

Forecast-In-A-Box


Leveraging EWC to run on-demand data-driven weather forecasts

European Weather Cloud (EWC) User Workshop 2025

Harrison Cook*, Vojtech Tuma, Corentin Carton de Wiart,
James Hawkes, Tiago Quintino, Ana Prieto Nemesio,
Matthew Chantry

*harrison.cook@ecmwf.int





Copernicus ERA5:
15 billion (one off)

Hersbach, H et al. (2020)

**Training dataset
Based on a physical model**

Physics-based Forecast:

180 000
per forecast

AI Model:




0.3
per forecast

~1000x
reduction in
energy
Forecast time
reduced from
~30minutes to
~3minutes


Unit: System Billing Units (SBU)



From Supercomputer to Super Convenient!



Forecast-in-a-Box: a Forecasting System on your Desk



New Paradigm:

Instead of shipping huge forecast data, we can send the model instead

Give them the “**box**” that makes the data



earthkit

anemoi



Forecast-In-A-Box – What is it?

- **Forecast, not only Model**
 - Includes all the the product generation and processing capabilities
 - Includes catalogue of outputs and products
 - Visualisation of graphical products
- **Packaged in a ‘Box’**
 - Easy to Install, Run and Use
 - Hardware agnostic - “Whatever works for you ...”
 - Workflow designed to work on resource constrained environments

Forecast-In-A-Box – What is it?

- **Product Driven**

- Focus on what products are needed for the usecase
- Work backwards from the users request the workflow graph
- Execute in the environment provided by the user

- **Customise**

- The box contains the knowledge of how to make the products
- Just produce what you need
- Change or tweak the model, and even the number of ensembles

Forecast-In-A-Box – “What’s in the box?”

- **UI + Catalogue**

- Provides control to the user, whilst keeping it accessible
- Web Interface

The screenshot displays the Forecast-In-A-Box web interface, which is divided into four main sections at the top: **Model** (Configure the Model), **Products** (Choose Products), **Environment** (Configure Execution), and **Confirm** (Execute the graph).

The **Model** section is currently active and contains several categories of options:

- Plots:** Display products as plots. Options: **Interactive Map** (selected), **Maps**.
- Standard Products:** Standard products. Options: **Deaccumulated** (selected), **Output**.
- Forecast Statistics:** Statistics over time for each member. Options: **Maximum**, **Mean** (selected), **Minimum**, **Standard Deviation**.
- Thermal Indices:** Thermal Indices. Options: **2 metre relative humidity**, **Apparent temperature** (selected), **Humidex**, **Normal effective temperature**, **Wind chill factor**.
- Ensemble:** Capture the distribution of the ensemble. Options: **Ensemble Mean**, **Ensemble Standard Deviation**, **Quantiles**, **Threshold Probability**.

The **Products** section shows the **SimpleMapProduct** configuration with the following options:

- Parameter ***: Dropdown menu.
- Level Type ***: Dropdown menu.
- Level List ***: Dropdown menu.
- Step ***: Dropdown menu.
- Reduce ***: **True** (selected).
- Domain ***: **DataDefined** (selected).

The **Environment** section shows the **Domain of the map** dropdown menu.

The **Confirm** section shows the **Domain of the map** dropdown menu.

The **Selected (3)** section shows the selected configuration items:

- Plots/Interactive Map**: param: ["2t","msl"], levtype: "sfc", step: [6,12,18,24].
- Standard/Deaccumulated**: param: ["tp"], levtype: "sfc", step: [6,12,18].
- Standard/Output**: param: ["tcw"], levtype: "sfc", step: [54,60,66,72], reduce: "True", format: "grib".

The interface also includes a **Quick Start** section with a clock icon and a **Standard** section with a bar chart icon. The **Quick Start** section includes a button labeled **ALPHA** and a description: "Quickly start with a simple interface to get you going." The **Standard** section includes a description: "Select from products, comprehensive".

Forecast-In-A-Box – “What’s in the box?”

- **Task Definitions**

- Constructed from Anemoi and Earthkit functions



Fetch Initial Conditions

Regrid and Interpolate

Execute ML model



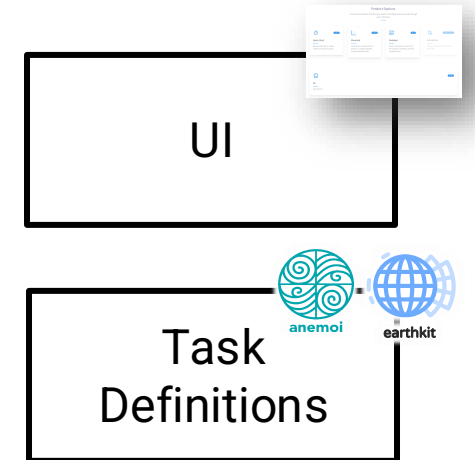

Compute statistics

Transform outputs

Visualise

Forecast-In-A-Box – “What’s in the box?”


- **Orchestration**
 - Task graph construction and scheduling



Forecast-In-A-Box – “What’s in the box?”

- **Execution Layer**


- Executes the task graph
- Hardware agnostic
- Heterogenous execution environment



Forecast-In-A-Box – “What’s in the box?”

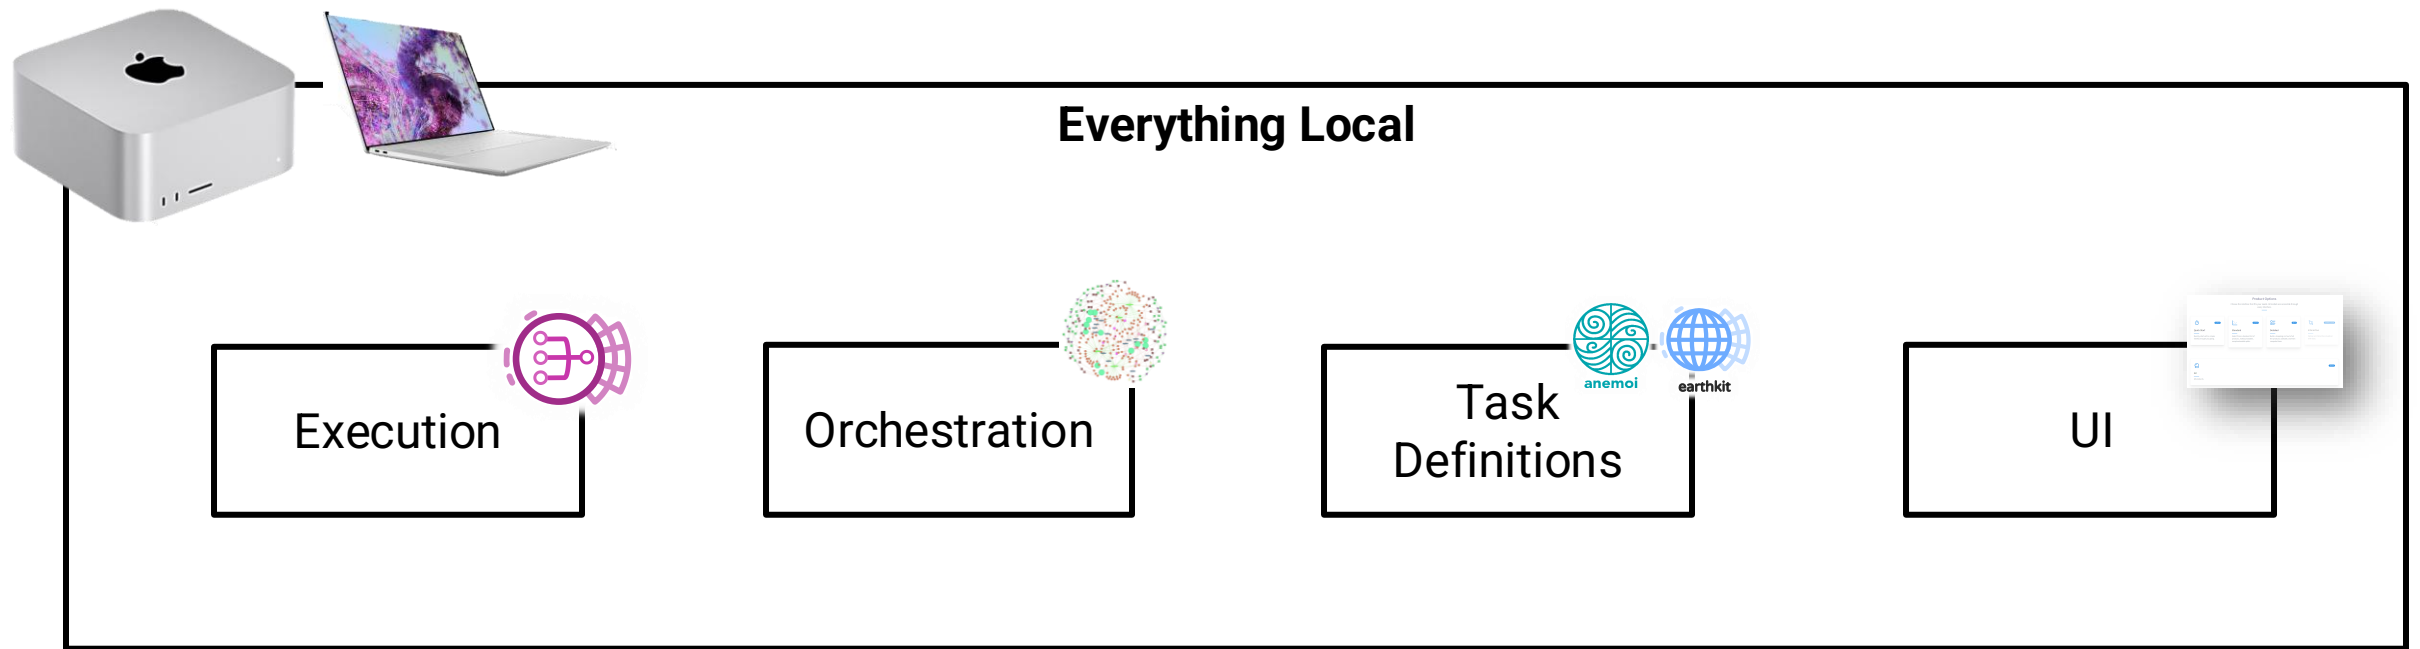
- **Putting it all together**

- A complete forecasting system
- Capable of running any Anemoi inference model
- Runs ECMWF’s full operational product and post processing software
- Provides ECMWF production quality graph and plotting software




Forecast-In-A-Box – “What’s in the box?”

- Can be deployed in many ways...



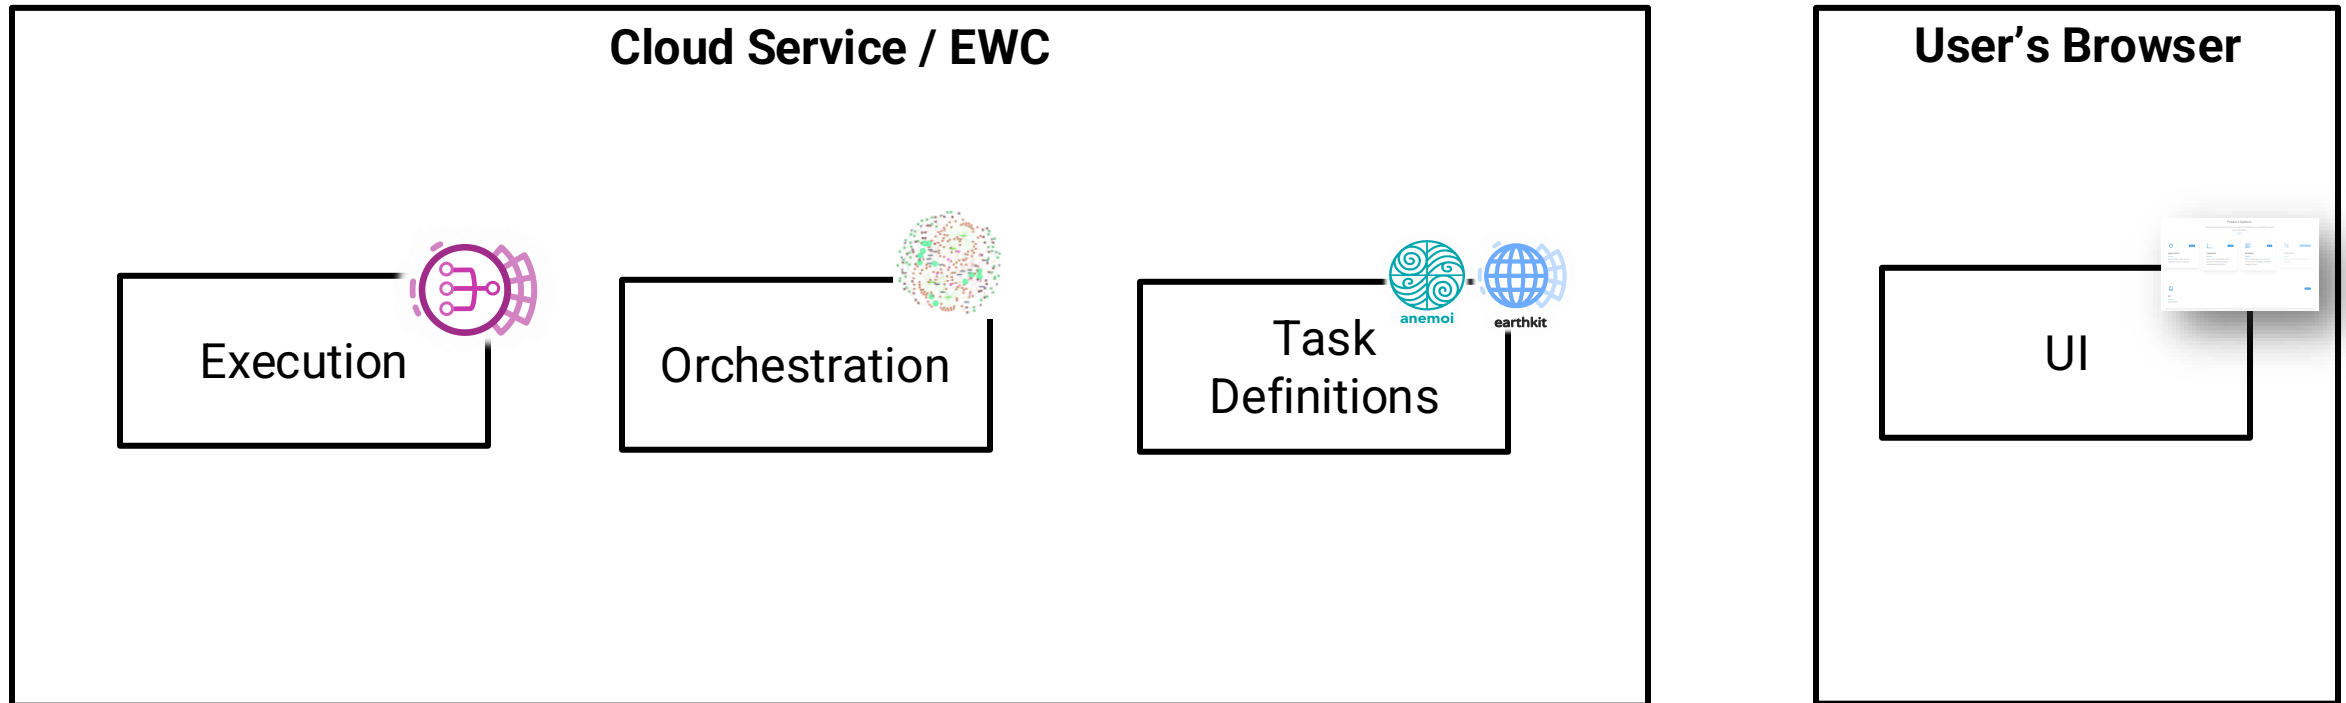
Forecast-In-A-Box – “What’s in the box?”

- Can be deployed in many ways...



Forecast-In-A-Box – “What’s in the box?”

- Can be deployed in many ways...



Forecast-In-A-Box – Implementation on EWC

EWC - VM

Cascade

Forecast

Rocky Linux™ FIAB ★

Running Env: Dev Type: Rocky GPU Plan: 16cpu-128gbmem-30gbdisk-40gbgpu

EDIT ACTIONS DELETE

✓	✗	99.948%	N/A	10%	12%	24%
HEALTH	LAST BACKUP	AVAILABILITY	RESPONSE TIME	MAX CPU	MEMORY	STORAGE

Summary Resources Runtime Storage Network Monitoring Backups History Console

Group: european-weather-cloud Cloud: [ecmwf-cci2-ecmwf-development](#) Date Created: 30/04/2025 01:37 PM

Owner: Harrison Cook Layout: rocky-9-gpu Version: 9

Cores: 16 Memory: 128.0GiB Total Storage: 1.1TiB

Source Image: Rocky-9.5-GPU-20250204155220 Provision Time: 2 minutes 20 seconds

Model Store

User DB

Forecast-In-A-Box – Implementation on EWC - Current

- Benefits
 - Collocated with ATOS
 - Initial conditions retrieval from MARS
 - Simple allocation of GPU's & VM's
 - Full control over system
 - Easy collaboration

Currently deployed on the EWC, and has been for months

Forecast-In-A-Box – Implementation on EWC - Future

- Use Kubernetes Clusters to run at scale
 - Dynamic Re-allocation of resources
 - Missing GPU's currently 😞
- Systems deployed on various hardware
 - Database doesn't need a GPU, etc...

Forecast-In-A-Box – Usage of the EWC



Easy
prototyping




Simple
deployment

<https://jellyseerr.eldrest.com>



Stable Interface
and system



Expandable to
meet needs

Forecast-In-A-Box – Projects

