

# The ECMWF research to operations (R2O) process

R. Buizza, E. Andersson, R. Forbes and  
M. Sleigh

Research Department

July 2017

*This paper has not been published and should be regarded as an Internal Report from ECMWF.  
Permission to quote from it should be obtained from the ECMWF.*



European Centre for Medium-Range Weather Forecasts  
Europäisches Zentrum für mittelfristige Wettervorhersage  
Centre européen pour les prévisions météorologiques à moyen terme

Series: ECMWF Technical Memoranda

A full list of ECMWF Publications can be found on our web site under:

<http://www.ecmwf.int/en/research/publications>

Contact: [library@ecmwf.int](mailto:library@ecmwf.int)

©Copyright 2017

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

Literary and scientific copyrights belong to ECMWF and are reserved in all countries. This publication is not to be reprinted or translated in whole or in part without the written permission of the Director-General. Appropriate non-commercial use will normally be granted under the condition that reference is made to ECMWF.

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

## Abstract

One of the key aspects of the European Centre for Medium-Range Weather Forecasts (ECMWF) business is the Research-to-Operations (R2O) process, which is followed to upgrade the IFS<sup>1</sup> cycle (i.e. the software used in forecast production). R2O starts with a strategic, 10-year vision, and ends with the cycles' implementations. It includes a series of actions that could be summarized in 6 activities: planning, development, testing, evaluation, communication and implementation. It involves the calculation of objective metrics to assess model performance, and communication with all key stakeholders.

In this document, we highlight the main phases and activities of the ECMWF R2O process. In Section 1 we present an overview of the 6 key activities of the R2O process, and then in the following Sections we discuss some of them in more detail. In Section 2 we discuss planning and development activities, and highlight their three main time horizons. In Section 3 we review the communication steps that are taken to ensure that all key stakeholders involved in the R2O process are properly informed. In Section 4, we analyze the last phase (the final 8-10 months) of the R2O process that leads to an IFS cycle implementation. In Section 5 we discuss the evaluation and diagnostic activities, and present the criteria, including the headline scores and other metrics used to assess forecast performance, and decide whether an IFS cycle is ready for implementation. In Section 6 we discuss how users' requests are taken into account in the whole process. Finally, on-going activities to review and improve the R2O process are discussed in Section 7, and a summary is presented in Section 8.

## 1 The 6 activities of the ECMWF R2O process

Figure 1 shows a schematic of a framework that can be used to understand and review the R2O process, based on 6 key activities:

1. Planning;
2. Development;
3. Testing;
4. Implementation (into operations);
5. Evaluation (and diagnostics);
6. Communication.

*Planning* starts with a high-level strategic discussion that identifies what are the main areas of development needed to deliver the strategy and achieve its main goals. During this phase, availability of people and computer power are analyzed, and taken into account.

Planning is followed by the *development* of IFS changes and testing. These two activities are strongly linked, with changes first tested in a modular way and then gradually merged:

1. Individual developments are tested stand alone in relevant configurations;
2. Developments are combined and tested in configurations closer to the operational suite;

---

<sup>1</sup>The Integrated Forecasting System (IFS) is used to generate the ECMWF forecasts: this includes the scripts and the main software of the Earth System used at ECMWF, including the suite configuration used in the forecast production (in assimilation and forecasting).

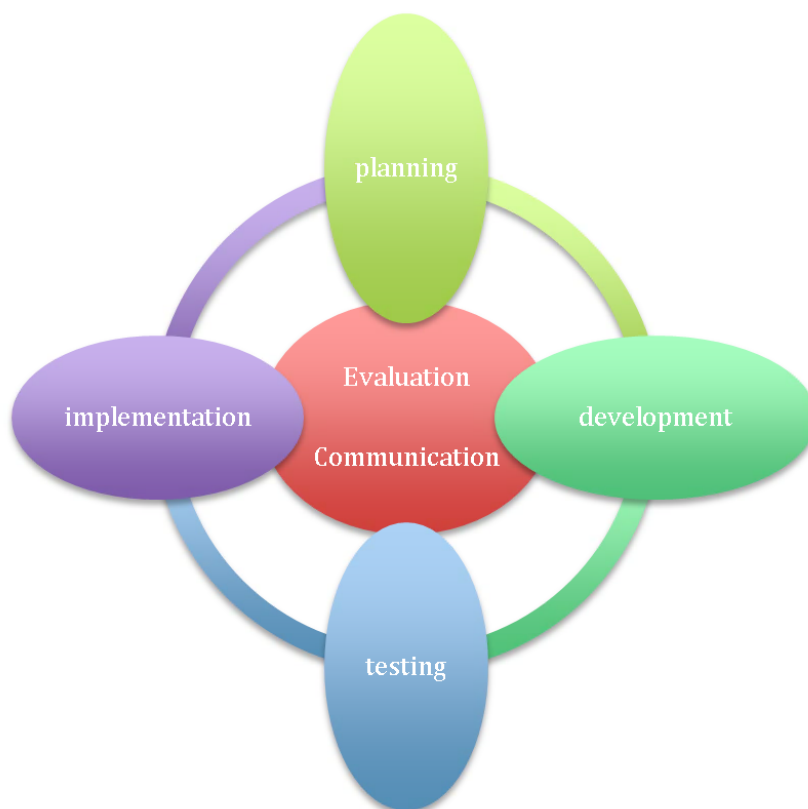


Figure 1: The ECMWF R2O process includes 6 key activities: planning, development, testing, implementation, evaluation and communication. Planning, development, testing and implementation follow each other in a circular and sequential way, while evaluation and communication activities are applied throughout the R2O process to support the other activities, to ensure that decisions are based on objective criteria and proper communication is applied.

3. All developments accepted for implementation are merged in the new IFS cycle and are tested in research experiments;
4. Once properly evaluated and finalized in the Research Department (RD), the new IFS cycle is handed over to the Forecast Department (FD).

In this final stage, the performance of the new IFS cycle is compared to the operational suite<sup>2</sup>, and the impact of the planned model changes on operational products is assessed.

This extensive testing phase is followed by the *implementation* phase, provided that results are judged to be positive overall.

*Evaluation* and diagnostics are key activities applied to assess in an objective way whether the planned changes lead to improvements or deteriorations. They are applied both in the development and testing phases, and routinely to monitor the forecast performance. In this routine phase, they can help to identify

<sup>2</sup>The term 'operational suite' (o-suite) includes the data acquisition, observation quality control and selection, analyses [the Ensemble of Data Assimilation, EDA, and the high-resolution analysis (4D-Var), and the ocean ensemble of analysis (ORAS5)], forecasts [the single high-resolution, and the coupled ensembles up to 46 days (ENS) and 13 months (S4)] and reforecasts (both for ENS and S4).

deficiencies in the forecast, triggering investigations that feed back to future model changes. One key aspect of this phase is the selection of the metrics that are used to assess performance, including tracking the impact of model changes throughout the years, and comparing the impact of model upgrades in the testing phase.

*Communication* activities support the other activities, to ensure that all key stakeholders are properly informed.

## 2 The three main time horizons of the *planning and development* activities: long-term strategic, medium-term and short-term

We can identify three main time horizons in our planning activities and developments, which are also reflected in the way plans and progress are reported to the ECMWF Committees:

- *Long-term strategic horizon*: for developments that may take many years of work, sometimes up to a decade, and might need substantial resources; they are discussed in the 10-year strategy at high-level and are presented in more details in the 4-year plan;
- *Medium-term horizon*: for development work over a 1-4 year time scale; these developments are discussed in the 4-year plan;
- *Short-term horizon*: this time range may include smaller developments or concludes long- and/or medium-term work; it could also conclude an evaluation and diagnostic process that was initiated to solve a specific problem with the forecast system, possibly in response to user feedback; it could also include developments triggered by changes in the computing environment, and/or adaptations to changes in software libraries.

To give some concrete examples, a long-term strategic planning and development activity was the coupling of the atmosphere/land model developed at ECMWF to a dynamical ocean model developed by the NEMO (the Nucleus of the European Modeling of the Ocean) Consortium, which also included an interactive sea-ice model (LIM, the Louvain-la-Neuve Sea Ice Model). This work required a long planning phase, the acquisition of relevant expertise not already present at ECMWF, and implementation of the ocean and sea-ice model components. It included the coupling of the two models to the atmosphere, substantial changes to the ECMWF software infrastructure, significant testing and evaluation, and eventually implementation.

An example of a medium-term strategic planning and development activity was the upgrade of the ECMWF microphysics scheme with the introduction of a 5-species scheme. This required a major development of a part of the ECMWF model that already existed, a reasonable amount of resources, and a few years of development and testing.

An example of a short-term planning and development activity was one of the recent changes in the configuration of the stochastic scheme used in the ENS to simulate model uncertainty. The need for such a modification was identified while testing a new model cycle, as a result of evaluation and diagnostic activities. Work that required a rather small amount of resources and testing (compared to the other two developments discussed above) was planned and conducted rather quickly, and changes were implemented after about 1 year from the time the problem was identified.

Development phases	Period	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Long-term (10-year strategy)	2006-2015	Yellow	Yellow	Yellow	Yellow	Yellow						
	2011-2020	Orange	Orange	Orange	Orange	Orange	Orange	Orange	Orange	Orange	Orange	
	2016-2025						Red	Red	Red	Red	Red	Red
Medium-term (4-year Plan)	2016-2019						Light Blue	Light Blue	Light Blue	Light Blue		
	2017-2020							Blue	Blue	Blue	Blue	
	2018-2021								Dark Blue	Dark Blue	Dark Blue	Dark Blue
	2019-2022									Very Dark Blue	Very Dark Blue	Very Dark Blue
	2020-2023										Very Dark Blue	Very Dark Blue
	2017							Grey				
Short-term (1-year plan)	2018								Grey			
	2019									Grey		
	2020										Grey	
	2021											Grey

Figure 2: Schematic of the links between the long-, medium- and short-term developments, which are presented to the Committees in the 10-year strategy, the 4-year plan, and the 1-year detailed plan.

There is a link between the three development phases and the documents that ECMWF presents to its Committees: the 10-year strategy, the 4-year plan and 1-year detailed plans (Fig. 2). Long-term developments are first identified and flagged, usually in rather general terms, in the strategic planning phases, then more specifically during the development of the 10-year strategy which is presented to the Committees<sup>3</sup> once every 5 years. The 10-year strategy is reviewed every 5 years, to allow for a smooth and effective evolution of the work between two consecutive strategic periods.

Strategic plans are then refined and defined in more detail in the 4-year plan, which is presented to the Committees every year. In the 4-year plan, long-term and medium-term developments are discussed, and the overlap between consecutive 4-year plans again guarantees continuity and a smooth evolution of the work. It is worth remembering that when the 4-year plan is reported, the same document discusses progress in the main areas of work, so that the Committees can keep track of developments that were presented in previous years and how they have been progressing.

Annual plans are defined in even more detail, with deadlines and milestones clearly identified and traceable. The 1-year-detailed plan discusses short-term developments, which are mostly linked with the medium-term developments presented in the previous years' 4-year plan. The 1-year detailed plan also includes developments required to address (known!) problems identified while evaluating and diagnosing the performance of the operational system.

### 3 Communication activities linked to the R2O process

Committees are informed about the plans and work progress regularly:

<sup>3</sup>ECMWF has a number of advisory committees to give external opinions and recommendations including the Scientific Advisory Committee (SAC) and the Technical Advisory Committee (TAC).

- The 10-year strategy is updated and presented to the Committees every 5 years;
- The 4-year plan of activities, partially updated annually, is presented to the Committees every year;
- The 1-year-detailed plan of activities is presented to the Committees every year.

Committees comment and review the ECMWF plans and progress, and ask ECMWF to report on specific issues deemed to be relevant and important. One example is given by the Scientific Advisory Committee (SAC) Special Topic papers, where ECMWF reports progress and plans on specific areas decided by the SAC. Another example is the Technical Advisory Committee (TAC) report on model performance, which discusses the performance of the operational system, verification in the Member States, and the impact of the most recent changes on forecast performance.

External communication relies also on public, peer-reviewed literature, and on ECMWF publications such as the Newsletter and the Technical Memoranda. These publications are also complemented by material accessible via the ECMWF web site. Examples for different areas are available following the links below:

- Support software: <http://www.ecmwf.int/en/computing/software>;
- IFS documentation:  
<http://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/ifs-documentation>;
- Characteristics of the operational suites:  
<http://www.ecmwf.int/en/forecasts/documentation-and-support>;
- Quality of the ECMWF forecasts and the Headline Scores: <http://www.ecmwf.int/en/forecasts/quality-our-forecasts>;
- Evolution of the ECMWF IFS cycles:  
<http://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model>;
- OpenIFS project (see Section 6): <http://www.ecmwf.int/en/research/projects/openifs>;
- Forecast changes and known issues: <https://software.ecmwf.int/wiki/display/FCST>.

Internal communication relies on a series of meetings and events, including:

- Informal and formal Team/Section meetings;
- Topical meetings (e.g. seminars held by people leading the development of IFS changes);
- Plan and progress report meetings, held at Section (once every 4-8 weeks) and Department (once every 3-6 months) level;
- Cross-department Quarterly Evaluation and Diagnostic (QED) meetings;

- Monthly meetings of the Senior Management Team (SMT: it includes the Directors, Deputy Directors and Lead Scientists), and quarterly meetings of the extended-SMT (SMT plus Head of Sections);
- 1-year annual plan meetings of the SMT.

Scientists/experts developing model changes and upgrades would use informal and formal Team/Section meetings to discuss plans, present progress and review work. Internal communication relies also on web-based tools and shared web pages. Team Leaders and Heads of Sections update the relevant Director and the SMT routinely on how work progresses (in particular, the SMT reviews progress monthly, and keeps track of whether the planned implementations are completed on time or not). Departments meet every 3-to-6 months to discuss and review progress, and to define the 1-year and 4-year plans.

Every IFS cycle implementation is managed by an Implementation Working Group (IWG). This group includes key people involved in the R2O process. It usually holds its 1st meeting about 2 months before the Research Department starts merging all code changes into a common repository. It discusses what should be included in the next cycle, reviews progress, assesses performance, and manages the implementation process. The way the IWG works is discussed in Section 4, where the last phase (the last 8-10 months) of the R2O process is discussed.

Model changes (past and forthcoming) are also communicated to the users during meetings, in particular at the annual User of the ECMWF Forecast (UEF) meeting, and Member State visits. Information is also reported on the ECMWF web site.

Developments and research achievements, and interesting studies are communicated to the scientific community in the peer-reviewed literature, and internal publications (workshop and seminar proceedings, Technical Memoranda, and Newsletter articles).

## 4 The last phase of the R2O process leading to *implementation*

Let us now assume that the planning activity has been completed, and that the development and the first set of testing has been completed. Let us also assume that the evaluation of this first set of experiments has suggested that the next round of planned changes could lead to improvements, and that communication activities have also been completed as expected. We have now reached the last stage of the R2O process, when the scientists/experts who have been developing possible changes are ready to include them in the next IFS cycle, start the final round of testing, and implement them into operations. This last phase takes usually between 8 and 12 months:

- To test all the new developments incrementally, so that it is easier to identify interactions, pinpoint the source of any problems and work on solutions;
- To experiment thoroughly first at lower resolution, so that more independent cases can be included in the tests, and then at full resolution;
- To evaluate and diagnose performance, applying clearly defined and objective metrics;
- To keep communication flowing as expected.

The first step of this last phase of the R2O process is to define who will be part of the IWG. In the past 24 months, the IWP has been co-chaired by one person of the Research Department (for the past 2



years, this person has been one of the two ECMWF Lead Scientists) and by one person from the Forecast Department (for the past two years, this person has been the Head of the Production Section). The IWG includes the Research Department (RD) Head of Sections, the Forecast Department (FD) Head of Production, and key people from RD and FD who are involved in the preparation and hand-over of the new IFS cycle.

To keep the process efficient and effective, IWG meeting attendance is defined depending on each meeting agenda. The meetings are attended by a wider audience only when the overall assessment of the new experimental suite (the e-suite<sup>4</sup>) is analyzed and key decisions are taken. This ensures that all the required expertise is 'in the room' to contribute to the discussion and help interpret the results, without over burdening everyone with unnecessary meetings.

The IWG manages the implementation process, advises the Directorate whether the new IFS cycle is ready for implementation, and suggests its implementation date. Based on the IWG advice, the Directorate decides whether to proceed as suggested.

In more detail, the IWG is tasked to:

- Initiate the production of control experiments based on the current operational cycle over a past period, to be used during the testing phase as a reference to assess the impact of each proposed change;
- Review all planned IFS changes, which are usually available as branches of a previous IFS cycle, and assess whether there is enough evidence that all of them are ready for merging into the new cycle, and whether changes have any impact on the computing resources required in operational production;
- Decide which branches are ready to be implemented in the e-suite, and how they should be incrementally merged;
- Define the time periods (usually the past two warm and cold seasons) on which the merged branches should be tested, and start the testing activities;
- Finalize the content of the IFS cycle and establish the content of the e-suite; start running the e-suite as an RD experiment and compare it to the reference; define a possible implementation date;
- Once a large number of cases have been run, assess whether the RD e-suite has been performing as expected and the cycle is ready to be handed over to the FD, so that they can start running it in parallel to the o-suite;
- Start a period of parallel real-time testing of the e-and o-suites, including product generation; provide member states with access to e-suite data;
- Task the FD Forecast Evaluation Team to carry out a further round of evaluation and diagnostics, independent from the developers;
- Assess the level of user impact, technically as well as meteorologically, and identify and coordinate the development of modifications to the supporting software and the meteorological application software, required to fully support the new cycle and to help the users adapt to the changes;

---

<sup>4</sup>The e-suite (experimental suite) is the corresponding forecast production system to the o-suite (operational suite) but based on the new IFS cycle.

- Freeze the new IFS cycle; confirm and/or revise the implementation date, and communicate it with enough lead time for the users;
- Implement the new IFS cycle into operation (i.e. switch off the operational cycle, and ‘label’ the e-suite as the o-suite);
- Throughout all the phases above, coordinate the preparation of communication material, both internally and externally, ensuring all stakeholders are informed on progress.

During these 8-10 months, the IWG meets regularly: initially about once a month, and then more frequently, depending on how work is progressing and how close we are to the implementation date.

Once the new IFS cycle is frozen, the IWG prepares an ‘IFS cycle report’, an internal document that summarizes the key content of the cycle and the expected impact on the forecast performance. This document is used by the FD analysts as a reference during their comparison of the e- and o-suites, to assess whether the results based on their parallel run of the e- and o-suites are consistent with the earlier RD experimentation (consistency is seen as evidence that the hand-over of the new software from RD to FD has worked as expected). This document is also used to prepare the external communication documents and the web pages that announce to the users the planned IFS cycle changes.

It is worth mentioning again the issue of computing resource availability, mentioned above. Since ECMWF aims to deliver its forecasts in a timely manner, care is taken that changes to the IFS cycle do not cause any delay in the products’ dissemination. This is one of the first points that is reviewed by the IWG when deciding which proposed changes can be implemented.

For example, at the first meeting of the 45r1 IWG meeting (this will be the cycle after 43r3, and it is planned for implementation in early 2018), which was held at the beginning of April 2017, the possibility was discussed of coupling the dynamical ocean model NEMO also in the high-resolution forecast model. Earlier tests indicated that this change brings improvements, but requires an increase in the computing resources. It was then decided to do more work to assess exactly how much more resources are required, and to see whether there are ways to compensate them by changing parts of the system (e.g. the I/O, the input/output software components). The results will be reviewed and a decision made as to whether this change can be included in this model upgrade or if more substantial work is required to make it affordable.

#### *Example: The R2O process for IFS cycle 43r3*

Figure 3 shows, for IFS cycle 43r3, how all the tasks and activities highlighted above have been linked together, the time that was assigned to each of them when the 43r3 IWG started working, and when key decisions were made. Note that, compared to the plan prepared in July 2016, the implementation was eventually delayed by about 3 months (11 July 2017 instead of 14 April 2017). This delay was caused by problems that were found in the RD e-suite test phase, in December 2016 and January 2017. They required the revision of some of the model changes that had been already merged, and further experimentation. The fact that these problems were identified before operational implementation, and solved, gives us confidence that the R2O process has been working as expected. One of the key areas of process improvement is to do more testing and evaluation earlier, so problems are easier to find and fix than if they were found at a late stage after substantial merging has taken place.

The 43r3 IFS cycle example indicates that the last phase of the R2O process took about 11 months from the time when the reference control experiments were made available (22 August 2016) to implementa-

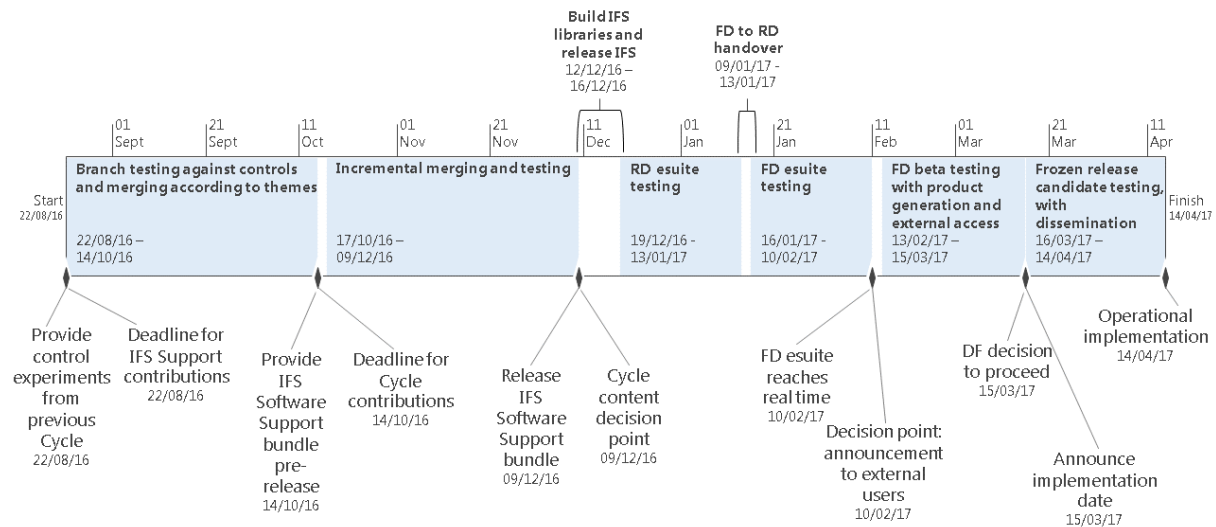


Figure 3: Gantt chart of the last phases of the R2O process prepared in July 2016 by the 43r3 IWG, applied to IFS cycle 43r3, implemented later than originally planned on 11 July 2017. Diamonds identify key decision points.

tion (11 July 2017). For the previous IFS cycle 43r1, this last phase of the R2O process was about 12 months. For the next cycle due to be implemented, 45r1, we are aiming to reduce the time further.

## 5 Metrics and decision criteria applied in the *evaluation* and diagnostic phase

Objective scores/metrics are used to assess the expected impact of potential model changes, and to monitor the forecast system performance. Different levels of granularity are applied in the different stages, with the full breadth of evaluation applied only in the last stages of the R2O processes.

Sudden changes of scores/metric indices are used to flag potential problems. These metrics include the 6 headline scores (of which 4 are supplementary scores) agreed by Council to monitor the ECMWF forecast performance. Time series of these indices are available on the ECMWF web site, and are routinely communicated to the Committees on an annual basis. Scorecards are used to summarize the performance of the e- and o-suites, and to identify variables and/or areas that need particular attention.

Figures 4 and 5 show two examples of scorecards, applied to assess the performance of IFS cycle 43r3 (these scorecards helped taking the decision whether the cycle was ready for operational implementation, and were included in the internal document prepared by the 43r3 IWG to describe the 43r3 cycle).

Average scores are supported by synoptic evaluations of e- and o-suite for specific case studies. Case studies are used to understand better features seen in average scores, and/or to investigate ‘potential’ impacts linked to model changes that might not be seen in average scores, e.g. because the changes affect only rare events, for which it is difficult to collect a large-enough sample. The interpretation of features seen in average scores is also helped by the application of diagnostic tools designed to investigate



Figure 4: Example of a scorecard for the single, high-resolution (Tco1279L137) forecast. It was built using 43r3 experiments, by merging the summer and winter experiments. For errors (ccaf and rmsef): upward green triangles indicate that 43r3 has (statistically significant at the 99.7% if solid and 95% if not) lower RMS errors (better scores) than the 43r1 (operation); downward red triangles indicate the opposite. For activity (sdaf): magenta indicates that 43r3 is more active than 43r1 and blue indicates that 43r3 is less active than 43r1.

specific aspects, e.g. the identification of model biases and random error components.

In the last phase of the R2O process, the analysts who monitor the operational forecast performance are also tasked to look at the e-suite performance. During the final parallel run, the analysts also have access to the whole range of products generated by both the e- and o-suites, and they comment in the Daily Report<sup>5</sup> when they see differences between the two suites (either positive or negative). At the end of the RD and FD e-suites, there is usually about 1 year of clean e- versus o-suite forecasts for the data-assimilation (Ensemble of Data Assimilation, EDA, and the high-resolution analysis) and the high-resolution forecast, and about 100 cases of medium-range/monthly ensemble (ENS).

If the results are judged to be convincing, robust and based on a large enough sample of cases, the IWG

<sup>5</sup>An analyst at ECMWF monitors aspects of the analysis and forecast daily, producing a Daily Report which discusses model performance for major meteorological events, problems or issues with the analysis or forecasts, comparison of different models etc.




Figure 5: Example of a scorecard applied to the 43r3 medium-range/monthly ensemble (ENS). It shows the difference of ranked probabilistic skill score between 43r3 and 43r1 experiments. Positive difference (dark blue or cyan) means that 43r3 outperforms while negative differences (red or yellow) indicate a degradation with 43r3.

advises the Directorate whether to implement the new cycle, and suggests an implementation date. Once the Directorate has decided whether to proceed, and has confirmed the implementation date, users are informed. Care is taken so that users have enough lead time before the operational change: usually the communication happens at least 4 weeks before implementation for a change that does not require major work from the users to be able to continue using the data. For major upgrades (e.g. when resolution is changed), users are informed at least 3 months before the implementation. Unless any unforeseen problem arises, the new IFS cycle replaces the operational one on the planned implementation date and the e-suite is stopped.

## 6 User requests and the R2O process

Users might have requests that need to be addressed in a new IFS cycle. For example, they might ask that a new model parameter is produced (e.g. a few years ago they asked that the wind field was also post-processed and produced at a 100 metre height for wind-farm applications); or they might ask that the dissemination schedule is changed.

All requests are taken into consideration, and the ones that are judged suitable for implementation are analyzed in more detail. The TAC takes an active role in prioritizing such requests. If they require changes to the model, then they need to be fed into the R2O process:

- If the changes require a ‘small’ amount of work, then they are passed to the scientists with the right expertise, who are asked to take them into account when they prepare a new model cycle.
- If instead these changes require substantial work (e.g. changes in the way the operational suite is organized), then the requests are discussed in the planning phase, when the 4-year and 1-year-detailed plans are prepared. If accepted, the required work is included in the 4-year or 1-year-detailed plan, and from then onward it follows the usual process.

At other times, users have spotted errors and problems in the operational forecasts. In this case, the issue is immediately discussed:

- If the error is severe and there is a quick solution (e.g. when it is clear that the error is due to a ‘bug’), the solution is tested on a number of cases, and if clean and parallel experimentations suggest that the fix leads to better results, it is implemented.
- If instead there is not a quick solution, or the parallel experimentation gives mixed results, RD discusses whether there is any chance to address the problem in the forthcoming IFS cycle. If this is the case, work progresses as discussed above, with the inclusion of the change in the next available/suitable cycle. If this is not possible, then the discussion is moved to the planning phase.

Scientists working with the ECMWF model, e.g. within the OpenIFS project (<http://www.ecmwf.int/en/research/projects/openifs>), can also influence model upgrades. This project provides research institutions with an easy-to-use version of some components of the ECMWF IFS: the forecast capability of IFS (no data assimilation), the Single Column Model (SCM) version, and supporting software and documentation. OpenIFS has a support team at ECMWF for technical assistance but limited resources for detailed scientific assistance. It was established a few years ago to increase the number of external academic users of the IFS, and thus to increase the number of collaborations between ECMWF and academic institutions, so that ECMWF can more easily integrate new developments from outside ECMWF into the IFS. An example of this type of feedback has been the inclusion in recent IFS cycles of changes that allow the model to run with single precision: this work was conducted in collaboration with Oxford University using the OpenIFS model version.

## 7 On-going review activities of the R2O process, including the IDEA project

The R2O process, and more specifically its last stage (the last 8-10 months of work), have been revised recently, and continue to be under review. This has been facilitated by the creation of an internal project, the IFS Development and Evaluation (IDEA) project, which was initiated in 2016 to improve the process of delivering IFS cycles. The IDEA project is a cross-department activity organized into three themes (Development and Testing, Evaluation and Diagnostics, Communication) with a remit to look across the IFS cycle development process to identify and action improvements to make sure we can do research efficiently that feeds effectively into operations. There are many different aspects, some of which have already been implemented and some still in progress. A few key outcomes include:

- To make the most efficient use of computational resources, we have identified a reduced resolution test-suite configuration for assimilation and forecasts, which is now used for the majority of the development testing in RD, with documented control experiments. This provides relevant results for evaluation at the same time as improving throughput on the supercomputer and reducing data volumes in the archive.

- The process of merging individual contributions has been clarified so that any potential interactions are identified early and are tested appropriately within Teams/Sections. Once ready, these tested pre-merged contributions are combined and tested in a structured way by the IFS Section so that there is traceability of the scientific and technical changes to the system in case of any unexpected results or interactions.
- A major improvement, introduced for implementation in cycle 43r3, is the application of the JIRA software project management tool for the IFS. This provides a central repository of all information for IFS developments; JIRA is a best-in-class system for documenting software developments, from bug fixes to major scientific improvements, in a systematic way that provides a widely available, transparent, searchable database. Each code change is logged in the system with information about the change and its interactions, the testing and evaluation that has been performed and other details such as dependencies on other parts of the R2O process. For example, an IFS code change for a new output parameter can be linked within JIRA to the Member State User Request for this parameter. It can also be linked to Data Governance to ensure the information is communicated through to FD and beyond, as a new operational product for external users. The new system will facilitate project-management and delivery of cycles, improve communication within and between ECMWF departments, Member and Co-operating States, and external collaborators; and centralize and streamline IFS-related documentation.
- One recommendation is for more technical testing to be done earlier on in the R2O process. With an increasing set of configurations, it is important that changes to one configuration do not cause problems in another. If tested early on, they can be diagnosed and fixed much quicker than later in the merging process. This can be facilitated by creating standard, fast-running technical tests with clear acceptance criteria. In the longer term these can be automated, forming the basis for a possible move to Continuous Integration (CI), a standard software development model whereby individual branches are continually merged into an aggregate development branch, which is frequently tested (e.g. daily). Such a move could shorten the feedback cycle considerably and reveal technical and scientific problems and conflicts as they arise, rather than after a considerable period of development in isolation.
- Improvements in the process of communication of information relating to an IFS cycle, such as earlier involvement of FD in IFS cycle meetings, and improved communication of evaluation results across departments.
- Improvements in the process of documentation of all aspects of the IFS cycle process, from testing procedures for code developers, linking updates to the full scientific IFS documentation more closely to the IFS cycle process, to more comprehensive documentation for users on operational output parameters.

These are just some of the recent activities that are looking at the R2O process, promoting communication and providing the tools and environment for an efficient and effective process. There is always a balance to strike between a good turnaround of IFS upgrades to ensure that new innovations become operational quickly, with careful testing and management of impacts to ensure robust and successful implementations for users.

The IDEA project is planned to end during 2017, once it has followed the implementation of two IFS cycles. The plan is to continue to review the R2O process in the future after each cycle, building on lessons learned and identifying any areas that require further improvement in areas such as software infrastructure, testing framework, merging and testing strategy, metrics, diagnostic tools, and communication.

## 8 Summary

In this document, we have highlighted the key stages of the ECMWF R2O process, and we have discussed how some of its key aspects have been organized:

- In Section 1 we have described the 6 main activities of the ECMWF R2O process: planning, development, testing, implementation (into operation); evaluation (and diagnostics) and communication.
- In Section 2 we have discussed the planning and development activities, how long-, medium- and short-term developments are planned, and how they are inter-linked. We have also discussed how these plans are communicated to the ECMWF Committees.
- In Section 3 we have reviewed the communication activities that are performed during the other 5 activities, from planning to operational implementation.
- In Section 4 we have highlighted the key phases of the last 8-10 months of the R2O process. We have discussed how these phases are managed by a working group, and have shown the timings of these phases. In particular, we have illustrated the timings of these phases for IFS cycle 43r3, implemented on 11 July 2017.
- In Section 5 we have briefly reviewed the performance and evaluation metrics used to assess the proposed model changes.
- In Section 6 we have shown how users' requests are taken into account in the R2O process.
- In Section 7 we have discussed some on-going ECMWF internal activities, including the IDEA project, that are continuing to review the R2O process and identify possible ways to further streamline and improve it.

Having a clear, resilient and effective R2O process is key to progress at ECMWF, and it is important to continue to evaluate and review the R2O process to identify opportunities for further improvements in the future.

## Acknowledgements

The authors would like to acknowledge Erland Källén for reviewing and commenting on the document.