

# The Unified Forecast System (UFS) Workflows Workshop Report

December 2020

Points of Contact: Mrinal K. Biswas<sup>1,2</sup>, Mike Ek<sup>1,2</sup> and Ben Cash<sup>3</sup>

## Organizing Committee:

Ligia Bernardet<sup>2,4,5</sup>

Benjamin Cash<sup>3</sup>

Arun Chawla<sup>4,6</sup>

Mike Ek<sup>1,2</sup>

Evan Kalina<sup>2,4,5,7</sup>

Louisa Nance<sup>1,2</sup>

Mariana Vertenstein<sup>1</sup>

<sup>1</sup>National Center for Atmospheric Research

<sup>2</sup>Developmental Testbed Center

<sup>3</sup>George Mason University/Center for Ocean-Land-Atmosphere

<sup>4</sup>National Oceanic and Atmospheric Administration

<sup>5</sup>Global Systems Laboratory

<sup>6</sup>Environmental Modeling Center

<sup>7</sup>Cooperative Institute for Research in Environmental Sciences



## Executive Summary

The [Unified Forecast System \(UFS\) Workflows workshop](#) was held virtually on 29 - 30 April 2020 to foster more discussion and learn about the different workflow efforts that are going on within and outside the UFS community. The workshop was attended by over 60 people. The workshop was preceded by the Configuration manager for Research and Operational Workflows (CROW) Review (28 April 2020), and the agendas for the two workshops were coordinated.

This workshop was an outgrowth of the Advanced Workflows working group in-person meeting held on 22 November 2019, followed by several discussions that then led to the planning of the Workflows workshop.

The presentations during the workshop included overviews of different types of workflows, e.g., UFS Global Model, Joint Effort for Data assimilation Integration (JEDI), Common Infrastructure for Modeling the Earth (CIME) for different applications, FV3-Stand Alone Regional (SAR) (now called Limited Area Model), Hurricane Analysis and Forecast System (HAFS), GEFS-Aerosol, and CONcepts of OPerationS (CONOPS) for WAM-IPE (Coupled Space Weather Forecasts).

Talks on workflow requirements, expectations, and Rocoto (a workflow management system) were also presented. There were two breakout sessions: 1) session A discussed the limitations, duplication of effort, cultural differences, barriers to move between applications, and issues supporting the general user community, and 2) session B discussed possible solutions to the issues explored in Session A.

Major problems identified:

- Too many workflows.
  - Users need to learn a completely different workflow system when switching between applications.
  - Significant duplication of effort required to solve similar problems.
  - Developers are not learning from other (parallel) efforts.
- Many key terms mean different things to different people.

- Communication needs to improve across the board.

How did we end up like this?

- Every application has its own specific requirements, so diversity is sometimes appropriate and necessary.
- Lack of communication:
  - Developers and end users often do not communicate with each other.
  - No established channel/forum/group/committee to talk to each other.
  - Lack of communication with National Centers for Environmental Prediction (NCEP) Central Operations (NCO), where standards set by NCO are often opaque to the transition-to-operations centers in NCEP, National Weather Service (NWS) and NOAA, as well as to the general community.
  - Lack of directive and corresponding standards.
  - No communicated development standards.
- No governing body to steer the flow of discussion between different groups.
- Success is measured by how well a workflow works for a particular application.

Solutions require:

- Creating a governing body.
- Establishing communication channels e.g., Slack channel, webinar series; there may be other possible options.
- Identifying a common set of tools that can be used by different applications, which will help to reduce duplication.
- Involvement of NCO in the developing, testing and adopting of workflow processes.
- Updating the [UFS glossary](#) to add workflow related terms where definitions need to be agreed upon by a broad cross-section of the community.

The summary of the workflows workshop was presented at a UFS steering committee meeting in May 2020 and at the [UFS Users' workshop](#) in July 2020.

# The Unified Forecast System (UFS) Workflows Workshop Report

**December 2020**

## 1. Introduction

The Unified Forecast System (UFS) Workflows workshop was held virtually on 29-30 April 2020. The organizing committee included UFS collaborators from the Developmental Testbed Center (DTC), George Mason University (GMU), the National Oceanic and Atmospheric Administration (NOAA), and the National Center for Atmospheric Research (NCAR). The logistics of the virtual meeting were provided by the DTC. The aim of the workshop was to gather information from different partners about their workflow systems, challenges, and how different groups can collaborate to solve similar problems.

There are several workflows that are used for different applications. But few of them are used by the community due to lack of portability, documentation, flexibility, and user support. However, there is a mandate to support workflows to the broader community for UFS applications. Several discussions preceded the UFS workflow workshop; a high-level initial meeting was held during summer of 2019 in Boulder, CO. During that discussion everyone felt greater coordination was necessary to bring all the stakeholders together to brainstorm ideas on how to leverage several efforts that were already underway for different UFS and other applications. A workflow did need to be selected for the Medium Range Weather application in order to meet the timelines of the release, but only needed to address model forecast requirements and not e.g. data assimilation, though something more encompassing would be needed in the future. At this time planning for a UFS Workflows workshop began (initially planned for Fall 2019, but subsequently postponed until Spring 2020).

On 22 November 2019, an advanced workflow telecon was held to discuss some different aspects of workflow, e.g. the Medium-Range Weather and HAFS applications, and use of the Common Infrastructure for Modeling the Earth (CIME). The role of CROW was also discussed in the context of HAFS. During this meeting it was realized

that, while HAFS had borrowed pre- and post-processing scripts from SAR, there was some duplication of effort between these applications, with possible divergence of workflows in the future. An evaluation was also needed of CROW; in particular, a decision needed to be made regarding the use of CROW in the HAFS application. The focus, format, and timing of the UFS Workflows workshop was then discussed during the meeting. A workflow workshop organizing committee was formed in January 2020 and it was decided that the timing and agenda of the Workflows Workshop and the [CROW review](#) would be coordinated because the attendees and topics would be of mutual interest.

The 29-30 April 2020 virtual meeting was attended by over 60 participants. The presenters included attendees from DTC, Joint Center for Satellite Data Assimilation (JCSDA), and Cooperative Institute for Research in Environmental Sciences (CIRES), NCAR, and GMU. There was also significant representation from NOAA, particularly Environmental Modeling Center (EMC) and Oceanic and Atmospheric Research (OAR),

This report provides an overview of all the presentations and the discussions that followed them. Before going into the in-depth discussions about workflow systems, it seems that readers should be clear on what the term workflow actually means. Based on the presentation by Arun Chawla on Day 1 of the Workshop this document uses the following general definition, at least in the context of numerical weather prediction (NWP): “A number of interdependent jobs that are run at regular intervals in a predefined order based on the application”. However, it is important to note that the UFS does not have a universally agreed-upon definition of workflow.

Section 2 of the report provides a high-level summary of the presentations. Section 3 summarizes the discussions during the breakout sessions. Section 4 briefly describes the steps planned to be taken after the workshop.

## 2. Highlights of the presentations

On Day 1 (29 April 2020), all [invited presentations](#) were made virtually. Welcoming remarks were given by Ben Cash (GMU/COLA), who noted that this was originally planned as an in-person event and thanking everyone for their interest and willingness to try what was then a very experimental meeting. The opening remarks were then given by DaNa Carlis (OAR), who noted that plans for a workshop of this kind had been in the making since July 2019. He established several of the major themes for the workshop, including the need to work across organizations, the need to take operational

requirements into account, and the need for evidence-based decision making for maintaining multiple workflows.

[Arun Chawla \(NOAA/NWS/EMC\) and Ligia Bernardet \(NOAA/GSL and DTC\)](#) provided a broad definition of workflow, overview of workflow requirements, expectations, and EMC's path in designing the workflow for the UFS applications. They defined the workflow manager as a system that consists of setting up, performing and monitoring interdependent jobs in a predefined order based on the application.

[Kate Friedman \(NOAA/EMC/EIB\)](#) presented an overview of the global workflow that drives the operational GFS. She stressed the need for using CROW to make the workflow more flexible and modular. However, the workflow is hardcoded to work only with the global configuration. She noted that there is no plan to merge the global and regional workflows at this time.

[Thomas Auligne and Yannick Tremolet \(JCSDA\)](#) presented the workflow for the Joint Effort in Data Assimilation (JEDI). This piece of the puzzle is critical because it will impact all NOAA, NASA, Air Force, NCAR and the community models that may use JEDI for their data assimilation. The objective is to build a generic configuration manager for all experiments.

[Mariana Vertenstein \(NCAR\)](#) provided an overview of the Common Infrastructure for Modeling the Earth (CIME). The CIME Case Control System (CCS) is an object-oriented python based workflow infrastructure targeted to creating a user-friendly environment for, configuring, building and running coupled earth system model experiments. The goal of the CCS is to facilitate research and remove the need for scientists to understand the entire complexity of the model in order to carry out an experiment. The CCS is used with the UFS Medium Range Weather application public release and is an integral part of the included Graduate Student Test which is a key measure of how successful the UFS project is in opening its code and development processes to the broader community.

[Jeff Beck \(CSU/CIRA at NOAA/GSL and DTC\)](#) presented an overview of the UFS Stand Alone Regional (SAR) model (now called Limited Area Model) configuration workflow. The workflow manager is Rocoto-based but can also be run with command line arguments. The workflow will be shared with the community including a user's guide along with the Short-Range Weather application public release slated for Fall 2020.

[Henry Winterbottom \(IMSG and NOAA/EMC, now at CIRES and NOAA/PSL\)](#) provided an overview of the Hurricane Analysis and Forecast System (HAFS), the hurricane application of the UFS model. The workflow currently supports the limited area model and global plus nest configurations.

[Li Pan \(IMSG and NOAA/EMC\)](#) presented the development of the Global Ensemble Forecasting System (GEFS) - Aerosol workflow. The GSD Chem Aerosol model is coupled with the FV3GFS and used as an ensemble member of the GEFS. This model is under active development.

[Raffaele Montuoro \(NOAA/ESRL/GSL and CU/CIRES\)](#) presented the development of the Coupled Space Weather Forecasts: Concepts of Operations (CONOPS) for the Whole Atmosphere Model (WAM) coupled to an Ionosphere, Plasmasphere, and Electrodynamics Model (IPE) workflow. The workflow is Rocoto-based and is under development.

[Christopher Harrop \(NOAA/ESRL/GSL and CIRES\)](#) presented the history, current capabilities and future plans for the Rocoto workflow manager, which is used by the majority of UFS applications. A number of challenges were also discussed, including the lack of dedicated funding for this critical work. While answering an attendee's comment, he mentioned that ecFlow is used in operations by NCEP and European Center for Medium-Range Weather Forecasts (ECMWF) but is not considered suitable for research because it needs a server that admins need to set up, though ecFlow can still be used but without being able to take advantage of all of its capabilities.

### 3. Breakout sessions

Day 2 (30 April 2020) of the workshop included two breakout sessions: Session A discussion was focused on analyzing the different workflows, with a focus on identifying areas where the workflow effort as a whole could be improved. Common problems identified by multiple breakout groups included the proliferation of workflows, challenges faced by newcomers to a given application, working with NCO, and various other topics. During Session B, participants brainstormed potential solutions identified in the first session and ways to improve the workflow effort in general.

Session A highlighted many issues identified by developers and experienced NOAA and UFS collaborators.

We have broken down the discussions into three parts: (i) discussion of the issues that are faced now, (ii) how did we end up like this?, and (iii) potential solutions to the broader issues.

### 3.1 Current issues

- Too many workflows
  - As reflected from the list of presentations on Day 1, it was evident that there has been a “workflow explosion” among UFS applications, where every application has its own workflow. While it was acknowledged that each of the varied applications do require unique capabilities, many of these needs could be addressed by assembling a common set of tools and modifying them to fit a particular application. Currently, users who work on multiple apps need to learn each from scratch (e.g. directory structure, workflow managers, scripts, etc). While there may be a need for certain specialized workflows, there should be evidence-based justifications for why they need to remain separate and unique given the associated costs in development and user training they require.
- NCO
  - The role of NCO is critical because they are the deciding body for operational code standards and are responsible for running the operational systems. It is worth pointing out that there was no representation from NCO at this workshop. A general consensus was that NCO is often resistant to new approaches (e.g. Python, GitHub), which at times is due to lack of familiarity related to logging and debugging. It was suggested that demonstrating a new capability (including language) to NCO can facilitate adoption of new infrastructure. There was a mixed response by workshop participants on when to involve NCO during the development process. However, there was a consensus that EMC should be the conduit to NCO.
- Lack of standard definitions of terms used in the workflow:
  - There was confusion on the definition and scope of the terms used in the workflow. Arun Chawla presented a broad definition of workflow but there was discussion that operations and research may have different views of the scope of workflow.
- Lack of leveraging other groups effort:
  - Workflows have been developed by different groups and many components of them serve the same purpose. However, due to the lack of communication and coordination, similar jobs within the workflow have



been written in different languages by different groups (e.g. regional uses bash library and HAFS uses bash and python), hence, there is duplication of efforts.

- Lack of access for University communities:
  - The future developers and users come from the universities. Many times the workflows are either incompatible with computational systems available to the university community or may be too complex to allow for easy use. Also, the applications and associated workflows require pre-existing software stacks that are a barrier for graduate students. Limited access to NOAA machines (especially for foreign nationals), where the workflows are usually tested and installed, is a serious roadblock for getting feedback from the community. Finally, not all data sources are available outside of NOAA machines, and existing workflows are not necessarily set up to handle varied or missing input datasets.
- NOAA's timelines:
  - While researchers outside NOAA use a given workflow system and want to contribute to development, it seems the community is always behind and the deadline to contribute is over. Innovation takes time, and with ever-changing codes, it is difficult to meet NOAA's timelines. However, it was suggested that if the community demonstrates the scientific merit, DTC can help them in transferring new innovations into operations.

## 3.2 How did we end up here?

- Every application has its own specific requirements, so diversity is sometimes appropriate
  - The development of different UFS applications happened simultaneously and respective groups felt the need to develop their own workflows which was best for them. Unification of workflows will make it a complex process, and may defeat the purpose of user-friendly applications for end-users.
- Lack of communication and governance:
  - Developers and end-users (in research and operations) often do not communicate during the development process. Codes are usually handed over with minimal instructions, and left to the users to learn how to use it for their purposes; at that time it is too late to go back and "reinvent the wheel".
  - There is no established channel/forum/group where developers and users can discuss their requirements and follow the development process, e.g.

the NCO standards are sometimes unknown to the research centers, and there is no mandate from NCO to engage with developers.

- There is no governing body to steer the flow of discussions between different groups. The developers are left to tackle challenges and find solutions on their own, often under pressure to meet deadlines.
- There is no standard w.r.t. language (e.g. Python vs shell, or a mix of the two), code standards, variable names, configuration formats, etc. There is no agreed upon infrastructure.
- Lack of Mandate
  - The metrics of success depend on how the workflow works with a given application, and not on how well the effort is made in coordination with other groups. There is no mandate to coordinate development among different groups.
  - Unifying duplicated portions of scripts and other workflow elements requires time and effort. These resources are often applied to advancing the application, rather than reengineering an existing capability within the workflow.

Session B focussed on potential solutions to the problems identified above, where we tried to focus on the high-level issues faced by developers and users.

### 3.3 Possible Solutions

- Create a governing body:
  - There was a consensus that a workflow governing body (or coordination group) can facilitate the discussions among different groups. Agreeing to use a common set of tools and calculating the resources saved by doing so will help in decision making. This approach will also motivate developers to share their ideas and challenges with other groups. Knowing what kind of tools/solutions are available in the community can help the development team decide whether it makes sense to take a divergent workflow path because of a unique need for a particular application.
  - The governing body can collectively decide on the coding standards and corresponding rules using evidenced-based decision making. There is a clear need to identify the requirements for each workflow. There should be a clear understanding of what each workflow *needs* to do, and why. For example, how much overlap is there between operational and research requirements? If it isn't possible to use one workflow for both research and operations then we should be able to demonstrate and document why.

- The governing body can decide on broad directions, and can help facilitate discussions between the developers and users before merging code, thereby fostering a collaborative environment in reviewing workflow issues.
- The body can create a mentoring program for future developers.
- Create communication channels:
  - There was an overwhelming consensus that there is a lack of communication between different groups (application developers, operations, NCO, and end users). This is detrimental for collaboration. Several communication channels (e.g. Slack channel, webinar series, etc) can be used to foster improved communication.
- Involvement of NCO in the developing, testing and adopting workflow processes:
  - NCO needs evidence-based decisions from the developers, helping them in making more informed decisions. But, this should be a two-way communication, where NCO should work with EMC and other developers (e.g. provide training) on acceptable [implementation standards](#), while keeping the communications open to modern software developments. This will help in making our software the best in our field of work.
- Updating the [UFS glossary](#):
  - There is an existing UFS glossary. During the course of this workshop it was felt that the workflow terminologies used can mean different things to different people. There is a need to collectively address this issue, which will help communication among workflow developers and users. We need to add workflow-related terms with definitions agreed upon by a broad cross-section of community.

## 4. Summary

The virtual workflow workshop provided a great opportunity to get acquainted with the various workflow solutions that encompass UFS and related applications. It was quickly recognized that there has been significant duplication in workflow development across the UFS that could potentially have been avoided if there had been better communication between different groups. The workflow developers could have leveraged other groups' efforts when trying to solve similar problems. Though the workflow is designed for a seamless end-to-end system, in the end-to-end channel of developers to researchers and finally to operations (here NCO), coordination and communication is not seamless and sometimes totally absent. The breakout sessions provided an opportunity to discuss these problems and challenges and begin to

consider solutions to mitigate towards, such as a commitment to unification (in some manner) where commonalities in workflows exist. It was felt that we need to come out of our own bubble and think of ourselves as a workflow team that spans all developers of NOAA UFS Applications (and not just at EMC, or Lab X, etc.) and include the broader community, instead of individual groups with their own insular workflow. To facilitate discussions among different groups, a [UFS workflows slack channel](#) has been created. However, a governing body is needed to steer the discussions between different groups for mutual advantage and in order to reduce duplicate efforts. Also, based on discussions during the workshop, the UFS glossary was updated to include definitions for certain missing terms. Finally, in order to make the workflow system successful, it is imperative to involve NCO and developers (e.g. EMC) early on in the development process. Subsequent actions following the workflows workshop and the release of this report should include engaging the relevant UFS working groups on where to go next, and planning a follow-on workshop sometime next year to see where we've come by then.

## 5. Acknowledgements

We are thankful to the speakers including Ben Cash, DaNa Carlis, Arun Chawla, Kate Friedman, Thomas Auligne, Mariana Vertenstein, Jeff Beck, Henry Winterbottom, Li Pan, Raffaele Montuoro, and Christopher Harrop. Thanks are also due to the breakout sessions organizers who helped in coordinating the brainstorming ideas. Thanks to Jennifer Bolton for providing logistical support in running the workshop. Finally, we want to thank all the 60 plus remote participants who have contributed their time and effort during the workshop and made this virtual event successful by engaging in useful and insightful discussions throughout the event.