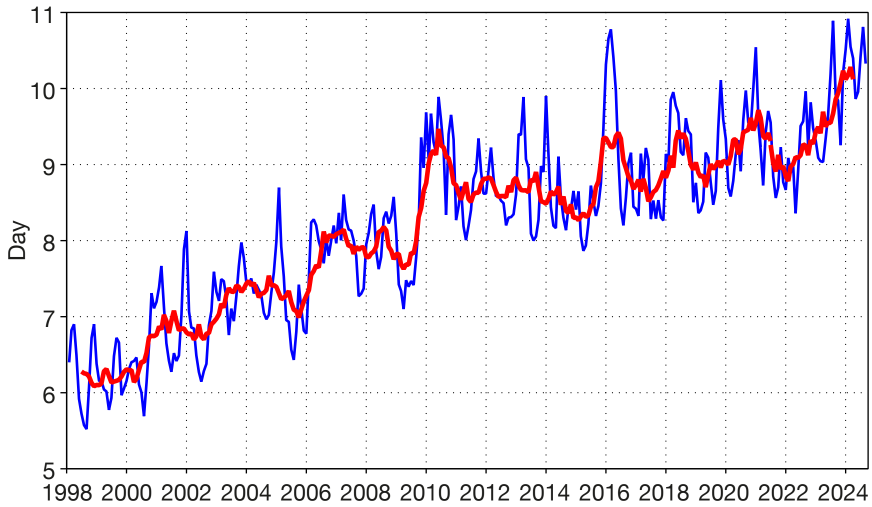

Upper-air ENS skill in Europe. Forecast lead time at which the Continuous Ranked Probability Skill Score (CRPSS) of 850 hPa temperature in Europe drops below 25%. Shown are 3-monthly values (blue) and the 12-month running mean (red). Verification is against own analysis.

ECMWF's experimental in-house machine-learning (ML) forecasts have progressed substantially in 2024, with the deterministic Artificial Intelligence Forecasting System (AIFS) now ranking above, or very close to (depending on the parameter), other ML models both for upper-air and surface forecasts. The first version of an ensemble AIFS, which started to run experimentally in June 2024, shows both forecast skill and spread-error match comparable to the IFS.

The sub-seasonal forecast from ECMWF ranks first overall when compared to other global sub-seasonal forecast models, although the statistical significance of the lead drops when going from week 2 to weeks 3 and 4.

ECMWF's seasonal forecast gave very good predictions of the transition in 2024 from El Niño back to more neutral conditions, both at seasonal and annual timescales. Record-breaking heat in parts of central and southeastern Europe in summer 2024 was well indicated, while the cold conditions in northwestern Europe, especially in early summer, were not so well captured.

Verification of Copernicus Atmosphere Monitoring Service (CAMS) ozone profiles against sondes in the northern extratropics has shown the highest scores so far. The latest World Meteorological Organization (WMO) quarterly air quality model intercomparison for North America (see https://hpfx.collab.science.gc.ca/~svfs000/na-aq-mm-fe/dist/) indicates that CAMS is leading among global air quality models for $PM_{2.5}$ and is approximately on a par for ozone and NO_2.

The complete set of annual verification results is available in ECMWF Technical Memorandum No. 918 on 'Evaluation of ECMWF forecasts', downloadable from https://www.ecmwf.int/en/publications/technical-memoranda.

The following are other sources of information about verification and forecasting system changes.

- Verification as part of ECMWF's charts page: https://charts.ecmwf.int
- WMO inter-comparison of global model forecast skill: https://wmolcdnv.ecmwf.int
- WMO ocean wave model inter-comparison results: https://confluence.ecmwf.int/display/WLW/WMO+Lead+Centre+for+Wave+Forecast+Verification+LC-WFV
- List of 'Known IFS Forecasting Issues': https://confluence.ecmwf.int/display/FCST/Known+IFS+forecasting+issues
- IFS cycle changes since 1985: http://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model

Assessment of ECMWF's Technical Advisory Committee, 10–11 October 2024

With regard to its overall view of the performance of ECMWF's operational forecasting system, the Committee:

a) congratulated ECMWF on the positive benefits of 48r1, noting that impacts of 48r1 are most marked in ENS surface parameters and upper air and least marked in HRES surface parameters, validating the decision to increase resolution;

b) noted 49r1 will bring further improvements with impacts on surface parameters in particular, especially 2 m temperature and 10 m wind speed, whilst recognizing some degradation at 50 hPa and in the tropics;

c) congratulated ECMWF on maintaining their lead over other centres in upper air scores, in particular noting a new high point in some scores due to the introduction of 48r1 and a clear signal for an increase in skill relative to ERA5;

d) noted lower summer 2024 scores compared to recent summers is linked to natural variability, with other centres and AI models showing a similar decrease in scores;

e) noted some improvement to ENS scores for surface parameters, for example in fraction of large errors in 2 m temperature and 10 m wind speed, brought about by 48r1;

f) welcomed improvements in scores for surface parameters in the shorter range, noting that the difference here between ECMWF and those centres who are leading in this respect has narrowed;

g) congratulated ECMWF on maintaining its long-term lead over other centres for significant wave height and its lead for peak period scores;

h) noted that although scores for tropical cyclone position were not as good as in 2023, comparison with ERA5 reveals this to be related to natural variability; additionally, the match between error and spread was good at days 3 and 5;

i) noted improved EFI verification scores for 10 m wind speed and 24 hour precipitation and encouraged further exploration of 2 m temperature skill in extreme situations;

j) welcomed the addition of AI models to this assessment;

k) congratulated ECMWF on the progress of AIFS, which has shown very marked improvement since summer 2023, and is now leading other AI models for some parameters whilst out-performing HRES for certain surface parameters with proposed development of additional surface parameters welcomed;

l) welcomed the inclusion of DestinE continuous extremes digital twin model evaluation, noting this model has been running routinely since August 2023;

m) noted DestinE shows very similar upper air skill to HRES and greater skill for surface parameters, recognizing that this should be expected given DestinE has a higher horizontal resolution than HRES;

n) welcomed the inclusion of hydrological model evaluation;

o) congratulated ECMWF on successful implementation of EFAS5 in September 2023, noting that in simulation mode this model performs considerably better than previous versions mainly due to increases in temporal and spatial resolution and an increase in the number of stations used in system calibration;

p) noted improvement in sub-seasonal scores relative to persistence for week 2 and limited clear signal for improvement in weeks 3 and 4;

q) welcomed the new comparison of sub-seasonal forecasts from different forecast centres, including introduction of a fair measure to adjust scores for ensemble size;

r) congratulated ECMWF on their lead in scores for sub-seasonal forecasts across all domains and parameters, noting sharp drop in skill in the extratropics from week 1 to week 2 and lead being less strong in weeks 3 and 4;

s) recognized that the seasonal forecast for the tropical Pacific captured the onset of last year's El Niño and its subsequent evolution to more neutral conditions well, even if the magnitude of the El Niño was overly amplified, noting this is a known bias; noted, also, the signal for El Niño onset was seen in the 13-month seasonal forecast plume, indicating longer-range predictability;

t) acknowledged that the seasonal forecast for winter 2023–24 highlighted warmer-than-average conditions over the northern hemisphere; that the persistent smaller-scale cold anomaly over Scandinavia was not captured is unsurprising given this was hard to predict at even two-week lead times;

u) noted that the seasonal forecast for summer 2024 captured the generally very warm conditions globally, for example showing some signal for particularly hot conditions over southern and eastern Europe but the cooler anomaly extending from the far northwest of Europe to the far west of Africa was not captured nor were some polar anomalies, though the latter is normal;

v) appreciated the continued development of new diagnostics and products and very good support ECMWF provided to Member and Co-operating States over the last year, with engagement via many mechanisms including online support, the annual UEF, online seminars, site visits and meteorological representatives at Member States; the 2025 UEF in Bologna celebrating ECMWF's 50th anniversary is eagerly anticipated.

Using ECMWF's observation processing (SAPP) in four countries in northwest Europe

Rónán Darcy (Met Éireann)

United Weather Centres West (UWC-West) is a consortium of the four national meteorological and hydrological services (NMHSs) of Denmark, Ireland, the Netherlands, and Iceland (DINI). Common numerical weather prediction (NWP) models have been running operationally since March 2024. The main configuration, DINI-EPS, is an ensemble prediction system based on HARMONIE-AROME, with hourly updates producing a 1+30 member ensemble, forecasting up to 60 hours ahead. The other common NWP configuration is the IG (Iceland–Greenland) domain producing 3-hourly deterministic forecasts for Iceland and Greenland. UWC-West utilises common high-performance computing facilities (HPCF) based in Iceland on the site of the Icelandic NMHS. In addition, high-availability clusters enable UWC-West to run ancillary products and services, including ECMWF's observation processing system SAPP (Scalable Acquisition and Pre-Processing). The observations assimilated by the NWP models are all processed by SAPP and include conventional observations, EMADDC (European Meteorological Aircraft Derived Data Centre) Mode-S aircraft observations, EARS (EUMETSAT ATOVS Retransmission Service) satellite radiances, and supplementary data streams from each of the four partner institutes.

History

SAPP is developed and used operationally by ECMWF. It handles the acquisition, processing, and extraction of meteorological observations of various types and sources into a format useable by NWP models.

In 2014, some NMHSs including Met Éireann, the Irish Meteorological Service, sent a request to ECMWF about the possibility of using SAPP for observation processing. Ireland's existing in-house processing system (Automatic Data Extraction, ADE) had become difficult to maintain given the ever increasing number and type of observations. Following a Member State visit from Enrico Fucile that year, ECMWF agreed to make SAPP available and development work began to replace ADE.

In December 2018, the ECMWF Council approved the Optional Programme supporting the provision of SAPP to participating States in response to interest in the NMHS community. Met Éireann became the first Member State to use SAPP operationally in November 2019. This participation in the Optional Programme has continued in UWC-West, where SAPP is used as the production system for providing observations in BUFR format to the NWP model for data assimilation. The ECMWF SAPP user forum has proved to be a useful source of help and information.

SAPP configuration in UWC-West

UWC-West currently uses the Virtual Machine (VM) version of SAPP v1.2 in a live/live configuration. This means that two VMs run in production mode, allowing a single VM to fail without any loss of data or any intervention required. An off-site VM is installed at Met Éireann, which provides UWC-West with a geo-resilient backup stream of observations should the two production clusters become degraded.

As the NWP model is run hourly, 24 extractions for conventional, satellite, Mode-S, and local

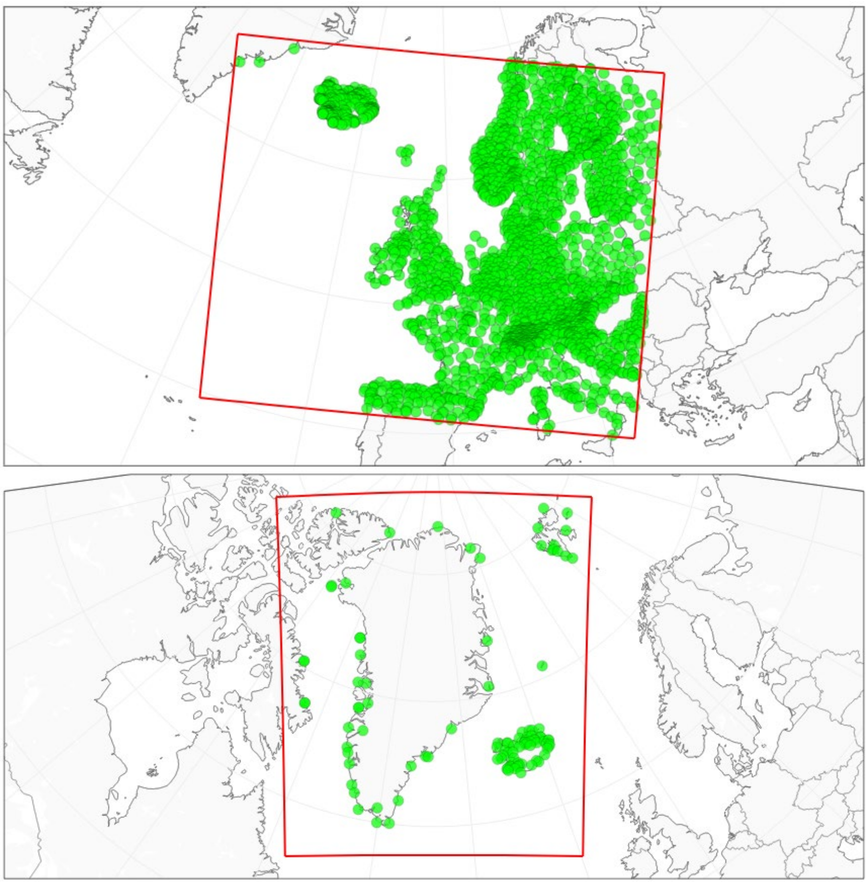

Two-metre temperature observations. Surface synoptic reports (2 m temperature) reaching the data assimilation step of the NWP model for the DINI domain (top) and the IG domain (bottom). The observation files are supplied by UWC-West SAPP.

observation streams are required daily, with a cutoff of 60 minutes.

The extraction scripts have been modified to perform a spatial check on the observations. Only observations contained within an area incorporating the NWP domains are used (see the figure). This has reduced the extraction time and improved the performance of the system for high-density observations types, such as EMADDC Mode-S aircraft observations.

CI/CD

As part of the requirements for UWC-West, all applications must be 'Continuous Integration and Continuous Deployment' (CI/CD) compliant. When any changes are made to the SAPP repositories, the user may trigger a deployment. This is sent to each of the SAPP virtual machines, where it is installed and SAPP services are restarted.

Future

A development VM has been installed using the latest version of the SAPP VM available from ECMWF, which is v1.3. This is based on Rocky Linux rather than CentOS and uses more up-to-date versions of Python 3 and ecFlow 5. Work is currently ongoing to carry out the customisations required for UWC-West, and it is anticipated that v1.3 will replace both v1.2 machines over the coming months.

ECMWF forecasts supported Dutch Solar Challenge team

Daan Nibbering (Solar Team Twente, the Netherlands)

With the help of ECMWF's weather predictions, Solar Team Twente from the Netherlands finished second out of fourteen during the Sasol Solar Challenge in South Africa in September 2024. This is a prestigious event, in which student teams from all over the world compete with their self-built solar cars, which are solely powered by the energy of the sun.

During the eight-day challenge, the teams try to cover as much distance as possible by completing additional loops on the route whilst driving from Johannesburg to Cape Town. The teams have to overcome the harsh and unpredictable environments that South Africa has to offer, from the scorching sun in the Kalahari Desert to wide mountain ranges in the Western Cape.

Desert trip. The Solar Team Twente car cruising through the Kalahari Desert. Credit: Gijs Versteeg

Solar Team Twente

Solar Team Twente is a student team based in Enschede in the eastern part of the Netherlands. For over a year, students from the University of Twente and Saxion pause their studies to design and build the most efficient solar car in the world. Each time, the students challenge the boundaries of existing technology and showcase the possibilities of sustainable energy and transportation.

Every other year, they compete with their newly built car in the Bridgestone World Solar Challenge in Australia. Over the years, Solar Team Twente has had a number of significant successes, including multiple second-place finishes in Australia and a win in the special edition held in Morocco in 2021. Last year, for the first time, the team spent a year innovating a car that had already been built, with the goal to finish first during the Solar Challenge in South Africa.

Importance of accurate weather predictions

Having an efficient solar car is not the only factor that is crucial for winning a solar challenge. With the help of ECMWF's accurate high-resolution ten-day forecast, the students devised the most optimal strategy. Weather factors such as wind and solar irradiation significantly impact the ingoing and outgoing energy of the solar car.

A carefully planned strategy is required to efficiently distribute the car's battery energy over the eight days. To determine the strategy, ECMWF's spatial and temporal weather forecasts are used as inputs to a mathematical model, which the race strategists developed.

During the race, the model and weather inputs proved to be accurate and reliable. In the northern part of the country, for example, the model suggested driving relatively slowly to conserve energy, enabling the team to drive faster during the

Solar Team Twente. The solar car team with their car after the race. Credit: Gijs Versteeg

final stages of the challenge along the western coast. Accurate long-term forecasts were crucial for calculating the impact of short-term decisions on later stages of the race, ultimately helping to maximise the distance covered.

Finishing second

On 13 September, Solar Team Twente started their challenge, competing against the first- and third-placed teams of the Bridgestone World Solar Challenge from Leuven (Belgium) and Delft (the Netherlands). During the first four days, the teams were very even, with each team in first place at least once.

Unfortunately for Twente, on the fifth day, the team from Belgium took a decisive leap by driving nearly 100 kilometres more. Despite catching up during the final days and maximising every aspect of the car and the team, the gap was too significant to overcome. Solar Team Twente ultimately finished second in the challenge, covering almost 4,200 kilometres.

Aim for the future

Solar Team Twente participates in worldwide solar challenges to showcase the potential of renewable energy and inspire future engineers to develop sustainable solutions. Through their innovative solar-powered vehicles, they demonstrate how cutting-edge technology can address global issues like climate change. Private transportation, a major contributor to greenhouse gas emissions, urgently needs a shift towards cleaner alternatives.
The mission of Solar Team Twente extends beyond competition, aiming to spark curiosity and drive progress towards a more sustainable future.

Training in support of Member and Co-operating States

Chris Stewart, William Becker

As part of its commitment under its Convention, ECMWF is dedicated to providing training that continuously evolves in response to user feedback, shifting priorities, and technological advances. The training covers advanced topics in numerical weather prediction (NWP) from medium to seasonal scales; transformative technologies, such as machine learning (ML); and training in the effective use of ECMWF products, tools, and computing resources. It is designed to meet the diverse needs of users across ECMWF's Member and Co-operating States, aligning with various strategic initiatives. ECMWF closely coordinates with global, regional, and national entities in carrying out its training activities. Recent updates include a shift towards online learning; new ML-focused activities; development of Massive Open Online Courses (MOOCs); and enhanced resource integration across departments and across programmes, such as the EU's Copernicus services we run and the EU's Destination Earth (DestinE) initiative, in which we participate. Looking ahead, ECMWF aims to expand its eLearning portfolio; further ramp up training in ML; launch MOOCs on topics related to climate monitoring; establish new training facilities at our Bonn site (Germany); and increase collaborative training efforts. Regular evaluations ensure training remains impactful, cost-effective, and scalable while addressing user needs efficiently.

ECMWF's training programme

ECMWF offers advanced NWP training to Member and Co-operating States, aimed at providing specialised knowledge and skills for scientists already working with NWP systems. These courses focus on both the foundational science and current research in various NWP topics.
On average, five week-long courses are held each year, now spread

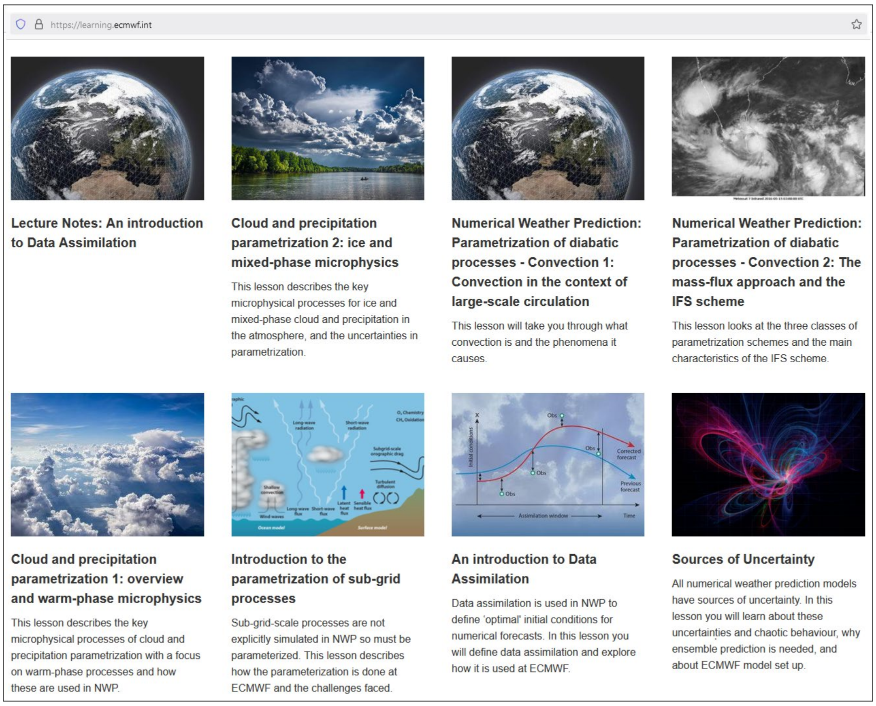

eLearning modules. ECMWF's new Moodle-based Learning Management System (LMS), which brings all ECMWF eLearning modules into a single platform (https://learning.ecmwf.int/).

throughout the year to avoid participant overload and better manage resources.

Training users on how to effectively use the products, tools and services offered by ECMWF, including in the context of Copernicus and DestinE, is another fundamental component of ECMWF's training programme. This includes the biannual 'Use and Interpretation of ECMWF Forecast Products' course, along with tailored training for key entities like EUMETNET; training at international symposia, for example organised by the European Meteorological Society (EMS); and collaborative training initiatives with organisations like the World Meteorological Organization (WMO) and the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT).

To help users engage with ECMWF's computing resources, we offer a variety of courses, including training on high-performance computing (HPC); on the European Weather Cloud (EWC) in partnership with EUMETSAT; and on computing services and software tools developed at ECMWF. These courses are well subscribed and evolve according to user needs and emerging technologies.

Given the increasing importance of machine learning (ML) in weather and climate, ECMWF has expanded the training we provide in this area. This will ramp up further in the scope of DestinE. The first MOOC on ML in weather and climate, which ran from January to April 2023, attracted over 9,000 participants globally. In-person courses on ML for weather prediction are held annually. Additionally, ECMWF's Machine Learning Pilot Project (MLPP) includes collaborative training with Member and Co-operating States, helping to build expertise in applying ML to weather and climate prediction.

Collaboration in training activities

ECMWF collaborates extensively with international, regional, and national partners to broaden the reach and effectiveness of its training programmes. Partnerships with the WMO include support for training in the context of several programmes. They include the Global Atmosphere Watch (GAW) and the Severe Weather Forecasting Programme (SWFP), as well as a fellowship scheme in which forecasters and researchers from national meteorological and hydrological services (NMHSs) of developing countries are hosted and trained at ECMWF. Collaboration with EUMETNET, through its training branch Eumetcal, ensures alignment with NMHSs. ECMWF also works closely with EUMETSAT, the European Space Agency (ESA), and others in joint training activities and in the context of partnerships, such as WEkEO (an EU service for environmental data). Tailored training at the national level is facilitated through Copernicus Climate Change Service (C3S) and Copernicus Atmosphere Monitoring Service (CAMS) National Collaboration Programmes (NCPs). These efforts enhance capacity-building globally by addressing diverse user needs and promoting best practices across the weather and climate domains.

Recent developments and future plans

ECMWF has significantly evolved its training activities in response to user feedback and changing technological and policy landscapes. Recent developments include improved strategic planning, which optimises the frequency and timing of training events to enhance efficiency.
The integration of learning resources into a unified online platform has streamlined access to ECMWF's eLearning materials (see the figure). Recent efforts have focused on developing Jupyter-based learning resources. This shift towards online learning has expanded course offerings. Additionally, ECMWF has introduced new ML-focused training initiatives and strengthened user engagement through direct collaborations with Member and Co-operating States.

Looking ahead, ECMWF plans to further develop its eLearning resources and offer more online courses, including further MOOCs. Within existing constraints, there will be more ML training courses and resources in response to growing demand. Collaboration with Member and Co-operating States will be intensified, incorporating hands-on learning in real-world scenarios, reflecting a shift toward 'training by doing together'. Regular evaluations of training activities will ensure continuous improvements in quality, scalability, and efficiency, while aligning with the evolving needs of users in the weather and climate community.

ECMWF and ESA start project to better exploit Earth system satellite data

Niels Bormann, Patricia de Rosnay, Sean Healy, Hao Zuo, Stephen English (all ECMWF), Filomena Catapano (ESA)

The exploitation of new satellite observations in ECMWF’s coupled data assimilation system has recently received a strong boost through the start of a project called Data Assimilation and Numerical Testing for Copernicus Expansion Missions (DANTEX). DANTEX is a three-year initiative funded by the European Space Agency (ESA) to pave the way for the advanced exploitation of upcoming satellite data, in particular from the Copernicus Expansion Missions (https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Copernicus_Sentinel_Expansion_missions). Four new ECMWF scientists participate in the project, which started on 18 November 2024 at ECMWF in Reading, UK.

DANTEX will develop novel ways to use so-called interface observations from a range of satellite instruments. These are observations that are sensitive not only to the atmosphere, but also to other aspects of the Earth system, such as the ocean, sea ice, snow, or land. Consistent exploitation of such observations across different parts of the Earth system in our coupled data assimilation system is a key strategic development at ECMWF. DANTEX is expected to become an important accelerator for such developments and to tackle a range of scientific and technical challenges, ensuring we get the best benefit from existing and upcoming future European satellites for numerical weather prediction and climate reanalyses.

Scope

Developments in DANTEX will target one existing and three future Earth Observation missions from ESA, adding new observations of unprecedented quality and breadth for Earth system analysis, modelling and prediction. Initially, efforts will target the cryosphere, land, and ocean waves. Cryosphere aspects will be enhanced through the Copernicus Imaging Microwave Radiometer (CIMR, a passive radiometer observing lower microwave frequencies) and the Copernicus polaR Ice and Snow Topography ALtimeter (CRISTAL) mission. The former will provide, for instance, high-resolution information on sea-ice concentration, whereas the latter will add information on the thickness of sea ice and snow. By combining the two, we expect to

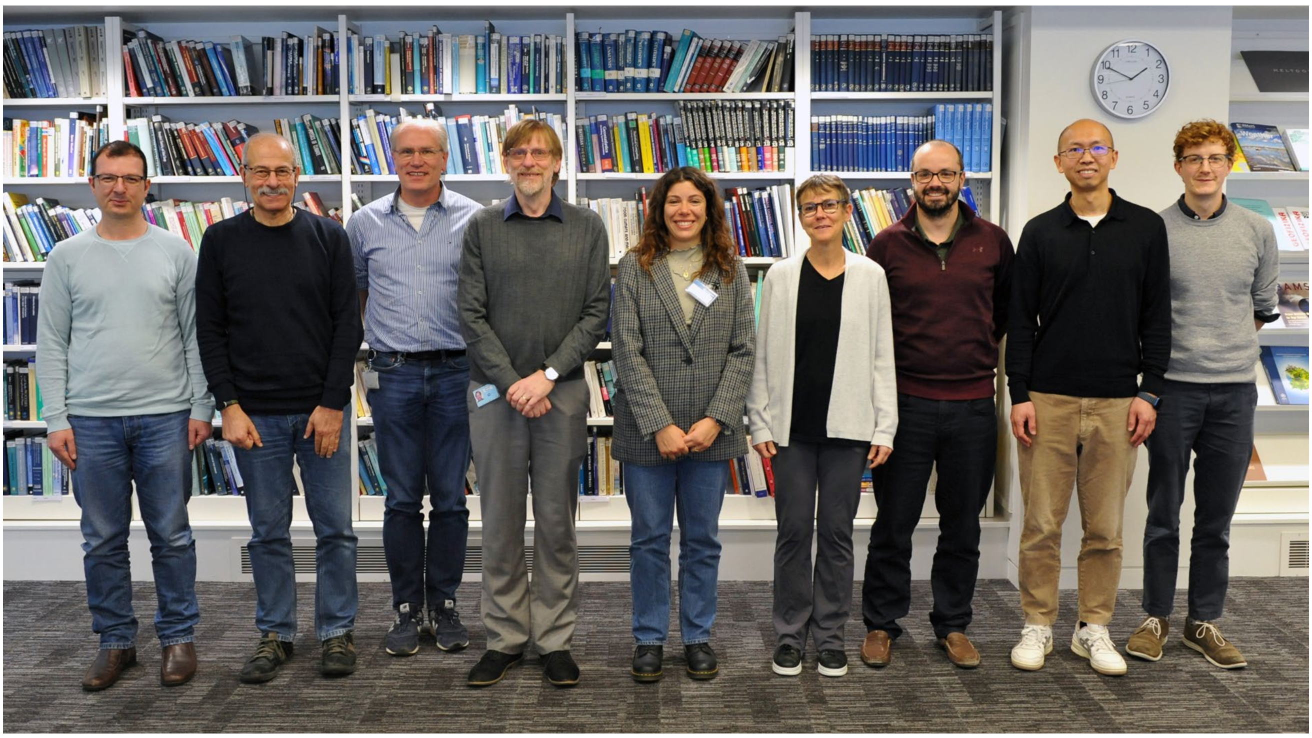

DANTEX and ECMWF. Staff involved in DANTEX get together for the kick-off at ECMWF in Reading on 18 November 2024.

see significant advances in our sea-ice analysis, as well as other cryosphere aspects. Our analysis of land temperature will benefit from the Land Surface Temperature Monitoring (LSTM) instrument, a passive radiometer observing the visible, near-infrared, and infrared parts of the spectrum. These three missions (CIMR, CRISTAL and LSTM) are Copernicus Expansion Missions. They are part of the Earth observation component of the European Union's space programme, expected to be launched late in the 2020s or early 2030s. In addition, DANTEX will enhance the representation of ocean waves and hence ocean/atmosphere interactions through better exploitation of data from the Synthetic Aperture Radar (SAR) on the existing Sentinel-1 mission (https://sentiwiki.copernicus.eu/web/s1-mission).

Raw satellite observations

A key aspect of DANTEX is that, for the Copernicus Expansion Mission instruments, we will use data as close as practical to raw observations ('level 1 data'), rather than retrieved products. This will ensure that we can maximise synergies between different observations and produce consistent initial conditions for the different components of our coupled Earth system model. To achieve this, we will develop and adapt forward models for these observations. These are models that map the Earth system model data to the satellite observations. DANTEX will employ state-of-the-art methods to facilitate this, including machine learning elements that are particularly promising to represent complex processes involved in surface radiative transfer. The coupled Earth system assimilation will build on the latest developments for sea ice and sea-surface temperature that will be implemented in Cycle 50r1 of ECMWF's Integrated Forecasting System (IFS). Developments in DANTEX will initially use existing observations as a proxy wherever possible. While these have significant weaknesses compared to the upcoming instruments, they nevertheless enable detailed development and testing of the scientific and technical capabilities prior to launch. The science developed and tested under DANTEX will enable fuller exploitation of all four missions by the weather and climate community, maximising their societal benefits.

High-altitude balloons: steering the future of navigation

David Brown (Harvard University, Urban Sky)

High-altitude balloons (HABs) offer a promising alternative to satellites for applications such as communication, environmental monitoring, and disaster response. While their ability to navigate using only altitude control has been demonstrated in practice, the underlying science has not been fully explored until now. New research published in *Nature Scientific Reports*, using ERA5 reanalysis data from ECMWF, investigates the feasibility behind their navigational capabilities. The study finds that navigating a HAB using the wind is possible to some extent, and that atmospheric conditions influencing this capability can be predicted as a function of season and geographic region.

HABs operate at a height between 15 km and 30 km, navigating by utilising the winds to steer rather than heavy and expensive lateral propulsion systems. These balloons rely on diverse wind layers within the stratosphere, which enable navigation by ascending or descending to wind layers flowing in different directions. The study asks a central question:

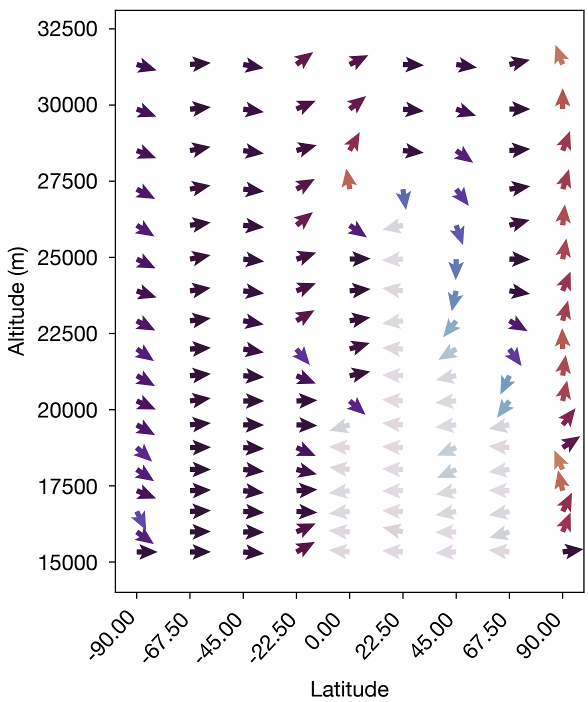

Wind direction at a specific longitude. Wind direction for wind columns at various latitudes at longitude 165 on 15 July 2021, according to ERA5. The arrow magnitude is normalised, and the colour and direction correspond to the direction of the wind layer at each altitude. This shows little variability in the southern hemisphere during austral winter, and wind diversity in the equatorial region and the mid-northern latitude region. (David Brown et al., *Nature Scientific Reports* **14**, 2024, http://creativecommons.org/licenses/by-nc-nd/4.0/)

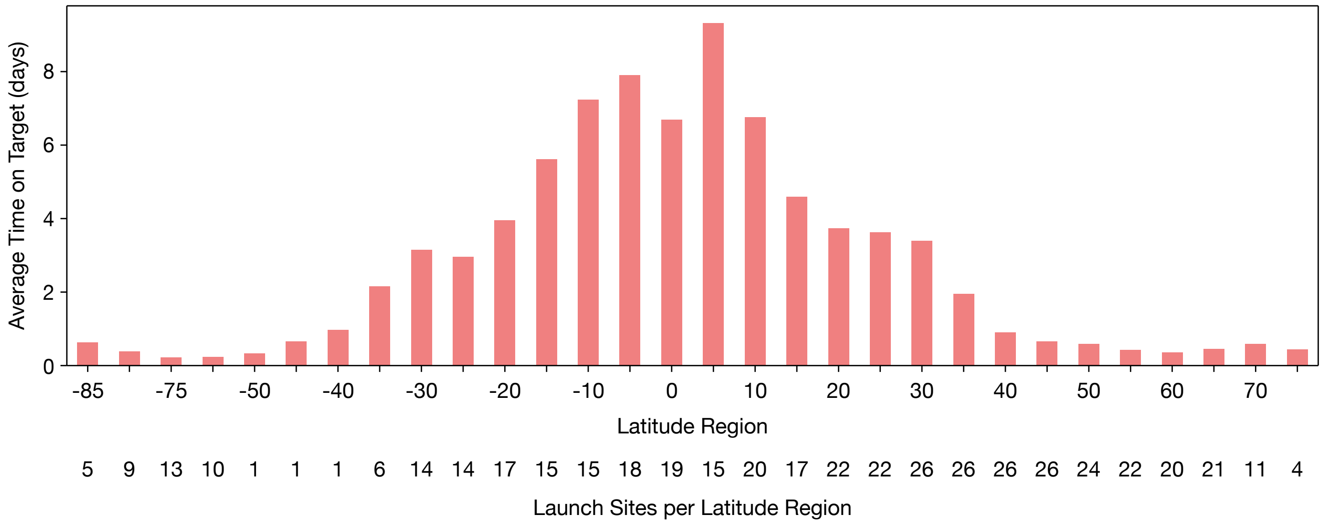

Time on target. Average time on target (within 50 km of target station) by latitude across all launch dates. The number of sites per latitude is listed below. (David Brown et al., *Nature Scientific Reports* **14**, 2024, http://creativecommons.org/licenses/by-nc-nd/4.0/)

Can you always steer a HAB just by adjusting its altitude? Simulations based on historical ERA5 wind data showed that, while some level of steering is always possible, true station-keeping – remaining within a close distance of a fixed location – is not universally achievable.

Modelling balloon navigability

The simulations utilised historical wind data from the ERA5 reanalysis dataset to model potential flight paths and determine how HABs could maintain station-keeping or navigate efficiently using only altitude control. The ERA5 dataset offers high-resolution atmospheric data on wind speed and direction, allowing for high-fidelity modelling of atmospheric conditions. The first figure shows the wind direction at a given time and longitude. The ability to keep station or navigate relies on varying winds in space and time, thus the hourly quarter-degree resolution that ERA5 offers is crucial for identifying these localised variations. The insights gained from this study rely on the detailed and accurate depiction of localised phenomena over a global scale. The study simulated a balloon traversing this reanalysis wind field for each sample and used a search algorithm to find optimal manoeuvring to keep the balloon in the target region.

Geographic and seasonal challenges

The study simulated thousands of balloon flights over five years across the globe using ERA5 reanalysis data to observe the effects of the geographic location and season on station-keeping performance (see the second figure for overall results across seasons). Key findings include:

- **Equatorial success:** Near the equator, diverse wind layers with opposing directions offer strong navigational ability.
- **Midlatitude struggles:** Strong, uniform winds in midlatitude regions limit station-keeping success, reducing overall navigational success.
- **Seasonal symmetry:** Hemispheric summer is more favourable than winter, and performance shows symmetries across opposing seasons.

Equatorial regions benefit from the Quasi-Biennial Oscillation, with alternating easterly and westerly winds offering the necessary wind diversity. Mean wind speeds are slower near the equator, and during the summer months these conditions extend into the wider tropics and the subtropics. These relatively weaker winds allow for both longer floating durations and greater ease of achieving significant deviations from the mean wind's direction.

Implications for HABs

The findings illustrate that systematic coverage cannot rely on operating single, independent balloons. The navigational limitations brought by the environment mean that the balloon will often be forced to drift away from its target coverage region. Thus, an approach using coordinated networks of smaller cost-effective balloons is required in unfavourable environments. These networks can operate similarly to satellite constellations, maintaining continuous coverage by distributing the problem among many balloons in regions where station-keeping is unfeasible.

High-altitude balloons have the potential to revolutionise the remote sensing industry. HABs can serve as an agile and cost-effective alternative to satellites. They operate at a relatively low altitude, enabling data to be of higher spatial resolution as well as offering faster data transfer. In addition, these systems can be deployed ad hoc, making them a strategic tool for responding to natural disasters such as wildfires or hurricanes.

Moreover, HABs can reposition by making simple adjustments to their altitude to catch different wind currents, offering more flexibility compared to satellites. Under suitable atmospheric conditions, they can loiter for periods of time, which no satellite can achieve. This capability makes HABs a powerhouse for monitoring evolving environments that need sustained observation.

The full study is available in *Nature Scientific Reports*. For more information, visit the publication: https://doi.org/10.1038/s41598-024-71445-9.

Five years of the CAMS Weather Room

Mark Parrington, Richard Engelen

The EU's Copernicus Atmosphere Monitoring Service (CAMS), implemented by ECMWF, has become a widely used and trusted source of information on atmospheric composition and air quality at the global and European scales since it became operational in 2015. The uptake of CAMS products by different users, and coverage in the global media in relation to different atmospheric pollution episodes, has also brought increased scrutiny of the data. This prompted the requirement to routinely monitor and evaluate CAMS products in near real-time in order to be able to quickly respond to specific queries from users and media, identify potential technical issues in the forecasts, and provide scientific inputs to CAMS communication and press activities. Taking inspiration from the physical Weather Room at ECMWF's headquarters in Reading (UK) and the daily reports, the CAMS Weather Room was established in November 2019 as a virtual space on the ECMWF intranet to compile information on specific cases using near real-time (NRT) products. During the past five years, the CAMS Weather Room has compiled information on more than 300 cases related to biomass burning emissions, desert dust storms, volcanic eruptions, the ozone layer, and global and European air quality in general.

Main use

NRT monitoring of CAMS products prior to 2019 had provided information on, for example, the large-scale Indonesian peat fires in 2015, covered in the Winter 2015/16 ECMWF Newsletter (https://www.ecmwf.int/en/elibrary/79901-newsletter-no-146-winter-201516), and the Arctic wildfires in 2019, which culminated in widespread global media coverage. However, these cases were dealt with on a more ad hoc basis. The CAMS Weather Room built on this existing monitoring to carry out more routine, day-to-day analyses of all NRT products, and to identify how the information could be shared for both technical developments and outreach activities. Cases are typically developed in response to anomalous features in CAMS operational products. Such features can, for example, occur in the biomass burning emissions from the CAMS Global Fire Assimilation System (GFAS), and Aerosol Optical Depth (AOD) and particulate matter (PM) forecasts at global and European scales. Forecast outputs are evaluated as much as possible against independent measurements available in NRT, using tools developed by ECMWF. During the past five years, these case studies have analysed the CAMS forecasting systems and identified potential developments of them. They have, for example, covered desert dust sources, the transport of atmospheric pollutants related to volcanic eruptions and wildfires, and large-scale anthropogenic air pollution episodes. In addition, many of these cases have provided essential inputs to a large number of website articles and press releases showcasing the availability and potential applications of CAMS products.

One of the first cases picked up by the CAMS Weather Room was to highlight a missing source of desert dust in Australia following reported air quality impacts in central parts of the country during January 2020 (see the first image). Subsequent updates to the source term led to an improvement in dust forecasts and

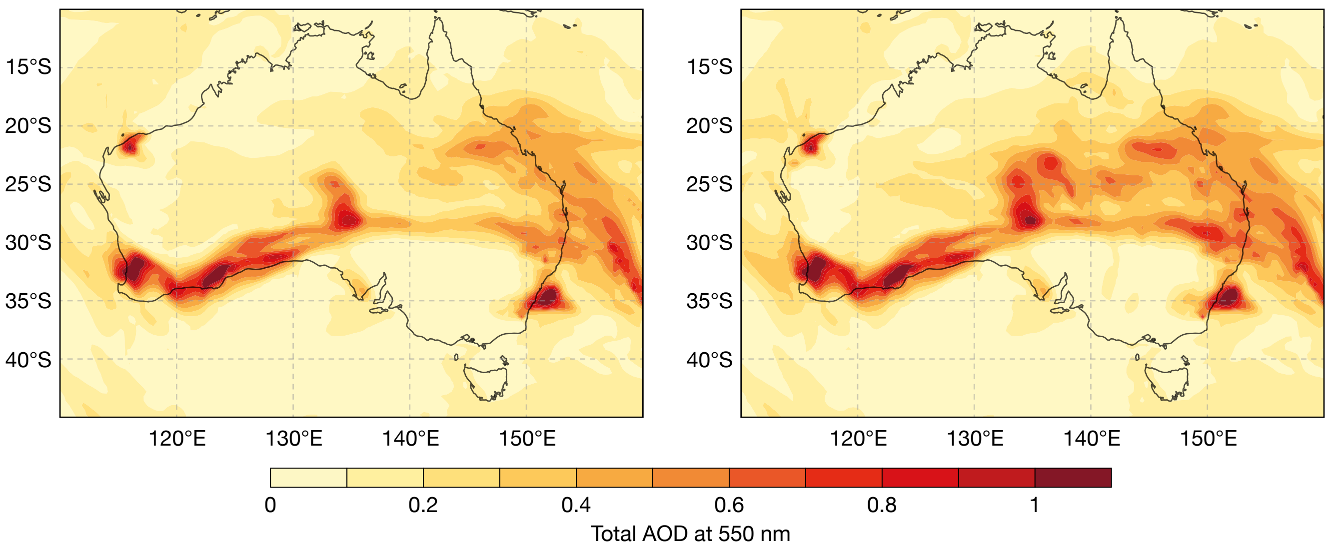

Missing dust source. CAMS global forecast of total aerosol optical depth at 550 nm initialised at 00 UTC on 11 January 2020 and valid for 12 UTC, with the original dust scheme on the left and the revised dust scheme, taking into account the missing source in central Australia, on the right.

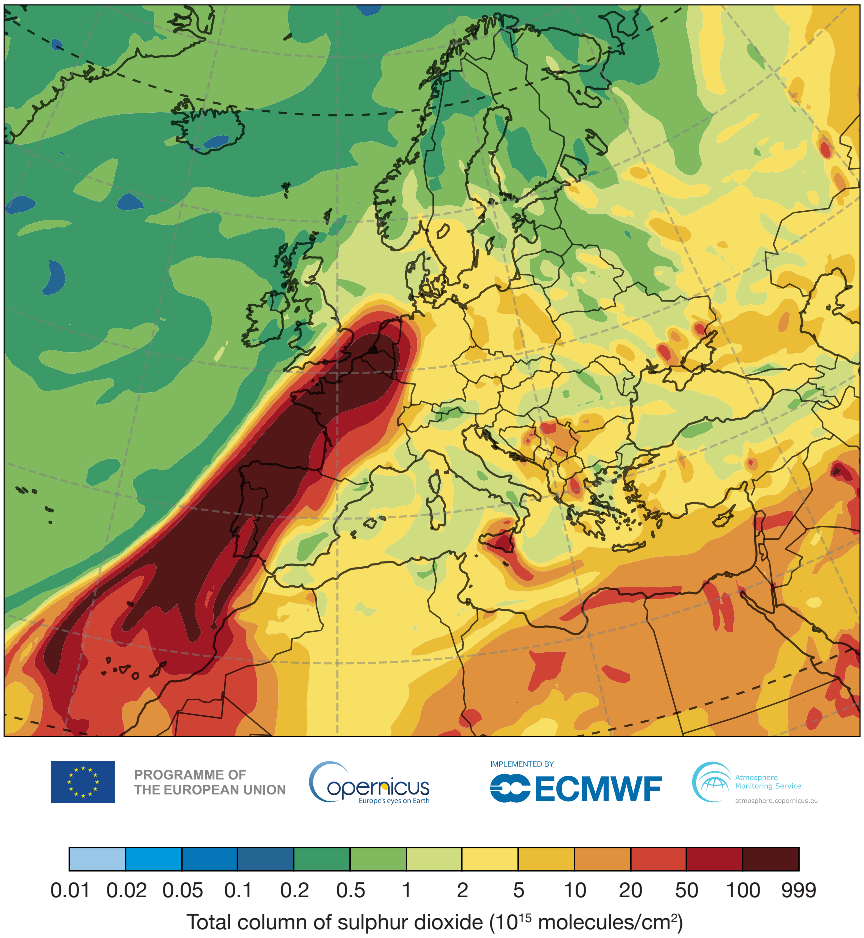

Sulphur dioxide plume. CAMS global forecast of total column of sulphur dioxide initialised at 00 UTC on 19 October 2021 and valid for 12 UTC showing the plume from the Cumbre Vieja eruption crossing Europe.

particulate matter forecasts for the region. One of the more significant episodes covered was the eruption of the Cumbre Vieja volcano on La Palma in the Canary Islands in September 2021 (see the second image). In this case, the information collected in the CAMS Weather Room was used to provide context in relation to possible air quality impacts, and the information was shared with Spanish colleagues and Copernicus representatives.

Expanded use

The scope has expanded over time to provide CAMS-related inputs and support internally to ECMWF (including the daily report, user support, user engagement, training, and climate intelligence) and CAMS contracts related to model development, validation and policy products. Information from different cases has been used in the European State of the Climate report (https://climate.copernicus.eu/ESOTC), in sections on European and Arctic wildfires. Case studies are also further developed in close collaboration with contractors working on quarterly Evaluation and Quality Control (EQC) of the CAMS forecasts and included in the EQC reports.

The CAMS Weather Room has become an important component of CAMS activities at ECMWF, and it is widely used for many of the different aspects covered by CAMS. Future developments will make selected cases more visible to external users to promote the use of CAMS data.

New observations October – December 2024

The following new observations have been activated in the operational ECMWF assimilation systems during October – December 2024.

Observations	Main impact	Activation date
SNOTEL snow data	Snow analysis, surface parameters	8 October 2024
SYNOP/METAR temperature	2 m temperatures in atmospheric 4D-Var (previously used in the land surface analysis only)	12 November 2024
Ground-based GNSS stations zenith total delay	Integrated water vapour	12 November 2024
NOAA-21 VIIRS AMVs	Tropospheric wind at high latitudes	10 December 2024
PlanetIQ radio occultation bending angles	Temperature and winds in upper troposphere/lower stratosphere	11 December 2024
Scatterometer winds from HY-2C	Near-surface winds over the ocean	11 December 2024

An update on AI–DOP: skilful weather forecasts produced directly from observations

Tony McNally, Christian Lessig, Peter Lean, Eulalie Boucher, Mihai Alexe, Ewan Pinnington, Patrick Laloyaux, Simon Lang, Florian Pinault, Matt Chantry, Chris Burrows, Ethel Villeneuve, Marcin Chrust, Niels Bormann, Sean Healy

In a previous Newsletter article (McNally et al., 2024a), we described how ECMWF research teams are embarking on a radical and ambitious project to investigate if weather forecasts can be made directly from meteorological observations, harnessing the power of machine learning (ML). We have called the method Artificial Intelligence–Direct Observation Prediction (AI–DOP). In this issue we report on progress and the first-ever skilful medium-range forecasts made purely from observations alone, without any use of a physics-based model, analyses, or reanalyses.

Here we briefly recall the rationale motivating this research. The initialisation of global physics-based forecast models is extremely challenging. This is because the majority of meteorological observations that we have from weather satellites do not directly measure the variables required by forecast models (e.g. temperature, humidity and wind), and they do not measure at the horizontal and vertical spatial scales required. To address this discrepancy, data assimilation systems blend information from the observations with fine-scale model grid information obtained from a previous (prior) forecast. For this blending process to be optimal, it requires a highly detailed and exacting knowledge of the uncertainty in the observations, as well as the uncertainty in the prior forecast state. As both of these uncertainties can be highly complex and variable (for example changing with the meteorology of the day), specifying these to the degree of accuracy required is extremely challenging and occupies substantial resources. In addition, to successfully blend observations with model states, we need to have a very accurate mapping between the quantities being measured (e.g. radiation being captured by a satellite sensor) and the geophysical variables of the physics-based model state. For some observations, such as cloudy infrared and visible reflectances from satellites, this mapping is so complex that we are currently unable to exploit these data in global numerical weather prediction (NWP).

Using artificial intelligence (AI) technology, we are exploring a completely different approach to using observations. Specifically, we have developed a system to enable Direct Observation Prediction (AI–DOP, see McNally et al., 2024b). Here, by applying ML to long historical datasets of observations, we train a neural network (NN) to forecast how the atmosphere evolves in time. Crucially, this forecast model operates directly upon the physical quantities that are actually measured by our meteorological observing systems. For example, it can predict the time evolution of radiances measured by satellites (that form the bulk of our observations), but also conventional observations of weather parameters, such as two-metre surface temperature or ten-metre wind. Formulated this way, the AI–DOP model can be initialised directly with values from the latest available observations without any need for data assimilation remapping to artificial grids or unmeasured quantities. This obviates the need for estimating large error covariances and allows the use of all observations, irrespective of the complexity of the measurement. The output of the model are predictions of observed quantities at future times. Owing to the design of the neural network and training procedure, the model can produce a forecast at any desired location, even where there may be no real input observations. Forecasts of weather-related variables, such as surface temperature or wind, are obtained by predicting future values of weather parameter observations, such as SYNOP weather station surface data or radiosonde data.

Curation of observation training data

Crucial for AI–DOP is a well-curated set of historical observations that can be used for training. For this, we extract observations from existing operational archives into special data formats suitable for ML training. While this is a laborious process, it is significantly eased by the archives containing standardised data representations (e.g. BUFR), and once this task is complete we anticipate the extracted data will additionally support Member State ML activity via Anemoi, a collaborative, open-source initiative to create ML weather forecasting systems.

At the time of writing this article, we have processed over 250 billion observations from the atmosphere, surface

 doi: 10.21957/tmi6y913dc

and ocean, covering several decades (some 2.5 TB of data – see Figure 1). This is of course a large volume of data. However, to put these figures into context, the ERA5 reanalysis, which underpins analysis-based data-driven forecast systems like GraphCast and our own Artificial Intelligence Forecasting System (AIFS), amounts to more than 6 PB of data. It is also worth noting that the historical data being curated for AI–DOP are typically at a higher spatial density than ERA5 data, and that they include some observation types which were never used at all by ERA5. The choice of which datasets to prioritise for inclusion in the training was made based on the contribution of each observation type to the current operational 4D-Var data assimilation system. During training, the AI–DOP neural network learns statistical correlations between different observation types. A particularly important relationship is that between satellite data (which have excellent global coverage) and sparser in-situ observations of weather parameters. Once correlations are learned at real weather station locations, they can be applied where there are no weather stations (e.g. over oceans) to enable global weather parameter forecasts. If required by users, these forecasts can even be specified on a regular grid.

Forecasts using different types of neural networks

We are currently experimenting with two different candidates for the type of neural network to be used in the production of forecasts, a transformer neural network (TNN) and a graph neural network (GNN). While the data curation process is still in progress, we present preliminary results with both networks trained on a subset of observation types. This subset comprises the main satellite-based systems (ATMS, IASI, SEVIRI, AMSU-A ASCAT, GPSRO) and in-situ conventional observations of weather parameters (from surface stations and balloons).

Both networks are successfully producing predictions of future observations many days in advance, where a highly realistic time evolution of weather patterns can clearly be seen in the predicted values. Both networks have demonstrated the ability to learn robust correlations between global radiance measurements available from satellites and the significantly sparser in-situ observations of weather parameters, in order to produce useful weather forecasts. Furthermore, results show that

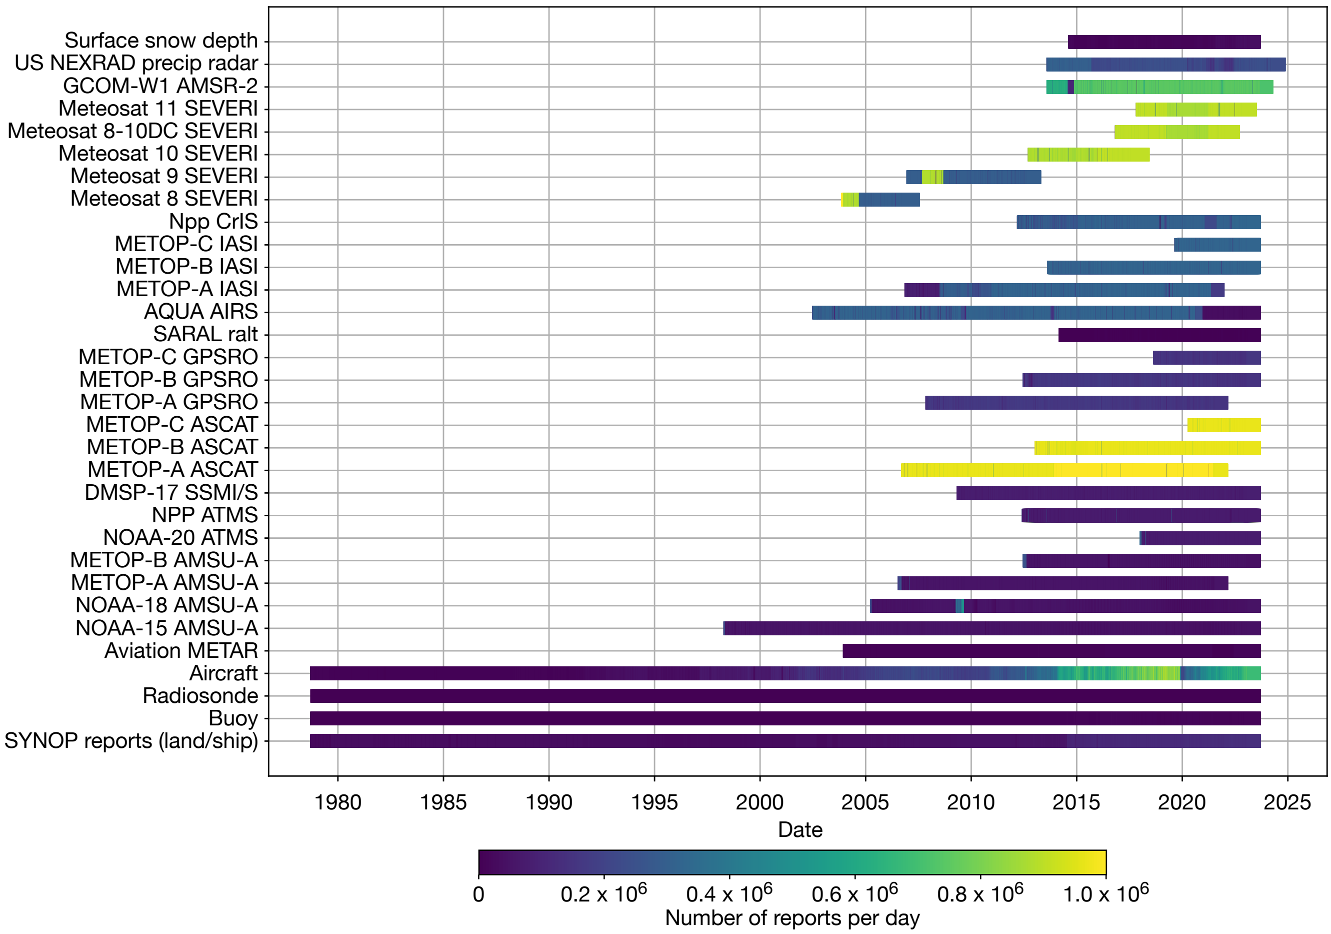

FIGURE 1 Summary of the different observation types currently included in the training dataset. They include both in-situ and satellite observations, including from EUMETSAT's Meteosat geostationary satellites and Metop polar-orbiting satellites. Satellite observations are generally indicated by satellite names and instrument names. Colours indicate the number of reports per day.

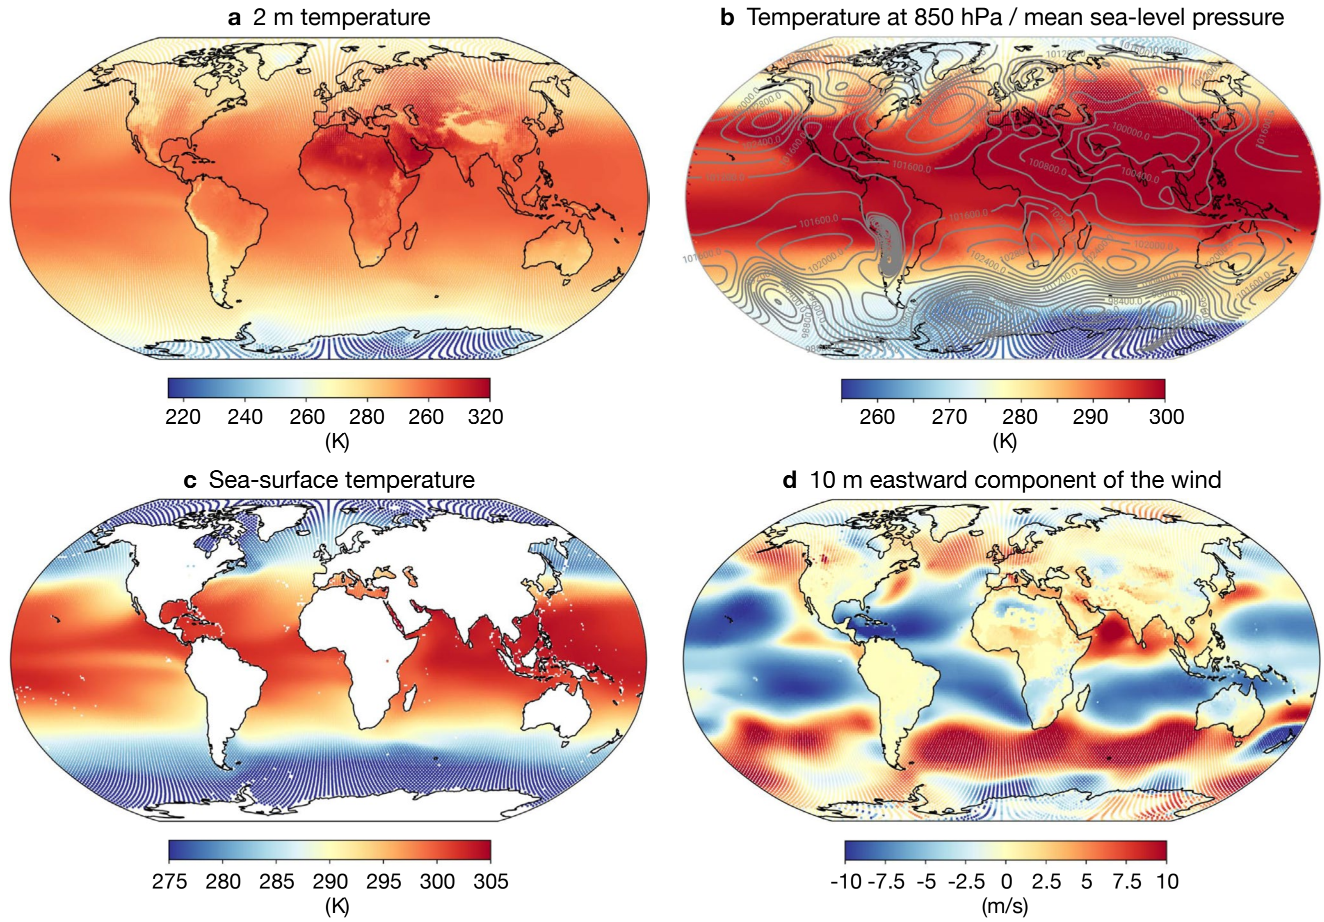

FIGURE 2 An example of gridded (O96) weather parameters from the AI–DOP network (TNN) forecasting day five (20 June 2022, 12:00 UTC). The figure shows the forecast for (a) 2 m temperature, (b) temperature at 850 hPa and mean sea-level pressure, (c) sea-surface temperature, and (d) 10 m eastward component of the wind. The projection used generates some plotting artefacts in high northern and southern latitudes.

these relationships (once learned at real weather station locations) can successfully be applied to arbitrary locations. In Figure 2, we show an example of predictions projected onto a regular user-specified grid.

Work is continuing to gain more insight into the relative merits of the two different network architectures, with a view to converging upon a single approach (or possibly a hybrid of the two) for further development.

Extending forecasts to the medium range

An immediate priority is optimising the process of forecasting into the medium range. Currently, both the TNN and GNN are trained to take 12 hours of real observations as input and predict observation values 12 hours in the future. To obtain (for example) a five-day forecast, this prediction is repeated 10 times, with the output of one 12-hour prediction fed recursively as the input to the next 12-hour prediction. In Figure 3, we can see that this so-called 'autoregressive roll-out' approach performs extremely well in the short range, but that there is a loss of performance beyond day two. Experience from the development of other data-driven forecast systems suggests that the skill of longer-range forecasts can be improved significantly by fine-tuning the network. This will involve feeding knowledge of the accuracy of the longer forecasts back to the training process to refine the learned correlations. Another area where we hope to achieve accuracy gains is in preferentially learning from tropospheric satellite radiances with a strong predictive correlation to weather parameters (and conversely down-weighting learning of stratospheric data) and guiding the network towards preferentially fitting observations known to be most reliable.

Finally, we are also exploring options for AI–DOP to produce probabilistic forecasts analogous to the ensemble forecasting systems of the physics-based Integrated Forecasting System (IFS) and the AIFS. Here we hope to build upon existing developments, such as diffusion- and score-based models designed for the AIFS (Alexe et al., 2024), which will make it possible to produce ensemble forecasts from AI–DOP. It is also expected that using a diffusion-based model will sharpen meteorological features in the forecast, which are prone to blurring with the current roll-out approach.

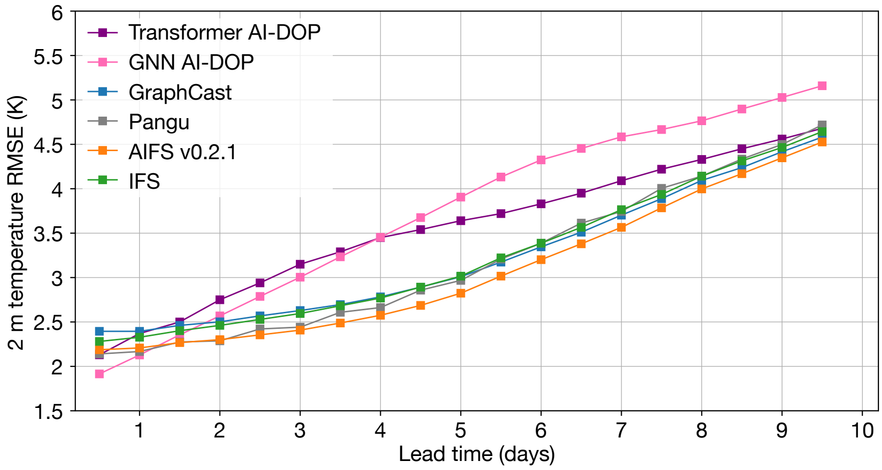

FIGURE 3 Root-mean-square error (RMSE) of AI–DOP forecasts (October–November 2022) of 2 m temperature (TNN in purple, GNN in pink) compared to the physics-based IFS and some state-of-the-art reanalysis-trained data-driven systems that rely on traditional data assimilation, such as Google DeepMind's GraphCast, Huawei's Pangu, and our own AIFS (October–November 2023). The different time frames are due to the AI–DOP observations dataset ending in early 2023.

Concluding remarks

The successful generation of medium-range weather predictions using only observations is a highly significant milestone in the field of AI data-driven forecasting. AI–DOP represents a radical departure from using observations in data assimilation to create initial conditions for physics-based models or analysis-based data-driven systems. It remains to be seen, of course, to what extent the skill of these new observation-based forecasts, either in the pure form described here or possibly hybridised with other approaches, will challenge other more conventional methods. This activity remains an extremely exciting area of research for ECMWF.

Further reading

McNally, T., C. Lessig, P. Lean, M. Chantry, M. Alexe & **S. Lang**, 2024a: Red Sky at night… producing weather forecasts directly from observations, *ECMWF Newsletter* **No. 178**, 30–34. *https://doi.org/10.21957/tmc81jo4c7*

McNally, A., C. Lessig, P. Lean, E. Boucher, M. Alexe, E. Pinnington et al., 2024b: Data driven weather forecasts trained and initialised directly from observations. *https://doi.org/10.48550/arXiv.2407.15586*

Alexe, M., S. Lang, M. Clare, M. Leutbecher, C. Roberts, L. Magnusson et al., 2024: Data-driven ensemble forecasting with the AIFS, *ECMWF Newsletter* **No. 181**, 32–37. *https://doi.org/10.21957/ma3p95hxe2*

Alexe, M., E. Boucher, P. Lean, E. Pinnington, P. Laloyaux, A. McNally et al., 2024: GraphDOP: Towards skilful data-driven medium-range weather forecasts learnt and initialised directly from observations. *https://doi.org/10.48550/arXiv.2412.15687*

Modernisation of the Integrated Forecasting System

Michael Sleigh, Andrew Bennet, Paul Burton, Paul Cresswell, Patrick Gillies, Adrian Hill, Zak Kipling, Michael Lange, Olivier Marsden, Ahmad Nawab, Balthasar Reuter

The Integrated Forecasting System (IFS) is mission-critical software for ECMWF. It fulfils our primary purposes of (a) developing a capability for medium-range weather forecasting, and (b) providing medium-range weather forecasts to our Member and Co-operating States, and it is also used for other applications. The IFS is, however, also extremely complex. The main driver in its development has been the need to improve meteorological quality and capability – an effort to which hundreds of people have contributed over decades. While very successful in terms of providing forecasts, maintaining and updating the IFS has also led to the accumulation of a great deal of technical debt in the system.
Now, because of growing demands on the system and increasing diversity of the environment, the rate of accumulation of technical debt is increasing. This leads to the need to modernise the IFS by adopting a modular design, a new representation of data, and an open-source approach.

The need to modernise

The growing demands on the system come from the expanding range of applications it is used in, and the growing population of developers and researchers who work with it. Since its launch, the IFS has evolved to support various configurations and applications: 4D-Var data assimilation, ensemble forecasting; ensemble data assimilation; full Earth system modelling; sub-seasonal and seasonal forecasting; atmosphere, ocean and land reanalysis; atmospheric composition and greenhouse gas forecasting and reanalysis; flood and fire forecasting; use in academic research/teaching (via OpenIFS); and most recently, with Destination Earth (Geenen et al., 2024), in climate and extreme weather at particularly high resolution.

Additionally, the environment in which it operates is becoming more complex and challenging. For a long time, all our operational and research work ran in an ECMWF-owned data centre and high-performance computing (HPC) facility, and it was able to use an industry-standard combination of central processing units (CPUs) and Message Passing Interface–Open Multi-Processing (MPI–OpenMP) parallelism. The first of these has ceased to be true, with Destination Earth pioneering the move to take advantage of external data centres and HPC machines we do not administer. On the second point, much effort has already been made to adapt to graphics processing units (GPUs), over several years. These developments together presage a much more diverse and fast-changing computing environment, which will include off-premises and cloud components and a wider range of vendors. The range of architectures has already broadened, and although we might not expect rapid change in the hardware market in the foreseeable future, at the very least it will become essential to utilise equally well those architectures that have already emerged. The rise of successful machine learning methods in weather forecasting also creates many opportunities for use in a complex, hybrid operational system, with the additional challenges that brings. As noted in the ECMWF Strategy from 2025 to 2034, a high degree of flexibility and agility is called for, to be able to respond to rapid changes.
The increasing need to be open in our science, data and software applies additional pressure. Furthermore, it will become more difficult to find software engineers who are willing to work on 'old' software frameworks and computing languages, such as Fortran.

The above factors – prioritising progress in performance and capability, continual increases in demands on the system due to diversifying applications, and a rapid increase in the complexity of the environment – have led to a system which is difficult to work with: less easy to understand, more brittle (easy to break), harder to test, larger and more complex than is necessary, and hence slower to improve. Like any debt that continually increases, there is a risk that at some point this becomes unsustainable. This might happen because of increasing complexity, gradually eroding our ability to make functional improvements in a realistic amount of time, at a time when the need to move quickly is at a premium.

This challenge points to an urgent need to modernise the architecture and infrastructure on which our forecasting systems are based; to clear much of the legacy of technical debt already inherent in the system; and to implement for the long term a development approach in which technical debt is explicitly recognised, documented, and dealt with continuously, alongside functional improvements. Much effort and progress have been made in recent years, for example through the Scalability Programme (Bauer et al.,

 doi: 10.21957/m9ad5hv72s

2020); a project to adapt the IFS so that it is ready for a hybrid CPU–GPU compute model (Hybrid 2024); and continual modernisation of our infrastructure and research-to-operations (R2O) tools and processes (Buizza et al., 2017; Buizza et al., 2018). But a more concerted, coordinated and exhaustive strategy is needed.

The necessary improvements can be achieved by the application of familiar and well understood software engineering techniques and processes to the design and the ongoing development of our forecasting-system software. The central part of this strategy is therefore the encapsulation, where appropriate, of compact, standalone components that can be individually tested and developed, at least to some extent, in isolation from the whole system, i.e. the separation of concerns. Efforts along these lines are not wholly new: many of them were implemented through the Scalability Programme, which introduced concepts of modularisation and separation of concerns in the form of an overarching plan to adapt the IFS to forthcoming HPC architectures. Hence, much of what we propose here builds on these initial successes and implements an overarching approach that takes the separation of concerns in a more coordinated manner to its logical conclusion.

We also propose additional actions that complement this approach, such as the definition of API (application programming interface) specifications and rigorous versioning. These ensure that integrated systems can be composed from the encapsulated components without descending into 'dependency hell'. New and more extensive standards will be introduced to guide developers, and tooling will be developed initially to detect and fix departures from the standards, and later to automatically enforce them. Standards will be of particular importance because, in addition to guiding developers, they will ensure the code is suitably structured to allow in-house tools to operate on it. An example of such a tool is Loki, which is an in-house source-to-source code translation tool. This will also facilitate the automatic extraction of OpenIFS, a supported and easily accessible version of the IFS provided for research and education, from the larger IFS source code. Standards and tools will also recognise the importance of code deprecation and removal.

Importantly, the strategy of devolving the code into components is not to move away from the idea of being integrated: we still intend to meet the same range of applications from a single source code base. Rather, it is to move away from an integrated forecasting system that is a monolith, to one that is a coherent 'ecosystem' of components from which forecast applications can be composed.

Component design

The overarching principle is the move towards a truly modular overall design, in contrast to the traditional monolithic approach that was used to create the core IFS code. Importantly, there will be a separation between technical infrastructure, individual scientific components, and the different variations of the overarching codes. This will ensure that technical and scientific changes can be adopted and migrated easily between the various supported cycles and configurations of the IFS. Model components are shared with Météo-France and the ACCORD consortium, and careful consultation will be undertaken to determine how best to separate code into a set of self-contained libraries that meets everyone's requirements. Additional testing infrastructure will be deployed to better track significant scientific and technical changes and coordinate the respective inclusion in release cycles across different organisations and projects.

A modular component structure makes it possible to improve test coverage and streamline change management. Individual module-level testing will facilitate greater test coverage. It will also enable the testing of non-operational features and model-specific code paths in jointly developed core modules, in addition to rigorous technical testing of software infrastructure components. Importantly, this does not diminish our current strong emphasis on scientific testing and evaluation in the R2O process. Here, components need to be integrated and tested together to understand their cumulative impact, no matter how well tested any individual component is in advance. We propose that the frequency of scientific release cycles will not be affected (with one major cycle per year on a fixed schedule), while their coordination will be improved.

A key change that such a refactoring will bring is that code ownership and governance can be applied per component. While scientific and infrastructure packages that include jointly developed code will require agreement and approval from all stakeholders, other components will be under local ownership. This does not preclude external contributions or use of specific subroutines in other modules. However, it makes the testing and technical management responsibility as localised as possible, with external contributions going through formal code review and integration testing before cycle synchronisation.

Data structures

A prime consideration in a more modular design is the representation of data, in particular gridded field data in grid-point and spectral space representations. Many of the modular components that will be developed in the approach described above will benefit from a harmonised representation of such field data. This will be provided by a new standalone library to enhance the modularity and uniformity of component interfaces.

A central representation of field data structures will also enable a clean separation between scientific abstractions used in forecast and data assimilation contexts, and technical concepts. The latter include concepts such as offload to accelerators and programming model compatibility. This separation can happen along the following lines:

- **Scientific API:** Here, further use of object-oriented Fortran and inheritance can separate high-level concepts, for example prognostic and diagnostic variables, and facilitate quick access to scientific metadata. Similarly, encapsulating higher-level concepts, such as tendencies or surface variable groups, and consolidating their use across the code base will be done in close collaboration with scientists.
- **Computational API:** Existing features in Atlas, a software library supporting the development of Earth system model components and data processing (Deconinck, 2018), and other libraries will be expanded to generally encapsulate technical features such as data storage backends, GPU-compatibility, parallel communication support, and I/O interface and utilities.

Moreover, a uniform, object-oriented representation of field data will serve two purposes:

- Provide a common API for field abstractions, compatible with Atlas, that encapsulates technical, computational and scientific functionality.
- Store IFS forecast states (model state, grid and fields) in unified data structures that provide a clean interface to OOPS (a framework for running different variational data assimilation formulations with a variety of forecast models) for forecast and data assimilation configurations.

A unified and individually tested data structure library will enable the graceful migration towards operational use of Atlas as the default memory data backend. The use of Atlas in the IFS will allow wider compatibility with novel packages, such as the FVM (a new nonhydrostatic dynamical core under development at ECMWF).

Workflow code

Separate to the compiled executables, the surrounding 'workflow code' divides into: (1) scripts which drive the model and manage data flow at run-time, and (2) 'suite builder' code, which runs once at deployment to generate the structure of the suite.

The envisaged workflow architecture is summarised in Figure 1. Multiple run-time configurations are combined with a suite generator tool which builds an ecFlow suite definition. ecFlow is a work flow package that enables users to run a large number of programs in a controlled environment (see *https://ecflow.readthedocs.io*). The suite generator tool also builds shell wrappers for its tasks, and it initiates a deployment of those, alongside required static component configurations and dependent code libraries. Control is then passed to the ecFlow suite, which orchestrates the execution of the task wrappers, which optionally generate on-the-fly state-dependent configurations and invoke relevant binaries and library scripts.

There are currently multiple sets of code for generating and deploying suites. All are written in Python, but they have substantial differences in both functionality and style, with varying levels of use of common frameworks. None has the required flexibility to cover all the needs of the complete range of research and operational forecasting suites. The strategy is to build a system which delivers the necessary flexibility. This is to be done by unifying the design and implementation of suites across applications, based on a common software stack such as pyFlow and Troika. Current differences in the use of suites should be harmonised as much as possible. It is also crucial to harmonise between research and operational environments. Differences here have often been the source of unexpected issues during the R2O process.

To begin, ECMWF is developing a set of common standards for suite design. Both the tooling for suite generation, and the design of the suites themselves, will be developed in line with these standards as they evolve. Suite design standards and their publication as part of invitations to tender (ITTs) will ensure consistency and optimal design for both internal and contracted-out work. The introduction of a full end-to-end forecast-suite-level test platform, which we refer to as the 'development suite' (d-suite), will demonstrate how new workflow standards and designs will work in an operational-like context.

All forecast systems are driven at runtime by scripts, which have grown organically to a high level of complexity, with a variety of different styles and approaches. This leads to a significant maintenance burden and difficulty in making changes and adding new features. It also adds substantial overhead to the migration to other HPC platforms.

A major effort to restructure and refactor the scripts is proposed, with the aim of making them cleaner, more robust and maintainable, and easier to understand and modify with confidence. The starting point for this will be rolling out a recently published IFS shell scripting standard, in three stages:

- Migrating away from the legacy Korn shell to Bash.
- Implementing a clean separation between ecFlow task wrappers, calling plain shell scripts to do the actual work which can be tested outside of ecFlow.
- Bringing scripts into line with the detailed provisions in the standard.

In parallel with this, a longer-term effort will be to

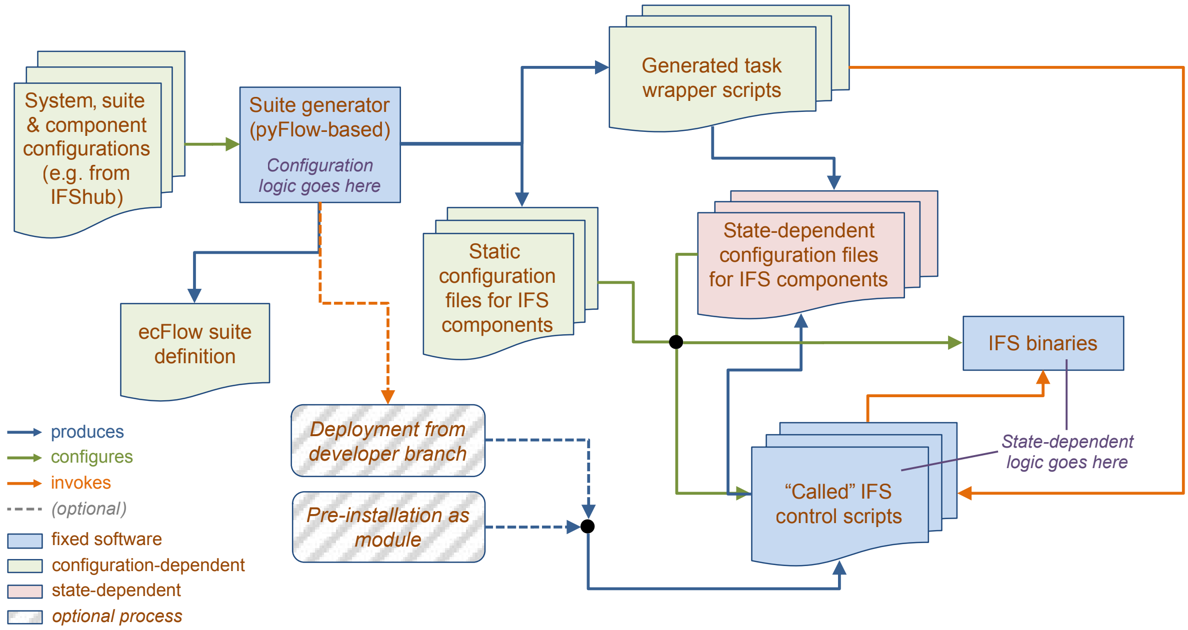

FIGURE 1 Future architecture of IFS workflow code, illustrating the interaction between suite generation, control scripts, configuration files, etc. Workflows are defined by the user in IFShub, a web-based interface for ECMWF and external users. Then, suite-generation code written in Python, and using the pyFlow library, generates both the ecFlow suite definition that defines the overall workflow (the set of tasks and how they interrelate and depend on each other), and the task wrapper scripts that control each task. Ultimately this leads to the production of forecasts via the IFS binaries. Care is needed to separate properly the static configuration defined once for the overall workflow, and the state-dependent configuration that evolves as the work proceeds.

modularise scripts, factoring out common repeated code and breaking down large scripts into smaller pieces that can be meaningfully tested in isolation. Where appropriate, functionality should be factored out into self-contained tools (e.g. Python packages) to be installed on the target platform. The primary motivation for this encapsulation is that it results in a collection of components that can be tested in isolation, outside the context of an ecFlow suite.

Adoption of open source and open development

IFS software is partially open-source, following a Council agreement to open-source selected components while keeping the system as a whole closed. The process of open-sourcing components aims to enhance collaboration with both Earth system and computer scientists. To date, a number of scientific kernels, such as the radiation package ecRad, and technical infrastructure libraries have been made public under the Apache 2.0 licence.

The short-term plan is to include these existing open-source components in the operational IFS build, and to remove the corresponding code from the central IFS source repository, so all scientists and developers will work directly in the public repositories. This will benefit the modernisation of the code.

In the longer term, the new ECMWF Strategy says that "ECMWF will build on the successful OpenIFS efforts and move to an open-source approach for the whole of the forecast model." Council decisions permitting, an initial open-source code will be based on OpenIFS, along with all the infrastructure to permit running 'out-of-the-box'. Since the initial open-source version of the IFS forecast model will effectively be OpenIFS, we suggest that the new public offering retain for the long term the OpenIFS identity.

Testing

Testing is of critical importance. Most recent improvements in IFS development workflows at ECMWF have been improvements in testing. These have included introducing a fast interactive test framework at the compiled IFS code level, and continuous integration (CI) testing which works at the full integrated-suite level. Building on this, much of the motivation for the separation of concerns/abstraction/modularity made above is that it enables more, earlier, and faster testing. The full integration-level and scientific testing in our current R2O process is still required, but there is more that can be done earlier in the process.

One limitation in current technical testing is that only the testing of changes expected to be bit-identical has been automated. We have no procedures to automatically test and accept meteorologically neutral technical changes, such as those related to GPU adaptation, optimisation, or code refactoring, without recourse to running long experiments of the order of a month and making a manual assessment. This is a particular barrier for purely technical code developers, and for HPC vendors, who are likely to need to make technical changes to tune the IFS to their systems as part of any procurement benchmarking. One approach already used by MeteoSwiss is a test ensemble, in which confidence intervals for output values are determined using a known good configuration. These are then used to detect problems in runs on different setups (different compiler, optimisations, accelerators) (see *https://github.com/MeteoSwiss/probtest*). We will explore this and similar ideas for the IFS.

It is a critical requirement to be able to identically replicate behaviour and results of operations in a sandboxed test environment. We will ensure the full software environment of an IFS run is controlled and replicable between research and operations. Being able to run the same code and configuration in both research experiments and operational suites helps ensure we fully test changes only once, and early in their development.

In addition to improved integration-level testing, we plan to significantly expand the scope of unit testing to verify the behaviour of individual components. This applies at both the coarse level of the modular components (for example the ECMWF ocean wave model, ecWAM), but also where possible to individual routines and scripts. Such an approach is already taken for suite generation code, but it will be extended to scripts and compiled code.

As previously mentioned, the introduction of a 'd-suite' (a fully functional, end-to-end clone of the operational suite) as an intermediate integration testing system between forecast-system development and operations will be a key mechanism to support enhanced testing, harmonise between research and development and operations, and allow us to move towards continuous delivery.

Conclusion and outlook

We have proposed a software strategy for our forecasting system, to complement the wider ECMWF Software Strategy (Quintino et al., 2023). This is intended to meet the immediate-to-medium-term challenge of ensuring the forecasting system that we have now – the IFS – remains sustainable on at least a ten-year timescale. It intends to make our forecasting software substantially more agile to meet the challenges of a rapidly changing system and environment. An implementation project called FORGE (Forecast-System Regeneration) is being launched to deliver the strategy.

While the more immediate software challenges have been thought through in detail so far, another strand of software strategy, which is less mature at this stage, looks towards and beyond the ten-year horizon. In particular, this concerns how the forecasting system might move wholly away from the current IFS and Fortran to a new FVM-based system written in Python, and exploiting domain-specific languages (DSL) to separate scientific from technical concerns. This will be the subject of future articles.

Feedback is actively being sought from developers and researchers within ECMWF and our Member and Co-operating States, and from our collaborators and other stakeholders, on the detail of the approaches presented here. This is done to take account of the feedback in the ongoing process of developing a final detailed strategy document, covering both the more immediate software challenges and the longer term, which will be presented to ECMWF's Committees in the future.

Further reading

Bauer, **P.**, **T. Quintino**, **N. Wedi**, **A. Bonanni**, **M. Chrust**, **W. Deconinck et al.**, 2020: The ECMWF Scalability Programme: progress and plans. *ECMWF Technical Memorandum* **No. 857**. *https://doi.org/10.21957/gdit22ulm*

Bonavita, **M.**, **Y. Trémolet**, **E. Hólm**, **S. Lang**, **M. Chrust**, **M. Janiskova et al.**, 2017: A Strategy for Data Assimilation. *ECMWF Technical Memorandum* **No. 800**. *https://doi.org/10.21957/tx1epjd2p*

Buizza, **R.**, **E. Andersson**, **R. Forbes** & **M. Sleigh**, 2017: The ECMWF research to operations (R2O) process. *ECMWF Technical Memorandum* **No. 806**. *https://doi.org/10.21957/m3r9bvg6x*

Buizza, **R.**, **M. Alonso-Balmaseda**, **A. Brown**, **S.J. English**, **R. Forbes**, **A. Geer**, 2018: The development and evaluation process followed at ECMWF to upgrade the Integrated Forecasting System (IFS). *ECMWF Technical Memorandum* **No. 829**. *https://doi.org/10.21957/xzopnhty9*

Burton, **P.**, **M. Martins**, **S. Siemen** & **M. Sleigh**, 2021: IFShub: a new way to work with IFS experiments. *ECMWF Newsletter* **No. 167**, 28–32. *https://doi.org/10.21957/bu599oxq27*

Deconinck, **W.**, 2018: ECMWF releases Atlas software library. *ECMWF Newsletter* **No. 155**, 12–13. *https://www.ecmwf.int/en/newsletter/155/news/ecmwf-releases-atlas-software-library*

Geenen, **T.**, **N. Wedi**, **S. Milinski**, **I. Hadade**, **B. Reuter**, **S. Smart et al.**, 2024: Digital twins, the journey of an operational weather system into the heart of Destination Earth. *Procedia Computer Science*, **240**, 99–108. *https://doi.org/10.1016/j.procs.2024.07.013*

Quintino, **T.**, **U. Modigliani**, **F. Pappenberger**, **S. Lamy-Thepaut**, **S. Smart**, **J. Hawkes et al.**, 2023: Software strategy and roadmap 2023-2027. *ECMWF Technical Memorandum* **No. 904**. *https://doi.org/10.21957/c6d7df0322*

Twenty years of the Framework for Member State time-critical applications

Paul Dando, Umberto Modigliani

In 2004, ECMWF's Council approved the 'Framework for Member State time-critical applications'. The Framework allows ECMWF Member and Co-operating State users to run time-critical work at ECMWF with varying levels of monitoring and support. Following the 20th anniversary of the approval of the Framework, we review why it was created, describe the applications that have been supported under it, and look at how it has evolved over the years.

Before the time-critical Framework

Since the introduction of the dedicated Member State Unix system in the mid-1990s (Dando et al., 2023), Member State applications running on ECMWF's computing systems have needed to access data produced by the latest ECMWF operational analysis or forecast in near real-time. Initially, to satisfy this requirement, a mechanism called 'job submission under SMS control' was developed. It ran on the Member States' general-purpose server (ecgate) and made use of SMS, the Supervisor Monitor Scheduler, to manage the job submission. Users provided a job script and indicated at which stage of the ECMWF operational forecast production they wanted the job to be executed.

This mechanism allowed users to make effective use of ECMWF by accessing data produced by one of the ECMWF models as soon as possible after production. Typically, the data were post-processed on ecgate to create tailored products, which were then transferred to the users' local systems. By 2004, about 300 jobs belonging to approximately 60 Member State users were being run each day under the mechanism.

Although this mechanism worked well, there were limitations. In particular, the service only provided a mechanism to submit the users' jobs. If the job failed for any reason, including temporary system issues such as a file system being full, then it was the users' responsibility to notice this, fix any issues with the script, and resubmit.

For more complex applications initialised by output from ECMWF models, which needed to run routinely and in a timely fashion on ECMWF systems, different ad-hoc technical solutions were used. In some cases, such as the prediction of ocean waves (which became a core ECMWF activity on 1 January 1999) and the Boundary Conditions (BC) for Limited Area Modelling, ECMWF Optional Projects were used. Others, such as COSMO-LEPS (Consortium for Small-Scale Modelling – Limited-Area Ensemble Prediction System), which started to run

a The Framework in detail

Option 1: Simple time-critical job submission monitored by ECMWF:

- Suitable for single jobs or simple jobs with multiple steps.
- Available to all registered users – no formal request needed.

Option 2: Member State suites monitored by ECMWF:

- Suitable for more complex applications comprising several interdependent tasks.
- Suites monitored by ECMWF with second- and third-level support provided by the Member States.
- Requested by the Technical Advisory Committee (TAC) representative of the relevant Member State.

Option 3: Member State suites managed by ECMWF:

- Application developed, tested and maintained by the Member State.
- ECMWF monitors the suite and provides second-level on-call support.
- It must be possible to test the application using ECMWF pre-operational (e-suite) data.
- Third-level support is provided by the Member State.
- Requested by the TAC representative of the relevant Member State.

 doi: 10.21957/p31dmu80sd

at ECMWF in November 2002 (see Montani et al., 2003, for details), ran under SMS suites in a similar way to ECMWF's operational suite.

As more requests were received – notably from the Norwegian Meteorological Institute to support a semi-operational run of its TEPS (Targeted Ensemble Prediction System) suite, and from the Hellenic National Meteorological Service to support its limited-area modelling forecasting activities during the Olympic and Paralympic Games in 2004 – it became clear that a more formal framework for managing time-critical activities was needed.

The Framework for Member State time-critical applications

Recognising the need to overcome the limitations of the job submission under the SMS control mechanism, and to offer a service better suited to the more complex applications being requested, a Framework for Member State time-critical applications comprising three options was proposed (see Box A):

1. Simple time-critical job submission monitored by ECMWF
2. Member State suites monitored by ECMWF
3. Member State suites managed by ECMWF.

The Framework was discussed by ECMWF's Technical Advisory Committee in October 2004 and approved by Council in December 2004. It has been in use ever since.

Time-critical option 1

Option 1 of the time-critical Framework (TC-1) provides a simple job submission mechanism, which replaced the 'job submission under SMS control'. TC-1 is available to all registered Member and Co-operating State users

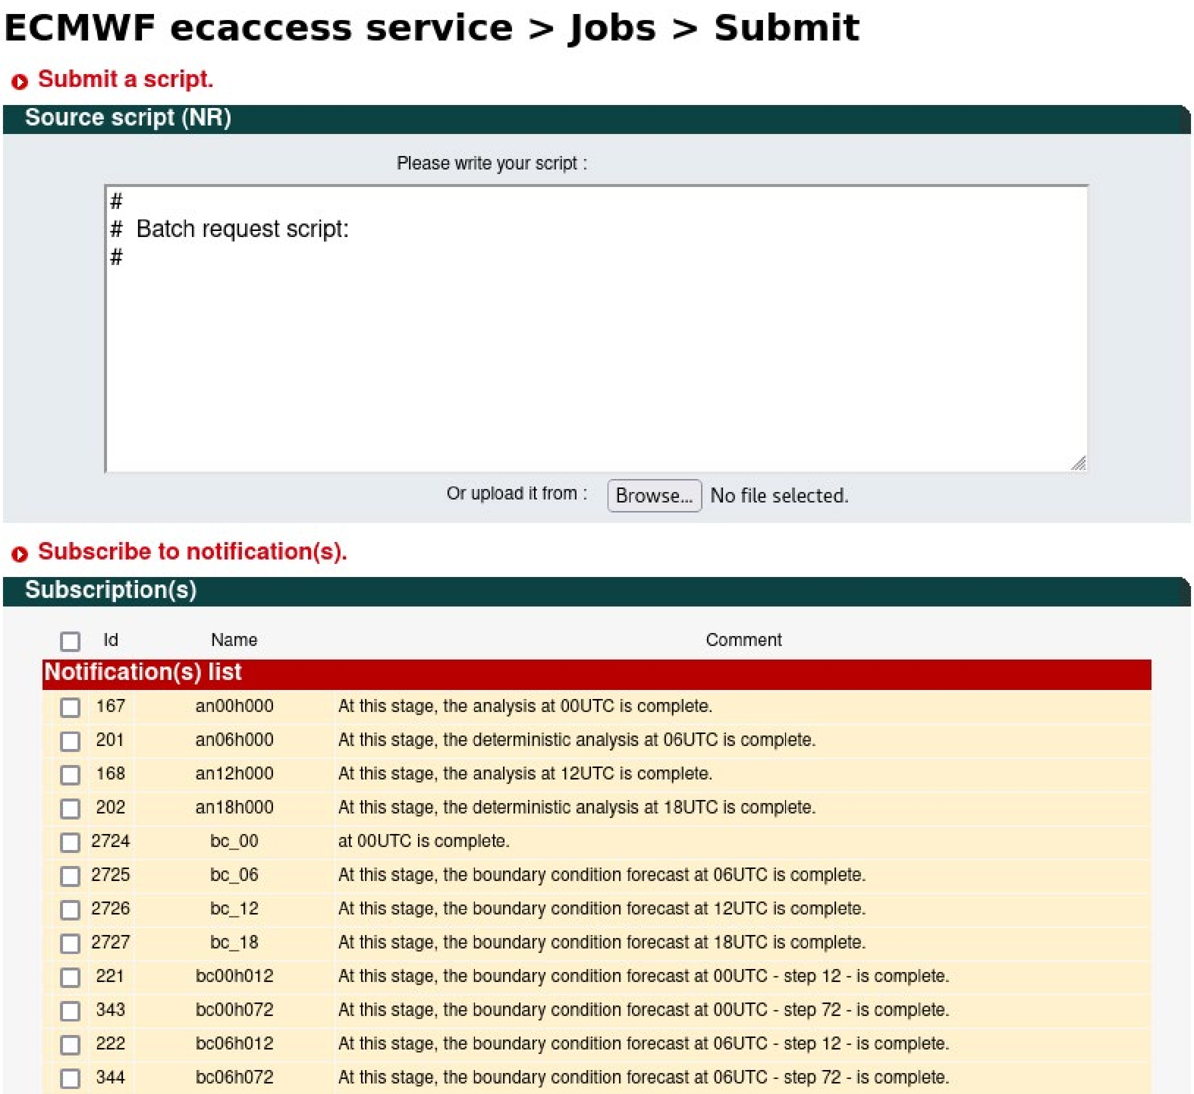

ECMWF ecaccess service > Jobs > Submit

Submit a script.

Source script (NR)

Please write your script :

```
#
#  Batch request script:
#
```

Or upload it from : Browse... No file selected.

Subscribe to notification(s).

Subscription(s)

	Id	Name	Comment
		Notification(s) list	
☐	167	an00h000	At this stage, the analysis at 00UTC is complete.
☐	201	an06h000	At this stage, the deterministic analysis at 06UTC is complete.
☐	168	an12h000	At this stage, the analysis at 12UTC is complete.
☐	202	an18h000	At this stage, the deterministic analysis at 18UTC is complete.
☐	2724	bc_00	at 00UTC is complete.
☐	2725	bc_06	At this stage, the boundary condition forecast at 06UTC is complete.
☐	2726	bc_12	At this stage, the boundary condition forecast at 12UTC is complete.
☐	2727	bc_18	At this stage, the boundary condition forecast at 18UTC is complete.
☐	221	bc00h012	At this stage, the boundary condition forecast at 00UTC - step 12 - is complete.
☐	343	bc00h072	At this stage, the boundary condition forecast at 00UTC - step 72 - is complete.
☐	222	bc06h012	At this stage, the boundary condition forecast at 06UTC - step 12 - is complete.
☐	344	bc06h072	At this stage, the boundary condition forecast at 06UTC - step 72 - is complete.

FIGURE 1 A screenshot of part of the ECaccess web interface for batch job submission. The user can enter the script to be run in the box at the top or upload from a file. The events to which the job should be subscribed can be selected from the notification list, part of which can be seen at the bottom of the page.

with access to ECMWF's high-performance computing facility (HPCF) and ECGATE Class Service (ECS). It is suitable for single jobs or jobs comprising multiple steps with straightforward interdependencies.

To support TC-1, an enhancement of the ECaccess batch system was implemented. The ECaccess command line tools could already be used to submit jobs to the batch systems on ecgate and the HPCF, and a retry mechanism was already available for file transfers using ectrans. This retry mechanism was extended to the batch jobs and a new concept of 'events' or 'notifications' was introduced.

The events are linked to the ECMWF operational dissemination schedule and are triggered by tasks in the ECMWF operational suite. Users subscribe their jobs to the events, and the jobs are held in standby mode until the operational suite sends an event notification. When the notification for an event is received, ECaccess submits all the jobs subscribed to that event and creates a new job entity in standby mode, ready for the next notification of the event. Users can submit their jobs to the TC-1 mechanism either from the command line or using the ECaccess web interface (see Figure 1).

To enable ECMWF 24x7 shift staff to monitor the TC-1 jobs, the web interface shown in Figure 2 was developed. The interface alerts staff when jobs fail and provides an option to try to re-run them.

Over the years, the 'Simple time-critical jobs' service has proved to be very popular. Today, there are about 150 different users running 1,800 jobs on average per day. Each job is subscribed to one of 100 defined events. Many use the service to access real-time data, using the Meteorological Archival and Retrieval System (MARS) to create bespoke products and charts; run

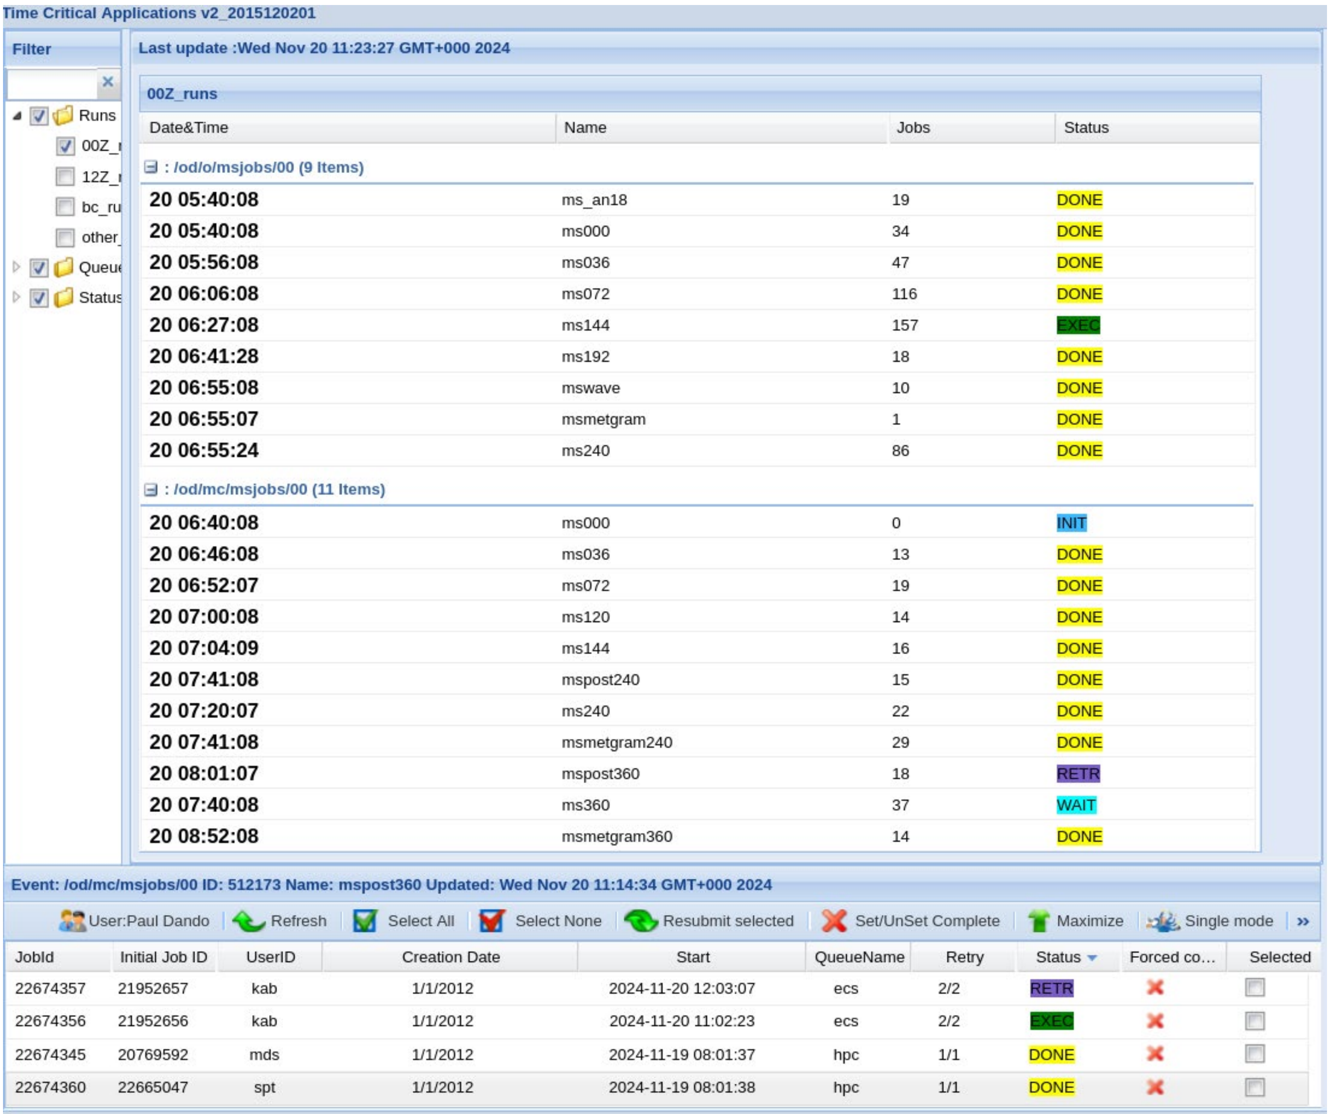

Date&Time	Name	Jobs	Status
: /od/o/msjobs/00 (9 Items)			
20 05:40:08	ms_an18	19	DONE
20 05:40:08	ms000	34	DONE
20 05:56:08	ms036	47	DONE
20 06:06:08	ms072	116	DONE
20 06:27:08	ms144	157	EXEC
20 06:41:28	ms192	18	DONE
20 06:55:08	mswave	10	DONE
20 06:55:07	msmetgram	1	DONE
20 06:55:24	ms240	86	DONE
: /od/mc/msjobs/00 (11 Items)			
20 06:40:08	ms000	0	INIT
20 06:46:08	ms036	13	DONE
20 06:52:07	ms072	19	DONE
20 07:00:08	ms120	14	DONE
20 07:04:09	ms144	16	DONE
20 07:41:08	mspost240	15	DONE
20 07:20:07	ms240	22	DONE
20 07:41:08	msmetgram240	29	DONE
20 08:01:07	mspost360	18	RETR
20 07:40:08	ms360	37	WAIT
20 08:52:08	msmetgram360	14	DONE

JobId	Initial Job ID	UserID	Creation Date	Start	QueueName	Retry	Status	Forced co...	Selected
22674357	21952657	kab	1/1/2012	2024-11-20 12:03:07	ecs	2/2	RETR	✗	☐
22674356	21952656	kab	1/1/2012	2024-11-20 11:02:23	ecs	2/2	EXEC	✗	☐
22674345	20769592	mds	1/1/2012	2024-11-19 08:01:37	hpc	1/1	DONE	✗	☐
22674360	22665047	spt	1/1/2012	2024-11-19 08:01:38	hpc	1/1	DONE	✗	☐

FIGURE 2 Web interface for monitoring time-critical option 1 jobs. The top part of the interface shows a summary of the status of all jobs for each event grouped by specific forecast runs. Here only those associated with the 00 UTC forecast runs are shown. The lower part shows the status of all jobs associated with a specific event. Job output can be viewed and any failed jobs re-run through the interface.

trajectory computations; make statistical analyses of ensemble forecasts (ENS); provide data for customers of national meteorological services; or even to trigger suites running under option 2 of the TC Framework. One particularly popular use is to download ENS meteograms for user-specified locations.

Time-critical option 2

Option 2 of the time-critical Framework (TC-2) is appropriate for running more advanced applications comprising several tasks with potentially complex interdependencies. Typically, TC-2 is used for running limited-area models (LAMs), such as HARMONIE, COSMO or ICON, with initial and lateral boundary conditions (LBCs) provided by ECMWF's Integrated Forecasting System (IFS). Workflows running as TC-2 are monitored by ECMWF's 24x7 shift staff, who provide first-line support. Jobs submitted have access to high-priority batch queues on the ECMWF HPCF to ensure they execute in a timely manner. Figure 3 shows currently running TC-2 suites viewed in ecFlowUI.

The first activities to run as TC-2 applications

In July 2005, COSMO-LEPS, which had been running on the ECMWF HPCF since November 2002, became the first suite to be declared officially as a TC-2 application. The application was developed and maintained by the weather service of the Italian Emilia-Romagna region, Arpae-SIMC, with the aim of providing high-resolution mesoscale ensemble forecasts of localised high-impact weather events. Initially, it ran a control forecast plus a ten-member ensemble based on the 'Lokal Modell' with initial conditions and LBCs from ECMWF ENS forecasts, selected via a clustering selection technique. Forecasts were produced twice per day to 120 hours, with 3-hourly output archived in MARS and disseminated to other COSMO partners. Over the years, COSMO-LEPS has been upgraded on several occasions. Today, it runs a 20-member ensemble producing forecasts to 132 hours.

Following COSMO-LEPS, other Member States started to run either regional or global ensemble numerical weather prediction (NWP) forecasts as TC-2 applications.

- The UK Met Office's MOTHS (Met Office Thorpex Suite, later renamed to MOGREPS-15) suite ran as a TC-2 application from 2007 to 2016. It produced a global 15-day, 24-member ensemble forecast based on the Unified Model, primarily as the UK Met Office's contribution to the Thorpex Interactive Grand Global Ensemble (TIGGE) archive.
- The Grand Limited-Area Model EPS (GLAMEPS – see Iversen et al., 2020) was a multi-model ensemble prediction system on a limited, pan-European domain

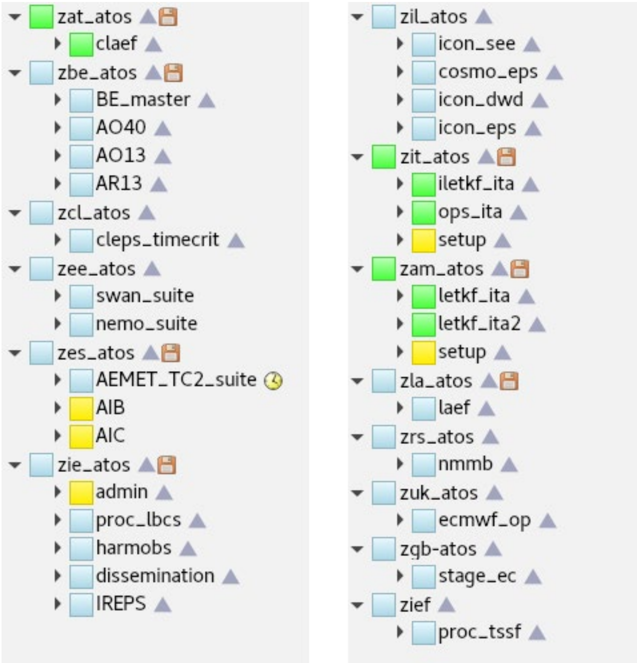

FIGURE 3 The TC-2 suites active in November 2024, viewed in ecFlowUI.

which ran as a TC-2 application on behalf of the HIRLAM and ALADIN consortia from 2011 to 2019. It aimed to provide reliable forecasts for the next two to three days with a higher spatial resolution than ECMWF's global ENS forecast at the time.
- In 2011, the Austrian national meteorological and geophysical service (ZAMG – now part of the GeoSphere Austria) started to run the ALADIN-Limited Area Ensemble Forecasting (LAEF) suite as TC-2. LAEF was based on the ALADIN/ALARO NWP model and was developed in the framework of the Regional Cooperation for Limited Area modelling in Central Europe (RC-LACE). In 2019, LAEF was replaced by C-LAEF (Convection-permitting LAEF), which still runs today. Output is made available to countries in the LACE community and also provides uncertainties for the Ensemble INCA nowcasting system at ZAMG.
- The German National Meteorological Service (Deutscher Wetterdienst – DWD) also ran its Boundary Condition Ensemble Prediction System - Multi Model (BCEPS MuMo) suite as TC-2 between May 2011 and March 2017. The suite interpolated global model data from ECMWF's IFS and forecasts from several national meteorological services to provide LBCs for DWD's COSMO-DE members.

TC-2 applications providing backups to operational NWP activities

Other Member States have used option 2 of the TC Framework to run backups of their operational NWP.

For example:

- The Instituto Português do Mar e da Atmosfera (IPMA) ran a backup of its ALADIN-Portugal LAM as a TC-2 application between November 2013 and May 2019.
- Between January 2015 and the end of October 2021, DWD ran a backup of part of its NWP activity as a TC-2 application in case of a catastrophic disaster at its main site in Offenbach.
- In May 2015, Spain's Agencia Estatal de Meteorología (AEMET) began running a backup of its operational HARMONIE suite as a TC-2 application. The model is run for two domains: the Iberian Peninsula and the Canary Islands.
- Since July 2019, the National Centre for Aerospace Meteorology and Climatology (CNMCA) of the Italian Air Force Meteorological Service (ITAF) has been using four TC-2 suites to run its NWP system, mostly for back-up purposes. The COSMO model and the NETTUNO ocean wave model are run for two domains and the ICON model for a single domain to provide both deterministic and probabilistic forecasts. A fourth suite runs the KENDA-LETKF atmospheric ensemble data assimilation system to provide initial conditions for the COSMO and ICON forecasts.

Other NWP activities run as TC-2 applications

In more recent years, several Member States have begun to use the TC Framework to run part or all of their operational NWP.

- In June 2017, the Irish Meteorological Service, Met Éireann, started to run its full operational NWP production as a TC-2 application. The original deterministic forecasts were based on the HIRLAM and HARMONIE models. By June 2019, these had been replaced by a new HARMONIE-based control plus ten-member ensemble system named IREPS. Nowadays, Met Éireann takes its main NWP products from the United Weather Centres – West (UWC-W) forecasts. However, a skeleton version of IREPS still runs as TC-2.
- Met Éireann also elected to run its operational Tidal and Storm Surge Forecasts (TSSF) under the TC Framework. In October 2022, TSSF was the first new suite approved to run as TC-2 on ECMWF's Atos HPC.
- With a lack of computer resources preventing an increase of model resolution or the production of ensemble forecasts, the Republic Hydrometeorological Service of Serbia (RHMSS) developed a suite to run its operational forecast model. The suite was approved as a TC-2 application in early 2018.
- Between May 2018 and October 2022, the Hellenic National Meteorological Service (HNMS) ran the COSMO model as a TC-2 application to provide high-resolution products for its various operations. Forecasts were run twice per day for two different domains over the Mediterranean Sea and Greece.
- In July 2020, A-LAEF (ALARO-LAEF), the RC-LACE consortium's common ensemble prediction system based on the ALARO LAM (see Belluš et al., 2022, for details), began running as a TC-2 application using resources provided by Croatia, Slovenia and Türkiye. Products are distributed to the LACE partners as well as Türkiye, while Poland receives data via servers at the Slovak Hydrometeorological Institute (SHMU).
- Since May 2022, the Belgian Meteorological Service (IRM/KMI) has run its operational NWP as a TC-2 application as three different suites: ALARO forced by ARPEGE data, uploaded by IRM to the ECMWF Atos HPCF, is run at 4 km resolution; ALARO and AROME are both run at 1.3 km resolution with initial conditions and LBCs based on IFS forecasts.
- Finally, in August 2024, two new TC-2 suites were approved for ESTEA, the meteorological service for Estonia. The suites provide wave forecasts in real time for a domain covering the Baltic Sea in the vicinity of Estonia. One suite makes daily runs of the NEMO-EST ocean forecast model, based on the NEMO-Nordic version of the widely used NEMO ocean model, while the other runs Delft University of Technology's SWAN ocean wave forecast model twice a day.

TC-2 applications processing ENS data

The UK Met Office has also used the TC Framework to run applications which post-process large volumes of ECMWF ensemble forecast data without the need to transfer these first to its Exeter headquarters.

The Site-Specific Post-processing Suite (SSPS) has been running as a TC-2 application since 2013. SSPS extracts site-specific information from ECMWF forecasts for locations around the globe. This is blended with output from UK Met Office models in its operational post-processing system to produce site-specific forecast data.

The UK Met Office has also implemented its Standard Gridding Engine (StaGE) as a TC-2 application which was accepted as a TC-2 suite in September 2024. StaGE processes ECMWF ENS forecast data and delivers products to the UK Met Office in a NetCDF format, ensuring they are in a standard gridded format for use in downstream UK Met Office systems. The primary goal is to make it easier to integrate ECMWF ENS output with that from the Unified Model to produce novel post-processed products.

Israel's activities under TC-2

Although not an ECMWF Member State, Israel has also been able to take advantage of the TC Framework to run operational weather forecasts as TC-2 applications (see Khain et al., 2022). This is thanks to the Optional Programme, which, since 1 January 2020, has allowed Co-operating States without the possibility to become Member States to use ECMWF's HPCF. The Israel Meteorological Service (IMS) has used resources provided under the Programme to run a computationally demanding COSMO ensemble system over the Eastern Mediterranean as well as a high-resolution ICON model over southeastern Europe. These provide weather forecasts for IMS and other weather services in the region. The ICON forecasts are also used within the World Meteorological Organization SEE-MHEWS-A (South-East European Multi-Hazard Early Warning Advisory System) project. In 2023, a third suite for running the ICON model, with initial conditions and LBCs from DWD, was accepted as TC-2, and this year a fourth suite to run an ensemble forecast based on the ICON model was approved.

b **Examples from Member States**

Trond Iversen, MET Norway: GLAMEPS (Grand Limited-Area Model Ensemble Prediction System) was a multi-model ensemble numerical weather prediction system on a limited, pan-European domain. The system was run as a time-critical facility (TC-2) at ECMWF from 2011 to 2019. GLAMEPS was technically ambitious and would hardly have been realised without the availability of TC-2. It was a challenge to obtain timely and frequent forecast updates without interfering with ECMWF's production, which was why six-hourly time-lagging was implemented in an upgrade in 2014. ECMWF staff were crucial for the operational implementation, including a hierarchy of real-time emergency procedures.

Andrea Montani, formerly Arpae-SIMC, now ECMWF: When COSMO-LEPS started to run at ECMWF in November 2002, no single country had the computing capabilities to run the application in a timely manner without interfering with routine operational activities. The cooperation and special efforts provided by ECMWF for the full implementation of the project under the time-critical framework had (and still has!) to be explicitly praised.

Time-critical option 3

Option 3 of the time-critical Framework (TC-3) is an enhancement of TC-2, where ECMWF both monitors and manages the suites. The TC-3 suites are closely integrated with ECMWF's operational suites. They are managed by the Integration Team of ECMWF's Production Section, who also undertake any testing of the suites whenever the IFS is upgraded.

One of the main applications being run as TC-3 provides initial conditions and LBCs based on ECMWF forecasts as Format Arpege files for the AROME, ALARO and ALADIN LAMs, as well as for the MOCAGE chemical transport model. The processing is based on the '903 configuration' available in IFS/ARPEGE. It was originally developed under a Special Project which started in 2006, led by Météo-France and supported by Hungary. Output is disseminated using the ECMWF Product Dissemination System (ECPDS) to members of the LACE community as well as Météo-France, Türkiye, Morocco and Bulgaria. Belgium and Austria also use the output from this TC-3 activity to initialise their forecasts running as TC-2 applications.

Several other TC-3 applications are being run by Météo-France, each of which was developed and run initially as a TC-1 job.

- The PERLE (Modèle de dispersion de polluants à échelle locale et régionale) suite runs on demand. It extracts 3D data from MARS on a domain and for a period of time included in a request file, which is sent to ECMWF via ECPDS. The data extracted are sent back to Météo-France, providing part of the input data used to run the PERLE dispersion model. The PERLE suite is triggered typically when there is a chemical incident, or possibly a volcanic eruption, anywhere around the French territories or elsewhere in the world.
- ASOT2MIFS (Adaptation Statistique à partir d'Observations de la Température à 2M pour IFS) produces three-hourly and extreme temperature forecasts from statistical adaptations of ECMWF forecast output.
- The ALPHA (Algorithmes et modèLes pour la Production Homogène globAle) suite post-processes ECMWF ensemble forecasts using an algorithm developed at Météo-France to create bespoke probabilistic diagnostics.
- FAME (Forecast of Airport Meteorological Elements) uses ECMWF forecast output to produce forecasts of variables of interest to the aviation industry at a list of airport locations worldwide.

Other applications currently running as TC-3 are:

- Jason2, which post-processes ECMWF analysis fields for EUMETSAT and the Centre national d'études spatiales (CNES - the French national space agency), and
- the UK Met Office's extratropical Cyclone Database (CDB) application. This application creates a web-based product that aims to represent, objectively and in a variety of ways, the location and behaviour of near-surface synoptic scale features in ECMWF forecasts. The features represented are those typically associated with adverse weather, such as warm and cold fronts, barotropic lows and frontal and diminutive waves.

The currently running TC-3 suites viewed in ecFlowUI are shown in Figure 4.

Using the TC Framework in emergency situations

On two occasions, the TC Framework has been used to support ECMWF Member States in emergency situations.

First, in August 2012, the Danish Meteorology Institute (DMI) experienced a series of problems with its power supply, resulting in severe damage to some of the central components of its HPCF, which consequently went out of service. DMI was already running a version of its HIRLAM model on the ECMWF supercomputer. It contacted ECMWF to ask if it would be possible to get high-priority access to sufficient HPC resources for a few days so that its NWP output could continue to be produced in a timely fashion while engineers worked on repairs to its HPCF. Access was quickly enabled to allow DMI to run in the higher-priority batch queues that are reserved for time-critical work under the TC Framework.

On a second occasion, a 5.3 magnitude earthquake on Sunday, 22 March 2020 severely damaged the headquarters of the Croatian Meteorological and Hydrological Service (DHMZ) in Zagreb. DHMZ quickly ported its NWP activities and some post-processing tasks to ECMWF's HPCF (see Abellan et al., 2020). Within a matter of days, the main components of the NWP system were running on the Centre's supercomputers, again making use of the higher-priority batch queues provided through the TC Framework to enable the forecasts to run in a timely manner.

Current status and future outlook

Twenty years after the TC Framework was approved, it remains very popular with Member and Co-operating

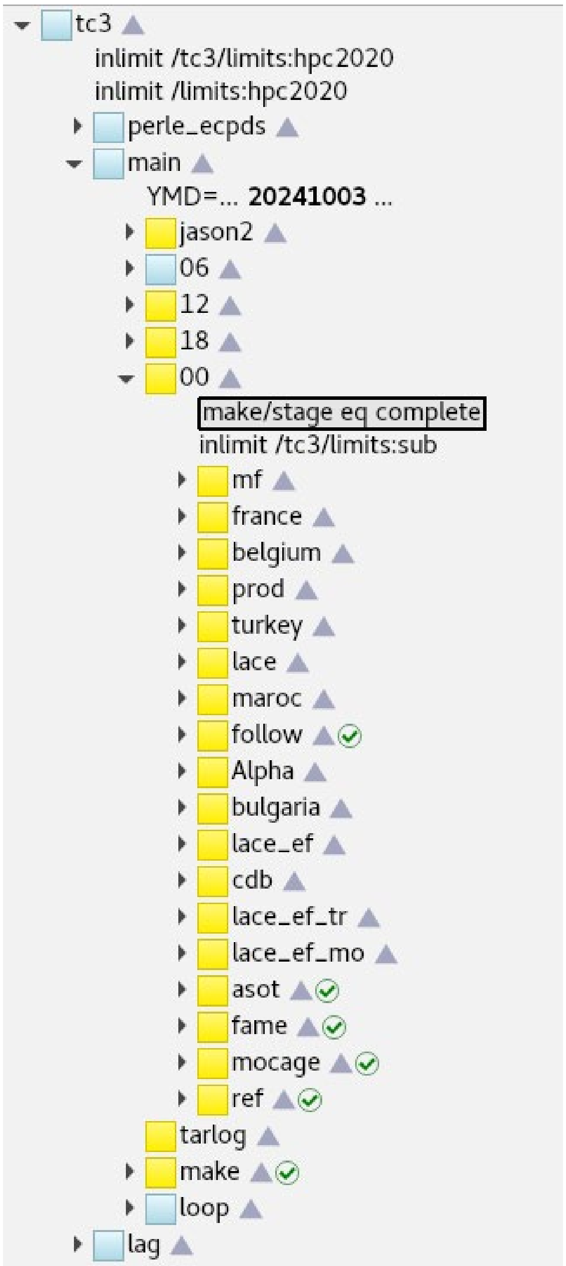

FIGURE 4 The TC-3 suites active in November 2024, viewed in ecFlowUI.

States. Four new TC-2 suites were approved during 2024, bringing the total to 23 suites being run by 11 different countries and consortia. Two further suites are being prepared and are expected to be approved as TC-2 within the next 12 months. There is also a major upgrade planned for the oldest TC-2 application, COSMO-LEPS, with the COSMO model being replaced by ICON.

There have also been three new TC-3 suites added during the last 12 months to provide initial conditions and LBCs as Format Arpege files for Morocco, Türkiye and Bulgaria. This brings the total number of TC-3 suites being run to 17 on behalf of 11 countries, organisations and consortia. For the future, Météo-

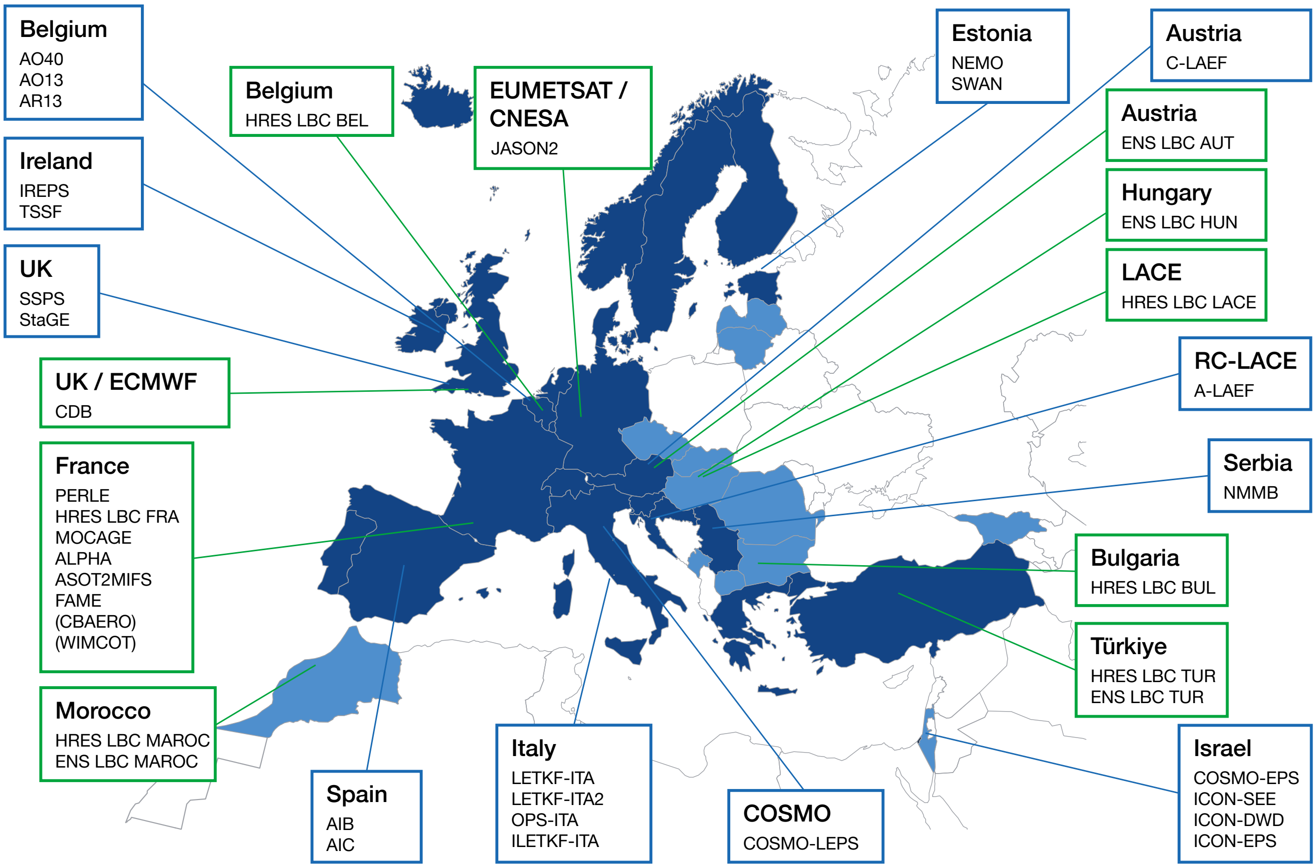

FIGURE 5 TC-2 (in blue boxes) and TC-3 (in green boxes) suites at the beginning of November 2024, grouped by country, organisation or consortia.

France is preparing two further potential TC-3 activities to support forecasts for commercial aviation. The TC-2 and TC-3 suites active in November 2024 are depicted schematically in Figure 5.

The TC-1 activity also remains very popular with Member and Co-operating State users. During August 2024, around 150 different users ran a total of almost 55,500 jobs (about 1,800 jobs per day on average), with each job subscribed to one of 102 different defined 'events'. In 2024, the Aviso service became operational (Iacopino et al., 2021). This provides users on the ECMWF HPCF or the European Weather Cloud with notifications when disseminated data or data in MARS are available, allowing them to trigger workflows at remote sites. However, Aviso is not a direct replacement of TC-1 events as it does not provide a monitoring facility or a re-run of failed jobs by ECMWF.

Recently a review of the TC Framework was undertaken to understand where updates and improvements can be made. The review took into consideration best practices and lessons learned from experiences with ECMWF's operational suites and with Member States' time-critical activities. It also looked at new technologies, such as the Atos HPCF, the European Weather Cloud, containerised workflows, and dedicated Virtual Machines for hosting the ecFlow servers. The review generated several straightforward and non-controversial suggestions for improving the Framework, which are gradually being implemented. Other recommendations, such as the possibility of using the European Weather Cloud to run certain applications, are still being explored.

Over the years, the TC Framework has enabled ECMWF Member States to run workflows which they could not easily run themselves. These include running high-resolution LAMs or ensemble forecasts or post-processing large volumes of data to produce novel products, in a routine and timely manner. Clearly, the TC Framework continues to provide great benefit to ECMWF's Member and Co-operating States, and we look forward to the next 20 years of the service.

Further reading

Dando, **P.**, **D. Lucas** & **U. Modigliani**, 2023: Thirty years of the ecgate service at ECMWF. *ECMWF Newsletter* **No. 177**, 29–33. *https://www.ecmwf.int/en/newsletter/177/computing/thirty-years-ecgate-service-ecmwf*

Montani, **A.**, **M. Capaldo**, **D. Cesari**, **C. Marsigli**, **U. Modigliani**, **F. Nerozzi et al.**, 2003: Operational limited-area ensemble forecasts based on the 'Lokal Modell'. *ECMWF Newsletter* **No. 98**, 2–7. *https://www.ecmwf.int/en/elibrary/78234-newsletter-no-98-summer-2003*

Iversen, **T.**, **I.L. Frogner**, **X. Yang**, **K. Sattler** & **A. Deckmyn**, 2020: The life of GLAMEPS. *ECMWF Newsletter* **No. 165**, 12–13. *https://www.ecmwf.int/en/newsletter/165/news/life-glameps*

Belluš, **M.**, **M. Tudor** & **X. Abellan**, 2022: The mesoscale ensemble production system A-LAEF. *ECMWF Newsletter* **No. 172**, 27–34. *https://www.ecmwf.int/en/newsletter/172/earth-system-science/mesoscale-ensemble-prediction-system-laef*

Khain, **P.**, **A. Shtivelman**, **H. Muskatel**, **A. Baharad**, **L. Uzan**, **E. Vadislavsky et al.**, 2022: Israel uses ECMWF supercomputer to advance regional forecasting. *ECMWF Newsletter* **No. 171**, 28–35. *https://www.ecmwf.int/en/newsletter/171/earth-system-science/israel-uses-ecmwf-supercomputer-advance-regional-forecasting*

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/news/croatian-met-service-backs-its-production-ecmwf-after-earthquake*

Iacopino, **C.**, **T. Quintino**, **J. Hawkes**, **B. Raoult**, **M. Partio**, **M. Grønlien Pejcoch et al.**, 2021: Aviso: ECMWF's data availability notification service. *ECMWF Newsletter* **No. 170**, 28–31. *https://www.ecmwf.int/en/newsletter/170/computing/aviso-ecmwfs-data-availability-notification-service*

ECMWF Sites: websites as a service

Manuel Martins

ECMWF's web presence is pivotal not only in disseminating its forecast products, but in collaborating and communicating research findings to a diverse audience, including scientists, policymakers, and the general public.

ECMWF Sites is a service that enables ECMWF staff and visiting scientists to effortlessly create and publish websites, providing a basis for collaboration, communication and fast prototyping. It ensures appropriate control and communication about the quality of ECMWF websites or publications outside our main website, *www.ecmwf.int*, meeting a minimum quality bar for our web presence. *ECMWF Sites* provides mainly private websites for use by our scientists, but websites can also be made publicly available, subject to an approval process.

The platform is built using Kubernetes, an open-source system for automating the deployment, scaling, and management of containerised applications (*https://kubernetes.io*). It leverages the use of the Kubernetes operator pattern to manage the lifecycle of a website. The platform and all websites are hosted in the ECMWF data centre, running in Kubernetes on a dedicated virtual infrastructure.

This article presents the service and provides an overview of the platform's architecture along with insights on the past, present and future of the platform.

Background and inception

At ECMWF, collaboration is ingrained in our operational framework: scientists, researchers, and engineers work together, leveraging their diverse backgrounds to improve weather forecasting models and climate prediction tools. ECMWF actively collaborates with external entities, including national meteorological services, academic institutions, and other international organisations. Web-based collaboration tools play a crucial role in facilitating communication and cooperation. These platforms enable real-time data sharing, collaborative editing of documents, and the ability to provide feedback on ongoing projects, regardless of geographical location.

ECMWF Sites is a web platform that was planned with this in mind. It provides a simple way for users to create a space, manage content and share it with others, internally or externally. It was initially designed as an HTTP hosting service for static content, with simple and straightforward requirements:

- enough disk space
- easy content management
- accessible with or without authentication.

These requirements were gathered throughout the years from real use cases. Some simple solutions were previously available, but they were complex to use, creating confusion and frustration for users. These solutions also did not reflect a modern service-oriented way of working for users.

Designed and released to pre-production by the end of 2019, as a proof of concept on Kubernetes native applications, *ECMWF Sites* served as the experimental model for a controlled service fully integrated with Kubernetes. This offered a solid understanding of the platform and provided insights into how the team could leverage the Kubernetes platform at ECMWF, running production containers at scale.

The use of containers for production services was first implemented at ECMWF in mid-2018, with the re-architecture of the Atlassian suite of products. These products were deployed in Docker containers, on dedicated virtual machine infrastructure, managed by *systemd* and provisioned with *Puppet*. We soon realised the need for a container orchestration platform. By the beginning of 2019, ECMWF was running the first production Kubernetes cluster and started to migrate small applications. A huge shift in technologies happened during this period, which meant some effort to ensure a smooth transition for all teams using the infrastructure.

ECMWF Sites was then released to production in April 2020, with less than 1,000 lines of Python code and around 25 websites published.

From the outset, the potential of the service was evident. Soon after, user feedback brought many additional requirements to light. The service became more customisable to the point that users can now build and run their own containerised applications as part of a website. Some of the requirements introduced later are:

- content retention control
- web analytics and reporting

 doi: 10.21957/m5gl27f96m

- multi-factor authentication
- support for caching, web robots and cross-origin resource sharing (CORS)
- running custom applications (containers).

Almost five years later, *ECMWF Sites* is about 10,000 lines of code, and it is written in Python, Golang and Lua. At the time of writing, it hosts around 300 websites (see Figure 1).

Overview of the service

ECMWF Sites is a web platform accessible through *https://sites.ecmwf.int*. Accessing this sends the user to the *Sites Hub*, which serves as the central entry point, providing a list of all available websites, along with their details and URLs. Each website is part of a 'space', typically the user's username, and has a unique name within that space. Users can view and configure their websites through the *Hub*'s site view page (accessible through the first column on *https://sites.ecmwf.int/hub/list/all/*). Websites are accessible via specific URLs, and administrators can manage all content through a web-based file browser or programmatically via a Representational State Transfer (REST) application programming interface (API).

Websites can be either private or public. Owners of private websites can share access with specific users or groups, who must authenticate to view the content. In addition to the web interface, users can interact with the platform using a Python Software Development Kit (SDK) or a Command Line Interface (CLI). A unique authentication token is generated for each website, enabling secure API interactions for automated tasks, without linking to specific user credentials.

Use cases

The service is being used for many different purposes by users with different backgrounds. Figures 2 and 3 show the distribution of websites per type and use case.

Verification and validation assessment plots are a very common use case of websites. The Iver private websites are used by scientists and researchers to access Iver verification tool results. This is done to assess the improvement their research provides over the current version of ECMWF's Integrated Forecasting System (IFS). During new IFS Cycle implementations, these websites are responsible for more than 1 TB of ingress traffic and 0.5 TB of egress traffic, per week. As ECMWF scientist Alan Geer describes:

> *"ECMWF Sites has revolutionised the way we share diagnostic plots with ECMWF scientists both within and outside ECMWF premises. My example is the Iver forecast verification tool, which has around 100 users across ECMWF and is one of the main ways in which the research department assesses the impact of upgrades to the IFS. When comparing the forecast and analysis quality of two experiments, Iver generates at least 14,000 plot panels, and these are summarised in a web page. Scientists run lots of experiments, meaning possibly hundreds of separate web pages. In the past, these web pages were only accessible within ECMWF, but now they are also accessible externally. To achieve this, each Iver user's workstation had to be turned into a mini web server to share the*

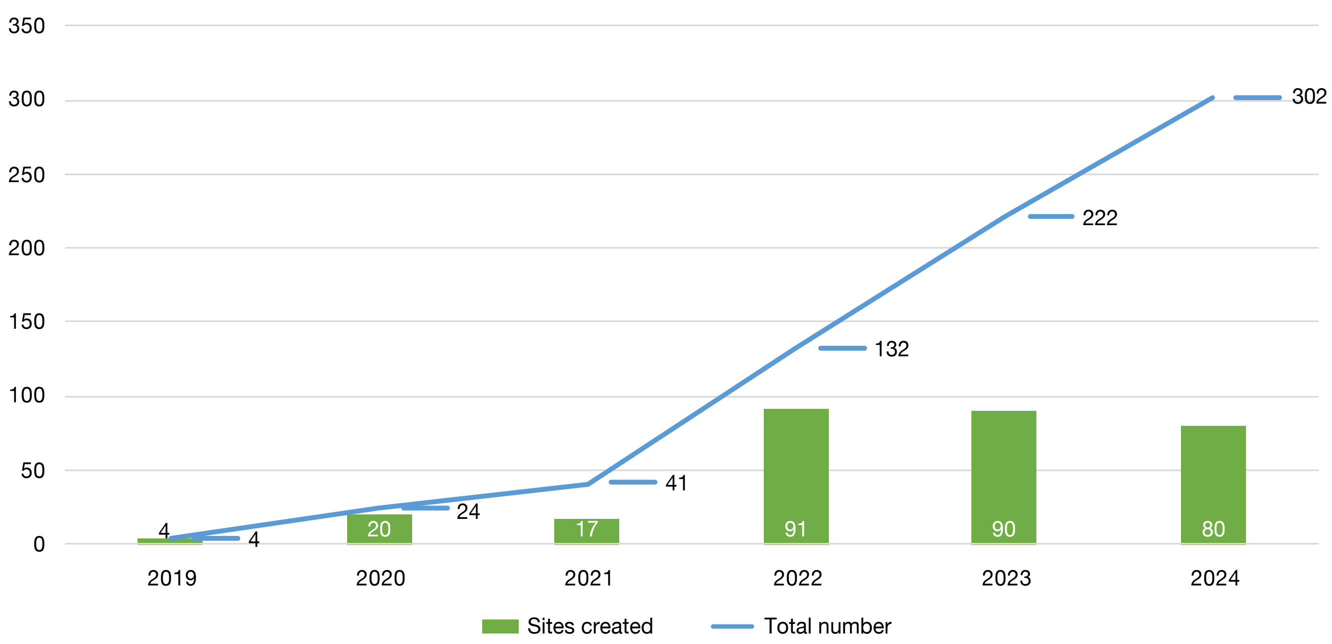

FIGURE 1 This chart shows the evolution of websites created and their overall number by the end of 2024.

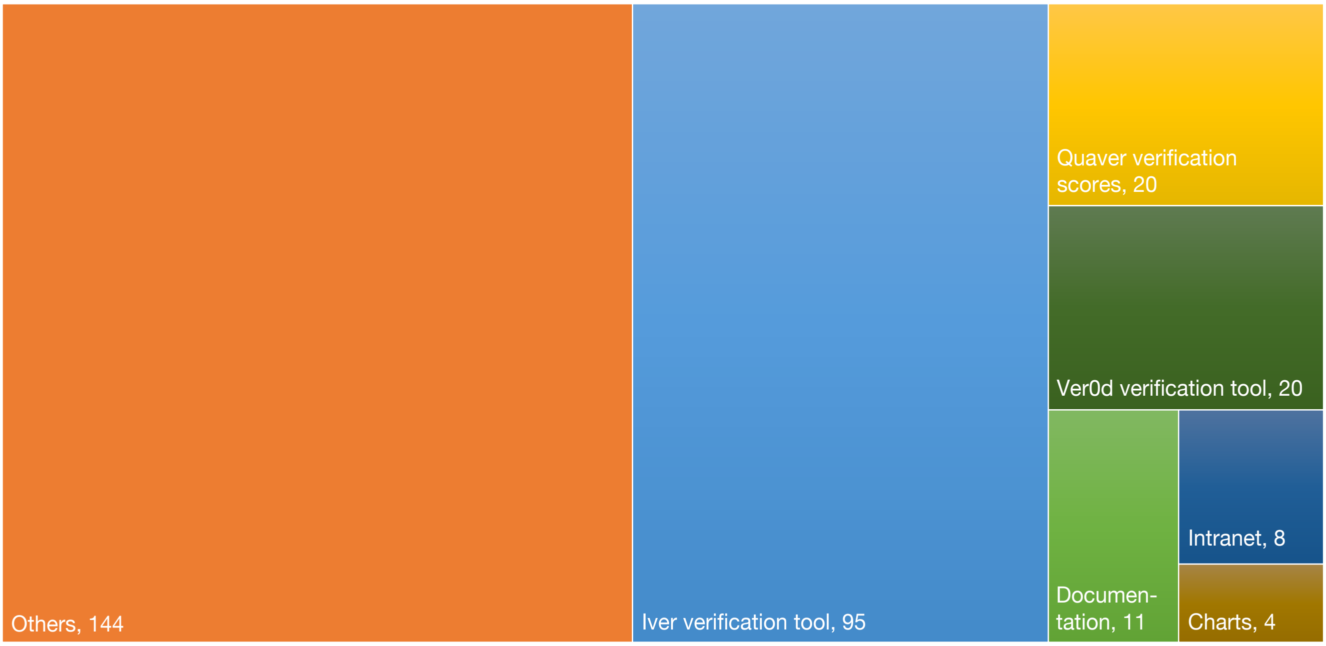

FIGURE 2 This chart shows the distribution of websites per use case at the end of 2024.

plots with others. Sites now deals with all of this and allows us to share these plots easily both inside and outside the organisation. It is comfortably hosting as many as 100 million plots from Iver users alone."

ECMWF Sites has been used as a fast-prototyping platform, allowing seamless integration with other web services. With its API-driven model, it allows easy uploads and it is often used as a repository for data, images, JavaScript and other static content, allowing other services to make these accessible, mainly to internal staff. As Helen Setchell, senior content architect and user experience coordinator at ECMWF, describes:

"As ECMWF's senior content architect and UX coordinator, I work with a variety of platforms to support our online presence, and while they each serve their purpose, they come with certain limitations – whether it's in development flexibility or optimising user experience. That's where Sites has made a real difference for ECMWF. We now have the freedom to create custom websites and experimental pages that break free from the constraints of our traditional platforms. It's given us a way to innovate without worrying about disrupting our core systems or forcing those platforms to do things they weren't designed for. One of the standout benefits is the ability to create solutions to share content seamlessly between platforms. This has allowed us to design user interfaces that truly cater to our audience's needs, instead of bending to the limitations of the platforms we use. Overall, ECMWF Sites has been a game-changer, offering the flexibility and control we've been looking for."

The *ECMWF Sites* platform supports many other public-facing web portals, such as *https://pulse.climate.copernicus.eu*, which uses data and charts constantly updated through the API. Furthermore, it is used to distribute monthly communication bulletins through media channels. As Julien Nicolas, a climate scientist at ECMWF, describes:

"Since 2023, ECMWF Sites has emerged as an indispensable tool for the C3S Climate Intelligence (CI) team, serving two key functions: file sharing with users external to ECMWF and back-end storage for web applications. For its monthly Climate Bulletins, the CI team relies on Sites to share a variety of graphics and data files. These are initially shared with web content collaborators (pre-publication) and later with journalists (post-publication). Additionally, each report includes embargoed content, for which access is securely restricted via password protection. ECMWF Sites has significantly streamlined this entire process, effectively replacing Dropbox, which was previously in use. The second major application of sites within the CI team is file storage for web applications. A notable example is Climate Pulse, where all maps, CSV, and JSON files are stored and updated daily using the sites API, driven by an ecFlow suite. This ensures that data is kept up to date seamlessly. The CI team's overall experience with ECMWF Sites has been excellent. Its combination of a file browser and API offers great flexibility, supporting both manual and automated workflows. Based on our feedback, several enhancements were introduced, making the platform even more user friendly."

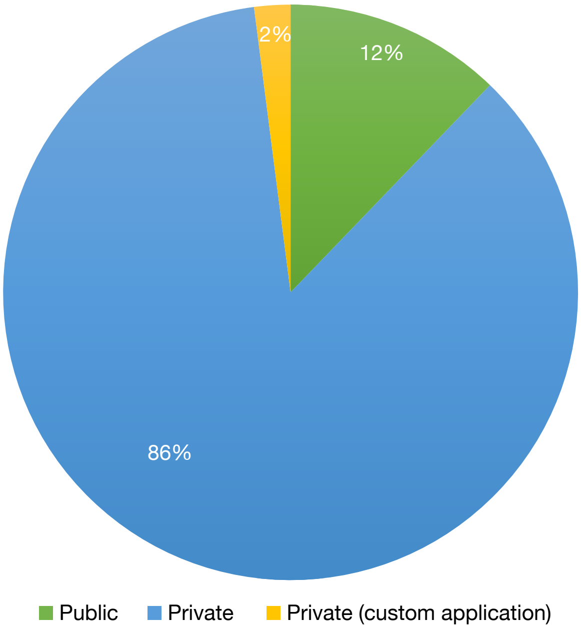

FIGURE 3 This chart shows the distribution of websites per type (private or public) at the end of 2024.

Custom Applications is a feature introduced in early 2022 and allows users to run a Docker container as part of their website. This is useful when users want to run some processing as part of their website, or access internal data from other systems. As Martin Janousek, an analyst at ECMWF, describes:

> *"I very much welcomed the implementation of ECMWF Sites as I had been calling for some sort of internal HTTP server to present various raw HTML documents since the early 2010s. Before ECMWF Sites, sharing of HTML documents was rather cumbersome. For example, one important product of every new model cycle evaluation, the scorecard, made as an HTML file with included JavaScript objects, was distributed either as an email attachment or put to a shared directory. That was rather impractical as staff started to work increasingly remotely. With ECMWF Sites a scorecard is shared as just a URL, conveniently accessible in users' local browsers. ECMWF Sites became even more useful when it offered an option to create dynamic applications. Step-by-step instructions and support from the team on how to build, test and deploy a docker container to ECMWF Sites were very helpful. It is also very important that, although website content can be conveniently accessed from anywhere, editing is managed by ECMWF Sites, limiting it to internal staff. I consider ECMWF Sites to be a key asset of ECMWF's IT systems, and I use it as a go-to solution for the implementation of future applications, in particular for data visualisation and access."*

A wide range of websites is available and can be accessed at *https://sites.ecmwf.int/hub/list/all/*.

Architecture

This section details the architecture of the main components of the system. With reference to the schematics shown in Figure 4, the system is divided into four main components: *Hub*, *Operator*, *Site*, and *Jobs*.

The *Hub* component allows users to interact with the system through a web interface or a REST API. It is essentially an enhanced proxy to the Kubernetes API server. It aggregates all the websites from all the users and allows users not only to explore the available websites, but to manage their own websites. Users can also access all sorts of information, such as configurations, server logs and web analytics. The *Hub* is open to all users with an ECMWF account, but only ECMWF staff and visiting scientists are configured to create websites. The *Sites Hub* is accessible through:

- Web UI (*https://sites.ecmwf.int/hub/*)
- REST API (*https://sites.ecmwf.int/hub/api/v1/spec/*).

The *Operator* component is the most important of the four. This component follows the operator pattern, which is a design approach that extends the functionality of Kubernetes. It does so by automating the management of complex applications, acting as a custom controller encapsulating the operational knowledge needed to deploy, manage, and scale the applications. Specifically, this component is responsible for ensuring the desired state of each individual website. A CRD (Custom Resource Definition) in Kubernetes is a way to extend the Kubernetes API by defining custom resources. The *Sites Operator* takes these *Site* CRD events and makes sure the necessary components composing a website are aligned with the CRD configuration. These resources allow the management of new types of objects, specific to an application or infrastructure needs, just like the built-in Kubernetes objects (like Pods, Services, etc.). Figure 5 shows an example of a *Site* CRD object.

The *Site* component is essentially a set of built-in Kubernetes objects that, when composed, will host and expose an individual website. These objects are specified by an individual *Site* CRD, and its lifecycle is managed by the *Sites Operator*. The Kubernetes objects that compose a website are:

1. One *StatefulSet* running four container applications, or five when users specify a custom application, where each of these applications is responsible for a specific task:

    a. Oauth2 Proxy – responsible for the authentication part and integrated with ECMWF's authentication system. Will ensure

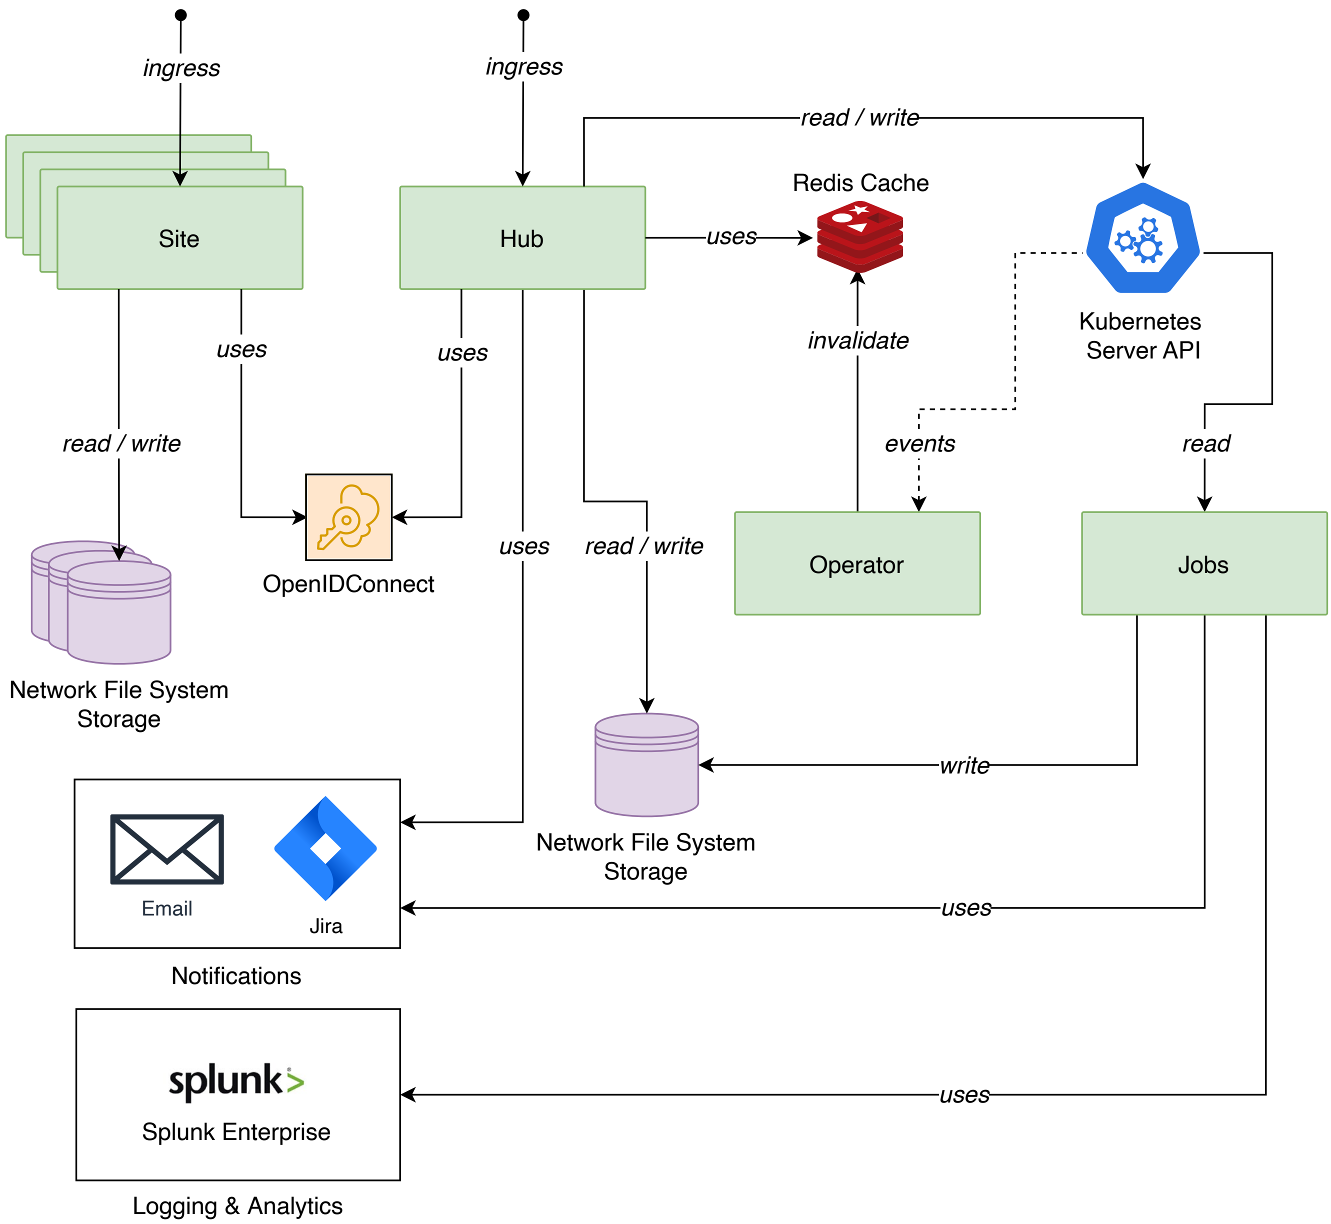

FIGURE 4 This diagram shows the high-level architecture of the platform with all the interactions between internal and external components.

users are logged in on private websites.

b. Nginx Proxy (OpenResty) – responsible for proxying requests accordingly and serving files as part of the web server. It is responsible for the authorisation as well, ensuring only configured users can access private websites.

c. Admin File Browser – this is the administration File Browser, which allows administrators to manage website contents.

d. Admin REST API – this is the administration REST API, which also allows administrators to manage website contents.

e. Custom Application – if users specified their own application to run as part of the website, then this will be served as the main application. It will not replace the Nginx Proxy. Instead, it will be proxied by the Nginx Proxy.

2. Two *ConfigMaps*, one with the global website configurations as environment variables and another with configuration files.

3. One Secret containing sensitive information to be used for authentication and authorisation purposes.

4. One *PersistentVolumeClaim* with the storage amount requested.

5. One Ingress to access the website through the main domain on a specific path: https://sites.ecmwf.int/*space/name/*

The *Jobs* component is composed by a set of packages to execute global administrative tasks such as: configurations backup, web analytics pre-loading, website storage quota, and retention date expiration notifications. These are essentially *CronJobs* running on intervals on the system.

ECMWF Sites is integrated with a few other systems to provide some of its functionality.

- All the components that are accessible to users, *Site* and *Hub*, use OpenIDConnect for authentication and authorisation.
- Network File System (NFS) volumes from TrueNAS are used as storage, and each individual website has its own volume.
- Redis is used for caching the *Hub* data, allowing users to interact with the system without overloading the Kubernetes Server API.
- Splunk is used for ingestion and extraction of log information, which is then used to produce web analytics.
- Jira and Email are used for notifications, such as when a website retention date expires, or a website storage is almost full, or to request access to a website, etc.

Security considerations

In order to ensure a high degree of compliance with security best practices, *ECMWF Sites* implements various strategies to minimise the attack surface of the platform and of other internal systems, and to prevent leakage of information.

In terms of infrastructure, the platform runs on a Kubernetes cluster which is deployed on Virtual Infrastructure within a specific network security zone. From within this security zone, it is only possible to access a set of standard ports running on other specific security zones, where other forecasting services run. Even though our internal Network Architecture has a high degree of segmentation and segregation, a set of *NetworkPolicies* further enforces internal traffic rules, blocking all traffic connections from and into each website, other namespaces and components within the Kubernetes cluster. This provides internal isolation between websites, and between other services running within the same Kubernetes cluster.

The Kubernetes API server is only accessible internally, from within a set of VLANs, within the ECMWF network. A *ServiceAccount* with a token is created for

```yaml
apiVersion: operators.ecmwf.int/v1
kind: Site
metadata:
  labels:
    app.kubernetes.io/name: site-symm-test
    sites.ecmwf.int/app: site
    sites.ecmwf.int/id: vx8vodk0flr1a43i
    sites.ecmwf.int/name: test
    sites.ecmwf.int/owner: symm
    sites.ecmwf.int/space: symm
  name: vx8vodk0flr1a43i
spec:
  configuration:
    applications:
      site-admin:
        image: eccr.ecmwf.int/sites/site-admin:2.0.0
      site-api:
        image: eccr.ecmwf.int/sites/site-api:2.0.0
      site-custom:
        extra_env:
          - SOME_VAR=VAL
        image: eccr.ecmwf.int/my-project/custom-app:1.0.0
      site-oauth2:
        image: eccr.ecmwf.int/sites/site-oauth2:2.0.0
      site-proxy:
        image: eccr.ecmwf.int/sites/site-proxy:2.0.0
    enabled: true
    headers:
      cors:
        enabled: true
        extra_domains:
          - site.copernicus.eu
      no_cache:
        enabled: true
        file_extensions:
          - js
      no_robots:
        enabled: false
    monitored: false
    quota: 1
    resource_limits: normal
    retention_date: 2025-08-20
    share:
      access_allowed_groups:
        - ecmwf_staff
      access_allowed_users: []
      administrator_groups: []
      administrator_users:
        - symm
      authentication_token: 0MQ4ODQzOTVhM2RjZDgzMw==
    space: symm
    twofa_enabled: false
    webdav_enabled: false
  created_by: symm
  creation_date: 2024-08-20
  description: This is an example website.
  id: vx8vodk0flr1a43i
  name: test
  owner: symm
  type: private
  updated_by: symm
  updated_date: 2024-08-21
status: {}
```

FIGURE 5 A sample Site CRD object definition.

the internal communication between the *Hub* and *Operator* with the Kubernetes cluster server API. Role Based Access Control (RBAC) is then configured for this *ServiceAccount*, allowing access to a few cluster-wide resources for the management of the system. The *ServiceAccount* token is not mounted on individual websites, meaning custom applications will not be able to use it to exploit the system. It is worth noting that the aforementioned *NetworkPolicies* would block this traffic anyway.

In terms of authentication and authorisation, all endpoints use OpenIDConnect for authentication, relying on the JWT *access_token*, using the *preferred_username* and *entitlements* for authorisation.
A 64 characters Hex token is randomly generated for each website. This token is used as an administration token and allows anyone holding it to manage all the contents of the website. It is useful for automated pipelines, and in case tokens are compromised they can be revoked by generating new ones. Only one token is usable at any time for a specific website, until it is revoked.

CPU and Memory are limited by a *ResourceQuota* object. This ensures a global limit for resources that the platform can take from the Kubernetes cluster. Every website is configured with a sensible number of resources that can be used. These are configurable per website within three resource levels, defined as Normal, Medium, and High. This is very relevant for websites running custom applications, since these run code out of our control; we can limit and protect the system from misbehaving containers:

- Normal – this is ideal for most use cases and is the default setting for all websites.
- Medium – this is ideal for websites that are integrated with other web services and require a bit more throughput.
- High – this is for high load traffic websites, including content management using the web file browser or REST API that might require additional memory or CPU.

Custom applications can only be configured on private websites. Custom application containers run in unprivileged mode and with a specific user ID and group ID. This means that users cannot configure Docker images where processes run with the root user, thereby decreasing the risk if using a compromised Docker image.

A set of *PriorityClasses* defines three priority levels: Normal, Medium, and High. These ensure that critical workloads receive the resources they need even in times of contention, maintaining the stability and reliability of essential websites and core components:

- Normal – by default, all websites get this priority class.
- Medium – websites that are monitored, usually public websites, get this priority instead of the low priority class.
- High – all the platform components, i.e.: Hub, Operator, Redis and all Jobs, get this priority class.

Websites can only be configured as public by *Hub* administrators, ensuring users need to go through an approval process. This decreases the risk of sensitive content exposure or non-compliant ECMWF design. Private websites can be configured with two-factor authentication, ensuring that users must enter a temporary one-time token (TOTP) before accessing the website, providing an extra level of security. Each website has a renewable retention date of one year from creation.

Ecosystem

ECMWF Sites offers a Command Line Interface (CLI) and a Software Development Kit (SDK).

These tools significantly enhance the platform ecosystem by providing users with powerful, flexible, and efficient ways to interact with the platform. They take full advantage of the *Hub* and the websites' REST APIs. Both the CLI and the SDK allow users to automate tasks, manage resources, and perform operations directly from the command line, streamlining workflows and improving productivity. Together, these tools empower users to fully leverage the platform's features, driving greater adoption and facilitating its usage.

Performance and load testing

To illustrate the performance of the system, a set of test scenarios was created. These tests consist in measuring the performance of a website in terms of content management, by uploading a set of static HTML and plot data of around 1.5 GB, and then accessing the website and navigating to a second location within that same website. The website is configured as Public and with Resource Limits set to High, in what we consider the ideal condition. These tests were executed from the three locations where ECMWF is based, Bologna (Italy), Bonn (Germany), and Reading (UK). They provide a good insight into the overall performance of the service across Europe.

The results of the test runs can be seen in Table 1 and Table 2, while Figure 6 and Figure 7 show the consumption of resources throughout the duration of these tests (circa 1 hour and 30 minutes per run). These results show that the website performs at the highest level under the tested load. In these performance tests, while the website is configured with High resources, CPU is not limited, hence leading to the highest

		Reference (Local)	Italy (Bologna)	Germany (Bonn)	UK (Reading)
50 Users	Apdex (Application Performance Index)	0.998	1,000	0.885	0.994
	90% (milliseconds)	84	99	892	237
	Throughput (requests/second)	73	74	69	76
25 Users	Apdex (Application Performance Index)	1,000	1,000	0.979	0.997
	90% (milliseconds)	32	77	275	177
	Throughput (requests/second)	76	83	57	64
10 Users	Apdex (Application Performance Index)	1,000	1,000	0.999	0.997
	90% (milliseconds)	13	73	96	157
	Throughput (requests/second)	65	51	40	32

TABLE 1 This table shows performance test results for the websites' content management under different load and different locations. Website content management performs very well, with 90% of the requests being served in 237 ms or less in the UK and in 892 ms or less in Germany. Content management is normally performed by a single user at a time, so these results show great REST API performance at scale.

		Reference (Local)	Italy (Bologna)	Germany (Bonn)	UK (Reading)
250 Users	Apdex (Application Performance Index)	1,000	1,000	0.999	1,000
	90% (milliseconds)	10	36	55	138
	Throughput (requests/second)	190	190	190	190
110 Users	Apdex (Application Performance Index)	1,000	1,000	0.999	1,000
	90% (milliseconds)	10	46	57	138
	Throughput (requests/second)	190	190	190	190
50 Users	Apdex (Application Performance Index)	1,000	1,000	0.999	1,000
	90% (milliseconds)	10	32	62	137
	Throughput (requests/second)	50	50	50	50

TABLE 2 This table shows the performance test results for website access under different load and different locations. Website access performs very well, with 90% of the requests being served in 138 ms or less in the UK and in 62 ms or less in Germany.

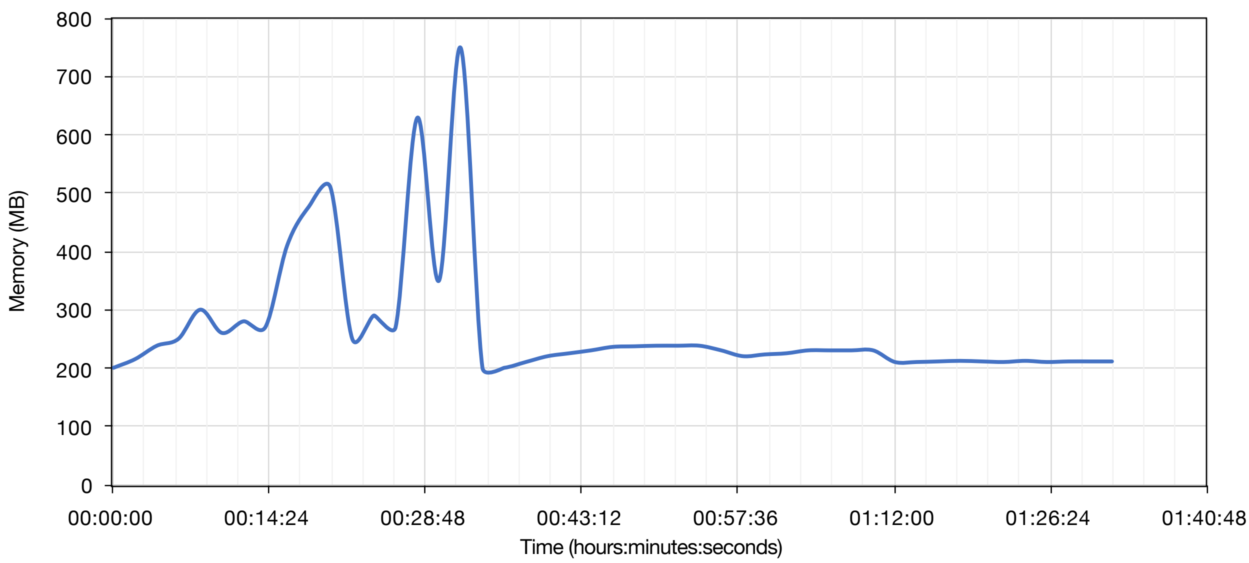

FIGURE 6 This figure shows the total memory usage throughout the performance tests. Memory peaks at 750 MB while performing content management, while during access to the website the peak reduces to only 238 MB.

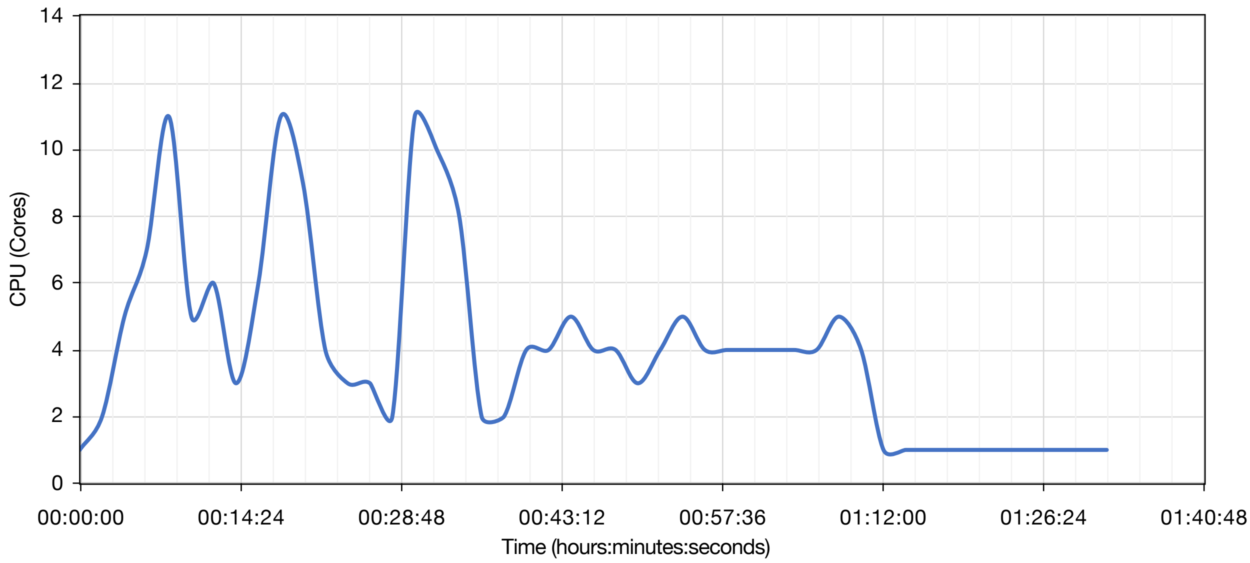

FIGURE 7 This figure shows the total CPU (Cores) usage throughout the performance tests. CPU usage peaks at 11 CPU Cores while performing content management, while during access to the website the peak reduces to 5 CPU Cores.

throughput. Memory consumption is very low for these operations, due to a great optimisation of the REST API. Usually, high memory consumption comes from the use of either custom applications or the web file browser open-source application: *https://github.com/filebrowser/filebrowser*.

These and other performance tests can be found in detail at *https://sites.ecmwf.int/performance/results/*.

Concluding remarks

In the past months, several developments have been made to improve the way the main components of the platform are distributed. The idea behind these changes is to open-source the project, decoupling the main components from ECMWF's specifics, and to make it even more configurable. Such a platform offers a robust system that empowers users to create, manage, and customise online spaces tailored to a diverse range of use cases and can be highly advantageous for other organisations or national meteorological services.

Several improvements on the Site REST API application were made while migrating from Python to Golang. This led to big improvements in throughput, while significantly reducing the memory consumption compared to the previous implementation. The latest version of the Site REST API supports streaming of large files, enabling *ECMWF Sites* to be used for sharing large datasets over HTTP. Looking ahead, the continued development of *ECMWF Sites* will focus on meeting evolving user needs, strengthening security measures, and furthering its potential as an open-source solution, ensuring it remains a reliable and adaptable platform for diverse use cases.

It is important to note that *ECMWF Sites* was introduced during the preparation for our data centre migration in 2022, providing a straightforward way to host many services without the need for custom integrations or involvement from other teams. This alleviated some of the pressure on the web development team within ECMWF's Forecasts and Services Department.

Further reading

Martins M., 2020: Introducing Sites: static websites as a service. *ECMWF Newsletter* **No. 164**, 17. *https://www.ecmwf.int/en/newsletter/164/news/introducing-sites-static-websites-service*

Varela D. & **Martins M.**, 2018: Re-Architecture of the Atlassian Collaboration Tools. *ECMWF Newsletter* **No. 157**, 17. *https://www.ecmwf.int/en/newsletter/157/news/re-architecture-atlassian-collaboration-tools*

Kubernetes, 2024: Concepts, *https://kubernetes.io/docs/concepts/*, accessed on 20 August 2024.

Kubernetes, 2024: Operator Pattern, *https://kubernetes.io/docs/concepts/extend-kubernetes/operator/*, accessed on 20 August 2024.

Kubernetes, 2024: Custom Resources, *https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/*, accessed on 20 August 2024.

ECMWF Confluence, 2024: User Documentation, ECMWF Sites - Websites as a Service, *https://confluence.ecmwf.int/display/UDOC/ECMWF+Sites+-+Websites+as+a+Service*, accessed on 20 August 2024.

Apdex (Application Performance Index), 2024: The Apdex Users Group, *https://ww.apdex.org*, accessed on 12 September 2024.

ECMWF Council and its committees

The following provides some information about the responsibilities of the ECMWF Council and its committees. More details can be found at:
http://www.ecmwf.int/en/about/who-we-are/governance

Council

The Council adopts measures to implement the ECMWF Convention; the responsibilities include admission of new members, authorising the Director-General to negotiate and conclude co-operation agreements, and adopting the annual budget, the scale of financial contributions of the Member States, the Financial Regulations and the Staff Regulations, the long-term strategy and the programme of activities of the Centre.

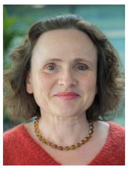

President Prof. Penny Endersby *(UK)*

Vice President Dr Roar Skålin *(Norway)*

Policy Advisory Committee (PAC)

The PAC provides the Council with opinions and recommendations on any matters concerning ECMWF policy submitted to it by the Council, especially those arising out of the four-year programme of activities and the long-term strategy.

Chair Ms Virginie Schwarz *(France)*

Vice Chair Prof. Dr Maarten van Aalst *(Netherlands)*

Finance Committee (FC)

The FC provides the Council with opinions and recommendations on all administrative and financial matters submitted to the Council and exercises the financial powers delegated to it by the Council.

Chair Mr Lukas Schumacher *(Switzerland)*

Vice Chair Ricardo José Squella de la Torre *(Spain)*

Scientific Advisory Committee (SAC)

The SAC provides the Council with opinions and recommendations on the draft programme of activities of the Centre drawn up by the Director-General and on any other matters submitted to it by the Council.
The 12 members of the SAC are appointed in their personal capacity and are selected from among the scientists of the Member States.

Chair Prof. Dr Thomas Jung *(Germany)*

Vice Chair Dr Susanna Corti *(Italy)*

Technical Advisory Committee (TAC)

The TAC provides the Council with advice on the technical and operational aspects of the Centre including the communications network, computer system, operational activities directly affecting Member States, and technical aspects of the four-year programme of activities.

Chair Ms Anne-Cecilie Riiser *(Norway)*

Vice Chair Mr Antonio Vocino *(Italy)*

Advisory Committee for Data Policy (ACDP)

The ACDP provides the Council with opinions and recommendations on matters concerning ECMWF Data Policy and its implementation.

Chair Ms Monika Köhler *(Austria)*

Vice Chair Anne Kristensen *(Norway)*

Advisory Committee of Co-operating States (ACCS)

The ACCS provides the Council with opinions and recommendations on the programme of activities of the Centre, and on any matter submitted to it by the Council.

Chair Mr Ilian Gospodinov *(Bulgaria)*

ECMWF publications

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

Technical Memoranda

924 **Pillosu**, **F.-M.**, **T. Hewson**, **E. Stephens**, **C. Prudhomme** & **H. Cloke**: Does a multivariate approach enhance univariate grid-to-point post-processed rainfall forecasts? A comparative analysis. *January 2025*

920 **Lopez**, **P.**: Solar Eclipses in the IFS. *January 2025*

925 **Ma**, **Z.**, **N. Bormann**, **K. Lean**, **D. Duncan**, **E. Berbery** & **S. Kalluri**: Forecast impact assessment of a potential ATMS instrument in the early-morning orbit using the EDA method. *December 2024*

923 **Molteni**, **F.** & **A. Brookshaw**: Indian Ocean teleconnections to the northern extra-tropics. *December 2024*

922 **Molteni**, **F.** & **A. Brookshaw**: Multi-decadal variability of tropical rainfall. *December 2024*

ESA or EUMETSAT Contract Reports

Salonen, **K.**, **P. Weston** & **P. de Rosnay**: Annual SMOS brightness temperature monitoring report - 2023/24. *January 2025*

Garroussi, **S.** & **J. McNorton**: D1- Towards enhanced fire fuel estimation with satellite-derived predictive models. *January 2025*

ECMWF Calendar 2025

Feb 3–6	Training course: Use and interpretation of ECMWF products
Mar 3–7	Training course: Parametrization of subgrid physical processes
Mar 17–21	Training course: Data assimilation & machine learning
Mar 24–28	Training course: EUMETSAT/ECMWF NWP-SAF satellite data assimilation
Apr 7–11	**ECMWF's 50th anniversary events in Bonn, Germany**
	Apr 7–11 Annual Seminar 2025
	Apr 8–9 The evolution of Copernicus Services at ECMWF: challenges and opportunities
	Apr 9–10 Workshop on ancillary data for land surface and Earth system modelling
	Apr 9–10 Workshop on surface process coupling and its interactions with the atmosphere
	Apr 9–10 Workshop on data assimilation: initial conditions and beyond
Apr 28–29	Finance Committee
Apr 30	Policy Advisory Committee (virtual)
Jun 2–4	8th C3S General Assembly
Jul 3–4	Council (virtual)
Sep 15–17	9th CAMS General Assembly
Sep 15–19	**ECMWF's 50th anniversary events in Bologna, Italy**
	Sep 15–18 Using ECMWF's Forecasts (UEF2025)
	Sep 15–19 21st workshop on high-performance computing in meteorology
	Sep 16–17 ECMWF DestinE Annual Meeting
	Sep 16–17 Code for Earth 2025
Oct 6–8	Scientific Advisory Committee
Oct 20–21	Technical Advisory Committee (virtual)
Oct 27	Policy Advisory Committee (virtual)
Oct 28	Finance Committee (virtual)
Dec 4–5	Council

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. 182** | Winter 2024/25
European Centre for Medium-Range Weather Forecasts
www.ecmwf.int