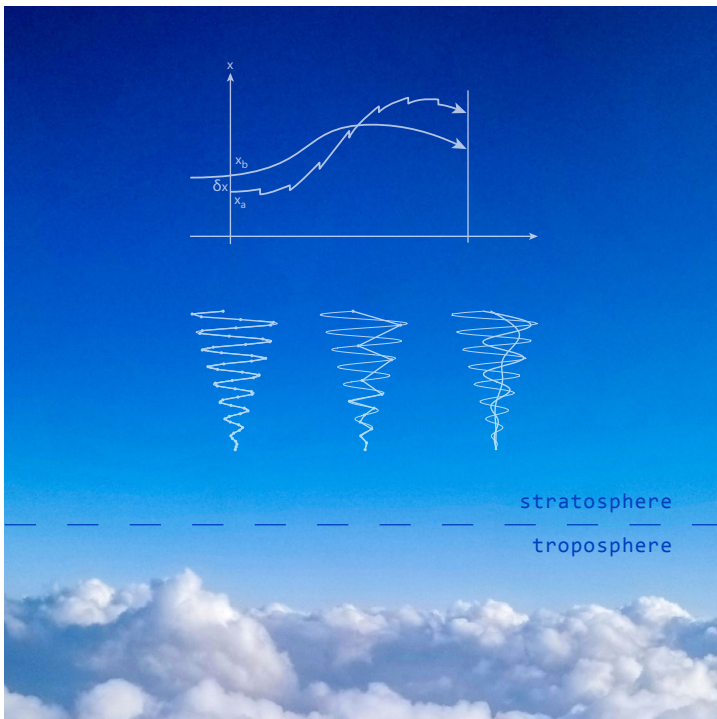


ECMWF Feature article

from Newsletter Number 163 – Spring 2020

COMPUTING

The new capabilities of ECMWF's product dissemination system



Cover photo: batuhanztirk / iStock / Getty Images Plus

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

doi: 10.21957/mz39cf41k7

This article appeared in the Computing section of ECMWF Newsletter No. 163 – Spring 2020, pp. 27–31

The new capabilities of ECMWF’s product dissemination system

Matthias Zink, Meghan Plumridge

ECMWF’s product dissemination system delivers approximately 240,000 post-processed files (about 45 terabytes) daily to ECMWF’s different users. Two important parts of the dissemination system are the Product Requirements Editor, in which users make data requests, and the Product Generation System, which tailors data in line with those requests. A few years ago, a major review of the product dissemination system identified key areas for improvement in these two components in order to better serve our Member and Co-operating States and other users in the future. Addressing these key issues has led to significant gains in computational efficiency, a more reliable and robust service as well as improvements in the user interface including new features. It has also prepared the ground for the migration of the product dissemination system to the new data centre in Bologna. Improvements in a third component of the product dissemination system, the ECMWF Production Data Store (ECPDS), were implemented earlier and are described in a previous Newsletter article (Gougeon, 2019). The new product dissemination system was implemented and is supported by a wide range of ECMWF staff, including user support, development, production and computing specialists.

Motivation

An important issue that motivated the review was scalability, especially when looking at the growth of produced and disseminated data volumes over recent years. Figure 1a shows the growth in dissemination volume over the last 13 years. Some of the bigger jumps can be attributed to model resolution upgrades, as in 2013 and 2016. Others, however, are simply driven by increasing demand for data from Member and Co-operating States and other users. One such jump was caused by making high-frequency products available to commercial users in 2018 (Box A). Overall, we can see exponential growth in the volume of data disseminated between 2007 and 2019. This growth is projected to continue in the future: it is estimated that the volume of meteorological data produced will reach the petabyte scale by 2026/27 (Figure 1b). The volume of data disseminated to users will increase accordingly. This growth, shown in Figure 1b, is an estimate based on expected model resolution upgrades only, without considering possible increasing data demand from users or growth in the user community.

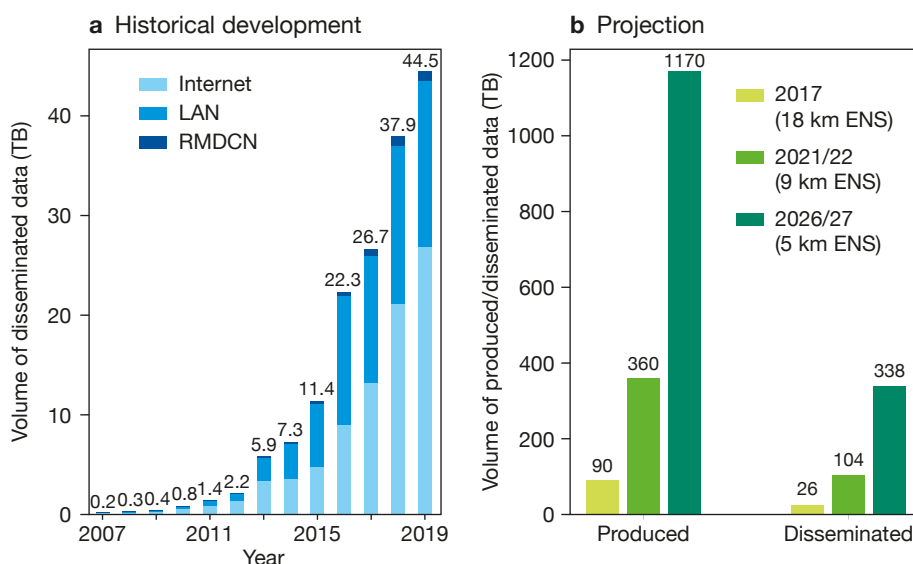


Figure 1 The development of (a) data volumes disseminated via different networks, i.e. the Internet, the Local Area Network (LAN), and the Regional Meteorological Data Communication Network (RMDCN), and (b) the projected volume of all data produced and disseminated by ECMWF’s different forecasts assuming an increase in the resolution of ENS forecasts expected to be implemented by 2026/27.

High-frequency products A

ECMWF’s high-frequency products are hourly weather forecast products that are generated four times a day (06, 12, 18, 00 UTC) 90 hours ahead in addition to the core products, which are generated at 00 and 12 UTC as 3- and 6-hourly data up to day 10 for the high-resolution forecast (HRES) and day 15 for the ensemble forecast (ENS). The high-frequency data come from ECMWF’s Boundary Conditions Optional Programme, which was established in 1991. The programme’s primary purpose is to make such data available to ECMWF’s Member and Co-operating States to use as boundary conditions for their limited-area models. Since October 2018, data from the programme have been available on request to all real-time data licence holders.

ECMWF’s product dissemination system

The previous product dissemination system was first implemented in 1999 (see Jokić, 1999, for details). It has evolved over the years to provide an increasing volume of numerical weather prediction products to a growing number of users. Today it consists of three main components: the Product Requirements Editor, the Product Generation System (pgen) and ECMWF’s Production Data Store (Figure 2). The Product Requirements Editor is a web application which allows users to define their real-time data requests (product requirements). These requirements are picked up before each run of the Product Generation System. Product generation is triggered by the progress of the respective forecast and tailors model outputs according to product requirements. The user-tailored products are then transferred to ECMWF’s Production Data Store (ECPDS), which disseminates them according to a fixed schedule via three available networks: the Regional Meteorological Data Communication Network (RMDCN), the Internet and the Local Area Network (LAN). Currently, we disseminate about 45 terabytes (TB) per day, of which 30 TB are pushed outside ECMWF while the remainder is stored locally. The data disseminated via LAN are usually data for Member and Co-operating States that post-process these products in-house.

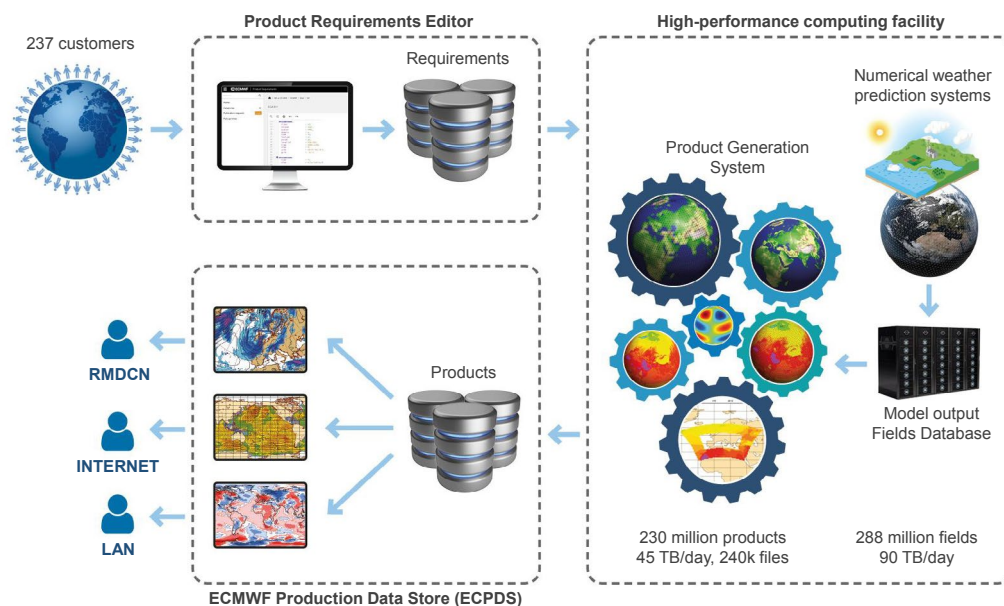


Figure 2 The main tasks of the product dissemination system are run on independent IT infrastructures. The Product Requirements Editor is launched on ECMWF web servers; the Product Generation System, comprising the workflow manager (ecFlow suite) and the product generation application, runs on the high-performance computing facility using the high-performance file system; and ECMWF’s Production Data Store (ECPDS) uses its own dedicated hardware.

The product dissemination system is run four times a day (at 00 UTC, 06 UTC, 12 UTC, and 18 UTC) for high-resolution and ensemble forecasts (HRES and ENS); twice a week (on Mondays and Thursdays) for extended-range ensemble forecasts and re-forecasts; and once a month for seasonal forecasts. The product dissemination system is in the time-critical path to ensure the timely delivery of products to our users. It is closely monitored by ECMWF’s operators and is supported on a 24/7 basis by analysts working on this system. To ensure resilience, the system can be easily transferred within minutes between ECMWF’s two high-performance computing clusters and between the two operational high-performance file systems (Hawkins & Weger, 2015). The functions of the main components of ECMWF’s product dissemination system and the reasons for upgrading them are summarised in Table 1.

Component	Function	Motivation for upgrading
1. Product Requirements Editor	Interface that allows users to configure their data requirements	Develop a comprehensive and robust tool for configuring user requirements easily
2. Product Generation System software	User-tailored post-processing of weather forecast outputs	Better scalability of the application, leading to a more efficient use of computational resources and improved robustness
3. Product Generation System workflow manager (ecFlow suite)	Orchestrates the components of ECMWF’s dissemination system	Easier to understand workflows, allowing ECMWF to respond as fast as possible in the time-critical operational environment

Table 1 Improved components of the product dissemination system.

The new Product Requirements Editor

Users of ECMWF real-time data can configure and maintain their product requirements via the recently developed Product Requirements Editor (PREd), which consists of a web interface and validation software (Figure 3). It has replaced the previous dissemination requirements tool, which had similar functionalities. The new Product Requirements Editor was released to users in February 2020. The PREd is available to ECMWF Member and Co-operating States as well as national meteorological and hydrological service (NMHS) licence holders and maximum charge commercial licence holders.

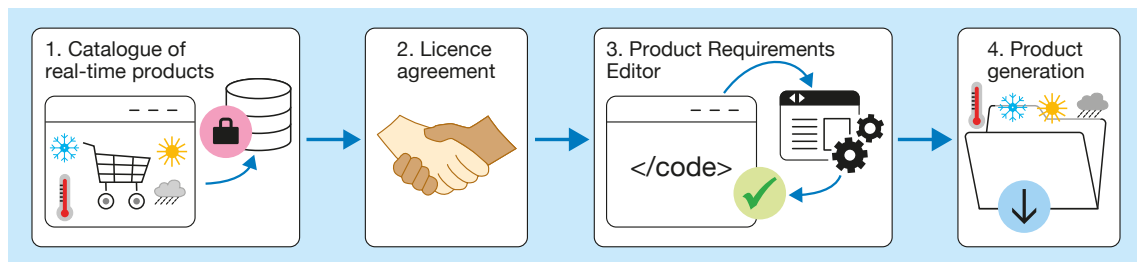


Figure 3 Schematic illustrating how users interact with the Products Requirements Editor. Prospective users can browse ECMWF’s real-time product catalogue. This web-based catalogue reflects the forecast products that are available for dissemination. Before getting access to the products, the user has to acquire a licence. The agreed data is configured in the PREd. The requirements are validated against the available products and cross-checked with the licence. For example, not all users have access to high-frequency products. The Product Generation System uses the product requirements to generate user-tailored data.

The PReD interface includes new features to assist users in the management of their real-time data requirements, including:

- autocompletion
- product requirement templates
- version history
- comparison between current requirements and previous versions
- publication requests for commercial customers.

The new validation software ensures that only accessible forecast data can be configured for real-time dissemination, preventing 'faulty' requirements from being fed to the Product Generation System. The validation tool checks the requirements against several catalogues, depending on the type of user. For example, a commercial customer may not have access to high-frequency products, whereas an NMHS non-commercial user could be licensed to access these additional products.

The Product Requirements Editor, together with future web applications, will streamline the real-time data requesting and delivery process for users and the Data Services Team at ECMWF by allowing users to modify their product requirements. This will improve the turnover for implementing modified or new product requirements in ECMWF's production system.

The new Product Generation System

The Product Generation System consists of two main components. The workflow manager (ecFlow suite) and the product generation software. ECMWF has rewritten both components from scratch, based on the experience gained with the previous system over the last 20 years. The task at hand was challenging as the Product Generation System is very I/O intensive and thus requires significant resources on the high-performance computing facility. The Product Generation System had to be thoroughly redesigned to interact with forecasts such that it will start generating products as soon as the respective forecast step is written to the Fields Database (Figure 2; for more details on the Fields Database, see Quintino et al., 2019). This had to be achieved without adverse effects on the I/O performance of the forecasting system.

The upgrade of the workflow management software involved redesigning the workflow itself and rewriting all scripts which are part of the workflow manager and job scheduler. We used a newly developed in-house python-based library (PyFlow), which handles the generation of the workflow and scheduling scheme (ecFlow suite) as well as the generation of the respective task scripts.

The new product generation software has been developed over the last few years. The aim was to replace the previous product generation software, which was based on the interpolation package EMOSLIB. This interpolation library was decommissioned in our production systems in early 2019. Its successor is the newly developed Meteorological Interpolation and Regridding (MIR) library (Maciel et al., 2017). The new product generation software takes full advantage of this new library and was designed to be more robust and performant. Major performance improvements were achieved by:

- introducing a producer–consumer pattern in the highly parallelised application
- streamlining the processing of data requests, as shown in Figure 2 of Maciel et al. (2017)
- reducing the number of I/O operations due to the implementation of parallelized writing of output files on dedicated CPUs.

These improvements made user-tailored post-processing of different forecasts more scalable and efficient. Comparing the runtimes between the old and the new product generation shows a big reduction in the time required to generate the same set of products, using the same computational resources (Figure 4). The new product generation runtimes for the core HRES and ENS forecasts (00 UTC and 12 UTC forecast runs) are less than a third of the old product generation runtimes. The improvement in runtime for high-frequency ENS products (00 UTC, 06 UTC, 12 UTC and 18 UTC forecast runs, labelled 'bc' in Figure 4) is smaller compared to all other products (all HRES and ENS core products) due to different write rates in the new product generation (write rates not shown). The write rates of the ENS high-frequency products are roughly half those of the latter products. This is likely due to competing demands on the high-performance file system: the ENS high-frequency products are generated in parallel to the HRES forecast and product generation, whereas the product generation of the core ENS products is delayed until the main workload of the HRES forecast is finished. This delay was deliberately introduced to avoid a too high load on the file system, which can potentially slow down all forecasts and product generation. However, runtime is still halved for high-frequency ENS products compared to the old product generation system.

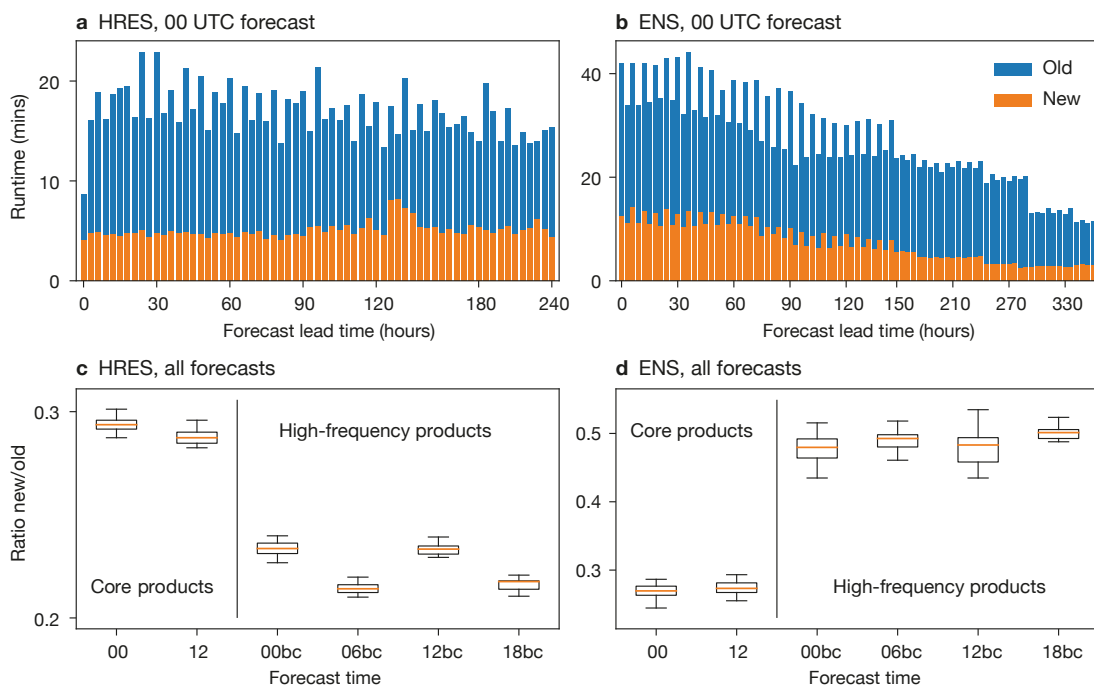
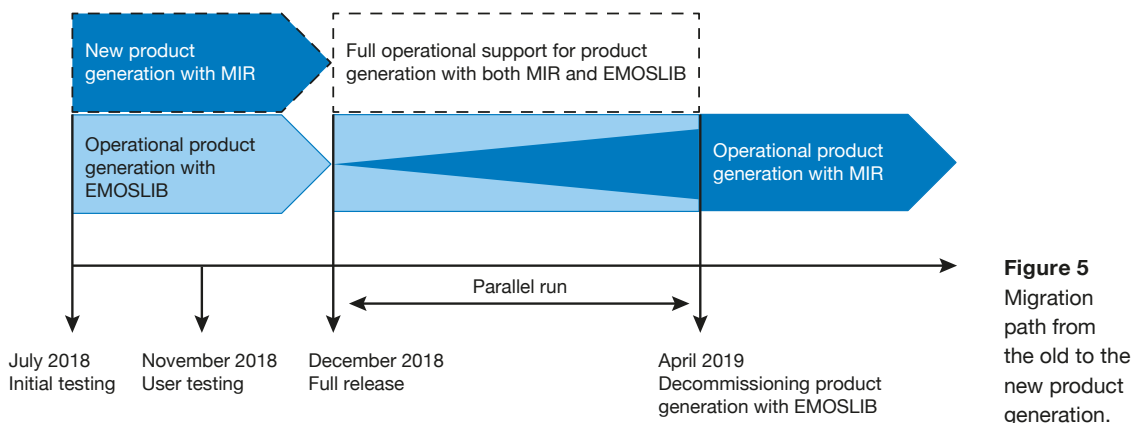


Figure 4 Product Generation System runtime comparisons between the new and the old system, showing (a) HRES product generation runtimes on a representative day (17 December 2018) for the 00 UTC forecast, (b) ENS product generation runtimes for the same day and forecast time, (c) the spread of the averaged ratios of runtimes for all forecast steps of the respective HRES forecast, e.g. 00 UTC as shown in panel a, over 24 days in December 2018, and (d) the same as (c) but for ENS forecasts. The boundary-condition programme/high-frequency data (see Box A) are marked 'bc'.

Seamless migration

An important design criterion for the new Product Generation System was to minimise the impact on users when migrating them from the old to the new system. For that reason, ECMWF dedicated significant resources to running the new and the old Product Generation Systems in parallel over a period of 10 months (Figure 5). The new Product Generation System was integrated into ECMWF's pre-operational environment for internal testing and scalability experimentation in July 2018. Selected users were given early access to the new products in November 2018 in order to assess potential impacts. The system became fully operational in December 2018 after positive feedback from our test users. By incrementally moving users from the old to the new product generation system over a five-month period, we ensured a seamless transition to the new system. Running both systems in parallel was technically challenging as both are heavily I/O intensive, as described above. Additionally, we facilitated parallel dissemination of data via a second instance of ECPDS, thus providing users with the infrastructure to test their systems with data from the new product generation. Upon successful testing and user migration, we phased out the production of backup data from the old product generation. The smooth transition to the new product generation was appreciated by our users.



Outlook

We expect that ECMWF's upgraded dissemination system will provide a more reliable and robust service, decreasing the number of delays in product dissemination compared to the previous system. We will continue to improve the quality of the dissemination service by applying state-of-the-art log management paired with data analytics. This will enable our analysts to identify potential issues early and to identify patterns. The new tools will facilitate a proactive problem-solving approach and minimise negative impacts on our users.

Further reading

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

Hawkins, M. & I. Weger, 2015: Supercomputing at ECMWF, *ECMWF Newsletter No. 143*, 32–38.

Jokić, D., 1999: ECMWF's new data dissemination system, *Seventh Workshop on Meteorological Operational Systems, ECMWF, Conference Paper*.

Maciel, P., T. Quintino, U. Modigliani, P. Dando, B. Raoult, W. Deconinck, F. Rathgeber & C. Simarro, 2017: The new ECMWF interpolation package MIR, *ECMWF Newsletter No. 152*, 36–39.

Quintino, T., S. Smart, B. Raoult, M. Fuentes, M. Zink, S. Villaume, J. Hodkinson, A. Mueller-Quintino, A. Bonet-Cassagneau, O. Treiber & C. Weihrauch, 2019: Upgrade makes Fields Database more resilient, *ECMWF Newsletter No. 160*, 5.

© Copyright 2020

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

The content of this Newsletter is available for use under a Creative Commons Attribution-Non-Commercial-No-Derivatives-4.0-Unported Licence. See the terms at <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

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