

## COMPUTING

### IFShub: a new way to work with IFS experiments



Cover Photo: roberto nova / iStock Editorial / Getty Images Plus

*This article appeared in the Computing section of ECMWF Newsletter No. 167 – Spring 2021, pp. 28–32.*

# **IFShub: a new way to work with IFS experiments**

Paul Burton, Manuel Martins, Stephan Siemen, Michael Sleight

Roughly 500 experiments on ECMWF's Integrated Forecasting System (IFS) run at any given time, comprising the research work of users from the Centre and our Member and Co-operating States. A host of different bespoke tools are needed to enable researchers to manage this experiment workload. Different applications are available to prepare and submit experiments, to monitor the status of running work, to manage the resulting data archived in our meteorological archive (MARS) throughout its lifetime, and to scientifically analyse and evaluate the results. Most of these tools are desktop applications, some with graphical user interfaces, and all are developed independently of each other by different teams. Each is unique in its language, interface, design, dependencies, and architecture.

In particular, PrepIFS is the tool that researchers use to create, check, and store experiment configurations, and to submit those experiments to the high-performance computing facility (HPCF). PrepIFS has been in use for several decades and has stood the test of time, but in recent years it has developed some maintainability problems. The application is almost unique now at ECMWF in being written in Java, in which there is very little expertise at the Centre anymore. This is because nearly all development has moved to Python and C++ in recent years, while the IFS itself is still written mainly in Fortran. Additionally, there is limited expertise in creating user interfaces amongst the scientists developing the IFS, and this task is better performed by application developers who have considerable experience in this area. These kinds of maintainability problems lead naturally to usability problems, as non-essential improvements and bug fixes become too time-consuming to implement and development necessarily reduces to the bare minimum. PrepIFS also uses a bespoke language to define logical rules that govern valid configurations, making it difficult to implement new or improved rules.

There has therefore been an aspiration for some years to replace PrepIFS with a new tool which solves many of these problems, and which is actively developed, tested and maintained. From this starting point, however, it is only a small step to thinking about the whole end-to-end workflow – from starting to prepare an experiment to having it fully evaluated and archiving its results – and how this might also be improved. To meet this challenge, the vision of IFShub was developed. This article will set out how this was done and explain where we stand in developing the new system.

## **Aims of the project**

Based on the shortcomings of existing systems, it was clear that a more coherent organisation of the overall model development workflow could help researchers in their day-to-day work. Especially for staff new to running the IFS, this would mean a faster learning process.

The IFS is increasingly run on different systems to support training and research activities, such as within the OpenIFS project and as part of a digital twin in the EU's Destination Earth initiative. A higher-level interface to the workflow will make it possible to hide some of the technical challenges of running experiments on different infrastructures while still keeping easy access to the outcomes. These developments should ease migrations to different systems and could be shared with other initiatives with similar challenges in the community.

## **Requirements gathering**

It was clear from the beginning that this would be a big multi-year project and involve many stakeholders across ECMWF. To ensure the success of such a project, it is vital to agree on and communicate a vision and its scope. It was decided to use external help, and the company Wavestone was chosen. In the first three months of 2020, two consultants spent time with staff at ECMWF to interview them and organise small workshops to gather requirements and develop a vision.



Figure 1 Impressions from the requirements gathering workshops held by Wavestone.

### IFShub vision

It became clear from user feedback that there is a multitude of tools currently in use with IFS experiments, and that it can sometimes be difficult for users to easily access all the available tools and learn how to use each one of them. Additionally, some tools require a user to be logged on to a specific platform and may require them to copy data across machines.

We concluded that users would benefit from a single access point to all tools which relate to IFS experiments and the data produced by them. These tools should be accessible via a consistent user interface on a web browser, which can be accessed easily from anywhere, across a range of devices.

The technical complexity of a tool's implementation can be hidden behind the web interface. Users will no longer have to concern themselves about which system they need to log in to or where their data must be copied from and to. If hardware platforms or data locations change in the future, users should be little impacted by such changes, as the web interface should remain largely the same. The adoption of the changing underlying platforms will therefore be hidden to the user.

Our vision for IFShub is to provide a web-based interface, accessible from anywhere, allowing ECMWF and external users to control the full workflow of running an IFS experiment, through a consistent and intuitive user interface. IFShub should facilitate and encourage strong cooperation and interaction between users, by allowing them to easily share configurations, data, analysis and associated documentation whenever they wish to.

### Wireframe mock-ups

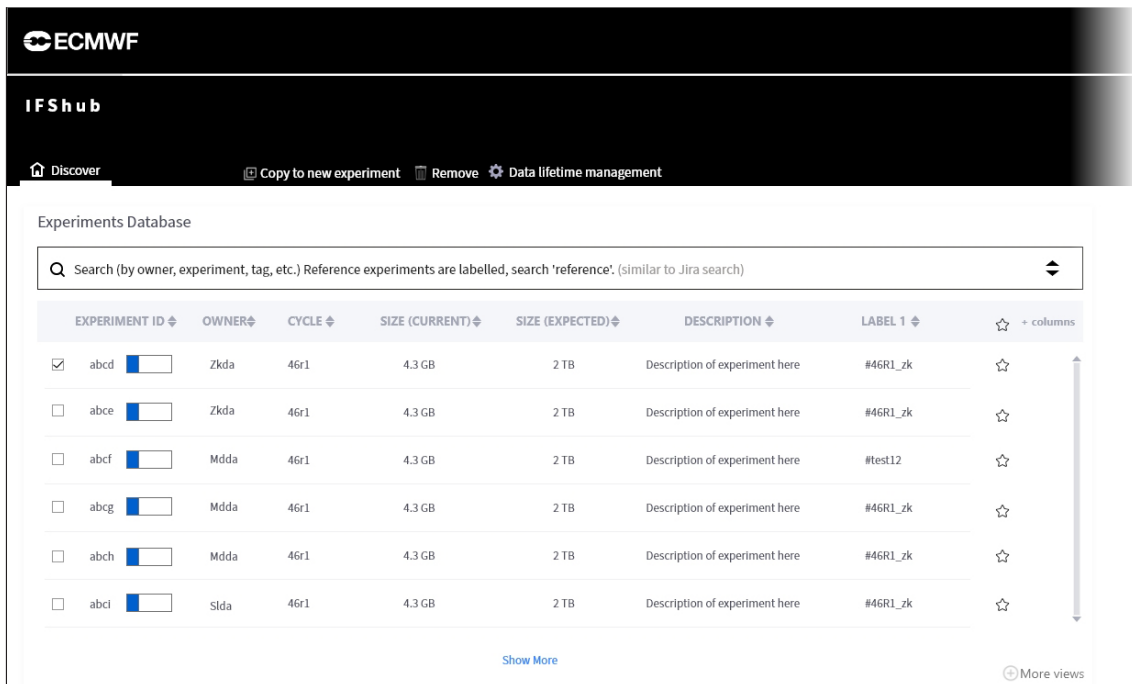
To make this vision more concrete, the next step was to produce some mock-ups of how the user interface could look. Together with our consultants, we identified the key areas of functionality we wanted to reproduce, and new functionality we could introduce. We then considered how this could be presented to users within a web user interface.

This was purely to provide a physical manifestation of the ideas produced from the requirements capture process – it was not a working prototype of a system. Despite this, it was still a very useful process to help us ensure we had captured all the major areas of functionality required in a way that was intuitive for users to interact with. It provides a very useful reference point for feedback and will be an invaluable aid to our developers as they start to develop a working user interface.

We present here two wireframes which illustrate some of the core concepts and functionality we plan to provide with IFShub. Of course, we are still very early in the development of IFShub, so they are still just ideas. They are very useful for generating comments and feedback as we work towards the final implementation, but they are not cast in stone and will no doubt be subject to revision as we develop and test the system and get feedback from our users.

### Experiment discovery

The ‘experiment discovery’ page is central to IFShub: this is where a user can browse and manage the full database of IFS experiments (both their own and those of other users), using a range of search tools (Figure 2). We aim to provide as much useful information about experiments as possible on this page, to enable a user to quickly and easily identify the experiment of interest to them. Once an experiment has been identified, the most common action will be to browse the experiment in more detail in the ‘experiment explorer’, but other useful actions from this page will include the ability to copy the experiment’s configuration to a new experiment, manage a user’s data assets, and remove old experiments which are no longer required.



**Figure 2** The experiment discovery page will give access to the full database of IFS experiments.

Additionally, this page will also include a summary of any useful information about ECMWF computer systems which may impact a user’s experiments – for example any current problems or planned system sessions.

### Experiment explorer

Once a user has identified an experiment of interest to them, they can open it in a new browser tab in the experiment explorer. From this page, the user can manage all aspects of the experiment. The main body of the page provides an overall summary of the experiment (Figure 3). This includes, for example, a progress monitor to quickly alert the user to any problems that need resolving; a summary of any plots produced from the experiment’s data; and an overview of the data assets of the experiment.

On the left side of the page is a menu to access further pages which allow the user to interact with the experiment, the main items being:

- **Configuration** – will open WebPrepIFS, a replacement for the existing PrepIFS application, presented within a web browser. WebPrepIFS will allow a user to comprehensively configure their experiment, check all the settings are sensible and consistent, and then submit the experiment to run on the HPCF.
- **Monitor progress** – will allow users to monitor the ecFlow suite relating to their experiment. ecFlow enables users to run a large number of programs in a controlled environment and is used at ECMWF to run all our operational suites across a range of platforms. Users can use this page to quickly identify any problems which require further attention. They will be able to view log files from failed tasks and restart failed tasks. We do not propose to replace ecflow\_ui with this tool but provide the most important functionality to allow users to monitor and manage running experiments.

- **Visualise data** – will provide a web interface to existing ECMWF tools, such as Metview and ecCharts, and provide the ability to view and manage the plots generated by these tools.
- **Evaluation tools** – provides a range of web-based tools (many of which will be new web interfaces to existing tools) used by researchers to evaluate and compare data from their experiments. Users are given the ability to manage the associated data and plots that these tools generate.
- **Notes** – will integrate with Confluence (the wiki software that hosts ECMWF’s intranet) to enable comprehensive notes and documentation to be kept for each experiment, with automatically added links to access other relevant experiment-related information.
- **Get data** – allowing the user to easily browse the available datasets generated by their experiment and download selections to their current machine.
- **Data lifetime management** – a tool allowing the user to easily view how much space their experiment is using on MARS or ECMWF’s File Storage (ECFS) system and how that data is being accessed. This enables redundant data to be easily identified and marked for removal.

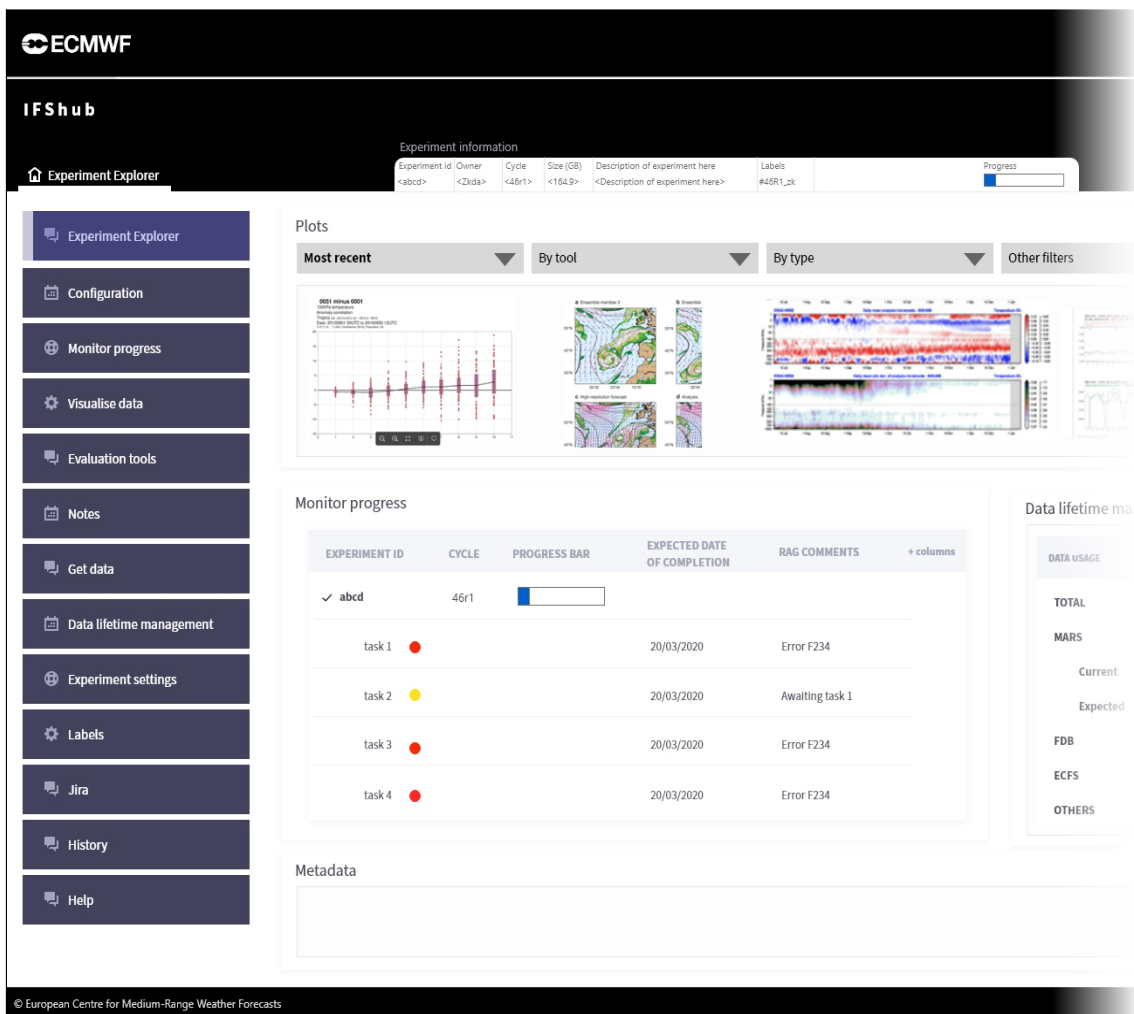
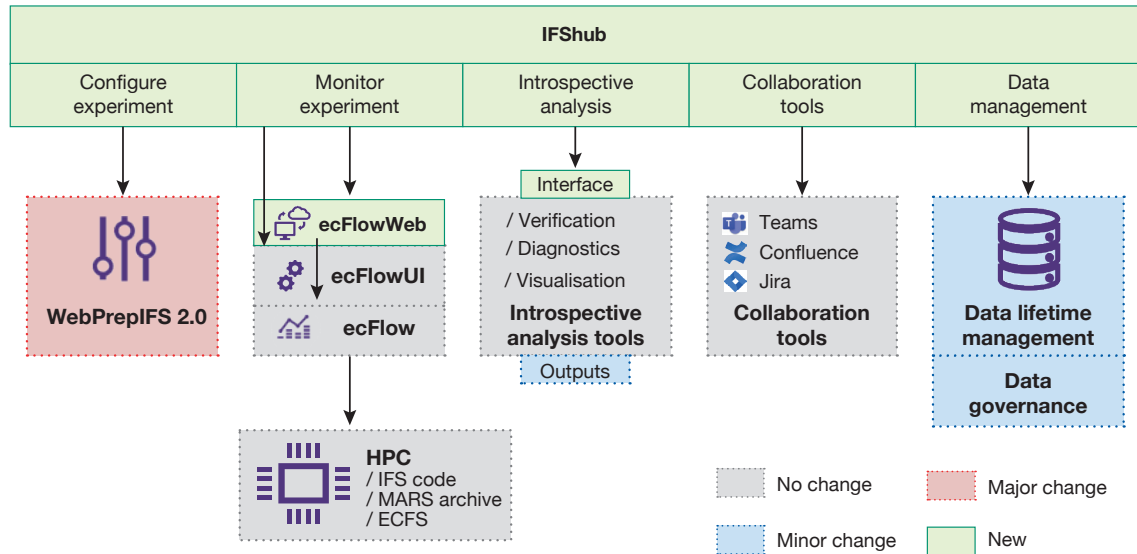


Figure 3 Overall summary of an IFS experiment that can be opened in a new tab. This ‘experiment explorer’ page will give access to all aspects of an experiment in the left-hand menu.

### Implementing the vision

As can be seen from the overall design of IFShub (Figure 4), most components already exist. The two major new developments are a loose web-based framework, which presents a unified interface to the different services, and a new application, WebPrepIFS, which allows users to define and manage their experiments. The latter will be developed in two steps.

The initial WebPrepIFS will be very similar in behaviour to the existing standalone PrepIFS application, so users should be able to easily migrate to this new web-based user interface as soon as it is available. We hope to be able to move all of a user’s existing experiments into the new user interface. This will allow users to have the advantages of the powerful PrepIFS system in an easily accessible web application within the IFShub framework.



**Figure 4** The different components which make up IFShub. Many parts already exist and might need only minor changes and integration.

Following on from this, we plan to start work on the successor, WebPrepIFS 2.0. Plans for this are at a very early stage, but the intention is to improve the user interface, making it easier to use, and to improve flexibility and maintenance aspects. One of the most valuable objectives will be to allow the creation of combinations of existing experiments. This will allow production analysts to use the same tool to set up their configurations, and research scientists to replicate an operational-like configuration, which can then be easily shared between research and operations, thus significantly simplifying and streamlining the research-to-operations handover.

### How will we get there?

With many key project participants busy with the migration of the computer centre to Bologna, it is unlikely that 2021 will see any major developments. Instead, it was decided to use the time to explore various aspects of the future system. For example, trials are underway on how information on experiments can be best stored and shared. As mentioned, the IFShub will link many existing web services. Work has started to explore how this can be achieved in a flexible way to allow services to evolve independently. In 2021, the management team of the IFShub project will also develop a long-term development road map for the next few years.

© Copyright 2021

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

The content of this document is available for use under a Creative Commons Attribution 4.0 International Public License. See the terms at <https://creativecommons.org/licenses/by/4.0/>.

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.