ECMWF

European Centre for Medium-Range Weather Forecasts

# Newsletter

**No. 170** | Winter 2021/22



**Publication policy**

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

The ECMWF Newsletter is not peer-reviewed.

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

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

# The role of computing

As I write this, work is under way on several fronts that underlines the importance of computing for ECMWF. In Bologna, Italy, we are beginning to take our new high-performance computing facility (HPCF) into operations; the European Weather Cloud, ECMWF's and EUMETSAT's cloud infrastructure that can be used to run workflows close to the latest data, has entered its pre-operational phase; we are one of the organisations implementing WEkEO, the European Commission's initiative to provide a single access point to all Copernicus data and information; and our EU-funded Copernicus Climate Change Service (C3S) and Copernicus Atmosphere Monitoring Service (CAMS) have developed their own databases, the Climate Data Store (CDS) and the Atmosphere Data Store (ADS), to access and process the data made available through those services. Many more activities are under way at ECMWF to make optimal use of computing resources. All of this underlines the central role of computing in everything that we do, including of course our core mission, numerical weather prediction (NWP): the two have always gone hand in hand. To be blunt, NWP would not be possible in today's shape if the last few decades had not seen huge advances in supercomputing.

Several articles in this Newsletter highlight the importance of computing. One is on our new data centre in Bologna. Work on the data centre is now complete, and testing of the new HPCF has begun. The new electrical and mechanical systems of the data centre have been designed to ensure resilience and reliability. Redundancy ensures that in the event of failures or maintenance activities, the power supply and cooling to the computer rooms is guaranteed. The data centre forms part of a wider Bologna Tecnopolo, which will host several new supercomputers and gives ECMWF enough flexibility to meet future expansion requirements. Another article explores the Aviso service for notifications of the availability of ECMWF data on the European Weather Cloud, including three workflows that have already been implemented by Member State users. A third describes new data exploration functions in ECMWF's Metview software.

Other articles highlight the collaborative nature of much of our work. They include an overview of our involvement in the South-East European Multi-Hazard Early Warning Advisory System (SEE-MHEWS-A) project, which aims to strengthen the existing early warning capacity in the region. There is also an article on ten years of OpenIFS at ECMWF, an initiative to provide and support versions of the IFS for research, education and training by Member and Co-operating States and academic institutes. What they make clear is that much of ECMWF's strength is derived from our active participation in the wider Earth system science communities and from our openness to new developments in that area.

**Florence Rabier**
Director-General

## Contents



**Editor** Georg Lentze • **Typesetting & Graphics** Anabel Bowen • **Cover** ECMWF's new data centre in Bologna, Italy

# Heavy snowfall in Denmark and Sweden at the start of December

Linus Magnusson, Timothy Hewson, Estíbaliz Gascón

The end of November and the beginning of December 2021 were much colder than normal in northern Europe. During the period, a sequence of cyclones brought challenging weather conditions for several countries. Over the British Isles, storm Arwen caused strong winds and snowfall, leading to long power outages locally. Here, we will focus on a cyclone that passed Denmark and southern Sweden on 1 December and brought up to 30 cm of snow in northern Denmark and inland parts of southern Sweden (see the snow depth analysis and observations).

## Precipitation forecasts

In terms of medium-range predictions, ensemble members started to pick up the risk of snowfall on 26 November, five days before the event. The signal became gradually stronger in forecasts from 28 November onwards, as can be seen in the Extreme Forecast Index (EFI) and Shift of Tails (SOT) figure. However, for northern Denmark the EFI for snowfall dropped in the forecast from 30 November before re-gaining strength again just before the event. As this dip occurred, the SOT forecast maintained relatively high levels (>2), indicating that the possibility of an extreme event was still very much there. This case can serve as a reminder to forecasters to not extrapolate trends in consecutive forecasts.

ECMWF provides various precipitation-type products based on the model output. Here we show the 24-hour high-resolution forecast (HRES) valid on 2 December at 00 UTC together with five-day precipitation-type meteograms for Aalborg in northern Denmark and Lund in southern Sweden, all starting on 1 December at 00 UTC. The products show, respectively, the diagnosed precipitation type in HRES and probabilities of each precipitation type from ENS, distinguishing in each case between rain, sleet, wet snow, snow, ice pellets and freezing rain. The ensemble product also includes three different intensity categories. During the cyclone passage, precipitation mostly fell as snow on the northern side, as was the case for Aalborg, where the forecast indicated wet snow for most of the time. On the southern side, the precipitation started as snow, then turned into rain in a warm sector of the cyclone, but then turned back to snow again later as the cyclone moved on.

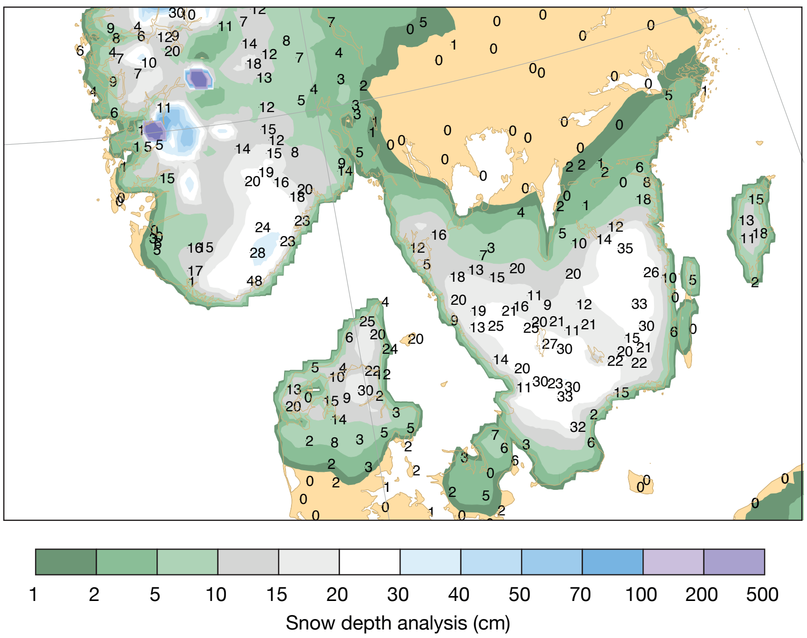


**Snow depth analysis and observations.** High-resolution (HRES) snow depth analysis (shading) and observations (numbers) for 2 December 06 UTC.

## Snow on the ground

A critical question for the impact is how much snow will accumulate on the ground. The complex evolution of the precipitation type going from snow to sleet to rain and then back to snow again makes the modelling of snow accumulation on the ground challenging. The ECMWF model has a known issue with the accumulation of snow on the ground when the precipitation type falls as sleet: the fraction of snow included in the sleet tends to erroneously accumulate on the ground in the model. A second known issue that makes this type of situation worse is that, once there is any snow on the ground in the model, it tends to melt too slowly. This problem is most acute when the total diagnosed depth is small. These issues are documented for users as items S2 and S3 here: https://confluence.ecmwf.int/display/FCST/Known+IFS+forecasting+issues. Both came to the fore in the forecast for Lund from 1 December 00 UTC. The model accumulated 5 cm of snow (10 mm of water equivalent), which then reduced to 4 cm only five days into the forecast. In reality, 2 cm was measured on 2 December at 06 UTC, and that all melted later that day. Such errors in accumulated snow can impact the 2-metre temperature further into the forecast range, especially during clear-sky situations leading to strong cooling over snow surfaces.

Of course, the prediction of accumulated snow on the ground is critical for users. How well did the short-range forecast do for predicting the precipitation and accumulated snow in Aalborg? The total 24-hour precipitation forecast for Aalborg for 1 December in the shortest-range

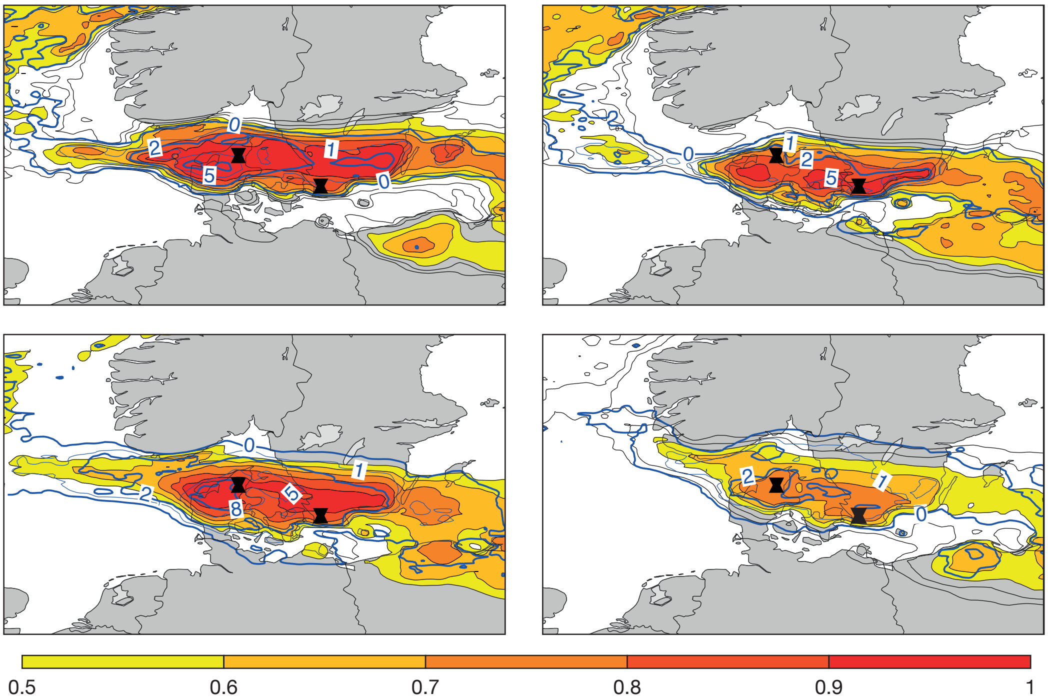


**Extreme Forecast Index and Shift of Tails.** The plots show the EFI (shading) and SOT (blue contours) for snowfall on 1 December 2021 in forecasts from 1 December (top-left), 30 November (top-right), 29 November (bottom-left) and 28 November (bottom-right), all at 00 UTC. The hourglass symbols indicate the location of Aalborg in Denmark and Lund in Sweden.

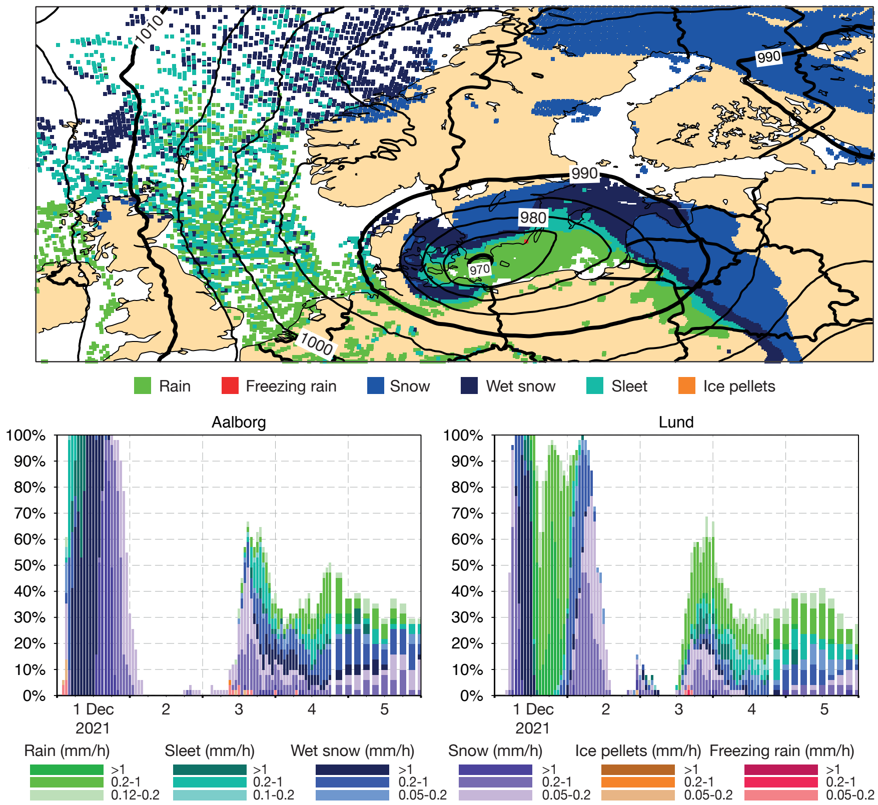


**Forecasts of precipitation type from 1 December 00 UTC.** The top image shows the HRES forecast valid on 2 December 00 UTC, including mean sea-level pressure, while the bottom images show ensemble forecast products for the probability of precipitation type in Aalborg (left) and Lund (right).

forecast (from 00 UTC that day) was 32 mm, of which 30 mm was predicted to be in snow form. The resulting snow depth is then a function of the accumulated snow water equivalent and the modelled snow density. The latter adds another dimension of complexity in the modelling. The modelled snow depth increase during the period was 16 cm, indicating a relatively high snow density for fresh snow. Meanwhile, nearby observations reported 17 mm of precipitation in 24 hours, and a snow depth increase of 20 cm. Thus, whilst the predicted total precipitation was almost twice the observed value, the predicted increase in snow depth seems to have been lower than observed. This discrepancy could be due to either overprediction of the snowfall accompanied by an overestimation of the snow density, or an under-catchment of snow in the rain gauge, resulting in a precipitation observation value that is too small. The latter is a common observation problem when heavy snowfall is accompanied by strong winds.

Cycle 48r1 of ECMWF's Integrated Forecasting System (IFS) is to include a major upgrade of the snow scheme in the land model. This will include improvements to the thermodynamics, which are expected to bring a number of benefits such as better melting of low snow amounts, although further work is needed to solve all issues related to snow accumulations.

# Metop-A satellite retired after 15 years of huge benefit to forecasting

Mohamed Dahoui, Tony McNally, Niels Bormann, David Duncan

In November 2021, EUMETSAT decommissioned the Metop-A satellite, which has brought tremendous benefits to numerical weather prediction (NWP) since its launch about 15 years ago. Metop-A was Europe's first operational polar-orbiting satellite and the first in a series of three satellites of the EUMETSAT Polar System (EPS) programme covering the mid-morning orbit. It was packed with a suite of novel and advanced instruments, such as IASI, ASCAT, GRAS and GOME-2, as well as valuable heritage sensors, such as AMSU-A, MHS, HIRS and AVHRR. With the combined use of all these instruments, Metop-A was arguably the most influential meteorological satellite of its generation. Follow-on Metop satellites (Metop-B launched in 2012 and Metop-C in 2018) offered extra benefits, useful resilience and new products based on the dual exploitation of satellite pairs.

### Metop instruments

| | | | |
|---|---|---|---|
| **IASI** | Infrared Atmospheric Sounding Interferometer | **AMSU-A** | Advanced Microwave Sounding Unit–A |
| **ASCAT** | Advanced Scatterometer | **MHS** | Microwave Humidity Sounder |
| **GRAS** | GNSS (Global Navigation Satellite System) Receiver for Atmospheric Sounding | **HIRS** | High-resolution Infrared Radiation Sounder |
| **GOME-2** | Global Ozone Monitoring Experiment–2 | **AVHRR** | Advanced Very High Resolution Radiometer |

## Improved forecasts

The combined use of data from the three Metop satellites represented nearly 45% of observations used at ECMWF and accounted for around 27% of the total impact of all observations assimilated by the Centre, as estimated by the Forecast Sensitivity to Observation Impact (FSOI) measure (see FSOI/data counts image). Having three Metop satellites provided additional impact and a healthy resilience in the observing system against instrument failure: it shielded the assimilation from intermittent outages affecting individual satellites or instruments. Two plots from 2013 show the big impact Metop-A and Metop-B had in terms of error reduction (see 'Impact by hemisphere' image).

Data from Metop helps to constrain many aspects of the atmosphere in the forecasting model: the vertical distribution of temperature, humidity and indirectly the wind (IASI, AMSU-A, MHS and GRAS); surface winds (ASCAT); surface moisture over land (ASCAT); atmospheric winds (AVHRR); and ozone (GOME-2). Derived products from IASI, GOME-2 and AVHRR are used to monitor and constrain many aspects of atmospheric composition.

## Novel instruments

Aside from heritage instruments such as the microwave sounders AMSU-A and MHS, Metop-A introduced a range of new instruments which significantly extended observing capabilities.

The **IASI** instrument was the first infrared interferometer sounder with high spectral resolution (8461 channels). It was designed to measure temperature and water vapour profiles in addition to the concentration of other infrared-absorbing constituents. Throughout its lifetime, the instrument has proven to be fit for operations with stable performance and a consistent

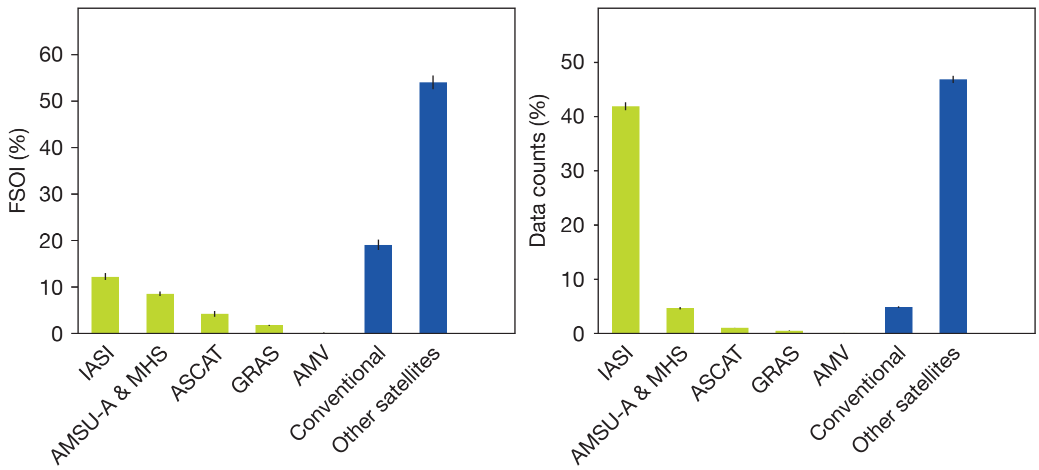


**FSOI and data counts in May 2020.** Contribution of Metop instruments (shown in green), conventional observations and the rest of the observing system to the reduction of forecast error as estimated by FSOI (left), and used data counts (right), with error bars.

positive impact. Due to the configuration of assimilation systems and available resources, only an evolving subset of the most informative IASI channels was used. In addition to using radiances for NWP, ECMWF uses several datasets derived from IASI for producing the Copernicus Atmosphere Monitoring Service (CAMS) global analyses, forecasts and reanalyses, including CO, $CO_2$ and $CH_4$ and, in test mode, ozone and $SO_2$.

**ASCAT** was a new-generation scatterometer instrument designed primarily to measure surface wind over oceans. The data are also useful to derive soil moisture over land. The instrument is equipped with two sets of three antennae allowing simultaneous measurements to be made from three directions to improve the wind direction ambiguity. Like IASI, ASCAT data started to be used not long after launch and has provided quasi uninterrupted service since then. For extended periods, ASCAT was the only instrument providing surface winds. Research is ongoing into the use of surface stress, which is closer to what the instrument measures but is not currently assimilated. Such work would allow ASCAT data to impact both the atmosphere and the ocean in the context of coupled data assimilation.

**GRAS** was the first operational European GNSS radio occultation instrument designed to measure temperature profiles and tropospheric humidity using refracted Global Positioning System signals. GRAS data are used by ECMWF in the form of bending angles and provide a consistent number of profiles. The consistent availability of GNSS-RO data helped fill large observation gaps in the stratosphere.

The **GOME-2** instrument is designed to provide information on total column and vertical profiles of ozone as well as other species, such as $NO_2$, $SO_2$ and aerosols – the latter in combination with information also from IASI and AVHRR (Polar Multi-sensor Aerosol optical Properties product, PMAP). Total column ozone is used in the ECMWF atmospheric analysis system. More products (ozone, aerosols and $SO_2$) are used by the CAMS analysis system.

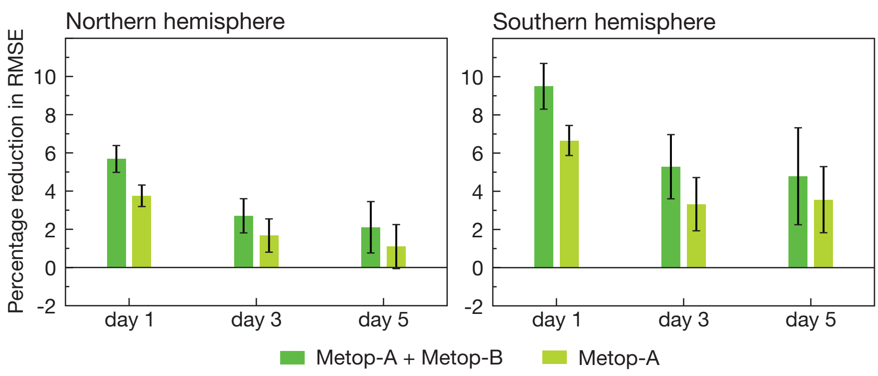


**Impact by hemisphere of adding two Metop satellites to the baseline system in 2013.** The plots show the percentage reduction in root-mean-square error (RMSE) of 500 hPa geopotential height at days 1, 3 and 5 in the northern and southern hemisphere extratropics as a result of assimilating just Metop-A and both Metop-A and Metop-B. The error bars are 95% confidence intervals and represent the statistical significance of the improvements with respect to the 'no Metop' baseline experiment.

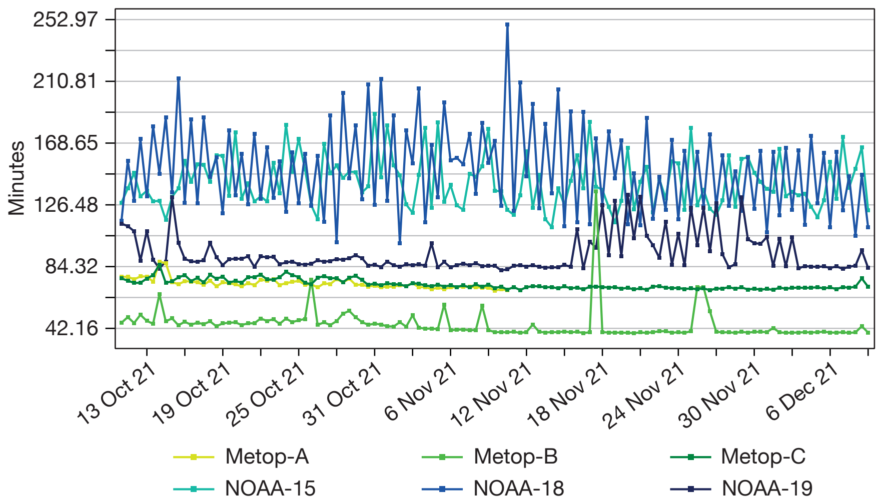


**Mean timeliness of AMSU-A data from various satellites.** The plot shows the mean timeliness of AMSU-A data from Metop compared to other satellites.

## High reliability

EUMETSAT provided operational support for Metop satellites, leading to very high availability as well as speedy communication regarding anomalies and maintenance operations. EUMETSAT's regional data service (EARS) and collaboration with the US National Oceanic and Atmospheric Administration (NOAA) to exploit the Antarctic dumping station at McMurdo allowed Metop data to be distributed to the user community in a timely manner, benefiting global and regional NWP models (see the AMSU-A data figure).

The dual and subsequently triple exploitation of Metop satellites was performed in a way to maximise the benefits for NWP. This was achieved by choosing complementary orbits and generating dual products, such as global Atmospheric Motion Vectors. EUMETSAT managed to extend the lifetime of the Metop-A satellite well beyond its life expectancy without noticeably compromising the performance of the instruments.

Metop-A was a very successful operational satellite that had a tremendous impact on NWP. Indeed, its success has been key to securing substantial investment to launch the second generation of Metop platforms in the mid-twenties. These will carry even more sophisticated technology to support operational NWP, but also to accelerate our understanding of Earth system science.

# Ten years of OpenIFS at ECMWF

Glenn Carver

The OpenIFS activity at ECMWF marked ten years in December 2021. It began in 2011, aiming to provide and support versions of the operational Integrated Forecasting System (IFS) for research, education and training in numerical weather prediction and meteorology by Member and Co-operating States and academic institutes.

## Motivation

The original motivation behind OpenIFS remains relevant today and is in line with ECMWF's strategy of close partnership with the meteorological community. A key objective was to deliver benefit to Member and Co-operating State users, the wider community and ECMWF. Over the last ten years, OpenIFS has: produced improvements in the IFS; established a growing user community amongst Member and Co-operating States; enabled new collaborations; increased research and training opportunities in universities; created an international scientific workshop series; and improved knowledge of the IFS and ECMWF in the wider meteorological community.

It was clear that ease-of-use, good support, and training were essential for external use of the IFS to be successful. These priorities have focused OpenIFS efforts to make the IFS, a complex operational numerical weather prediction model, easy to use externally to ECMWF. They have also led to significant engagement with the meteorological community to provide training and support. OpenIFS provides online user guides, public forums for users to ask questions, a dedicated support team at ECMWF, and training events.

## OpenIFS compared with the IFS

Updates to operational versions of the IFS occur at least yearly, but OpenIFS releases are less frequent, on a timescale typical of university research projects, and focus instead on producing releases to coincide with significant upgrades of the IFS.

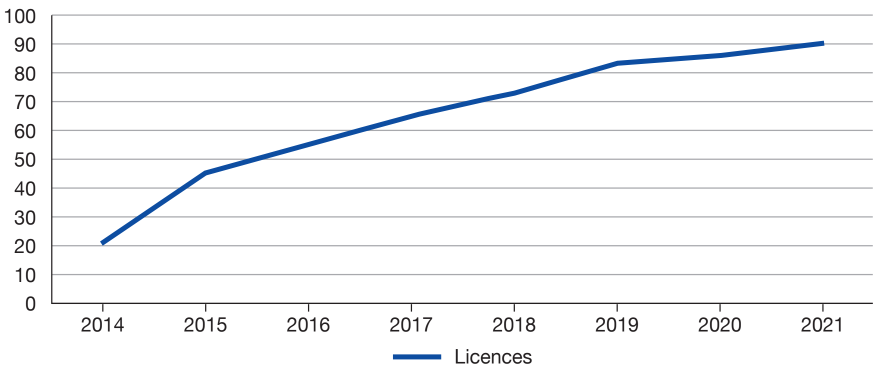


**Growth in the number of OpenIFS licensees in recent years**. Part of the increase is driven by the steady adoption of the OpenIFS model by EC-Earth consortium partners.

To date, three versions of the IFS have been made available at upgrade Cycles 38r1, 40r1 and 43r3. The OpenIFS model has the exact same deterministic and ensemble forecast capability as the operational IFS. It includes the wave and surface components but does not include the data assimilation nor the observation handling components of the IFS. This makes the model easier to provide externally and simpler to use. At the current time, the ocean model NEMO is also not included, although it is hoped that this may be possible in the future.

## Research and teaching applications

Overall, the number and range of research and teaching applications using OpenIFS over the last ten years are impressive. As expected, it has been used in a broad range of numerical weather prediction (NWP) studies, such as high-impact weather event and diagnostic studies. Many of these studies have been published in the peer-reviewed literature or as ECMWF Newsletter articles. OpenIFS has also been the focus of several Earth system modelling groups. The EC-Earth consortium (www.ec-earth.org) includes 30 research institutes from 12 European countries to collaborate on the development of an Earth system model. They have chosen the OpenIFS model as the atmospheric component of their next model release, whilst other research groups are working to build their own Earth system models based on OpenIFS (see ECMWF Newsletter No. 164).

OpenIFS has also played a key role in collaboration with the University of Oxford on the use of reduced precision in the IFS. This led to significant savings in computing resources at ECMWF for operational forecasting. OpenIFS was ideal for the proof-of-concept development as its reduced code base allowed rapid prototyping and testing.

An objective of OpenIFS is to contribute towards meeting education and training needs both at ECMWF and for licensed users. Whilst some use of OpenIFS in teaching was anticipated, the number of universities using OpenIFS in advanced-level teaching is more than originally expected. OpenIFS has clearly enabled a bridge between theoretical teaching and the practical experience of using a world-leading forecast model.

## OpenIFS user workshops

An OpenIFS workshop series was created to provide a forum for OpenIFS users to meet, to learn about OpenIFS, and to highlight ECMWF research. To date, there have been five such workshops (Helsinki, Stockholm, ECMWF, Trieste, Reading), which have rapidly grown from small local workshops to five-day international workshops with invited speakers. Participants use the OpenIFS model to explore interesting case studies on the

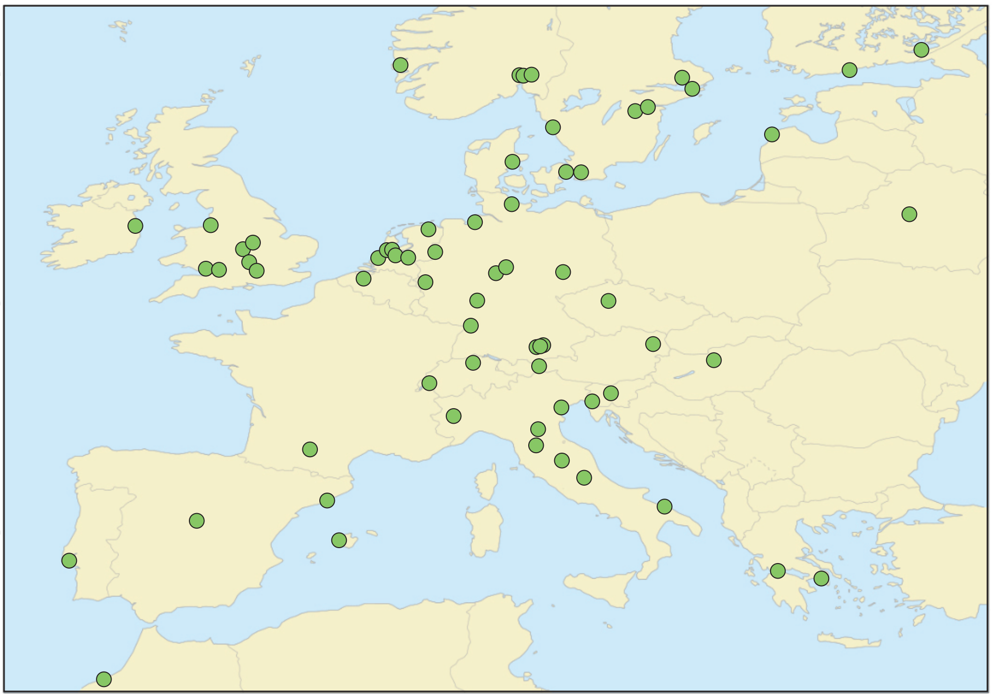

**Locations of OpenIFS licensees in Member and Co-operating States.** The chart shows the locations of most OpenIFS licensees. There are a few more licensees in other parts of the world, which are not shown here.

theme of the workshop. The case studies in the OpenIFS workshop series provide tutorials and practical examples where the model configuration is altered to examine its impact on the case study forecast. Some universities have adapted this material for their own use.

## Future

The biggest hurdle to using OpenIFS has been the generation of the model's initial conditions. This is best done at ECMWF, where the system 'knows' how best to produce balanced states from the various analyses available. A new web facility is under development, which will allow users to generate their own initial states rather than requesting them from the OpenIFS support team (see ECMWF Newsletter No. 168). Other plans include investigating the addition of the NEMO ocean model as used in ECMWF operational forecasts, to support seasonal and longer timescales, and incorporating the atmospheric composition capability from the EU-funded Copernicus Atmosphere Monitoring Service (CAMS) operated by ECMWF.

Over the last ten years OpenIFS has established a growing community of users, generated impact at ECMWF, universities and other research institutes, and established an international workshop series. Given these achievements, together with ongoing plans, we are sure OpenIFS has an exciting future and can continue to provide a valuable link between ECMWF and the wider community to the benefit of both.

# Newsletter readership survey results

Georg Lentze

A questionnaire that was run online from October to December 2021 shows that the appearance and content of the ECMWF Newsletter are much appreciated by its readers. The survey has also provided some valuable insights for possible improvements. Many thanks to all the readers who took part!

Almost two thirds of participants view the PDF version of the Newsletter, while over a third read the newer web version. The print version is now only available in limited numbers in the three outlets of the Centre in Reading (UK), Bonn (Germany), and Bologna (Italy). The results of the survey show that it continues to be a good idea to produce a PDF version to maximise access to the Newsletter.

As can be seen in the illustration, most participants are satisfied or very satisfied with the appearance of the Newsletter, its overall content, and its readability and timeliness. The articles considered most important are the ones discussing changes to ECMWF's forecasting system ('very important' for 58.9% and 'important' for 29.9%), but general news items are also rated highly ('very important' for 37.4% and 'important' for 52.3%). The range of readers is diverse, which is why computing articles will continue to feature ('very important' for 18.7% and 'important' for 39.3%). A list of new ECMWF publications and a calendar of events, which are currently only included in the PDF version and hard copy, are regarded at least as 'somewhat important' by most participants (85.1% and 75.7%, respectively).

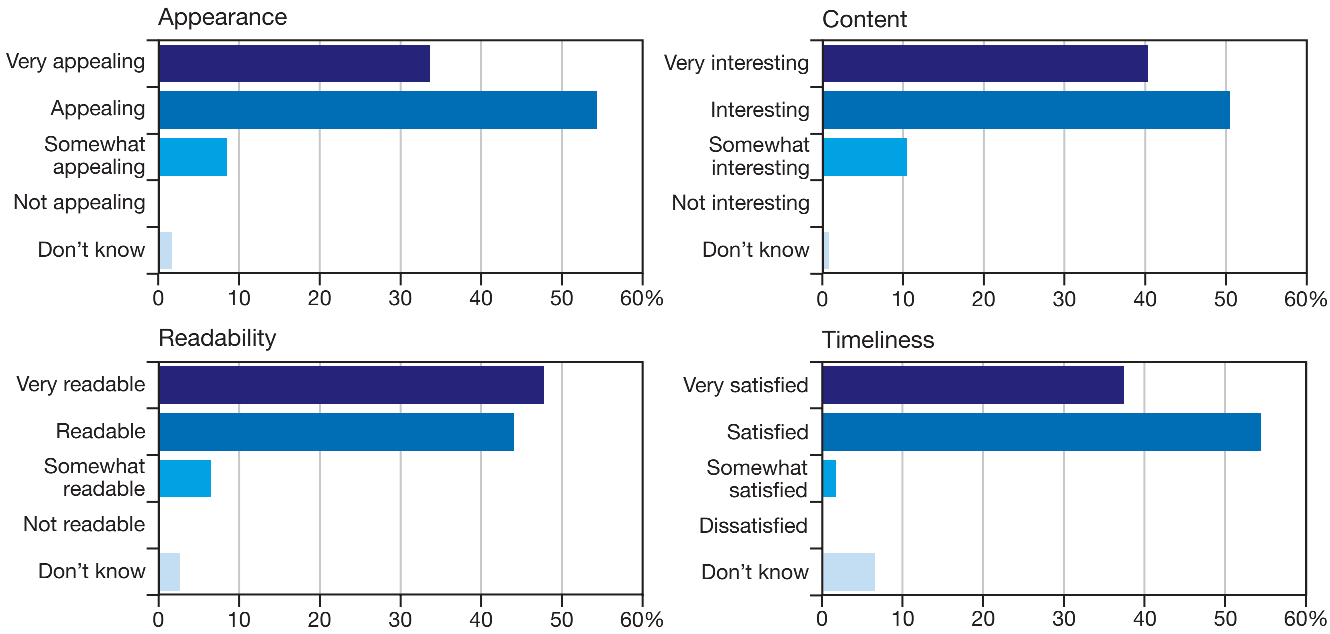


**Some results from the readership survey.** The results show a high level of satisfaction with the appearance, content, readability and timeliness of the ECMWF Newsletter. They are based on 105–109 responses.

In addition to providing much positive feedback, free text boxes were used for some constructive proposals, such as calls for articles about the use of ECMWF products, interviews with senior ECMWF scientists, and more articles about the Centre's activities relating to the EU's Copernicus Earth observation programme. We appreciate all the responses and will aim to incorporate ideas where possible into future issues of the ECMWF Newsletter.

For tips on how to submit an article or content idea, please consult the 'Guidance for Newsletter authors' at https://www.ecmwf.int/en/about/media-centre/media-resources or contact *Georg.Lentze@ecmwf.int*.

# ECMWF holds atmospheric composition training with ESA and EUMETSAT

Chris Stewart, Stijn Vermoote

Almost 790 participants from 107 countries registered for the third joint ECMWF, European Space Agency (ESA) and European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) training in atmospheric composition from 6 to 17 December 2021. The event was chiefly organised by ECMWF and a format was chosen to accommodate a large number of participants from across the globe.

The leading role of ECMWF in the domain of atmospheric monitoring is reflected in the number of lecturers and trainers from ECMWF who contributed to the training. These include Anna Agusti-Panareda, Johannes Flemming, Chris Stewart, Miha Razinger, Antje Inness, Mark Parrington, Martin Suttie, Erik Andersson and Julia Wagemann. Gabriele Pfister, an ECMWF Fellow from the US research centre UCAR/NCAR, also contributed, as did many scientists from a wide variety of respected institutions. A large part of the training made use of data and services provided by the Copernicus Atmosphere Monitoring Service (CAMS, https://atmosphere.copernicus.eu/), entrusted to ECMWF by the European Commission. We have recently entered the second phase of CAMS, paving the way for continued service delivery for the next seven years (2021–2027).

## Training elements

The online training took place as a series of webinars distributed over two weeks. Each webinar began with a short icebreaker quiz, followed by morning lectures given by leading atmospheric composition scientists. The afternoons comprised demonstrations of how to process data from ECMWF, ESA and EUMETSAT, using free and open source tools in the form of Jupyter notebooks. At the end of each practical demonstration, participants were given assignments to complete in the days between the practical sessions. The results of these were discussed in dedicated sessions, and prizes were awarded for the best results.

The distributed programme, combined with the use of online tools for managing questions and sharing results, and the fact that all lectures and practical demonstrations were recorded and available on the course website soon after each session,

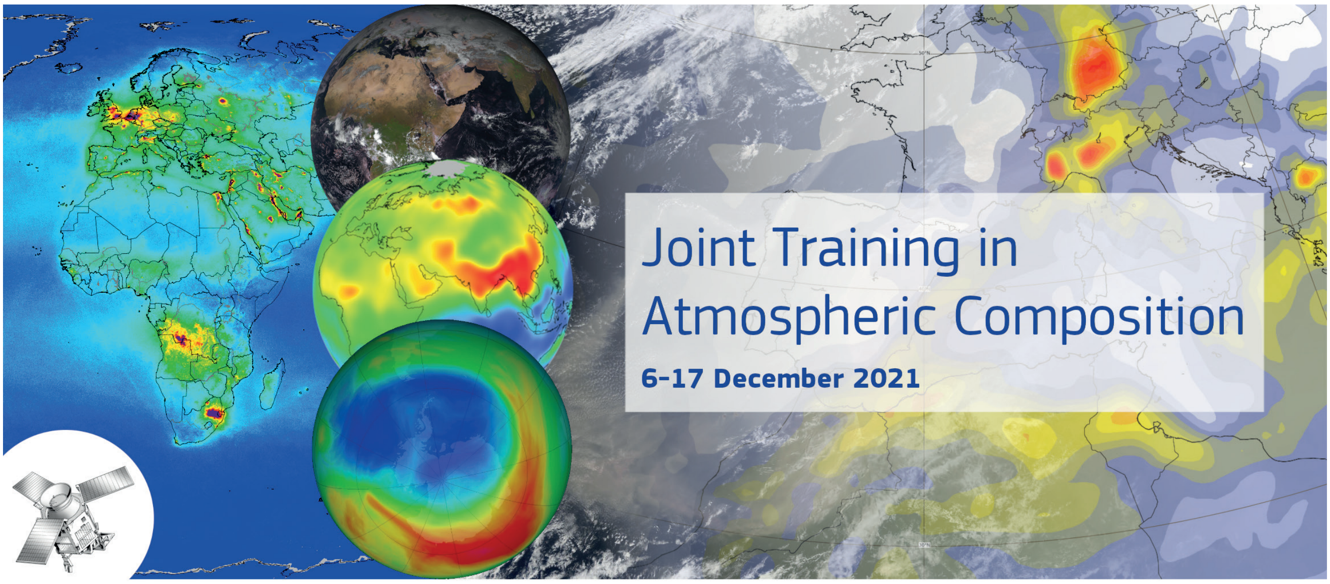


enabled students to actively participate regardless of their time zone. The emphasis on the use of freely available data and tools made it possible to increase the level of participation, and to prepare students to directly apply what they learned in the real world without any need for costly environments or cumbersome installations.

While the course was open to all, the target audience included young scientists at PhD and post-doc level. Given that ECMWF led the organisation, the focus was on modelling, but other aspects of atmospheric composition were also included. The first day provided an introduction to atmospheric processes. The second and third days focused on modelling and observations respectively. The fourth day covered the service element of atmospheric composition monitoring and the 'data value chain' from services such as CAMS to downstream businesses. The last day focused on future developments in modelling, future satellite missions, and future services, including the development of a European $CO_2$ monitoring and verification support capacity as part of CAMS.

Participants were given assignments each day, which were based on the practical sessions. These began with simply accessing CAMS data and analysing plots and animations of atmospheric composition events, such as wildfire activity, $SO_2$ emissions from the Cumbre Vieja volcano, or dust transport across the Atlantic. Later assignments included calculating the Air Quality Index and applying averaging kernels to compare satellite observations with model data. The latter was done using the ADC Toolbox developed by Alba Vilanova as part of the ECMWF Summer of Weather Code (ESoWC) programme.

> Value statements from participants
>
> *"Thank you for this amazing training. Probably the best one!"*
>
> *"Please continue using this online format that allows more young scientists to learn from the best."*
>
> *"Thank you all very much for giving me a chance to attend this excellent course. I have learnt a lot."*
>
> *"It was amazing. I gained so much I couldn't even imagine when I was registering."*
>
> *"Please keep it up. It is very, very useful."*
>
> *"I do not know who has been the visionary behind this training effort, but I wholeheartedly want to thank him/her."*

## Excellent feedback

The training received excellent evaluation from the participants. Eighty-three per cent gave the highest rating of 5 in their overall evaluation of the course. Many stated their intention to begin using for the first time, or more actively, the data, tools and services provided by the three organisations.

Taking advantage of the high participation rate, attendees were invited to complete a survey to capture information on their experience with ECMWF, ESA and EUMETSAT data and resources, their needs, and any suggestions they may have to help the three organisations better serve their respective user communities.

This was the third in a series of joint ECMWF, ESA and EUMETSAT trainings in atmospheric composition. Next year ESA will lead the organisation of the training, which will be more focused on satellite and in-situ monitoring of atmospheric composition.

For more details on the training, and to view the recordings and presentation files, which are available for all lectures and practical sessions, please visit https://atmosphere.copernicus.eu/3rd-eumetsatesaecmwf-joint-training-atmospheric-composition.

# WMO symposium on education and training

Sarah Keeley, Becky Hemingway, Chris Stewart

The first virtual meeting of the international World Meteorological Organization (WMO) symposium on Education and Training, and the 14th session of the symposium itself, took place from 22 to 25 November 2021. This WMO symposium is held every four years to provide recommendations to the WMO, policy-makers and governments, international organisations, and the education and training community itself. The meeting was attended by trainers and educators from all over the globe and had lots of lively discussion on how to develop education and training of current staff and the next generation of meteorologists and environmental scientists.

## Remote teaching

The symposium discussed a range of topics, which covered developments in the skills required by meteorologists and hydrologists and in turn those training and assessing them. Sarah Keeley led the thematic working group on the technological barriers that need to be overcome when remote teaching is needed. It was clear that the approach to any remote learning needs to take into account the available technology and the reliability of internet connections.

During the conference ECMWF presented a poster on how we had responded to the pandemic and rapidly adjusted our training delivery methods to maintain continuity of service, as well as the lessons we have learned over the past year to continue to improve them. Many attendees appreciated what we shared about our experience and the developments concerning our online classroom experience.

## Focus on new areas

Many of the new challenges for the training community reflect changes in the scope of ECMWF's work, where we are already making advances, but also in our working ethos. It was noted that the community could achieve more by collaboration and the pooling of resources and shared expertise. ECMWF has pioneered and championed probabilistic forecasting for decades, and the importance of training for the understanding and communication of probabilistic predictions was one key area that was highlighted. Other areas are gaining in importance throughout the meteorological community as we move to an Earth system approach and more predictions of impacts and climate-related variables are required. Changes we have already undergone at ECMWF include an Earth system approach for our forecast systems and research; machine learning research; and delivering and contributing to the EU's Copernicus services. These developments have naturally brought about changes in the training we offer as well, with new and adapted courses covering these areas. It was encouraging to see that this is supported by community requirements.
The symposium recognised the need to train more people in applying advances in numerical weather prediction, the use of big data, and machine learning.

## Outcomes

The meeting produced a roadmap for education and training, which aims to help increase the quality and capacity of meteorological training around the world. It endorsed the WMO Global Campus initiative to strengthen international and regional collaboration. The recommendations from the meeting aligned closely with the plans we have within our own Strategy for delivering and supporting training across the broad spectrum of ECMWF activities, which will also support our Member and Co-operating States and the wider community in meeting their training needs. For more details, visit https://public.wmo.int/en/media/news/education-and-training-period-of-rapid-change.

# WMO agrees on hourly global surface reporting

Bruce Ingleby

Most countries make hourly weather station reports (SYNOPs) for national use, but many only distribute 3- or 6-hourly reports for others to use. With 4D data assimilation systems, ECMWF and other global numerical weather prediction centres can make good use of higher frequency surface data, and they have been calling for the exchange of hourly SYNOP data. The World Meteorological Organization (WMO) also recognises that more exchange of existing data is a particularly efficient way of improving the global observing system. At a meeting in October, the World Meteorological Congress (representatives from all national meteorological services, the governing body of the WMO) agreed on improvements to data exchange. The Global Basic Observing Network (GBON) was agreed in outline in 2018 and an important detail was added in October: from January 2023, hourly surface reports and high-resolution radiosonde reports should be exchanged globally (https://public.wmo.int/en/events/constituent-bodies/cg-ext2021).

The figure shows the frequency of SYNOP reports currently received at ECMWF. Hourly reports are available from most of Europe and from Australia and Brazil following recent upgrades. The blue colours show 3- and 6-hourly reports. Some stations (grey) are even less frequent. The orange dots show that 11% of stations are not providing BUFR reports (the migration from alphanumeric codes was supposed to be complete in 2014). Unfortunately, some countries are providing no surface data at all. A few, notably the USA, provide more data via METARs than SYNOP reports. The impact on the forecast of reports near the end of the data assimilation window is larger than the impact of reports near the start of the window. The exchange of more frequent data is a potentially important step forward.

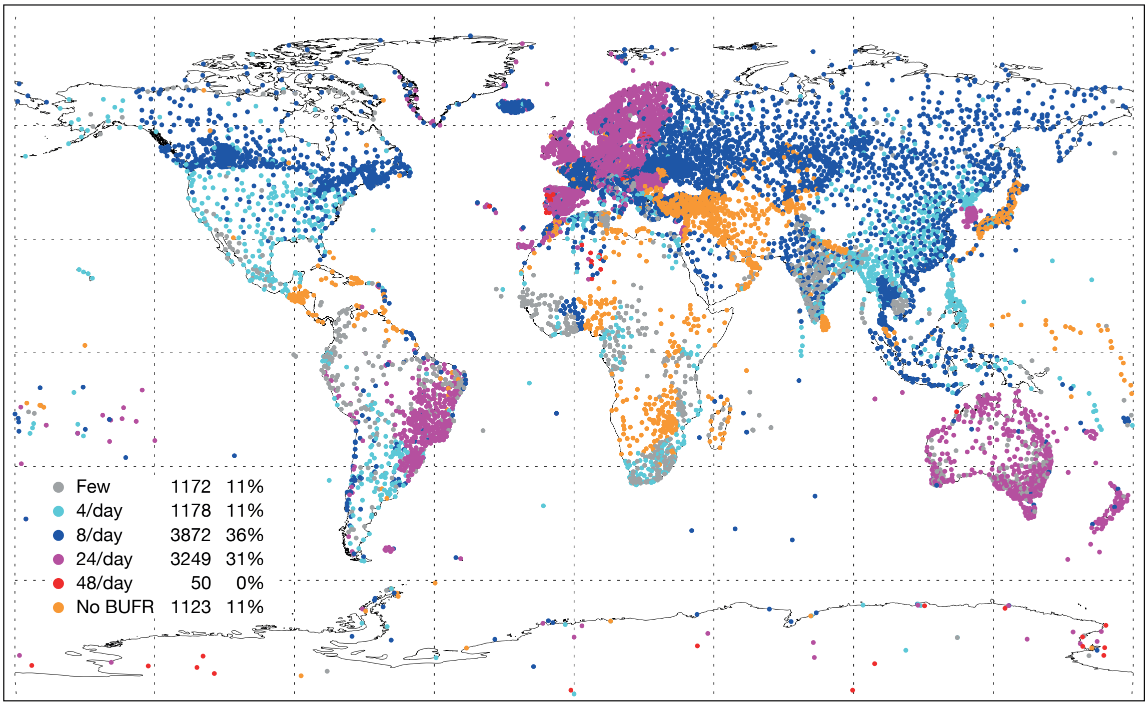


**SYNOP report frequency.** Frequency of BUFR reports received at ECMWF in November 2021. Updates of this map will be available via *https://confluence.ecmwf.int/display/TCBUF/Data+availability*.

## New observations since October 2021

The following new observations have been activated in the operational ECMWF assimilation system since October 2021.

| Observations | Main impact | Activation date |
|---|---|---|
| Sentinel-5P $NO_2$ retrievals | $NO_2$ in the CAMS system | 12 October 2021 |
| Radio occultation bending angles from Sentinel-6a and GRACE-C | Temperature and winds in upper troposphere/lower stratosphere | 13 January 2022 |

# ECMWF Representatives Meetings in 2021

Becky Hemingway, Carsten Maass, Esperanza Cuartero, Xiaobo Yang, Emma Pidduck, Maartje Kuilman

ECMWF Representatives in our Member and Co-operating States play a vital role in managing users and the use of ECMWF products and services, the dissemination of information, and providing feedback to ECMWF. During November and December 2021, virtual meetings were held with Computing Representatives, Meteorological Representatives and Catalogue Contact Points. The meetings were an opportunity for ECMWF to present updates and discuss a variety of topics, listen to user feedback, and understand user issues in Member and Co-operating States. The meetings also provided a forum for attendees to share experiences and network with their fellow representatives.

There were 51 attendees across the three meetings. The virtual nature of the meetings ensured a diverse attendance across Member and Co-operating States, which was beneficial for discussion and feedback. All participants said they found the meetings 'Very useful' or 'Useful', and we received positive feedback on the topics covered and the meeting structure. ECMWF looks forward to continued engagement with its Representatives and thanks them for the important role they play.

More information on each meeting can be found below.

## IT User Forum

Twenty-two Computing Representatives from Member and Co-operating States attended the IT User Forum 2021. The meeting focused on updates relating to our new data centre in Bologna, Italy, and discussion around actions which will need to be taken by Member and Co-operating States to ensure minimal impacts to users during the migration to the Atos supercomputer in 2022. Attendees were also given a video tour around the data centre in Bologna to show its progress and the scale of the centre.

Presentations were given by ECMWF on the European Weather Cloud, the Support Portal, software upgrades, and future working practices at ECMWF. Discussions included the use of tokens for access to ECMWF services, ecCodes and data dissemination during the migration to Bologna.

Attendees provided updates from their organisations and countries via recorded presentations, showcasing their advances in forecasting and technology over the past year. Numerous exciting projects were highlighted, including many collaborations across the meteorological community.

## Meteorological User Forum

Member and Co-operating State Meteorological Representatives were invited to the Meteorological User Forum 2021. This was the second time such a meeting was held. The forum welcomed 16 attendees and featured discussions and presentations from ECMWF staff on new ECMWF products, ecCharts and OpenCharts, our new data centre in Bologna, the Forecast User Forum, the Support Portal and the next Using ECMWF's Forecasts meeting (UEF2022).

A session on the recent upgrade of ECMWF's Integrated Forecasting System to Cycle 47r3 provided valuable feedback on model performance. It was complemented by recorded attendee presentations, which included various case studies and more feedback on forecasts and areas for development. Member and Co-operating State liaison visits were also discussed, with feedback being used to improve liaison visits in the current and future visit cycles. Other discussions included ecCharts and OpenCharts and how these services could be improved to provide more useful output to users, login issues, the Support Portal, and data volumes

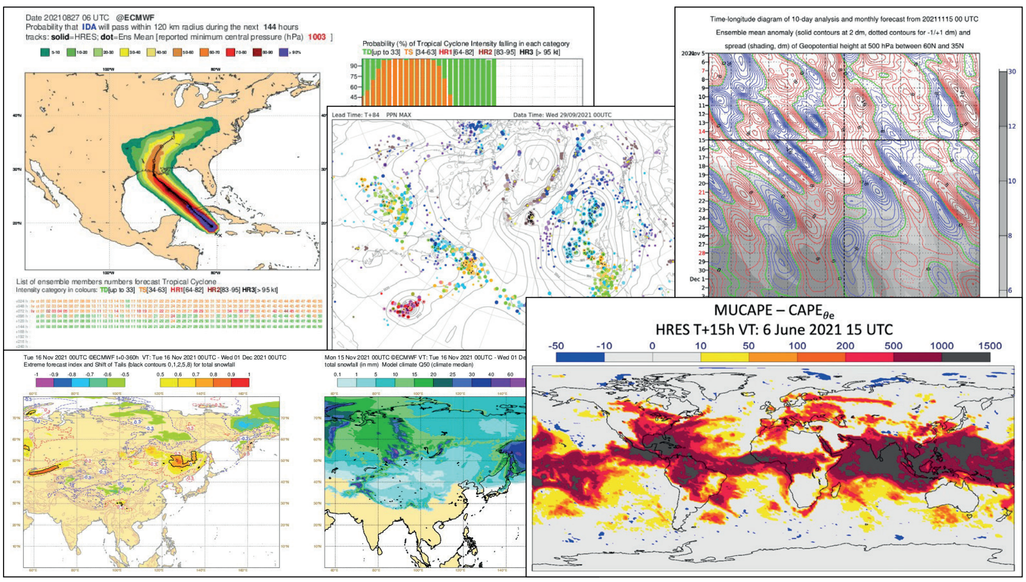


**Presentation of products.** Many new and updated ECMWF products were presented at the Meteorological User Forum 2021, including tropical cyclone products, longer period Extreme Forecast Index/Shift of Tails fields for snowfall, and improved Hovmöller diagrams for the extended range.

in the future, especially given the planned upgrades to model resolution and ensemble member numbers.

To facilitate networking between the representatives, a 'Social Break' was held. This involved groups of three or four participants coming together for an informal chat for ten minutes. The break was positively received, with attendees liking the opportunity to chat with others.

### Catalogue Contact Points Meetings

The Catalogue Contact Points (CCP) Meetings were attended by 13 CCPs from Member and Co-operating States, and a number of new CCPs were welcomed to the group. ECMWF provided updates on the Advisory Committee for Data Policy (ACDP) 2021 and June Council outcomes implementation, future developments, and research access policies. Participants presented information on open data, delivery services and policy changes within attendee countries. Feedback was also sought on the data redistribution policy, and discussions on the new data service charge model and open data took place.

A 'Social Break' also took place between attendees, with three or four participants per group for an informal chat. These were positively received.

# Nobel Prize winner Klaus Hasselmann and ECMWF

Luigi Cavaleri (ISMAR, Italy), Peter Janssen (ECMWF)

The Nobel Prize in Physics was awarded in 2021 to Syukuro Manabe, Klaus Hasselmann and Giorgio Parisi in acknowledgement of their scientific work, thus giving a clear message on the relevance of the climate issue. Out of the three, Klaus was interested in a number of problems which closely intersected with activities at ECMWF. What follows is a short account of his role in the activities of the Centre relating to both practical and fundamental questions.

**Klaus Hasselmann.** A wide range of his work, from ocean-wave dynamics and coupled modelling to climate-related work, had a direct impact on ECMWF. *(Courtesy of Susanne Hasselmann)*

## Ocean wave dynamics

With a degree in theoretical physics, the first contact Klaus had with the practical world was in oceanography, in particular wind waves. In 1960, he developed the concept of the energy balance equation, which is still the basis of all present wave forecasting models. In 1962, he shocked the oceanography community by showing that resonant four-wave interactions amongst gravity waves, which he represented by a perturbation expansion, play a key role in ocean wave dynamics. Four years later, he showed how the same problem can be solved using Feynman-type diagrams. He thus provided a sound explanation of why, during the evolution of ocean waves, their energy moves progressively to lower and lower frequencies. In 1973, the JONSWAP experiment gave the first empirical evidence of how wind waves evolve, followed one year later by the first quantification of the energy loss by white-capping.

All these theoretical and practical developments were ahead of time because computing capacity was very limited. But as soon as Cray computers became available, Klaus started interacting with Lennart Bengtsson, then Director of ECMWF, and a small group of wave modellers was formed trying to implement a sophisticated wave model at ECMWF, using the classroom as an office. However, the restriction was that only 10% of operational resources could be spent on wave forecasting, and even vector machines such as the Cray were not powerful enough to produce a reasonably accurate wave forecast. Once again, the solution was provided by Klaus and his wife Susanne, by developing an efficient approximation to the nonlinear interaction calculation, which is called the Discrete Interaction Approximation (DIA). Even after almost 40 years, the DIA is still the essential ingredient that makes wave forecasting possible at meteorological centres all over the world.

## Coupled modelling

But Klaus's interest was not restricted to wave modelling alone. His vision was to develop a comprehensive coupled ocean-wave and atmosphere model that would use global surface data, such as wind speed and direction, and sea-state parameters, such as wave height and wave spectrum, for initialisation of the coupled system. These data could only be provided by polar-orbiting satellites, so Klaus started interacting with the European Space Agency (ESA). He played an important role in the development of instruments such as the Scatterometer (surface winds) and the SAR (2D-spectrum of long waves) on board the first European Earth observation satellite ERS-1, which was launched in 1991. The coupled ocean-wave and atmosphere model was introduced in June 1998, and the wealth of wind and wave data has improved the performance of the ECMWF forecasting system ever since. At the same time, satellite data were instrumental in the validation of wave forecasting results, and many forecasting centres have obtained considerable improvements in the forecast skill of wave parameters by improving the ocean wave model.

**The authors of the book on the 'Dynamics and Modelling of Ocean Waves' in 1994.** Klaus Hasselmann, Gerbrand Komen, Peter Janssen, Susanne Hasselmann, Mark Donelan and Luigi Cavaleri (left to right).

## Climate-related work

All this could easily lead you to consider that Klaus was mainly a wave man, but this would be far from the truth. His dominant interest (think of the Nobel Prize) was climate, and this is where, from a different perspective, his work is again related to ECMWF. In many of the climate models used in the 1960s and 70s, the atmosphere was not explicitly included. Instead, it was placed in the model's statistical-diagnostic regime and was represented only through temporally averaged terms. However, in his seminal paper on 'Stochastic climate models' published in *Tellus* in 1976, Klaus pointed out that the atmosphere's influence is not limited to these temporally averaged terms and that its variability must also be considered. This results in differential equations for the slow components of the climate system, which include stochastic forcing terms. Analogous to Brownian motion, these short-term atmospheric variations cause long-term fluctuations in slow subsystems, which explains the observed red spectrum of slow climate variables. The theory of Brownian motion has been discussed in many applications since Einstein's paper in 1905, but it had not yet been applied to geophysical systems, such as the climate system. In two follow-up 1977 papers, the applicability of the concept was demonstrated through an analysis of sea-surface temperatures and thermocline variability and with a global energy balance model. A large variety of different applications of this stochastic approach followed in subsequent years.

Having framed the physics and the limits to the predictability of the climate system, Klaus addressed the problem of detecting a meaningful signal amidst the noise of the past. This was achieved by introducing the concept of 'detection and attribution'. First, in the detection step, the relevant change is examined to see if it falls within the range of natural variability. Having obtained a successful rejection, the change is compared with one or several theories. If a good fit is found, then the conclusion is drawn that the change can be attributed to the relevant factor(s). The concept was used successfully for investigating whether an external signal, as suggested by climate model simulations, would be detectable in the observational record of global temperatures. The technique was adopted in various corners of the academic world and eventually became a cornerstone in reports by the Intergovernmental Panel on Climate Change (IPCC): the fingerprint of human activity in a changing global climate.

From a general point of view, the concept of a possibly very wide range of frequencies interacting in a non-linear way forms the basis of various processes or timescales ECMWF deals with. An obvious example is the daily variations in weather details and the predictability of seasonal or longer-term forecasts. An application at the other, high-frequency end of the timescale is the effect of gustiness in wind fields on the average wind input to ocean waves. Klaus's climate work is also important to the activities of the EU-funded Copernicus Climate Change Service (C3S), which has been run by ECMWF since 2014. C3S offers information on the past, present and future climate, and it builds on Klaus's insights, especially as under the C3S umbrella a coupled ocean–atmosphere reanalysis is being developed.

## A fundamental role

The list of Klaus's achievements is very long, especially when you realise that most of his work has been in fields different from meteorology. In our daily work for better results, with a steady look to the future, it is easy to forget that we stand on the shoulders of the greats of the past. Although his main interest was in climate, with later work on the economy and a unified theory of particles and fields, Klaus had a fundamental role in shaping some of the activities at ECMWF. Knowing that we are working on the shoulders of a Nobel Prize winner is, no doubt, a privilege and a great pleasure.

# ECMWF's support for the SEE-MHEWS-A project

Fredrik Wetterhall, Umberto Modigliani, Bojan Kasic

**In 2016, the World Meteorological Organization (WMO) initiated the South-East European Multi-Hazard Early Warning Advisory System (SEE-MHEWS-A) project, which aims to strengthen the existing early warning capacity in the region. The project finished its first phase in 2017, which delivered an implementation plan. The project is now finalising the second phase, which is the implementation of a prototype of an operational system. ECMWF has been supporting this work, and several of our services have been used in the implementation. See Box A for more information on SEE-MHEWS-A.**

Since 2019, ECMWF has been involved in the project by providing expertise and support in: (i) observational data handling; (ii) setting up limited-area numerical weather prediction (NWP); (iii) setting up hydrological models over the area; and (iv) supporting the development of a web-based Common Information Platform (CIP).

The pilot phase of the project has so far been a success. Several countries in the region have agreed to share more of their observational data within the project. The data are handled by the Scalable Acquisition and Pre-Processing (SAPP) system at ECMWF and are now used in data assimilation for ECMWF's forecasts. Furthermore, the four NWP models in the project have all been implemented by ECMWF Member and Co-operating States and are running in real time, some as time-critical applications. Three hydrological models have been calibrated over the region and now run quasi-operationally. A web-based platform for disseminating the products is being developed and is hosted on the European Weather Cloud (EWC). The EWC also hosts a nowcasting component. Figure 1 shows a conceptual diagram of the tools used at ECMWF to deliver the system.

## Acquisition, pre-processing and use of observations

One of the goals of the pilot phase of the SEE-MHEWS-A project was to acquire and pre-process additional meteorological observations not yet distributed on the Global Telecommunication System (GTS), or having higher reporting frequency, and make them available for data assimilation. ECMWF achieved this goal by utilising the ECMWF Production Data Store (ECPDS) and the Scalable Acquisition and Pre-Processing (SAPP) system.

The ECPDS acquisition service/portal is used for automatic data retrieval (relying on ftp, sftp, and https protocols); both public and private accounts were

**a** **Scope and aims of SEE-MHEWS-A**

The SEE-MHEWS-A project is supported by the following WMO national meteorological and hydrological services (NMHSs) in southeast Europe: Albania, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Greece, Hungary, Israel, Jordan, Lebanon, North Macedonia, Republic of Moldova, Montenegro, Romania, Slovenia, Turkey, and Ukraine. Of these 17 countries, ten are ECMWF Member or Co-operating States. The project's aims are to:

- strengthen regional cooperation by leveraging national, regional and global capacities to develop improved hydro-meteorological forecasts, advisories and warnings to save lives and limit economic losses
- strengthen national multi-hazard early warning systems making tools and data available to the participating countries and other beneficiaries
- implement impact-based forecasts and risk-based warnings to support governments, disaster management authorities, humanitarian agencies and NGOs in their decision-making
- harmonise forecasts and warnings in trans-boundary areas
- provide training on operational forecasting for staff of the NMHSs participating in the project.

For more details, see: *https://public.wmo.int/en/projects/see-mhews-a*.


doi: 10.21957/fb92s3aj7w

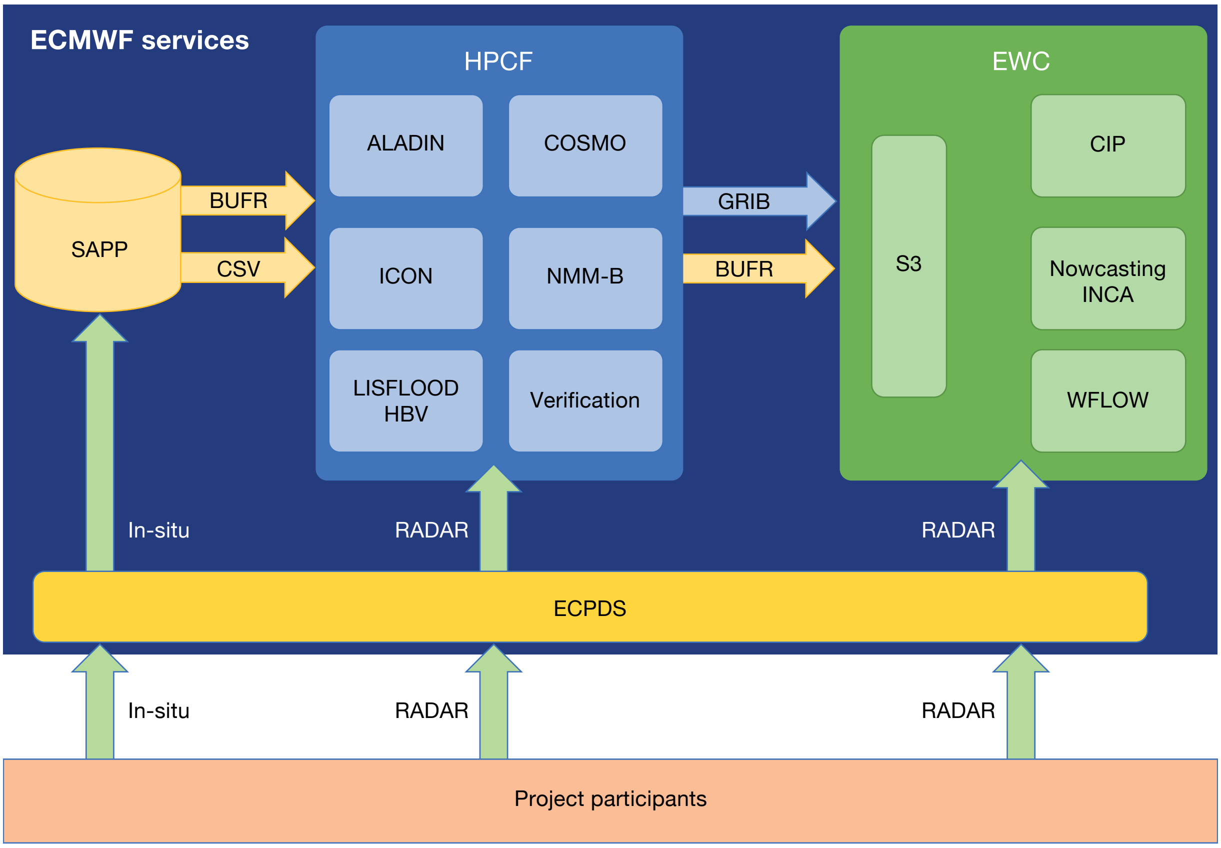


**FIGURE 1** This is a high-level view of the SEE-MHEWS-A main components and the workflow implemented at ECMWF. In-situ data from the project participants feed into ECMWF's Scalable Acquisition and Pre-Processing (SAPP) system. Together with radar observations, they are then directed to ECMWF's high-performance computing facility (HPCF). Here, they go through different limited-area weather forecasting systems (ALADIN, COSMO, ICON, NMM-B) and the LISFLOOD and HBV hydrological rainfall/runoff models. The output is then transferred to the European Weather Cloud (EWC) with the Common Information Platform (CIP) as well as the INCA nowcasting and the WFLOW hydrological model. A shared Amazon S3 bucket is used to store the data on the EWC, and ECMWF's Production Data Store (ECPDS) is used to retrieve observations from project participants.

created to support participating countries' needs to transfer public and/or reserved data. Files acquired by ECPDS are sent to the SAPP data driven processing system, responsible for observation decoding, initial format and quality control, and conversion to a consolidated ECMWF BUFR format before extractions for data assimilation. One of the advantages of using the SAPP system is the ability to configure new processing chains for the newly developed encoders for the project.

To facilitate provisioning of additional meteorological observations that could not be provided in BUFR format, a custom SYNOP CSV data exchange format and an ecCodes-based CSV to BUFR converter was developed. This includes handling of local and national identifiers for stations not yet registered in the WMO OSCAR database. The SAPP system has also been used to validate and process sub-hourly BUFR SYNOP observations that Slovenia started to circulate on the GTS in June 2021. Furthermore, SAPP extractions have been configured to feed SEE-MHEWS-A data assimilation, hydrological verification, and CIP for display purposes.

Thanks to this joint effort, it was possible to start exchanging many additional observations not available on the GTS (Tables 1 and 2). As of August 2021, Bosnia and Herzegovina, the Republic of Croatia, Cyprus, Hungary, Montenegro, the Republic of North Macedonia, the Republic of Turkey and Ukraine are all submitting SEE-MHEWS-A observations regularly in near-real-time (NRT). Israel and Slovenia are providing all their in-situ observations on the GTS. ECMWF provided SAPP technical support for BUFR processing and extraction to participant countries such as Turkey, which utilise SAPP in their local environment.

The additional measurements provided as part of the SEE-MHEWS-A project started to be assimilated operationally in November 2020 by ECMWF's upper air 4D-Var system and the Land Data Assimilation System (LDAS). The ECMWF operational analysis (both 4D-Var and LDAS) and the high-resolution forecast (HRES) are run on ECMWF's HPCF. ECMWF analysis and HRES forecasts are then used to provide initial conditions (IC) and lateral boundary conditions (LBC) used by the limited-area models (LAM) and nowcasting system run as part of the SEE-MHEWS-A project. The required

| Country | Number of stations – GTS | Number of additional stations through SEE-MHEWS-A | Total number of reports/day – GTS | Total number of additional reports/day through SEE-MHEWS-A |
|---|---|---|---|---|
| Bosnia and Herzegovina | 14 | 63 | 222 | 9072 |
| Republic of Croatia | 40 | 4 | 683 | 576 |
| Cyprus | 4 | 37 | 43 | 5328 |
| Greece | 44 | 11 | 282 | 82 |
| Hungary | 30 | 90 | 496 | 2160 |
| Montenegro | 6 | 4 | 120 | 80 |
| Republic of North Macedonia | 17 | 11 | 113 | 1584 |
| Republic of Turkey | 122 | 144 | 1640 | 3456 |
| Ukraine | 36 | 127 | 288 | 1016 |
| **Total Number** | **312** | **491** | **3887** | **26354** |

**TABLE 1** GTS stations and additional stations provided through the SEE-MHEWS-A project.

files are produced using ECMWF's product generation (PGEN) system and then transferred, using ECMWF's dissemination system (ECPDS), to the required file system. This is either on the HPCF for the LAM runs or on the object data store (S3 buckets) for the nowcasting system run on the European Weather Cloud (Figure 1).

In addition to IC and LBC from ECMWF's Integrated Forecasting System (IFS), the ALADIN-ALARO LAM system uses its own data assimilation based on three hourly cycles with 3D-Var for assimilation of atmospheric variables and Optimal Interpolation for assimilation of soil fields. The main forecast is run twice per day, at 00 and 12 UTC.

The main observations assimilated are collected and processed by the common pre-processing system of the RC LACE consortium (OPLACE) hosted by the Hungarian Meteorological Service (HMS). The system is pre-processing surface, aircraft, radiosonde and wind profile observations as well as atmospheric motion vectors, satellite radiances and scatterometer data.

Similar SAPP extractions of the SEE-MHEWS-A observations in CSV format to the HPCF have been set for the needs of one of the hydrological models used (HBV).

## Meteorological modelling

The limited-area NWP component of the SEE-MHEWS-A project consists of four models (see Table 3) with the aim of providing high-resolution, short-range weather forecasts, but also initial and boundary conditions for hydrological models and nowcasting.

Three models are configured to share the same domain (from 9.4°E to 45°E, and from 28.5°N to 52.5°N) and a spatial resolution of 2.5 km (Figures 2 and 3). Only COSMO is run on a slightly coarser resolution of 4 km because the system was configured before the start of the SEE-MHEWS-A project and some time is needed for

| Country | Pressure – 4D-Var | 2-metre relative humidity – 4D-Var | 2-metre temperature – LDAS | 2-metre relative humidity – LDAS | Snow depth – LDAS |
|---|---|---|---|---|---|
| Bosnia and Herzegovina | 12 | 1 | 63 | 47 | 0 |
| Republic of Croatia | 4 | 4 | 4 | 1 | 1 |
| Cyprus | 9 | 3 | 30 | 3 | 0 |
| Greece | 8 | 3 | 9 | 9 | 1 |
| Hungary | 21 | 19 | 87 | 86 | 0 |
| Montenegro | 2 | 0 | 3 | 3 | 2 |
| Republic of North Macedonia | 9 | 2 | 10 | 10 | 2 |
| Republic of Turkey | 114 | 46 | 130 | 130 | 7 |
| Ukraine | 122 | 110 | 126 | 126 | 120 |
| **Total Number** | **301** | **118** | **462** | **415** | **133** |

**TABLE 2** Additional SEE-MHEWS-A stations assimilated by 4D-Var and LDAS.

| Model | Responsible centre |
|---|---|
| ALADIN-ALARO (Aire Limitée Adaptation dynamique Développement InterNational) | ARSO (Slovenian Environmental Agency) |
| COSMO (Consortium for Small-scale Modeling) | HNMS (Hellenic National Meteorological Service) |
| ICON (Icosahedral Nonhydrostatic model) | IMS (Israel Meteorological Service) |
| NMM-B (Nonhydrostatic Multi-scale Model on B grid) | UoB (University of Belgrade) |

**TABLE 3** Limited-area NWP models used in SEE-MHEWS-A.

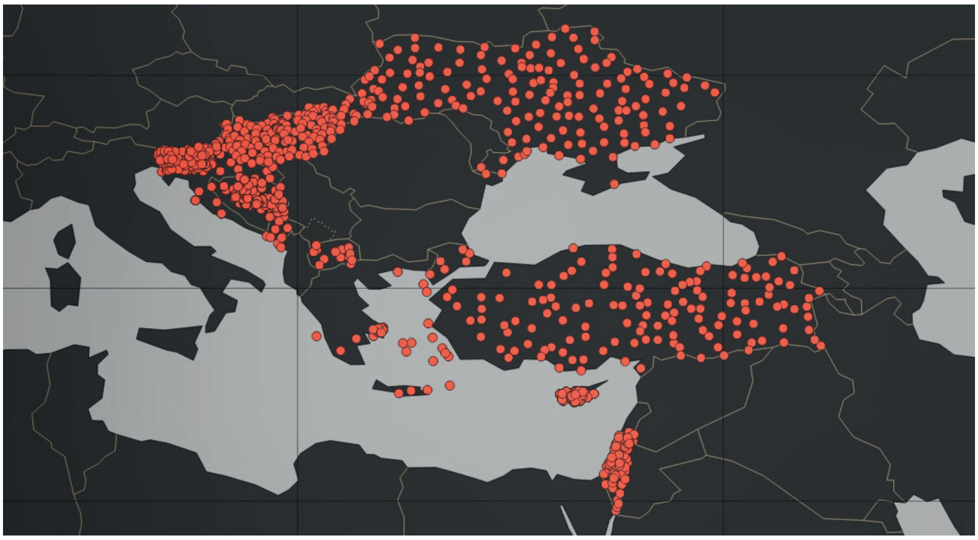

**FIGURE 2** Station locations from which SYNOP data were extracted for the NWP component of the SEE-MHEWS-A project during a time window of one hour from 11:30 on 9 August 2021.

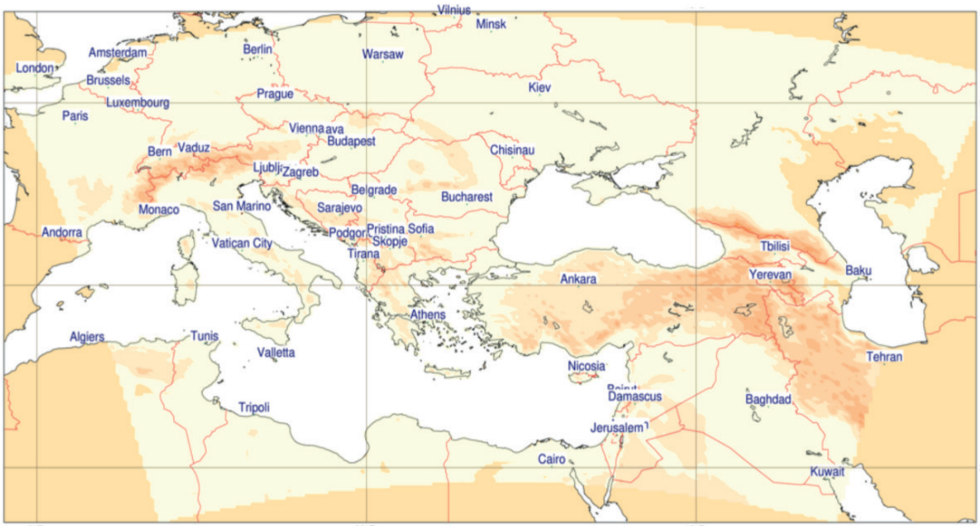


**FIGURE 3** The common NWP domain for SEE-MHEWS-A.

| Model | Correlation coefficient | Bias (mm/24 h) | Mean absolute error (mm/24 h) | Root-mean-square error (mm/24 h) |
|---|---|---|---|---|
| ALADIN | 0.76* | 2.55* | 3.83* | 8.74 |
| COSMO | 0.58* | 0.35 | 3.45 | 9.66 |
| ECMWF | 0.68 | 0.70 | 2.92 | 8.09 |
| ICON | 0.69 | 0.84 | 3.00 | 8.82 |
| NMM-B | 0.60* | 0.15 | 3.38 | 8.71 |

**TABLE 4** Example of conventional verification results for ALADIN, COSMO, ECMWF, ICON and NMM-B for 24-hour cumulative precipitation data over Croatia. The values significantly different from the ECMWF forecast (0.05 significance level) are marked with an asterisk.

adjustments to the project requirements. Each model is run twice per day, with starting times of 00 and 12 UTC, using ECMWF's HRES for IC and LBC over a forecasting period of 72 hours.

Meteorological outputs from NWP models are consistent between models, and include outputs for hydrological modelling, verification, nowcasting, and the CIP. They include commonly used surface and standard pressure level fields. All models are managed by their corresponding ecFlow suites. The ecFlow servers are run on a dedicated virtual machine or on ecgate with all suite components hosted by the ECMWF HPCF. IC and LBC are disseminated via the ECPDS directly to the HPCF.

While COSMO and ICON run as Time-Critical option 2 suites, ALADIN-ALARO and NMM-B run in near-real time at this stage of the project. The model outputs are kept on the HPCF for one month, so that the latest runs are quickly available to other project components, such as hydrology. The same outputs are archived in ECMWF's file storage system (ECFS) to make historical runs available to the project participants. Outputs for CIP and nowcasting are disseminated via ECPDS to the S3 object store in the European Weather Cloud, where these project components run.

## Verification of NWP output

An important activity of the SEE-MHEWS-A project is the verification of the performance of the NWP models. A comparison of the results in a structured, region-wide forecast verification procedure is a necessary component of the agreed process to establish a multi-hazard early warning advisory system in southeast Europe. The application of the verification methodology agreed in the project aims to evaluate the performance of several NWP models. The five models available for verification are: ALADIN-ALARO, COSMO, ECMWF IFS HRES, ICON, and NMM-B.

The 00 UTC initialisation model runs are used for the verification of the 24 h cumulative precipitation forecast. To avoid the model spin-up period, only the data from 06 UTC to 54 UTC is used, forming a forecast range up to two days ahead. The ALADIN, ICON and NMM-B model resolutions are ~2.5 km, COSMO's is ~4 km, and the ECMWF-IFS model resolution is ~9 km. For comparison with observations, model forecast data are extracted for the nearest model grid point over land. The verification methodology is applied on 24 h cumulative precipitation data gathered from the available observing stations in each country for the period from 22 September 2020 to 22 November 2020. As part of the project, by applying the above verification methodology, a set of country-specific verification reports has been produced by DHMZ, the Croatian Meteorological and Hydrological Service (Table 4).

## Nowcasting

Nowcasting is becoming increasingly important for warnings of high-impact weather and numerous other applications. As part of the SEE-MHEWS-A project, the Austrian Zentralanstalt für Meteorologie und Geodynamik (ZAMG) is piloting the multivariable analysis and nowcasting system INCA for a specific area in Bosnia and Herzegovina (Figure 4). INCA provides near-real-time analyses and forecasts of

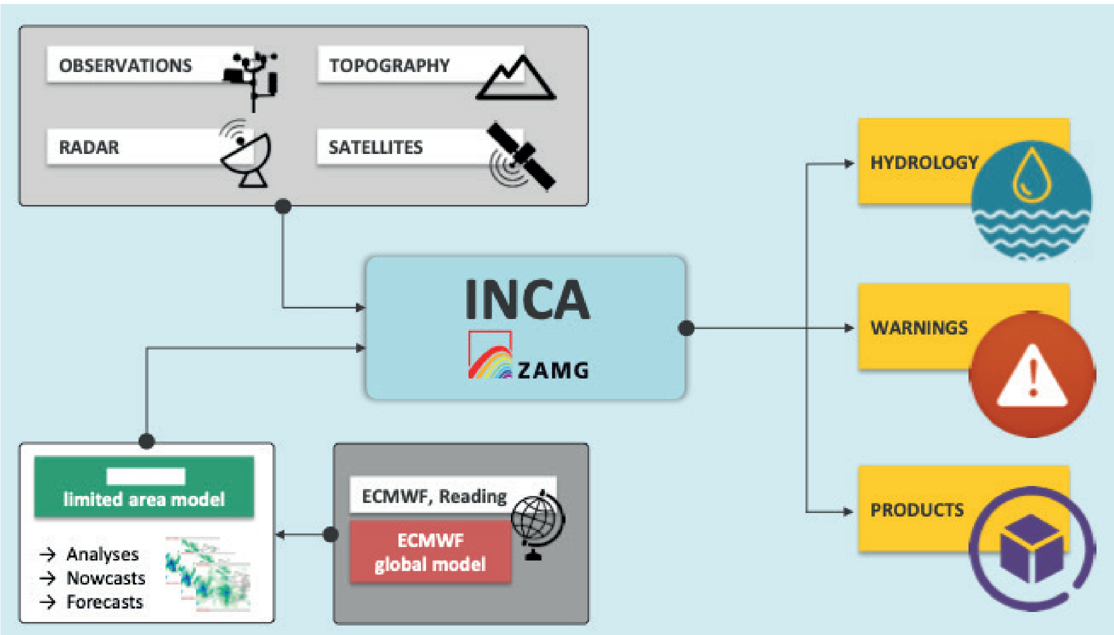


**FIGURE 4** Diagram showing the INCA set-up at ZAMG in Austria. Within SEE-MHEWS-A, INCA is run on the European Weather Cloud at ECMWF. Observations, radar, topographical data and satellite data as well as NWP data as the main input data are also provided within the cloud. The limited-area model used is ICON as run by the Israel Meteorological Service, with boundary conditions provided by ECMWF.

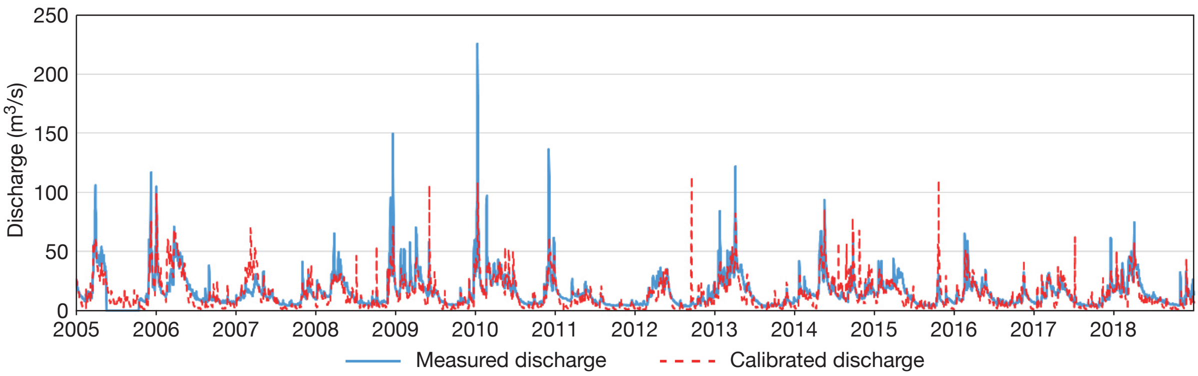


**FIGURE 5** Measured (blue line) vs. calibrated discharge (red line) at Daljan, the most upstream hydrological station on the Vrbas catchment river.

several surface fields. Its objective is to improve numerical forecast products in the nowcasting (0–4 h) and very short (up to about 12 h) ranges (Haiden et al., 2011). INCA has a very high-resolution spatial and temporal system. It can currently provide weather analyses and nowcasting for a 100 $m^2$ to 1 $km^2$ grid and has a rapid update cycle, running every five minutes to one hour depending on data availability, weather parameters and application requirements.

INCA analysis merges, in real time, the available observations of automatic weather stations, radars and satellites, forecasts of NWP models, and very high-resolution geographical and topographic data. As a multivariable system, INCA provides real-time analyses and nowcasting for precipitation amounts and types as well as for temperature, humidity, wind, cloudiness, snow fall line, ground temperature, convective parameters, global radiation, and so on.

ZAMG has implemented a version of INCA on the European Weather Cloud at ECMWF, using all available in-situ measurements from the area, including additional SYNOP and radar observations provided as part of the SEE-MHEWS-A project. The INCA configuration as of 31 August 2021 is tailored to the specific requirements of the project and to the current data availability. The current domain is 16.5°–17.9°E, 43.7°–45.3°N, with a spatial resolution of 1 km, covering the Vrbas river catchment area in Bosnia and Herzegovina. Despite the high temporal resolution of the rain gauge observations (10 minutes), the update frequency of INCA precipitation is set to 20 minutes because the radar is currently provided at 20-minute intervals. INCA was implemented in the European Weather Cloud using the latest Python-based version of the code at ZAMG. Adjustments and optimisations due to improved or additional data provision and/or feedback by local experts can be easily undertaken within the project.

## Hydrological modelling

The hydrological models used in the project are LISFLOOD (the hydrological model used in the European Flood Awareness System, EFAS; Smith et al., 2016), HBV and WFLOW. The HBV and LISFLOOD hydrological models were in the pilot phase set up over the Vrbas catchment in Bosnia and Herzegovina, and the WFLOW model was set up over the Vardar catchment in North Macedonia. Setting up the models entails adjusting the river network and making sure all static maps are in place, as well as calibrating the models using hydrological and meteorological observations from the catchment that were gathered and processed for this purpose. It is planned that the models be expanded to larger regions in the next phase of the project.

The LISFLOOD model was calibrated using similar tools to those developed for EFAS. Quite a substantial amount of development work was needed to adjust the tools to work over a single catchment such as the Vrbas. The calibration resulted in an improvement of performance of the stations in the catchment. An example of the calibration is shown in Figure 5.

LISFLOOD is implemented on the ECMWF HPCF using ecFlow. The workflow is that the model picks up observational data (precipitation and temperature) from the EFAS data flow and calculates the hydrological initial conditions. Then it retrieves the NWP data produced over the region and runs them through the hydrological model. Finally, products for the web are generated and are sent to the CIP for dissemination.

The LISFLOOD and HBV models were both set up to run on ECMWF's HPCF as part of our support for the project. WFLOW is set up using FEWS, which is a piece of open software that handles data flow, calibration and verification. The output of the hydrological models is transformed to JSON files for a flexible integration of the information on the CIP.

## The Common Information Platform (CIP)

The SEE-MHEWS-A CIP consists of four components: a user interface (Figure 6), an image server, a database, and an intermediary element between users and the resources they intend to obtain (REST API).

A prototype was developed and implemented on the European Weather Cloud (EWC) using Kubernetes (k8s) as a cluster solution. A virtual server has also been installed within the cloud environment to provide Network File System (NFS) server functionality. The whole application is based on Django apps and hosts a catalogue of available products, a viewer, and a dashboard to monitor points of interest. The application also deals with user management, such as user registration, verification and authorisation. The solution for the product generation in the CIP depends on the product. For NWP products, GRIB files are converted to web map services (WMS). For other products, such as hydrological reporting points and river discharge exceedance, the hydrological modelling chain creates images and xml files which are then converted to a web service at the CIP backend. This is done using Mapserver and a set of endpoint services made available through the Django REST API app.

The main technologies used in the CIP for visualising data are following Open Geospatial Consortium (OGC) standards. The technologies used are enabling interaction with the user, and plots can be adjusted on the fly. Several user interaction events, including dedicated workshops, were held to understand the needs of users to ensure that the prototype is fit for purpose.

## Conclusions and lessons learned

The use of the additional observations brought by SEE-MHEWS-A in ECMWF data assimilation is benefiting all of our Member and Co-operating States through the provision of improved initial and lateral boundary conditions. The additional observations are also used by the data assimilation part of the ALADIN contribution to SEE-MHEWS-A and the INCA Nowcasting system. A subset of them, targeting a specific area, will be used by the hydrological components. In January 2021, the SEE-MHEWS-A Project Steering Group agreed to share these additional observations with WMO Global Data-processing and Forecasting System centres, starting with the World Meteorological Centre DWD Offenbach.

The experience gained in the acquisition and processing of additional observations not shared via the GTS is proving useful and beneficial in the recently established EUMETNET Federated Data Coordination Mechanism (FDCM) programme.

The benefits for ECMWF and its Member and Co-operating States can be summarised as follows:

- The development of the pilot SEE-MHEWS-A system strengthens the early warning capacity in the region for the benefit of all national meteorological and hydrological services (NMHS) participating in the project.

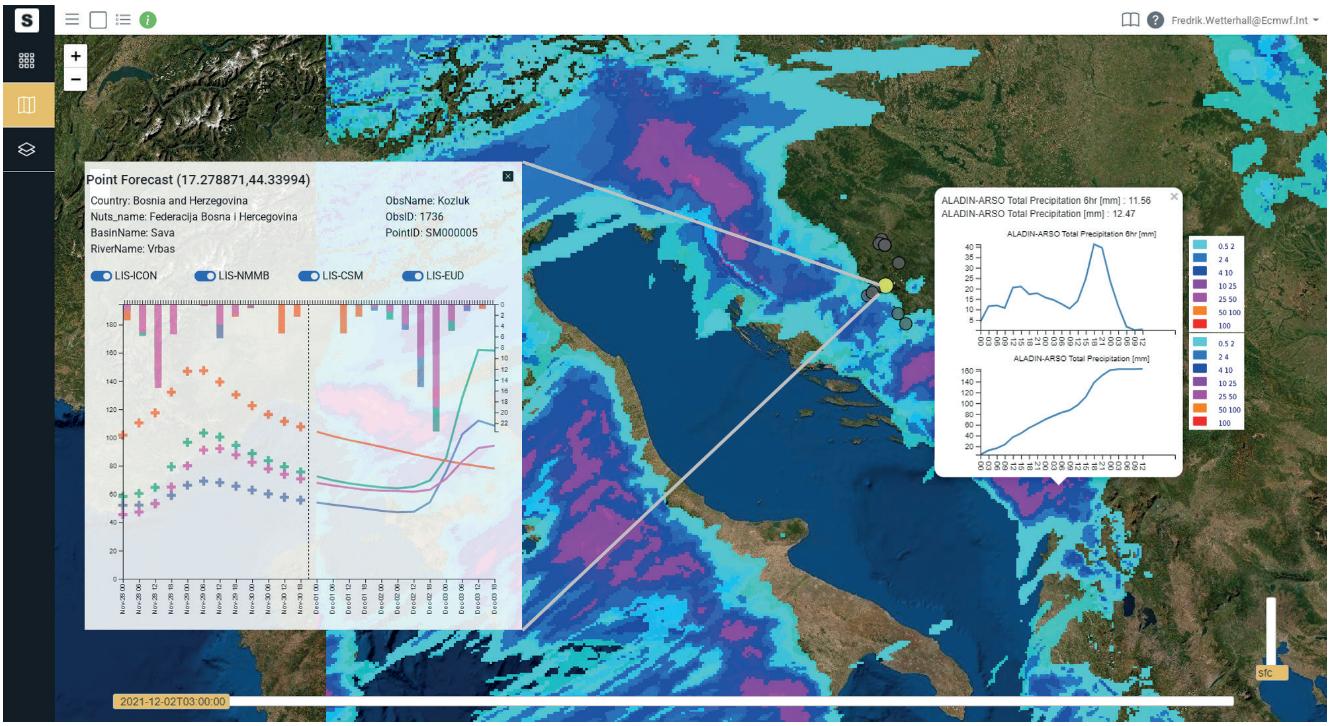


**FIGURE 6** Screenshot of the SEE-MHEWS-A CIP. The example shows a map of the forecast of total accumulated precipitation, as well as precipitation in 6-hour intervals visualised every 3 hours, with the ALADIN model forecast for the next 72 hours, issued on 2 December 2021 at 00 UTC. The graph on the left shows the discharge forecast in $m^3/s$ for the hydrological model LISFLOOD initiated on 1 December 2021 at 00 UTC, driven by NWP models.

- ECMWF has intensified its collaboration with several NMHS.
- There is increased visibility in other forums related to observational exchange (e.g. EUMETNET Supplementary Data Hub, other WMO activities).
- ECMWF has potential opportunities to convince countries in the region of the benefits of joining as Co-operating States.
- Additional observational data (especially temperature and precipitation) become available from stations in the region.
- The use of SAPP is improved, including at the Turkish Meteorological Service for its own data assimilation and the provision of observations to ECMWF as part of the project.
- The Member State share of HPCF resources is more fully used due to several NWP suites running at ECMWF, including time-critical ones.
- The project allows the testing of new technologies in web services as part of the development of the CIP.

The SEE-MHEWS-A project has been successfully implemented by the different participating countries, with ECMWF support. The various components of the workflow have been implemented in the most suitable ECMWF facility and re-using ECMWF software and systems such as ECPDS, SAPP, and ecFlow.
The project has demonstrated how a fairly complex Member and Co-operating State application can be developed and operated using all services available to ECMWF countries.

The third and final phase of the project, the implementation of a fully operational system, is planned to start in 2022. The project has proven to deliver a system that provides good output, but there is still a lot of work to do to make it a fully fit-for-purpose system for forecasters, especially regarding the development of a decision support tool that will meet forecasters' demands.

## Acknowledgements

The implementation has been made possible through the interaction of several NMHS in ECMWF Member and Co-operating States. The project is led by Milan Dacic and Sari Lappi from the WMO (*https://public.wmo.int/en/projects/see-mhews-a*) and is supported by the World Bank, the U.S. Agency for International Development/The Office of U.S. Foreign Disaster Assistance (USAID/OFDA), the European Union and the Global Facility for Disaster Reduction and Recovery (GFDRR). From ECMWF, in addition to the authors of this article, Volkan Firat, Christopher Barnard, Cinzia Mazzetti, Damien Decremer, Francesca Moschini and Xavier Abellan were involved.

## Further reading

**Haiden**, **T.**, **A. Kann**, **C. Wittmann**, **G. Pistotnik**, **B. Bica**, & **C. Gruber**, 2011: The Integrated Nowcasting through Comprehensive Analysis (INCA) System and Its Validation over the Eastern Alpine Region, *Weather and Forecasting*, **26**(2), 166–183. doi.org/10.1175/2010WAF2222451.1

**Smith**, **P.J.**, **F. Pappenberger**, **F. Wetterhall**, **J. Thielen Del Pozo**, **B. Krzeminski**, **P. Salamon et al.**, 2016: On the Operational Implementation of the European Flood Awareness System (EFAS), in: *Flood Forecasting: A Global Perspective*, Academic Press, 313–348. doi.org/10.1016/B978-0-12-801884-2.00011-6

# ECMWF's new data centre in Italy

Matteo Dell'Acqua, Jeremy Thomas, Michele Toni, Andrew Gundry

**After an international tender, on 22 June 2017 representatives of ECMWF Member States approved the proposal by the Italian Government and the Emilia-Romagna region to host ECMWF's new data centre in a refurbished former tobacco factory in Bologna, Italy. The Emilia-Romagna region provided the land, the existing building and the design, and it promised full delivery of an operational data centre as part of the regeneration and refurbishment of the former tobacco factory site. Following a tender process, the region appointed a consortium of Italian contractors, who started work on site in December 2018. By the end of 2021, all construction work had finished and a new Atos BullSequana XH2000 high-performance computing facility (HPCF) had been installed on the premises (for details on the HPCF, see Hawkins & Weger, 2020).**

Construction work included refurbishment of the old buildings to house five diesel rotary uninterruptable power supplies and the electrical infrastructure as well as data halls and offices; erecting a new building, known as the L2 building, to house the chillers and pumps; digging two tunnels connecting this building to the main building; and constructing a new electrical sub-station. Towards the end of 2020, all the main infrastructure work to the data centre site had been completed and the building services and systems almost finalised despite the COVID-19 pandemic. By the end of April 2021, most tests were completed, allowing ECMWF to start the deployment of the network and the installation and connection of servers.

In June 2021, partial handover from the region to ECMWF took place. This enabled the Centre to have control of systems and processes, ensuring managed resilience to the data centre, whilst ECMWF continued its installation of information and communications technology (ICT) systems and the start of the HPCF build in Bologna. During September, the data centre was formally opened by representatives of ECMWF and the Italian authorities.

## Mechanical and electrical infrastructure

The electrical and mechanical systems have been designed to ensure resilience and reliability in accordance with ECMWF tender requirements and operational needs. Each system and each component are redundant, so that in the event of failures or maintenance activities, the power supply and cooling to the computer rooms is guaranteed.

**FIGURE 1** The new data centre in Bologna.

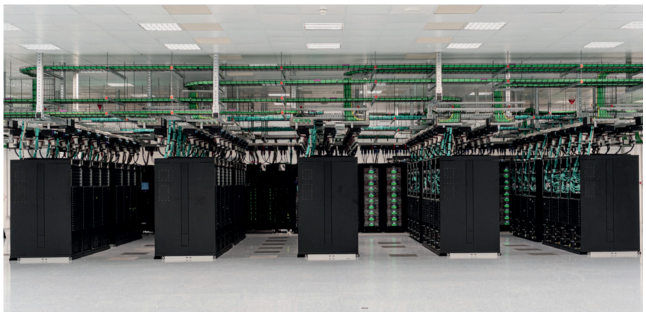

**FIGURE 2** The data centre houses ECMWF's new HPCF.

As for electricity, the site utility supply allows for two fully rated electrical 10 MW circuits. Two points of delivery, a main one and a spare one, are powered from two different lines. They are connected within a closed loop to five transformer substations, which feed the five 2MW diesel rotary uninterruptible power supply (DRUPS) units. The outputs of the DRUPS are in turn connected to an 'isolated ring': in this way, the data centre has five power supply branches, and if one of them fails, the other four support the load of the site automatically. From the five DRUPS, through the critical distribution boards, busbars and cables, the energy reaches the computer rooms. Each element of ICT equipment is powered by at least two of the five branches except the HPCF racks, which are powered by three branches, giving extra resilience.

The data centre cooling system relies on three sources, operating alternately and together, based on the season: geothermal wells, adiabatic dry coolers and chillers. To minimise the environmental impact and reduce operational costs to ECMWF, the system uses ground


doi: 10.21957/mb69s17x3d

| Project meetings | Monitoring meetings | Handover documents | Administrative acts | Regional laws | National law | National Decree |
|---|---|---|---|---|---|---|
| 136 | 39 | 7,685 | 144 | 4 | 1 | 2 |

| Companies | Workers per day – on average | Workers per day – peak | Electrical cables | Concrete | Steel |
|---|---|---|---|---|---|
| 80 | 180 | 382 | 50 km | 48,000 tons | 2,000 tons |

**TABLE 1** Some numbers illustrating the size of the project.

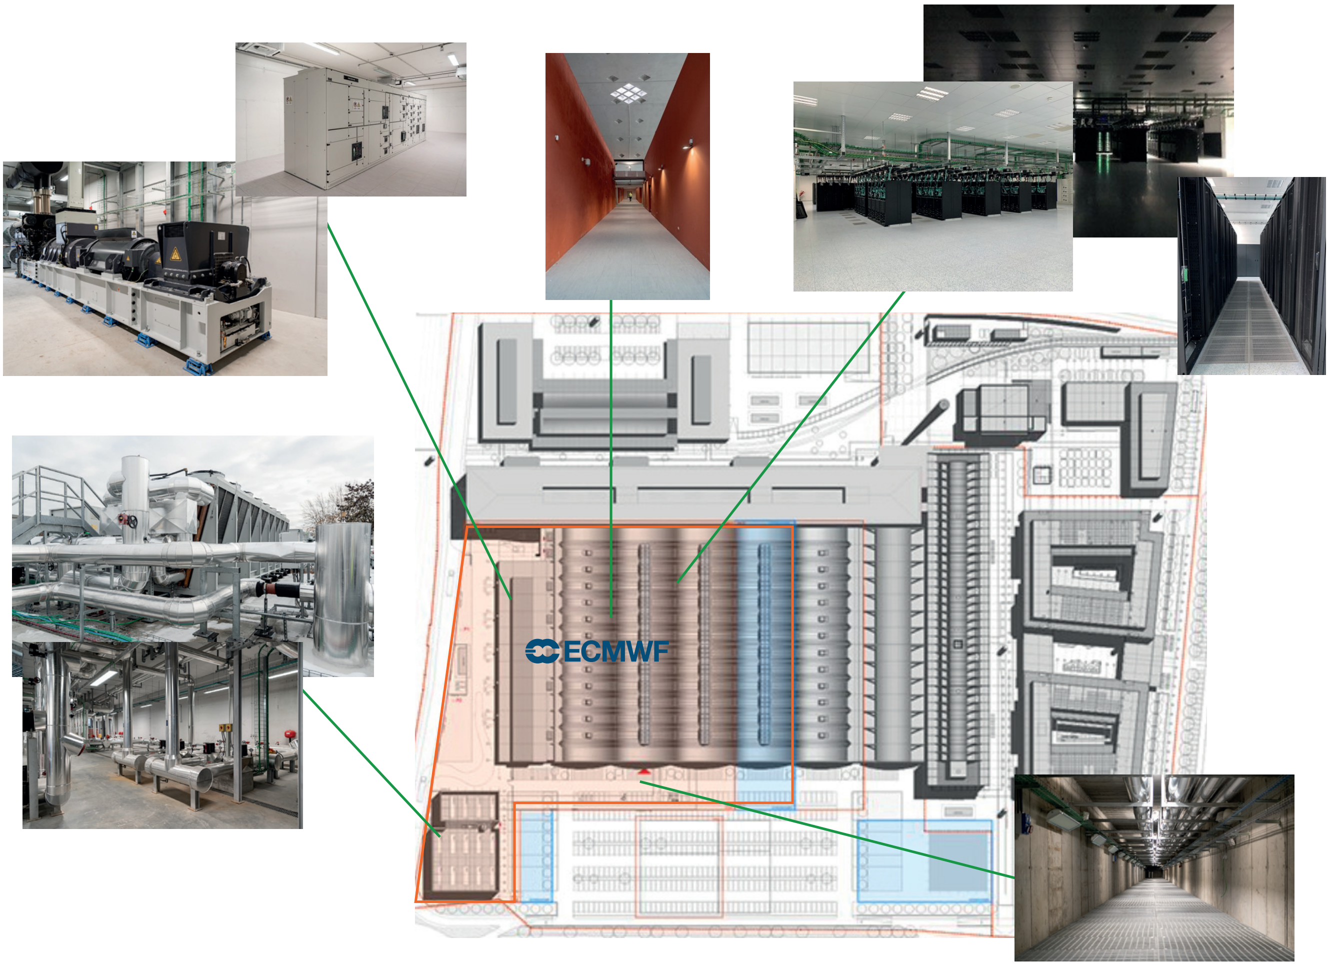


**FIGURE 3** Starting on the bottom left, the images show the L2 building with chillers on the roof and the mechanical plant inside; followed above by the newly refurbished technical block showing the DRUPS and one of the switchboard rooms; to the right a 100 m long corridor and the HPCF in the data halls; and to the bottom a tunnel bringing in cold water to the data halls.

water to support the data centre cooling for part of the year. So far, the wells have been dug and the pumps installed, with final wiring and connections almost finalised. Geologists are assessing the wells' performance before issuing ECMWF with the necessary licence to bring them online.

The site has two thermal power plants, with four chillers, two dry coolers and two geothermal well systems each, as well as a reserve chiller as a further degree of redundancy. The two mechanical plants, through two pumping stations and two underground tunnels, bring chilled water from the L2 building to the computer rooms. If needed, the plants can be connected, so that one also cools the thermal load of the other providing further flexibility.

| | Reading | Bologna |
|---|---|---|
| **Power available** | 5.6 MW @ 430 V | 10 MW @ 15 kV |
| **Total cooling capacity** | 7,500 kW @ 42°C outdoor | 7,800 kW @ 45°C outdoor |
| **DRUPS** | 4 DRUPS 1400 kW each | 5 DRUPS, 2 MW each |
| **Chillers** | 10 chillers: 2x700 kW, 1x850 kW, 7x750 kW | 9 chillers: 865 kW each |
| **Data halls** | 1x1000 $m^2$, 1x500 $m^2$ | 2x1000 $m^2$ |
| **Data storage** | 1x250 $m^2$ | 2x560 $m^2$ |

**TABLE 2** Brief comparison of the old ECMWF data centre in Reading, UK, and the new one in Bologna, Italy.

To optimise energy efficiency and therefore reducing consumption, the system has three operating stages: during the winter season, only the adiabatic dry coolers will work. When the external temperatures become too high, the geothermal wells will start operating and will gradually take over the entire plant. In the hottest season, chillers will be used to respect the groundwater withdrawal limits of the Italian environmental authority and provide the cooling capacity for the data centre.

The infrastructure is equipped with security systems for early warning fire detection and extinguishing (both inert gas and water mist, based on the characteristics of the room being protected), access control, anti-intrusion system and security cameras.

Systems are monitored and managed via a Building Management System (BMS), which can monitor all the main components and provide periodic reports of consumption. Aspects of environmental sustainability have not been neglected, with the data centre having been certified at Platinum level according to LEED (Leadership in Energy and Environmental Design). LEED is a set of rating systems for the design, construction, operation and maintenance of green buildings.
The certifications come as Silver, Gold and Platinum, and the highest Platinum certification indicates the highest level of environmentally responsible construction with efficient use of resources.

## The Bologna Tecnopolo

The Tecnopolo former tobacco factory in Bologna is the most recent node of the Regional Network of technopoles, which also includes ten other places that host research, innovation and technology transfer laboratories. It covers an area of over 100,000 square metres entirely dedicated to science and technology and their applications. The conditions of use make particular reference to supercomputing and data and their use for advancing knowledge, the improvement of the lives of citizens, and the competitiveness of the economic system. Bologna Tecnopolo, closely connected with the academic system and Italian and European research, offers enough flexibility to meet future ECMWF expansion requirements.

In addition to the ECMWF data centre, the Tecnopolo will host from 2022 the pre-exascale EuroHPC supercomputer Leonardo managed by Cineca and the new Data Centre of the National Institute of Nuclear Physics (INFN), which needs a new location to process the huge amount of data from European experiments on nuclear physics.

The construction of the headquarters of the National Agency for New Technologies, Energy and Sustainable Economic Development, as well as that of the Competence Center Industry 4.0, dedicated to the use of Big Data in industry, are also planned at the Tecnopolo site.

Finally, a new building is being built specifically for international research centres in the fields of supercomputing, climate, Earth observation and human development that would like to have their headquarters at the Tecnopolo. These include an advanced stage of development of the project to host an office of the United Nations University (UNU) on Big Data and artificial intelligence for the management of human habitat change.

ECMWF's new data centre in Bologna is the first project to be completed at the Tecnopolo site. It will soon be part of a broader structure at the forefront of technologies in supercomputing, Big Data and artificial intelligence (AI). This will support ECMWF on its path towards its long-term strategic goals.

*ECMWF would like to thank the Italian government and the Emilia-Romagna region for their role in building the new data centre. A video on the Bologna Tecnopolo becoming a leading place for supercomputers, in Italian with English subtitles, is available here: https://www.youtube.com/watch?v=96TfXHCWxf8*

## Further reading

**Hawkins**, **M.** & **I. Weger**, 2020: HPC2020 – ECMWF's new High-Performance Computing Facility, *ECMWF Newsletter* **No. 163**, 32–38.

# Aviso: ECMWF's data availability notification service

Claudio Iacopino, Tiago Quintino, James Hawkes, Baudouin Raoult (all ECWMF), Mikko Partio (FMI), Martin Grønlien Pejcoch (MET Norway), Tomas Karlsson (SMHI)

**To ensure the timely exploitation of its data, ECMWF has developed a system, called Aviso, designed to notify users when real-time forecast data or derived products are available, and to trigger user-defined workflows in an automated fashion. The Aviso service has been available since the beginning of 2021 to users of the European Weather Cloud (EWC). A few ECMWF Member States have already started exploring its advantages by executing time-critical workflows directly in the EWC. This article gives an overview of the Aviso service and showcases some of these workflows. Aviso is openly available on GitHub: *https://github.com/ecmwf/aviso*.**

## Motivation

Recent investments in cloud-based platforms have attracted a growing number of consumers of ECMWF data. An example of these initiatives is the EWC, an ECMWF–EUMETSAT cloud infrastructure built on ECMWF premises. EWC users wish to run automated, real-time tasks or workflows close to the latest data produced by the model in the ECMWF high-performance computing facility (HPCF), thus avoiding costly data transfers out of the data centre. This trend is likely to increase with the exponential growth of weather forecast data.

The real-time forecast stream currently produces about 120 TiB of raw weather data every day, while derived products amount to about 45 TiB of disseminated data per day. It is expected that in the next few years, taking into account resolution upgrades and more complex model physics, the raw forecast data will exceed a petabyte per day.

From an operational perspective, such a large amount of data and users requires an efficient way of synchronising the users' workflows with ECMWF's forecast schedule. A mechanism is needed to timely notify the consumers of specific data availability in a scalable manner and provide the capability to automatically trigger their workflows based on this data. ECMWF's current systems allow users to access the data only in specific time periods defined by the forecast schedule. Moreover, they are not designed to provide notifications to a large number of users. Aviso allows users to synchronise their workflows directly with the availability of the data.

## Overview

Aviso is a scalable notification system designed for a high throughput of notifications. It is developed by ECMWF with the following objectives:

- Notifying data availability events
- Domain-specific – speaks ECMWF meteorological language
- Triggering user workflows
- Supporting a semantic **when** <this> … **do** <that> …
- Persistent history and ability to replay events
- Independent service from HPC or cloud environments
- Protocol-agnostic
- Highly reliable and built for time-critical applications.

Aviso is developed with the intention of being generic and applicable to various domains and architectures. Notifications are designed to provide a lightweight message. ECMWF uses Aviso to provide notifications about data availability. Each notification is composed of a set of metadata required to identify and address the data availability event and a field to inform about the data location (e.g. URL). Moreover, these metadata are expressed using ECMWF Meteorological Archival and Retrieval System (MARS) language. This provides a richer experience to the users as they can reuse the same MARS keys used to define products and access data to subscribe to events of data availability.

Aviso is a client-server application. We refer to the notification server as *Aviso Server* and to the client application as *Aviso Client*, or simply Aviso. Aviso Server is a micro-services application deployed on the ECMWF Kubernetes infrastructure. This guarantees high levels of reliability, elasticity and scalability. The solution is based on a key–value store implementation where data availability events are persisted: the key represents the event metadata, while the value contains the data location. Unlike queue-based technologies, this key–


doi: 10.21957/c59dk312pi

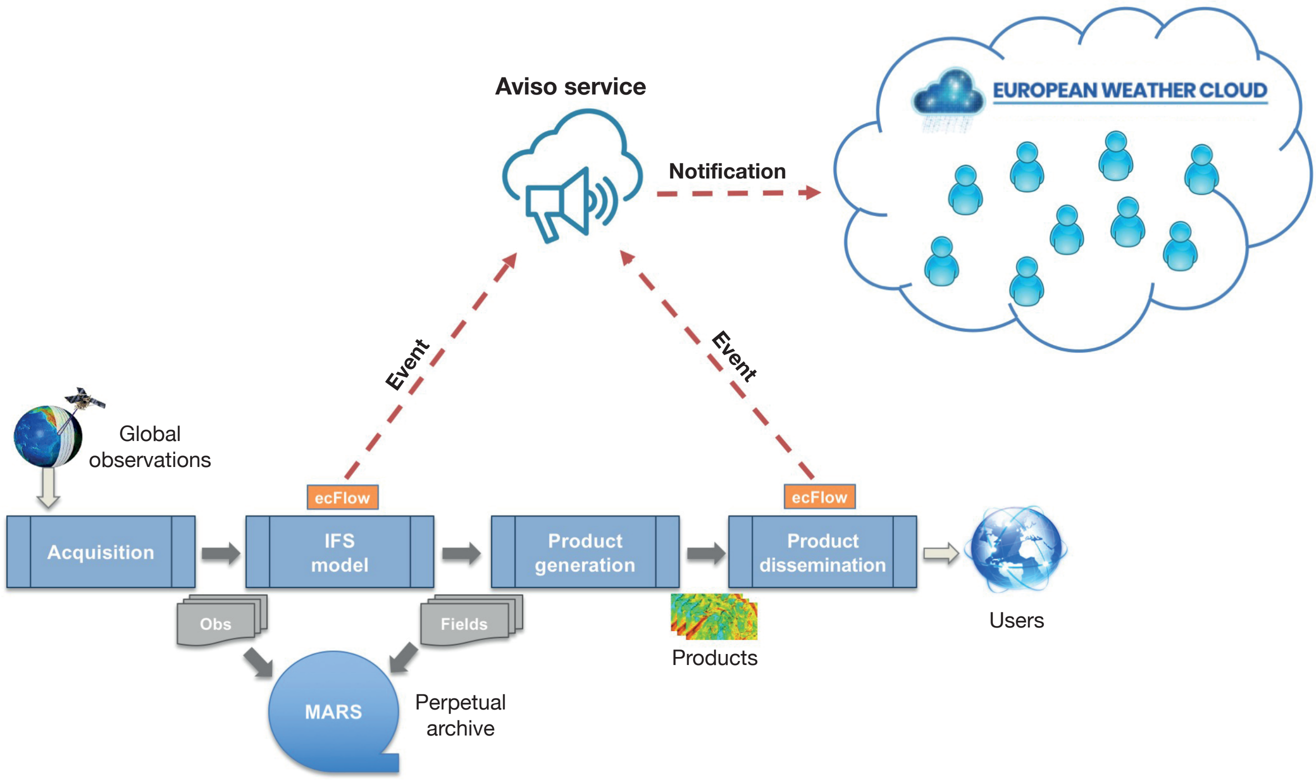


**FIGURE 1** Events submitted to Aviso from the ECMWF data flow.

value store makes it possible for notifications to be replayed, and it enables catch-up mechanisms which are more flexible for consumers and greatly increase the reliability of the system.

The remainder of the article focuses on Aviso Client, which is more relevant to the user community.

## Notification events

ECMWF has deployed Aviso as a notification service for the availability of the data produced by the Centre. Figure 1 shows the ECMWF data flow. It starts from the data assimilation of observations, followed by the generation of model output. The latter is a time-critical step relevant for users' workflows, and therefore Aviso is notified when this data is available in the HPC fast storage (MARS/Fields Database (FDB)). The data flow continues with the generation of derived products that are then disseminated via ECMWF's dissemination system. The delivery of these products to their individual destinations is also notified to Aviso as users depend on custom products for their downstream applications.

Aviso thus currently offers notifications for the following types of events:

- User products delivered via ECMWF's dissemination system. Each event of this type is related to one product disseminated to the user destination. The event contains the URL to access the product notified together with a set of metadata identifying the product.
- Availability of data in the ECMWF HPC fast storage (MARS/FDB). Each event of this type is related to one step of the forecast model output. The event carries all the metadata identifying these model step data. It does not contain its location because users will have to access it via the conventional MARS API.

## How to use Aviso

This section gives a short introduction on how to use Aviso Client. Detailed user documentation is available on Read the Docs: *https://pyaviso.readthedocs.io*

Aviso Client is a Python package available and installable from PyPI: *https://pypi.org/project/pyaviso/*. It can be used as a Command Line Application (CLI) or as a Python API, to integrate into users' own tools. In both cases users must subscribe to an event and program a trigger associated with the event. This is done through a listener configuration. A listener is composed of the following elements:

- **event –** the type of event to listen to
- **request –** a dictionary of keys identifying the specific events to be notified
- **triggers –** a sequence of processes to be executed once a notification is received

```yaml
listeners:
  - event: dissemination
    request:
      destination: FOO
      stream: enfo
      step: [1,2,3]
    triggers:
        - type: command
          command: my_script.sh
          environment:
            STEP: ${request.step}
```


**FIGURE 2** Aviso listener configuration defining the event to listen to and the trigger to execute.

Figure 2 shows a YAML file used as an example of this configuration. This is a basic example of a listener to dissemination events, identified by the keyword *dissemination*. In Aviso, we use event types to distinguish the nature of the events. The currently available types of events are *mars* and *dissemination*. A *mars* event notifies when a new step of the forecast model output is available in the HPC fast storage (MARS/FDB). A *dissemination* event notifies when a user's product is delivered via ECMWF's dissemination system to the predefined user destination. The event type determines how Aviso will interpret and validate the request section.

The *request* section allows users to define the *When* condition. It contains specific keys that are a subset of the ECMWF MARS language. The MARS key describes the data to which the event is related. Aviso delivers only the notifications that match the keys defined. Users can narrow or expand this selection by adding or removing keys, respectively.

The *triggers* section allows users to define the *Do* condition. It contains a list of *triggers*. Each *trigger* defines an independent action executed every time a notification is received. A few types of triggers are available, including logging, execution of a Linux shell command as an independent process or posting the notification to an external location. The latter makes it possible to submit events in other cloud ecosystems, such as Amazon, Azure, etc., enabling workflow execution across different clouds or data centres. The notification can be sent to Amazon Simple Notification Service (SNS) topics or to HTTP endpoints implementing the CloudEvents standard. In the example, the script *my_script.sh* will be executed.

Using Aviso as Python API is intended for users who want to integrate it into a bigger workflow written in Python, or who simply have their trigger defined as a Python function. Figure 3 shows an example of a Python script that defines a function to be executed once a notification is received, creates a listener that references this function, and finally passes it to Aviso to execute. The function in the example, called *my_mars_request*, performs a MARS request to retrieve the data notified. Note how the notification metadata that Aviso provides is then used to build the MARS request. Aviso can achieve this seamless integration by virtue of being a domain-specific notification system for meteorological data.

## Catching up on missed notifications

Before listening to new notifications, Aviso will by default check what was the last notification delivered for

```python
from pyaviso import NotificationManager
from ecmwfapi import ECMWFService

# define function to be called
def my_mars_request(notification):
    mars_server = ECMWFService("mars")
    request = notification["request"]
    request.update({
        "type": "fc",
        "levtype": "sfc",
        "param": 167.128})
    mars_server.execute(request, "my_data.grib")

# define the trigger
my_trigger = {"type": "function", "function": my_mars_request}

# create an event listener request that uses that trigger
my_request = {"stream": "enfo", "date": 20190810, "time": 0, "step": 0}
listeners = {"listeners": [{"event": "mars",
                            "request": my_request,
                            "triggers": [my_trigger]}]}
# run it
aviso = NotificationManager()
aviso.listen(listeners=listeners)
```

**FIGURE 3** Example of using Aviso Python API to define a function to be executed once a notification is received.

that user, and it will then return all the notifications that have not yet been delivered. It will then carry on listening and deliver new ones. This behaviour ensures high reliability because, even if the listening process is interrupted due to networking faults or system unavailability, the users will not miss any notifications.

Users can also choose to listen only to new notifications or replay past notifications where available. Replaying past notifications is also a useful mechanism for testing the listener configuration with real notifications.

## Pilot workflows

This section showcases three workflows implemented by Member State users in the EWC. These workflows are pilot applications. They aim to explore and demonstrate the potential in using Aviso as a notification system and triggering mechanism for downstream applications of weather and climate data that exploit the data proximity of the EWC.

### *MET Norway workflow*

MET Norway has always been keen to make ECMWF data available as soon as possible for downstream processing and utilisation. In 2016, it developed a system broadcasting an event every time a model time step is received from the dissemination system, triggering further actions. Dissemination goes through the Secure File Transfer Protocol (SFTP), and MET Norway uses the *inotify* Linux subsystem to detect a change. With Aviso, this task has become easier. Aviso provides these events, and ECMWF disseminates the data via an Amazon Simple Storage Service (S3) object store in the EWC.

Figure 4 shows the new set-up. Aviso Client is used in a proof-of-concept application, which converts GRIB files to NetCDF and serves them using an OPeNDAP server (THREDDS). Systemd is the orchestrator used to run Aviso and schedule conversion jobs using a language-agnostic background job system called Faktory (*https://github.com/contribsys/faktory*). The notifications include a full S3 path, which is reachable from wherever the conversion system is running.

This set-up makes use of templating in the configuration file to structure commands for the processing system. This feature makes it possible to parametrize post-processing programs and does not pose any requirements on the languages or libraries used. The implementation of the pipeline described is available at *https://github.com/metno/agnc*.

To leverage the flexibility of S3, MET Norway is planning to add a conversion to the Zarr format and distribute the data further by pushing it back to the S3 object storage.

### *Finnish Meteorological Institute workflow*

The Finnish Meteorological Institute (FMI) is planning to use Aviso to integrate an existing Amazon Web Services (AWS) processing infrastructure with a data processing pipeline deployed at the EWC. The overall workflow serves the following purposes:

- Converting GRIB data to a format used by an FMI Weather Information and Forecast Production System (SmartMet) workstation
- Exporting data to Geoweb as Web Map Service (WMS) layers
- Enabling product development through an interactive platform.

Figure 5 shows the overall concept of the pipeline at the EWC. A preliminary version of this pipeline is deployed as a Kubernetes cluster hosted by the EWC. Aviso is used to signal whenever a derived product has been disseminated to the FMI S3 bucket at the EWC. From this notification, a JSON message is sent to an Amazon SNS topic at the AWS cloud. By doing this,

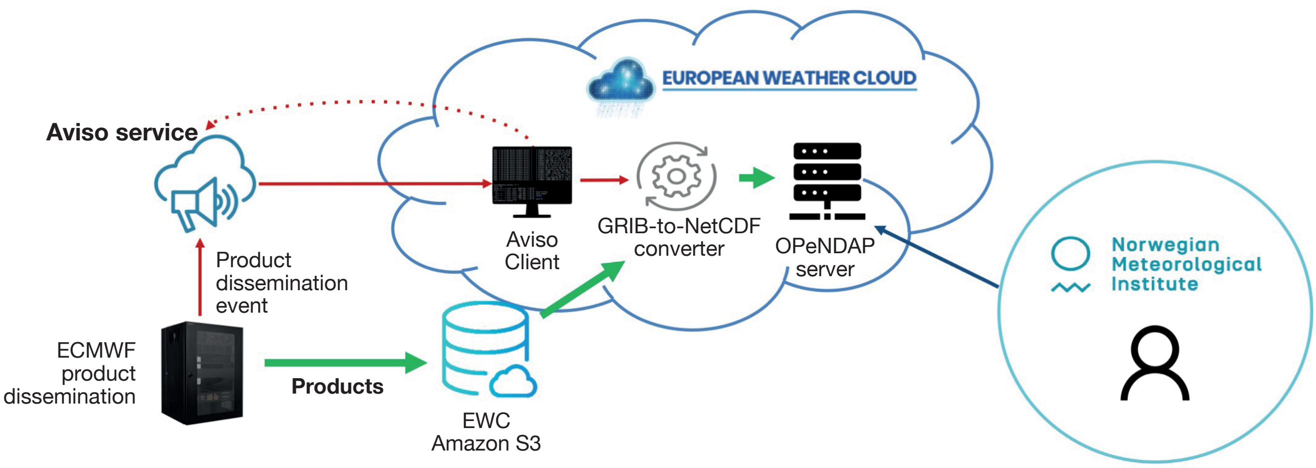


**FIGURE 4** MET Norway workflow using Aviso with the EWC.

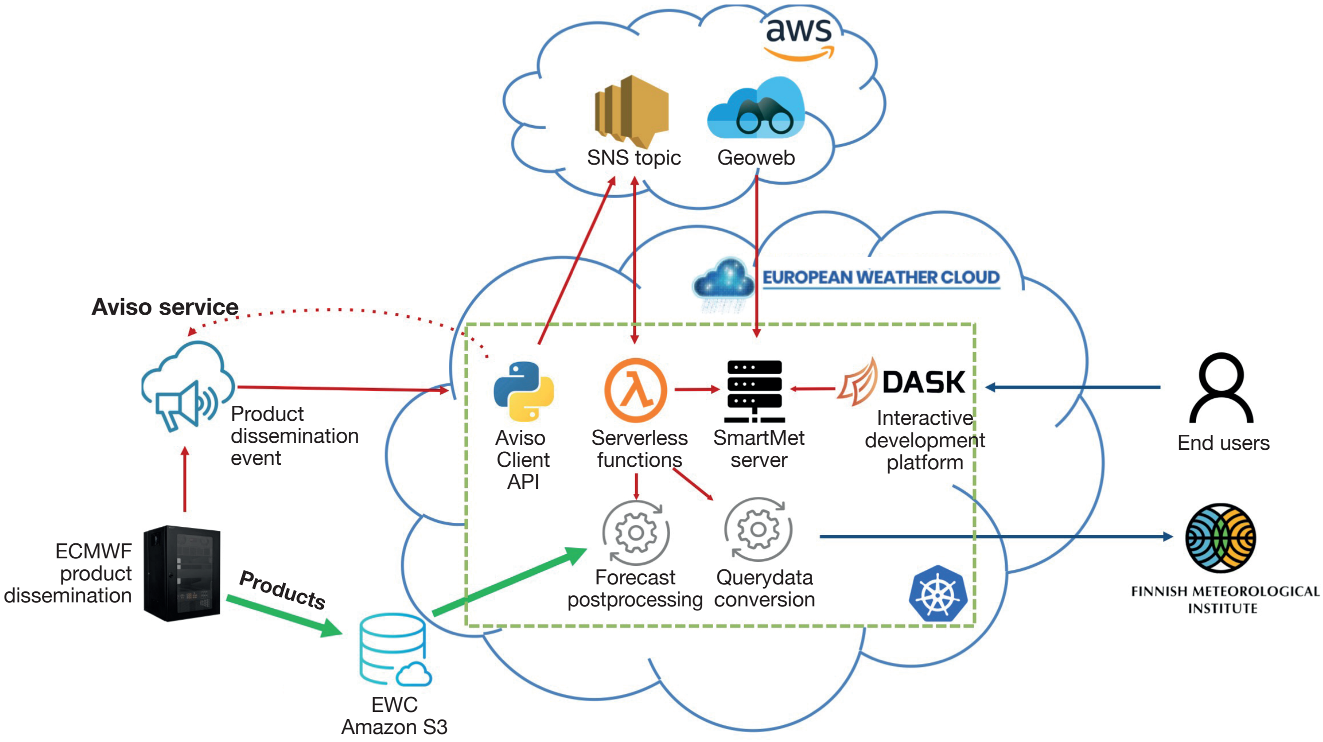


**FIGURE 5** Finnish Meteorological Institute workflow using Aviso with the EWC.

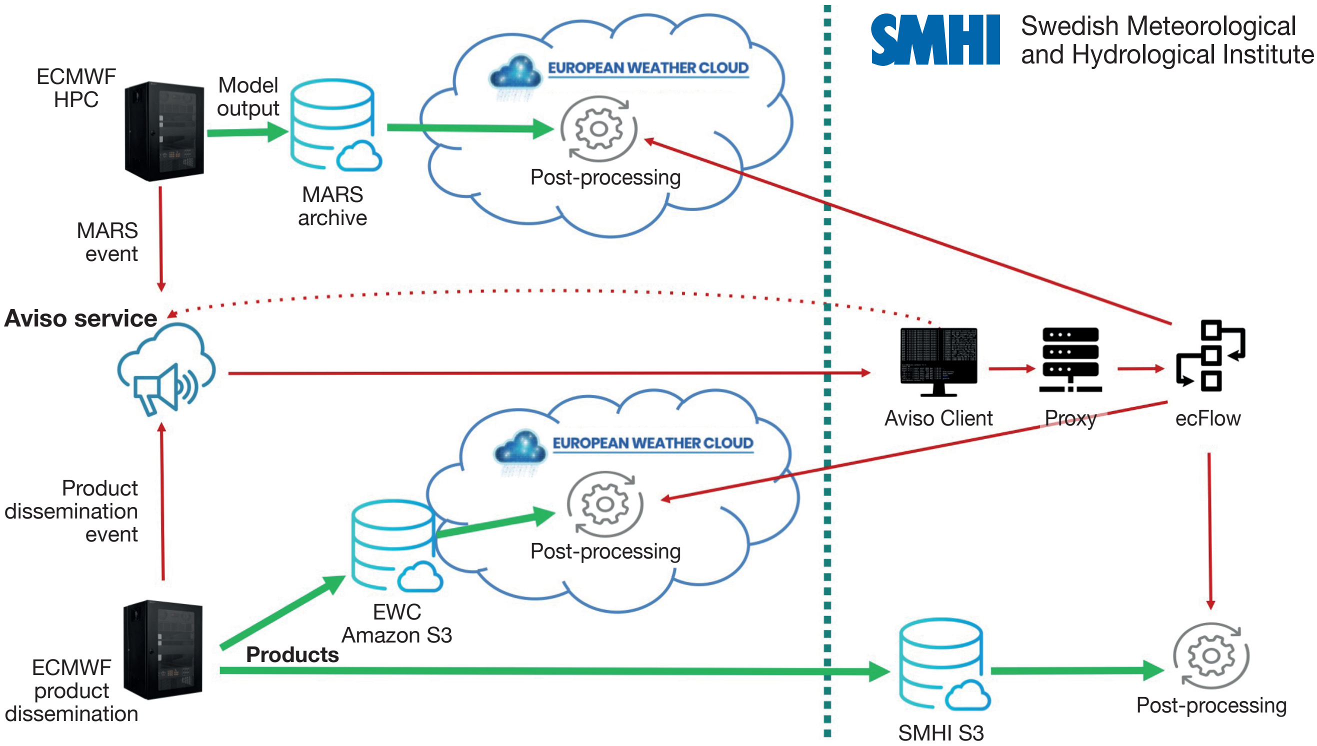


**FIGURE 6** Swedish Meteorological and Hydrological Institute (SMHI) workflow using Aviso with the EWC.

FMI is connecting the EWC pipeline with an AWS processing infrastructure previously developed; the SNS topic is part of that infrastructure. The message is then delivered back to the EWC, where a serverless function implemented by OpenFaaS is reading it and starting a subsequent processing step, dependent on the message contents.

One important task is forecast postprocessing, for example producing probability products with specific thresholds from ensembles. Another task is to convert GRIB data to a file type supported by the SmartMet workstation used by FMI meteorologists. As part of the FMI pipeline, a SmartMet server instance is running continuously at the EWC, providing WMS layers directly from GRIB data. These layers are read by a Geoweb instance running at AWS.

Another service is offered by a Dask instance running at the EWC, which provides an interactive development platform in the form of Jupyter notebooks to enable development of new products and algorithms. Resulting data is exported to the SmartMet server and visualised with Geoweb.

### *Swedish Meteorological and Hydrological Institute workflow*

The Swedish Meteorological and Hydrological Institute (SMHI) currently uses a central ecFlow workflow package to control and supervise all data production pipelines that involve the retrieval and processing of derived products from the ECMWF Production Data Store (ECPDS) and model output data from the MARS archive. In this architecture, Aviso can fill a void as currently there is no control flow support for new data available in MARS or ECPDS. At present, active polling is used to trigger the various pipelines, which puts unnecessary load on the servers. Figure 6 illustrates how Aviso would be used instead. Some preliminary tests have been conducted in running Aviso Client at SMHI and to use events from Aviso to set events in SMHI's ecFlow service. The plan is to trigger postprocessing jobs in both the EWC and in the SMHI private cloud from the central ecFlow at SMHI. This would make it possible to use the current infrastructure to control and also supervise production pipelines located in external cloud infrastructures like the EWC.

## Outlook

Aviso is currently pre-operational, and it is receiving nearly 400,000 notifications a day of both model output data and derived product availability. Users of the European Weather Cloud have started to explore the potential of this new service by building responsive data processing pipelines. Aviso is, however, not confined to the EWC infrastructure and can offer notifications across different data centres and cloud environments thanks to the adoption of cloud standards such as CloudEvents and AWS SNS topics.

Aviso is planned to acquire operational status in 2022, together with IFS Cycle 48r1. Several activities are taking place to achieve this target, including the set-up of the service support teams, lines of communication, operational procedures (including disaster recovery), training and the implementation of monitoring and analysis infrastructures. With respect to users, the intention is to gradually open the service to a larger audience while also integrating it with other ECMWF operational systems.

*This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreements No. 825532 and 824115.*

## Further reading

**Pappenberger**, **F.** & **M. Palkovič**, 2020: Progress towards a European Weather Cloud, *ECMWF Newsletter* **No. 165**, 24–27.

**Gougeon**, **L.**, 2019: The ECMWF Production Data Store, *ECMWF Newsletter* **No. 159**, 35–40.

# New Python data exploration functions in Metview

Iain Russell, Sándor Kertész, Marcus O. Köhler

**For 30 years, ECMWF's Metview software has been providing straightforward ways to access, process and visualise ECMWF forecasts and related data. Many workflows involve the exploration of large datasets. This can include data inspection, filtering and plotting. This type of work has traditionally been performed using Metview's Graphical User Interface (GUI), but in recent years new environments such as Jupyter notebooks have become popular, especially for teaching in virtual environments. This article describes new features in Metview's Python interface designed to meet these needs.**

## Working with multiple GRIB files

Grid-based and spectral output fields from numerical weather prediction (NWP) models, such as ECMWF's Integrated Forecasting System (IFS), are produced in binary GRIB format. When a GRIB file is read into Metview using the Python programming language, the data object returned is called a *Fieldset*. Metview has always had features for combining multiple GRIB files into a single Fieldset, but recent developments have made it even easier. The Fieldset constructor can now take a list of paths to GRIB files, or even a wildcard. The following example shows two ways of loading multiple GRIB files into a Fieldset:

```
import metview as mv

my_gribs = ["/path/data_a.grib", "/path/
data_b.grib"]

fs1 = mv.Fieldset(path=my_gribs)

fs2 = mv.Fieldset(path="/path/*.grib")
```

The features described in the rest of this article apply to all Fieldsets, whether obtained from single or multiple GRIB files. They also apply to *datasets*, described later.

## Exploring the contents of GRIB data

It is not immediately obvious exactly what data lie in a collection of GRIB files. Metview's GUI provides a tool to quickly examine the contents of a GRIB file. In a non-GUI Python environment, such as Jupyter notebooks, new functions provide these capabilities.

The most straightforward one is the ls() function. This prints details of each and every field in a Fieldset, with customisable columns (see Figure 1). For larger datasets, the describe() function groups the GRIB messages by parameter and gives a summary for quick inspection. It can also be used to delve deeper into an individual parameter (see Figure 2). In both functions, the output is customised to take advantage of the formatting capabilities available in a notebook environment.

## Plot widget for Jupyter

When using Metview's GUI, data are plotted by default into an interactive plotting widget called the Display Window, which provides many features to navigate through the different fields. In a Jupyter environment, a new plotting widget was developed to allow the user to easily navigate the fields. Plot frame selection is available via the slider control, and continuous playback

```
fs[:4].ls()
```

| Message | centre | shortName | typeOfLevel | level | dataDate | dataTime | stepRange | dataType | gridType |
|---|---|---|---|---|---|---|---|---|---|
| 0 | ecmf | t | isobaricInhPa | 1000 | 20111215 | 0 | 0 | fc | regular_ll |
| 1 | ecmf | z | isobaricInhPa | 1000 | 20111215 | 0 | 0 | fc | regular_ll |
| 2 | ecmf | t | isobaricInhPa | 850 | 20111215 | 0 | 0 | fc | regular_ll |
| 3 | ecmf | z | isobaricInhPa | 850 | 20111215 | 0 | 0 | fc | regular_ll |

**FIGURE 1** The ls() function displays columns of per-field information; extra keys can be specified by the user.



doi: 10.21957/st60e7fj9k

**a** Entire Fieldset

```
fs.describe()
```

| parameter | typeOfLevel | level | date | time | step | paramId | class | stream | type | experimentVersionNumber |
|---|---|---|---|---|---|---|---|---|---|---|
| q | isobaricInhPa | 100,250,... | 20111215 | 0 | 0,6,... | 133 | od | oper | fc | 0001 |
| t | isobaricInhPa | 100,250,... | 20111215 | 0 | 0,6,... | 130 | od | oper | fc | 0001 |
| tp | surface | 0 | 20111215 | 0 | 0,6,... | 228 | od | oper | fc | 0001 |
| u | isobaricInhPa | 100,250,... | 20111215 | 0 | 0,6,... | 131 | od | oper | fc | 0001 |
| v | isobaricInhPa | 100,250,... | 20111215 | 0 | 0,6,... | 132 | od | oper | fc | 0001 |
| z | isobaricInhPa | 100,250,... | 20111215 | 0 | 0,6,... | 129 | od | oper | fc | 0001 |

**b** Individual parameter

```
fs.describe('t')
```

| | |
|---|---|
| **shortName** | t |
| **name** | Temperature |
| **paramId** | 130 |
| **units** | K |
| **typeOfLevel** | isobaricInhPa |
| **level** | 100,250,300,500,700,850,1000 |
| **date** | 20111215 |
| **time** | 0 |
| **step** | 0,6,12,18,24,30,36,42 |
| **class** | od |
| **stream** | oper |
| **type** | fc |
| **experimentVersionNumber** | 0001 |

FIGURE 2 The describe() function can give (a) an overview of an entire Fieldset or (b) a more detailed overview of an individual parameter.

and looping are also possible (see Figure 3). For the purpose of integrating the plots into a saved version of the notebook, the widget can be deactivated.

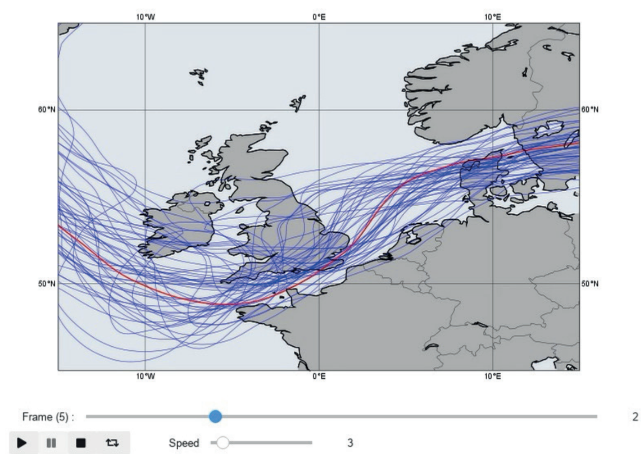


FIGURE 3 The plot widget provides animation controls for Metview's inline plots in a notebook environment.

## New plotting functions

Metview offers a rich plotting interface with fully customisable graphical end products. However, setting up plots like the ones showcased in the Metview Gallery requires several steps, including data management, layout generation and most importantly styling, resulting in lengthy code for some plots. This approach is fully acceptable when a new product is being developed, but in a training environment or when we just want to explore the data, it can be a big overhead. To overcome this difficulty, a high-level plotting interface was added to Metview's Python interface. It comprises a set of functions, each capable of generating a certain type of plot in a single line of code using predefined settings. These include plot_maps(), plot_diff_maps(), plot_xs(), plot_stamp(), plot_rmse(), plot_cdf(). All these functions take their default style settings from the newly developed style library, which includes predefined map areas, map styles and data visualisation styles (contouring, wind, symbol and

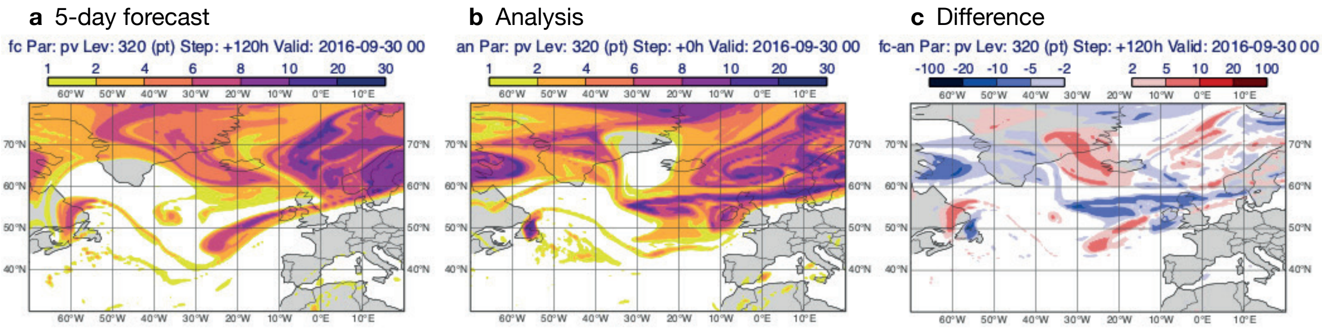


The code to create the difference plot is:

```
import metview as mv
fc = mv.read("fc_pv.grib")
an = mv.read("an_pv.grib")
mv.plot_diff_maps(fc, an,
area=[-30,-70,80,20])
```

**FIGURE 4** The plot_diff_maps() function provides an easy way to plot two Fieldsets and the difference between them. In this case the plots generated by the function are (a) a 5-day forecast of potential vorticity at a potential temperature of 320 K, valid on 30 September 2016, (b) the corresponding analysis, and (c) the difference between them.

storm track plotting) allowing for specialisation per plot type. For map-based plots, the contouring uses the ecCharts style when it is available for the given parameter. However, like all the other style settings, it can be overridden by the user.

Figure 4 shows how plot_diff_maps() can be used to generate a fixed-layout plot showing the difference between a potential vorticity forecast and the analysis using default settings.

## Selecting fields

Some training environments, such as OpenIFS user workshops, require the handling of many thousands of GRIB fields. In these cases, Metview's traditional GRIB filter using the read() function can struggle to select subsets of fields in acceptable time, especially if run multiple times over the same data. To handle these large cases, a new Python-based GRIB filtering function called select() was written. It uses pandas DataFrames to internally index the GRIB fields, and the cost of running it multiple times over the same data is virtually the same as the cost of running it once. An additional difference to the read() filter is that any ecCodes keys can be used in the filter query, making it usable with a wider variety of data than that which comes from the MARS archive.

Another new feature is a convenient shorthand notation for selecting fields from a Fieldset. The Python indexing operator can be used to select fields based on parameter, level and level type. The following pairs of lines are equivalent:

```
z = fs.select(shortName="z")

wind = fs.select(shortName=["u", "v"],
level=500)
```

and

```
z = fs["z"]

wind = fs["wind500"]
```

## Datasets

Metview has been extensively used in OpenIFS training courses to help participants carry out case studies and evaluate their simulations. Gaining from this experience and building on the new Python components discussed in this article, the dataset concept was introduced to Metview. A dataset is a collection of data files packaged together with a tailored style library, which, combined with a set of Python scripts or notebooks, can form a curated environment for a training course. The GRIB data in a dataset is always pre-indexed, offering optimised access to the fields. We can use describe() to inspect the contents and select() to extract data from it. When a dataset is loaded, its style library automatically becomes the current one and high-level plot functions like plot_maps() will automatically pick up the settings from it, guaranteeing that the plots look right by default and do not require further customisation. Datasets are not distributed with Metview but can be regarded as external resources that can be freely created, distributed and shared.

## Use in new NWP training course

In 2021, ECMWF offered a new training course: A Hands-on Introduction to NWP Modelling (Keeley et al., 2021). The practical modelling exercises were to a large part based on a case study featuring a storm event which took place in 2016. The participants were given the task to apply changes to some of the model's physics parametrizations and to investigate how this would affect the simulated development of the storm and its associated local weather impacts.

In group work, the participants were encouraged to use their own forecasting knowledge to select appropriate meteorological parameters to assess the storm development as modelled by OpenIFS. Correspondingly, a wide range of model fields and diagnostics needed to be written out by the model to offer the opportunity to look at different aspects. Given the time constraints during the training course, it was necessary to be able to process and visualise all parameters and diagnostics in a quick and efficient way and allow for comparison between different forecast experiments. Also, due to the virtual delivery of the course within a Jupyter notebook environment, and having limited resources for each training lab instance on the European Weather Cloud, the visualisation needed to be Python-based, easy to use and quick with low computational requirements, and it had to produce inline figures in the Notebooks.

The new Python-based Metview functions were used by each course participant to combine the output from their own reference forecast and sensitivity experiments into a Metview dataset object. Additional GRIB Fieldsets from the operational analyses, for the same forecast time period and interpolated to the same grid resolution, were added to the participant's OpenIFS model results to allow comparison. When additional forecast experiments were carried out later, they could also easily be added to the existing dataset.

The plotting functions developed for the dataset came with reasonable default presets. Without any prior experience with Metview, a wide variety of plots of instantaneous and accumulated model fields could be quickly created with different styles (regional mask, map projections, selected time steps for animation). Participants who wished to experiment with the style presets had the option to modify the default settings by editing a small number of configuration files. Difference plots between experiments highlighted changes in the model's representation of the storm caused by altered model physics. The time slider allowed the animation of the storm evolution within the Jupyter notebook inline figures.

## Future

With greater availability of ECMWF data through an open data policy, it is more important than ever to have straightforward facilities to inspect, filter, process and visualise data within a standard workflow. Raw model output data, which can be of greater volume and spatial complexity, should also fit into the same workflow. By evolving developments such as those described above, we will continue to support the Python community to make the best use of all our data.

## Availability

Metview is installed on ECMWF's computing platforms, including the Member and Co-operating State server ecgate. For use external to ECMWF's computing platforms, Metview is also available as a binary installation on the conda platform, available through the conda-forge channel with one of these commands:

```
conda install metview -c conda-forge
```

```
conda install metview-batch -c conda-forge
# no GUI installed
```

The Python interface can be installed with either of these commands:

```
conda install metview-python -c conda-
forge
```

```
pip install metview
```

Metview's Python documentation has been moved to readthedocs and can be accessed here: *https://metview.readthedocs.io/en/latest/*. Jupyter notebooks illustrating the concepts described in this article can also be found there.

## Further reading

**Keeley**, **S.**, **X. Abellan**, **G. Carver**, **M. Köhler**, **S. Kertész** & **I. Russell**, 2021: Making training mobile, *ECMWF Newsletter* **No. 168**, pp. 12–13.

# 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** Dr Daniel Gellens *(Belgium)*

**Vice President** Prof. Penny Endersby *(UK)*

## 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** Mr Eoin Moran *(Ireland)*

**Vice Chair** Ms Virginie Schwarz *(France)*

## 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** Dr Gisela Seuffert *(Germany)*

**Vice Chair** Mr Lukas Schumacher *(Switzerland)*

## 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** Dr Inger-Lise Frogner *(Norway)*

**Vice Chair** Dr Alain Joly *(France)*

## 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** Dr Sarah O'Reilly *(Ireland)*

**Vice Chair** Ms Anne-Cecilie Riiser (Norway)

## 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** Mr Paolo Capizzi *(Italy)*

**Vice Chair** Ms Monika Köhler *(Austria)*

## 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 Nir Stav *(Israel)*

**Vice Chair** Dr Kornélia Radics *(Hungary)*

# ECMWF publications

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

**Technical Memoranda**

888 **Magnusson**, **L.**, **S. Majumdar**, **R. Emerton**, **D. Richardson**, **M. Alonso-Balmaseda**, **C. Baugh et al.**: Tropical cyclone activities at ECMWF. *October 2021*

# ECMWF Calendar 2022

| | |
|---|---|
| Feb 7–10 | Training course: Use and interpretation of ECMWF products |
| Feb 14–18 | Radio Frequency Interference 2022 workshop |
| Feb 28–Mar 4 | Training course: Data assimilation |
| Mar 7–11 | Training course: Satellite data assimilation |
| Mar 14–18 | Training course: High-Performance Computing – Atos |
| Mar 21–25 | Training course: Predictability |
| Mar 28–31 | CAFE workshop |
| Mar 28 | MAELSTROM Dissemination workshop (virtual) |
| Mar 28–Apr 1 | Training course: Parametrization |
| Mar 29–Apr 1 | Machine Learning Workshop (virtual) |
| Apr 12 | Advisory Committee for Data Policy |
| Apr 25–29 | Training course: Advanced numerical methods for Earth system modelling |
| Apr 26–27 | Finance Committee |
| Apr 27 | Policy Advisory Committee |
| May 3–6 | Training course: Machine learning for weather prediction |
| May 9–12 | Workshop on model uncertainty |
| May 16–20 | Training course: A hands-on introduction to NWP models |
| Jun 7–11 | Using ECMWF's forecasts (UEF2022) |
| Jun 13–17 | Online computing training week |
| Jun 28–29 | Council |
| Sep 12–16 | Annual Seminar 2022 |
| Oct 3–5 | Scientific Advisory Committee |
| Oct 3–6 | Training course: Use and interpretation of ECMWF products |
| Oct 6 | Advisory Committee of Co-operating States |
| Oct 6–7 | Technical Advisory Committee |
| Oct 20–21 | Finance Committee |
| Oct 21 | Policy Advisory Committee |
| Oct 31–Nov 4 | Sixth WGNE workshop on systematic errors in weather and climate models |
| Dec 1–2 | Council |

# Contact information

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

Telephone National 0118 949 9000

Telephone International +44 118 949 9000

Fax +44 118 986 9450

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. 170** | Winter 2021/22
European Centre for Medium-Range Weather Forecasts
www.ecmwf.int