

D6.2 Operational workflows in the CDS and prototypes of the workflows

November 2025

Charalampos Karvelis (ECMWF)
Matthew Menary (ECMWF)
Carlos Delgado Torres (BSC)
Chiara Cagniazzi (ECMWF)

Document information

D6.4 Data Management Plan	
Grant Agreement number	101081460
Project title	Adaptation-Oriented Seamless Predictions of European Climate
Project acronym	ASPECT
Project start date	1 January 2023
Project duration	48 months
Work Package	WP6
Deliverable lead	European Centre for Medium-Range Weather Forecasts (ECMWF)
Author(s)	Charalampos Karvelis (ECMWF) Matthew Menary (ECMWF) Carlos Delgado Torres (BSC) Chiara Cagniazzi (ECMWF)
Type of deliverable* (R, DEM, DEC, other)	OTHER
Dissemination level** (PU, CO, CI)	PU
Date of first submission	20 Dec 2025
Revision n°	-
Revision date	-

Please cite this report as: Karvelis, C., Menary, M., Delgado Torres, C., Cagniazzi, C., (2025), Operational workflows in the CDS and prototypes of the workflows, D6.2 of the ASPECT project.

Funded by the European Union. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them.

* **R**=Document, report; **DEM**=Demonstrator, pilot, prototype; **DEC**=website, patent filings, videos, etc.; **OTHER**=other

** **PU**=Public, **CO**=Confidential, only for members of the consortium (including the Commission Services), **CI**=Classified

Table of Contents

Executive Summary	4
About ASPECT	5
1. Introduction	6
1.1 Purpose and Scope.....	6
1.2 Acronyms and Abbreviations.....	6
2 Applications and Operational Workflows	7
2.1.1 ERA Explorer.....	7
2.1.2 Thermal trace.....	8
2.2.1 Pensions case study.....	9
2.2.2 Humanitarian (BRC) case study.....	10
2.2.3 Agriculture case study.....	10
2.2.4 Humanitarian (Save The Children) case study.....	11
3 Technologies and design principles.....	13
3.1 Technologies.....	13
3.2 Accessibility, Mobile Optimisation, and Quality Assurance.....	14
3.3 Data lifecycle management.....	14
5 Overview of Operational and Prototype Workflows	14
5.1 Operational workflows.....	15
5.1 Prototypes workflows.....	15
6 Future Enhancements	16
References.....	16

List of tables

List of figures

Executive Summary

ASPECT designs and develops an advanced global ensemble prediction system covering seasonal to decadal timescales, tailored to user needs for climate adaptation. The project also produces various derived information, such as metrics for evaluating the predictive skill of averages and extremes (WP2), end-user-focused outputs (WP4 and WP5), temporally blended prediction–projection datasets (WP3), a catalogue of extreme events (WP3), statistically downscaled data (WP3), and specific impact-relevant indicators (WP4).

This document presents an overview of the applications, their technologies, and the workflows that support both operational and prototype activities within the ASPECT project. It outlines how the project develops user-oriented applications, implements robust operational workflows, and experiments with prototype approaches to ensure continuous improvement and readiness for future climate services.

About ASPECT

ASPECT aims to establish and demonstrate a **seamless climate information system** providing predictions with a time horizon of up to 30 years. This project is supported by cutting-edge research and by the practical application of climate information to various socio-economic sectors.

The project's goal is to enhance existing prediction systems and integrate their outputs across different timescales, combining them with climate projections to create a unified, seamless source of climate information. This integrated framework will serve as a standard for sectoral decision-making and climate adaptation. While the primary focus is on European climate information, ASPECT also extends its scope globally wherever there is substantial policy relevance, such as in disaster preparedness.

Coordinated by the **Barcelona Supercomputing Center (BSC)**, ASPECT brings together a consortium of eight partners and three UK-associated partners, combining world-leading expertise across the physical and social sciences as well as IT. The partnership covers areas such as climate prediction and projection, impact assessment and application, user engagement and co-development, communication, climate service design and delivery, and data management.

The consortium includes **BSC, CMCC, ECMWF, MPI, RHMZ, SMHI, UKMO**, the **University of Leeds**, the **University of Oxford**, the **University of Zagreb**, and **Raventós Codorníu**.

1. Introduction

1.1 Purpose and Scope

The purpose of this document is to describe the applications implemented in the Copernicus operational space and to present the corresponding prototype workflows that support their development. Furthermore, additional ASPECT-related applications and prototype workflows that facilitate visualisation, downloading and utilisation of the new ASPECT data will also be showcased.

The new applications developed within the scope of the ASPECT project provide web interfaces for accessing climate services, and their corresponding workflows ensure the reliable processing of these data.

The objective of deliverable D6.2 is to develop and deliver the users' applications and prototype workflows. Accordingly, this document offers a brief overview of the applications and workflows developed in ASPECT as part of deliverable D6.2, dedicated to the provision of these applications and prototype workflows.

1.2 Acronyms and Abbreviations

Acronyms and abbreviations used in this document are defined in the following tables.

Institutions and organisations

Acronym	Definition
ASPECT	Adaptation-oriented Seamless Predictions of European Climate
BSC	Barcelona Supercomputing Center
CMCC	Euro-Mediterranean Center on Climate Change
C3S	Copernicus Climate Change Service
ECMWF	European Centre for Medium-Range Weather Forecasts
MPI	Max-Planck Institute for Meteorology
RHMZ	Republic Hydrometeorological Service of Serbia
SMHI	Swedish Meteorological and Hydrological Institute
UKMO	UK Met Office
WMO	World Meteorological Organisation

Acronym	Definition
---------	------------

BRC	British Red Cross
CDS	(Copernicus) Climate Data Store
ESGF	Earth System Grid Federation
ERA	ECMWF reanalysis
FAIR	Findable, Accessible, Interoperable & Reusable
GRIB	GRIdded Binary or General Regularly-distributed Information in Binary form
NetCDF	Network Common Data Form
WP	Work Package

2 Applications and Operational Workflows

Several applications were created as part of the ASPECT project, targeting various audiences. These applications are a mixture of direct super-user driven applications with defined scope and parameters, and more general applications, co-developed with ASPECT WP4, which also serve as technological demonstrators.

2.1 General audience

General audience applications focus primarily on the general public with limited experience accessing and using climate data. The climate information is provided straightforwardly, offering an easily accessible, efficient and robust visualisation of the user request and the associated data.

In the ASPECT project, two main applications for the general public were developed, ERA Explorer and Thermal Trace.

2.1.1 ERA Explorer

The ERA Explorer can be used to explore historical climate data from anywhere on Earth. It is open and free, and the application can be found here:

<https://era-explorer.climate.copernicus.eu/>

The ERA Explorer implements an end-to-end workflow that provides users with seamless access to climate information. The workflow begins with the ingestion of ERA datasets from the CDS, followed by automated processing and quality-control steps to ensure consistency using modern technologies, such as close to the latest version of JavaScript and Python libraries. The application follows established DevOps practices that ensure the continuous use of up-to-date libraries and dependencies.

The ERA explorer is aimed at the general public with an interest in weather/climate, as well as policy-makers and journalists who need quick access to contextualised weather

information. Moreover, the ERA Explorer currently relies on ERA5 data to illustrate the capabilities of the ERA5 data products to a broader audience.

ERA Explorer can be used by a user to explore the vast ERA data and to calculate useful climate information supporting a mobile-friendly environment, and usage beyond traditional desktop settings. For example, to explore what the hottest month is in a specific location, when are the rainy seasons, or to answer question supporting a variety of sectors like for example agriculture (does a place often have frost in January), energy (which way does the wind usually blow and what is typically the wind power for a particular period over a place), etc.

To achieve this, new technologies have been used, making the application very powerful and responsive. For example, each plot is created on the fly in around two seconds. Another important aspect is that, because nowadays between 50 and 65% of internet traffic is on mobile devices, the application has been designed to work equally well on mobile. It is built modularly, so its components can be reused while ensuring consistency with other Copernicus and ECMWF applications.

Finally, continuous monitoring, alerting, and failure-handling strategies ensure operational reliability, with automated checks validating data integrity, workflow execution, and system performance.

2.1.2 Thermal trace

Thermal trace is another application developed by ASPECT WP6 to explore extreme temperatures, alongside other environmental factors that can impact human health.

The thermal trace application can be found here: <https://thermaltrace.climate.copernicus.eu/> and is divided into cold and heat stress sections, allowing users to select whether they are interested in exploring cold or heat extremes from anywhere on Earth. It provides maps and charts related to feels-like temperature and thermal stress and offers historical context to explore changes through time. It provides an extensive scientific discussion of the basis for thermal stress categories and how to use and interpret the data.

Thermal Trace has been designed and developed in a similar way to ERA Explorer, using the same technologies.

The key quantities displayed are feels-like temperature and heat/cold stress categories. A user can answer questions such as: which locations have experienced recently extreme heat stress, or what was the hottest feels-like temperature during the heatwaves last summer. Also, the application is relevant to a vast number of sectors like agriculture, energy, humanitarian, etc, where a user can find information about the number of days of extreme heat stress, additional days of heat stress on average compared to previous years, or the frequency of cold stress days changing over time.

As ERA Explorer, Thermal Trace is continuously being monitored to ensure operational reliability, while various strategies, such as dynamic pod autoscaling, load balancing, and resource optimisation within the Kubernetes cluster, are implemented to scale up during high application load.

2.2 Super users applications

Super User applications are tailored for a specific ASPECT user group, known as Super Users in WP4. These Super Users are central to the project and form a smaller, closely collaborating group with the project team to develop case study applications. All the relevant applications are detailed further below.

To date, four applications have been created. These include one for the pensions case study and another for the humanitarian (British Red Cross) case study. Within the ASPECT project, these tools are known as “scrolltelling” applications because they use a guided interface to explain the story behind the ASPECT data and how users can benefit from it. In addition, two dedicated interactive data applications built with Shiny have been developed, one supporting the agriculture case study and another designed for the humanitarian case study in collaboration with Save the Children.

In other words, the aim is to demonstrate how the climate information can be applied in real decision-making processes. Through visualisations and step-by-step explanations, the applications help users understand not only the data but the reasoning behind its interpretation and use.

The “scrolltelling” applications are hosted on the Copernicus infrastructure, using technology similar to that of the general public applications in terms of deployment and observability. This ensures that those applications benefit from the same operational robustness, scalability, and monitoring capabilities.

The Shiny applications are hosted locally by the institution that developed them, particularly BSC.

Finally, the primary aim of these Shiny applications is to support internal project activities and advanced users with prior experience in climate prediction. Nevertheless, they can also be used by a broader audience when supported by an experienced user (i.e. WP4). We have already identified the complexity of these applications and, as a mitigation measure, plan to explore developing a dedicated web-based application that presents the data in a simpler, more straightforward way, specifically targeting non-expert users.

2.2.1 Pensions case study

The web application developed within the pension sector is a website where users can easily access climate information about climate risk. It is designed to demonstrate how climate data, including extremes, trends, and associated risks, can be converted into relevant and actionable information for pension funds and long-term financial planning.

The web app aims to assist pension-sector stakeholders who may not have in-depth knowledge of climate science by helping them explore and understand climate risk. It also simplifies the complexity of climate science, providing clear visualisations and summaries tailored for users who are not experts in the field. It serves as the basis for further user engagement through questionnaires and surveys conducted in WP5.

The application is available under the link:

<https://apps.copernicus-climate.eu/aspect-pensions/>

2.2.2 Humanitarian (BRC) case study

Similar to the pensions case study, another scrollytelling web application was developed for the humanitarian sector, particularly for the British Red Cross. It is another website where users can find climate information on the impacts of extreme heat and support them in decision-making.

It acts as a “scrollytelling” application, combining climate data, contextual narrative, and user-oriented outputs (indicators, risk metrics, maps) to tell the story behind the data and support decision-making in humanitarian and disaster risk contexts.

The application is available under the link:

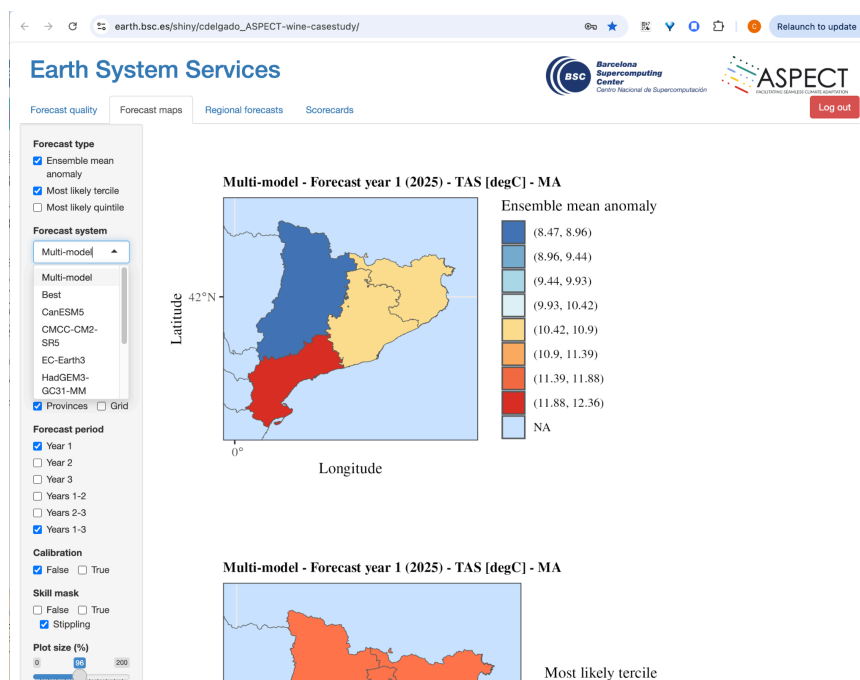
<https://apps.copernicus-climate.eu/aspect-humanitarian/>

2.2.3 Agriculture case study

In order to support the agriculture case study with Codorníu as Super User, an application was developed with the R Shiny framework.

The application is available under the link:

https://earth.bsc.es/shiny/cdelgado_ASPECT-wine-casestudy/ (the application is not open publicly, but access can be granted upon demand).



This Shiny application provides an interactive environment for exploring forecast quality and climate forecast products across multiple regions, variables, and forecast horizons. Also, users can analyse a range of skill metrics to assess forecast performance.

The application offers forecast maps, regional views, scorecards, and detailed metric visualisations. As the application focuses on the agriculture case study, users can select Catalonia, which supports a number of regional forecasts.

The application also supports only ERA5 land as reference datasets, while the forecast period range includes Year 1, Year 2, Year 3, and combined multi-year windows. Through this combination of inputs and the visual outputs, the Shiny app can give a detailed evaluation of forecast quality and facilitate a deeper understanding of multi-model seasonal and decadal prediction performance.

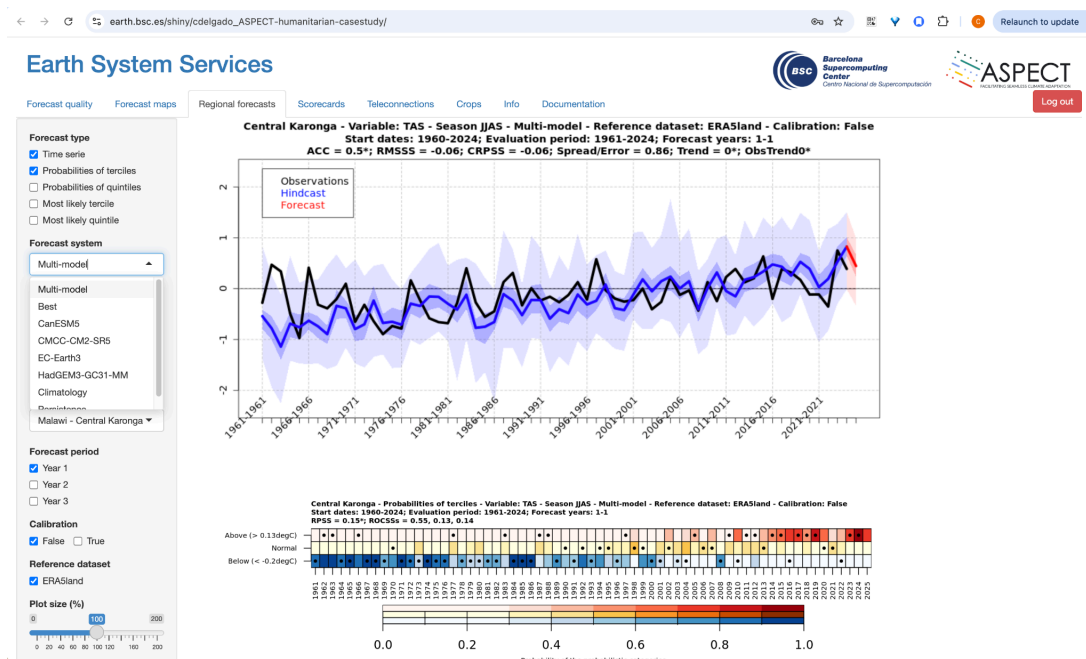
2.2.4 Humanitarian (Save The Children) case study

Similar to the agriculture case study, another Shiny application was developed to support the humanitarian case study with Save The Children as a Super User.

Similar to the previous application, this one provides access to multi-annual climate predictions and their evaluation in Malawi and Niger, with a focus on livelihood zones.

The application is available under the link:

https://earth.bsc.es/shiny/cdelgado_ASPECT-humanitarian-casestudy/ (the application is not open publicly, but access can be granted upon demand).



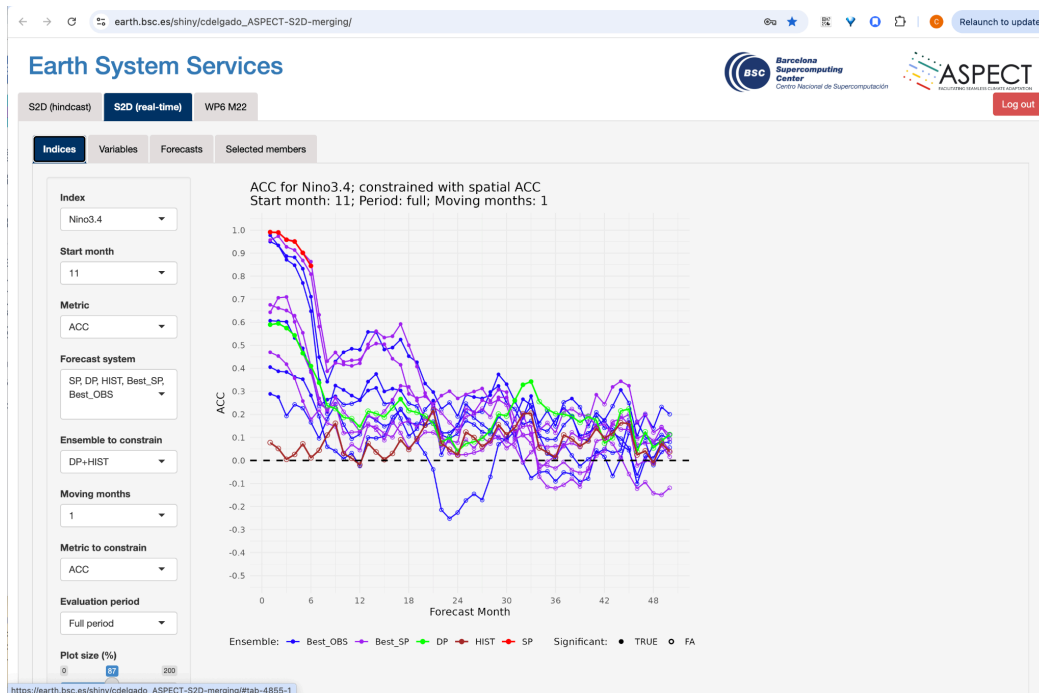
2.3 Additional applications

Under the ASPECT project, additional applications were developed, primarily to highlight the benefits of ASPECT's developments.

Those applications were developed with the R Shiny framework, so they are all called Shiny applications.

A couple of Shiny applications were developed:

- The online catalogue of decadal prediction graphical products, which is available here: https://earth.bsc.es/shiny/cdelgado_ASPECT-WP6-M22/
- The Seasonal to Decadal temporal merging. https://earth.bsc.es/shiny/cdelgado_ASPECT-S2D-merging/ (the application is not open publicly, and access can be granted upon demand).



All the above Shiny applications are delivered via a dedicated Shiny web application in R, providing an interactive, user-friendly environment for exploring climate information generated by the new ASPECT simulations and methods.

The web-based Shiny applications support systematic monitoring of the forecast by presenting a set of forecast quality metrics and products, enabling users to assess performance across forecast systems, regions, and time horizons, as well as to view the most up-to-date predictions.

All the Shiny applications developed in the ASPECT project include comprehensive forecast quality maps based on multiple skill metrics, along with ensemble-mean anomaly maps and most-likely tercile categories. The applications are complemented by documentation, ensuring users can interpret the products effectively and access methodological guidance as needed.

Finally, the additional apps can be used as training material on different types of variables and indicators, how to interpret plots, and the various types of quality scores (e.g., Relative Operating Characteristics (ROC) scores).

3 Technologies and design principles

3.1 Technologies

Copernicus applications

All the Copernicus Applications have been developed and deployed using a state-of-the-art technology stack, meaning that it ensures responsiveness, clear design, focus on high performance, and support for mobile use. The selected software components and frameworks were chosen for their scalability, modularity, and efficiency. Before beginning the design of the applications, several available technologies were surveyed, and their pros and cons were analysed.

At the end of the process, the following key technologies were chosen:

- **React:** a JavaScript framework to make the website
- **Mantine UI:** UI framework to style the website
- **React-Redux:** a state management framework to enable the user to interact with the website, and for various parts of it to talk to other parts
- **Leaflet:** Fast map-based visualisations
- **Flask:** lightweight, flexible Python web framework that provides the core backend functionality of ERA Explorer and Thermal Trace.

Furthermore, these web applications achieve high reliability and scalability through their cloud-native deployment on a Kubernetes cluster on Copernicus infrastructure. Running within a Kubernetes environment ensures that application components are automatically monitored, restarted, and load-balanced, providing continuous availability even during demand fluctuations.

This containerised, distributed setup improves resilience, fault tolerance, and operational robustness, ensuring that the ERA Explorer and Thermal Trace web applications can support a growing user base with consistent performance and minimal downtime.

Finally, these web applications use an ARCO (Analysis Ready, Cloud Optimised) Zarr archive, which is a chunked, compressed, cloud-friendly data format, to access and process large multidimensional datasets directly on the server. This structure allows selective loading of only the needed chunks, enabling high-throughput computation and very fast response times. This technology and archive format allow fast on-the-fly data processing and user interactivity.

Shiny application

For the additional applications, the technology used was the R Shiny framework, allowing users to explore data and visualise results through a browser-based interface without needing to write code themselves.

A Shiny application is an interactive web application built using the R programming language. It enables users to explore data, inspect visualisations directly through a web browser, without requiring local software installation. Shiny combines data processing, statistical computation, and visualisation in a single environment, allowing complex workflows to be exposed through intuitive user interfaces. The technology is widely used in scientific and

data-driven contexts, as it supports seamless integration with existing R-based analysis pipelines.

Shiny apps combine reactive elements (such as sliders, dropdowns, maps, and plots) with R's analytical capabilities, making them ideal for dashboards, data exploration tools, and scientific demonstrators.

3.2 Accessibility, Mobile Optimisation, and Quality Assurance

All applications were developed with a strong focus on accessibility and user experience. The applications are fully responsive, particularly those hosted on Copernicus infrastructure, ensuring seamless functionality across desktop, tablet, and mobile devices, reflecting the fact that a significant portion of internet traffic is mobile-based. The user interfaces were designed for clarity and ease of use, allowing a broad audience—from the general public to sectoral stakeholders—to explore complex climate data intuitively.

To ensure the highest standards of reliability and correctness, Copernicus applications underwent an extensive internal review process. This included rigorous code and workflow validation, UI/UX testing, and performance benchmarking. Additionally, external Evaluation and Quality Control (EQC) was conducted to verify further accessibility, scientific rigour, and adherence to Copernicus standards.

In addition, the Shiny applications were reviewed following the common ASPECT procedures applied to deliverables, ensuring consistency, quality, and alignment with project standards.

3.3 Data lifecycle management

All applications use structured data lifecycle management principles to ensure that both ERA and ASPECT datasets remain consistent and up to date. The underlying services operate on well-defined schedules aligned with the release of new ERA or ASPECT data, offering timely updates on the web interface.

Dependency management within the applications ensures that data stays current by coordinating tasks that enforce strict execution order, so each processing stage begins only after its prerequisite steps are successfully completed.

Also, to maintain high reliability, the system employs automatic recovery mechanisms that detect and respond to service interruptions, such as failed data retrievals, any kind of delays, or container restarts, by retrying failed operations or restoring application components through Kubernetes' native self-healing features. This method minimises operational disruptions without requiring manual intervention.

4 Overview of Operational and Prototype Workflows

In the ASPECT project, several operational but also prototype workflows are implemented as Jupyter notebooks.

4.1 Operational workflows

The operational workflows are the recipes behind the ERA Explorer application and are available to users as Jupyter [notebooks](#). A Jupyter Notebook is an interactive, web-based document that allows users to combine executable code, documentation, and visualisations in a single workflow. All the Jupyter notebooks come as a Jupyter-Book, which is a framework for creating structured, web-based documentation from Jupyter Notebooks.

In other words, the operational workflows consist of a set of Jupyter notebooks that provide clear, reproducible steps for accessing and processing ERA data and visualising it, in a similar way to ERA Explorer.

These notebooks serve as an interface for users in a guided environment, allowing workflows to be executed reliably while not requiring the user to be very familiar with the data.

Each notebook encapsulates a complete data workflow from downloading inputs and applying standard processing steps to producing all the available maps and plots in the web applications, ensuring consistency across users and use cases.

These Jupyter notebooks are designed to maintain reliability and reproducibility. The steps are similar to the operational applications, ensuring that users can run their analyses in a predictable and transparent manner.

The notebooks are also created following best practices in documentation and code clarity, making them easy to follow and adapt according to users' needs. Each notebook undergoes the same rigorous internal review process as the front-end components.

The operational notebooks can also act as a bridge between research and application. They harmonise all processing steps, providing a straightforward recipe, standardise visualisations, resulting in users working with climate data consistently and scientifically.

The Jupyter notebooks from ERA Explorer can be found under the link:
<https://era-explorer.climate.copernicus.eu/notebooks/intro.html>

At the moment, the Jupyter Notebooks for the Thermal Trace application are not publicly available due to licensing restrictions on the underlying data. Once the necessary permissions are obtained, the notebooks will be released openly, following the same approach used for the ERA Explorer. However, all the values are available for download in CSV format if a user wants to analyse further.

4.1 Prototypes workflows

The ASPECT prototype workflows are also implemented as Jupyter notebooks, designed to guide users in downloading and processing ASPECT data in a reproducible way. These prototypes serve as a testbed for exploring new datasets, evaluating their structure, their usability, and scientific relevance, possibly before integration into more operational workflows. Additionally, the notebooks support testing new standards and methods used in ASPECT.

Moreover, these prototype workflows could help refine methodologies, validate technical choices, and ensure that future operational workflows are robust and aligned with community best practices.

The prototype workflow designs in ASPECT offer a structured framework for experimenting with ASPECT datasets and processing methods. These prototypes are created with objectives, such as validating and testing ASPECT data and providing an initial opportunity for users to explore it.

Compared to the operational approach, the prototypes intentionally adopt a more flexible and lightweight design and may not yet be suitable for production environments. They use standard community tools, ensuring transparency and reproducibility.

All the prototype workflows are available in the GitHub repository:
<https://github.com/c3s-applications/aspect-notebooks>

For the moment, the GitHub repository is not open publicly, and access can be granted upon request. The prototype workflows are kept restricted because ASPECT is currently working to build and engage a community of users. By controlling access at this stage, the project can monitor interest and maintain a clear understanding of who is using the prototypes as they continue to evolve.

5 Future Enhancements

Future developments for the operational web applications, such as ERA Explorer and its supporting workflows, will focus on operational robustness, adding new functionality, and ensuring long-term alignment with Copernicus requirements.

Planned updates to operational workflows include optimising preprocessing routines, and supporting new datasets like the future ERA6. In other words, as C3S data evolves, workflows will be adapted to ensure compliance with updated formats, metadata and retrieval approaches. So both updates on the data and enhancements on the applications will drive the next phases of development, ensuring that the system remains scalable, interoperable, user-oriented, and ready to integrate new datasets and capabilities.

The user cases application will continue to be maintained and updated to meet the needs of WP4 and the specific requirements of the case studies.

References

European Commission, Research & Innovation, Participant Portal H2020 Online Manual, http://ec.europa.eu/research/participants/docs/h2020-funding-guide/cross-cutting-issues/open-access-data-management/data-management_en.htm

ASPECT Grand Agreement, Project: 101081460 — ASPECT — HORIZON-CL5-2022-D1-02

ASPECT Description of Action (DoA) Project: 101081460 — ASPECT — HORIZON-CL5-2022-D1-02