HiDALGO

Building Cloud-Based Data Services to Enable Earth Science Workflows across HPC Centres

John Hanley, Milana Vuckovic, James Hawkes, Tiago Quintino, Stephan Siemen, Florian Pappenberger

John.Hanley@ecmwf.int

Overview

Introduction to ECMWF

The Data Challenge

HiDALGO & ECMWF

European Centre for Medium-Range Weather Forecasts

- Established in 1975.
- Intergovernmental Organisation
  - 22 Member States | 12 Cooperation States
  - 350+ staff
- 24/7 operational service
  - Operational NWP centre
  - Supporting NWS (coupled models) and businesses
- Research institution
  - Closely connected with researchers worldwide
- Operates two Copernicus Services
  - Climate Change Service (C3S)
  - Atmosphere Monitoring Service (CAMS)
- Supports Copernicus Emergency Management Service (CEMS)

What do we do?

ECMWF

Short-range weather forecast	Medium-range weather forecast	Long-range weather forecast	Climate prediction
Very high resolution Regional models	**High resolution** Global models	Predicts statistics of weather for coming month or season	CO_2 doubling and other scenarios
1-2 hour production schedule	6-12 hour production schedule	1-8 times a month production schedule	

What do we do?

Operations – Time Critical

- HRES 0-10 day, 00Z+12Z, 9km @ 137 levels
- ENS 0-15 day, 00Z+12Z, 18km @ 91 levels
- BC 06Z and 18Z, 0-5 days hourly
- 100 TiB, 85 Million products
- Real-time Dissemination, 200 destinations world-wide

Research – Non Time Critical

- 100s Daily active experiments to improve our models
- Reforecasts, Climate reanalysis, etc

Meteorological Archive

- > 300 PiB of data @ 5000 daily active users
- 250 TiB added per day

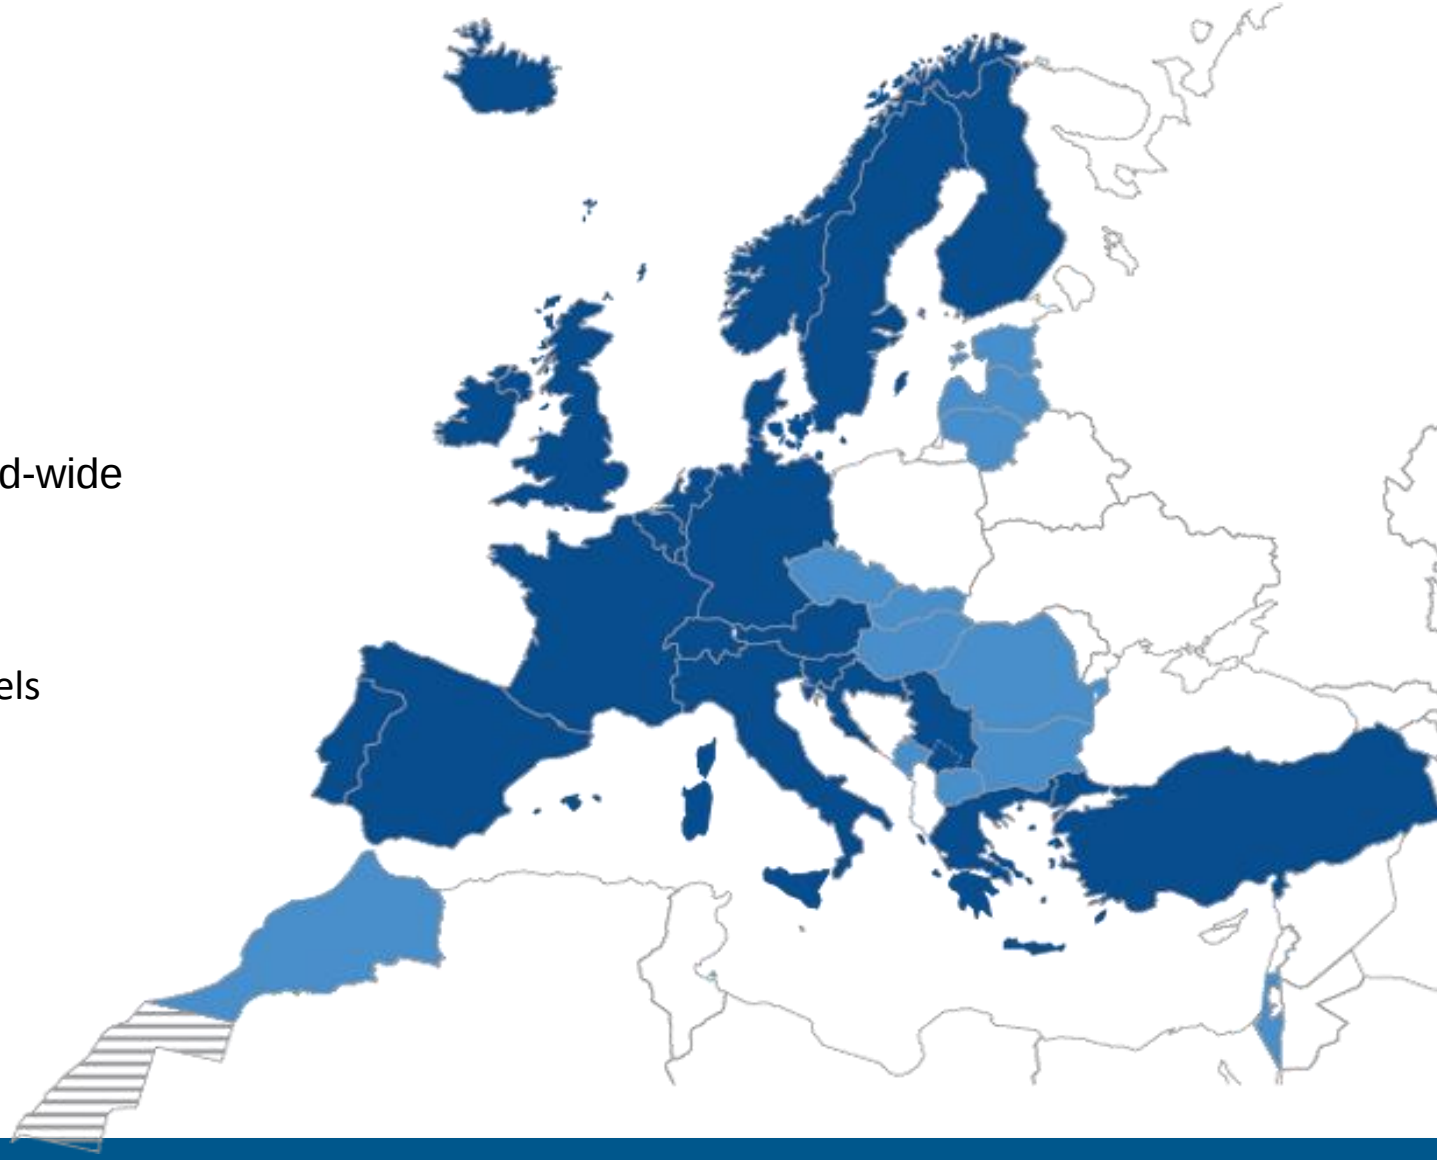

ECMWF's Facilities

HPC

2x Cray XC40 HPC

2x 129,960 cores Xeon EP E5-2695 Broadwell

2x 10 PiB Lustre PFS storage

Top500 42nd/43rd

Cloud

- **[SaaS]** Copernicus Data Storage (CDS) – Operational
- **[PaaS]** European Weather Cloud – Pilot currently being setup
- **[XaaS]** WEkEO www.wekeo.eu

Archive

Largest Meteo archive 4x Oracle (Sun) SL8500 tape libraries

~ 140 Tape drives

+ 100 TiB / day operational + 150 TiB / day other

Introduction to ECMWF

The Data Challenge

HiDALGO & ECMWF

ECMWF Data Growth – History and Projections

Historical Growth of Generated Products

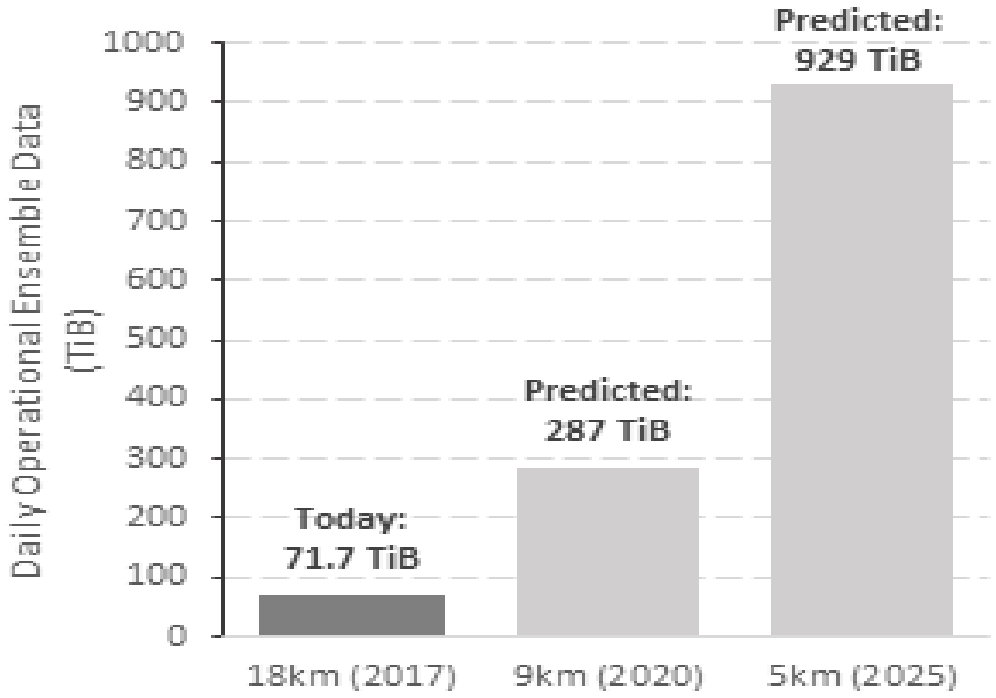

Model Output Projected Growth

- Data archival and retrieval system for ECMWF data
  - > 300 PB primary data

- Largest meteorological archive in the world
  - Direct access from Member States
  - Available to research community worldwide

Types of data growth

➔Reliability

Ensembles

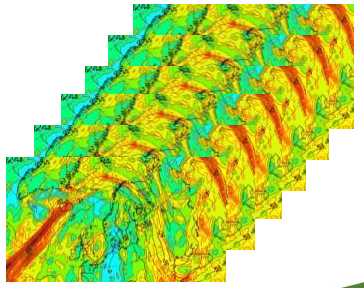

Traditional weather science domain

➔Range

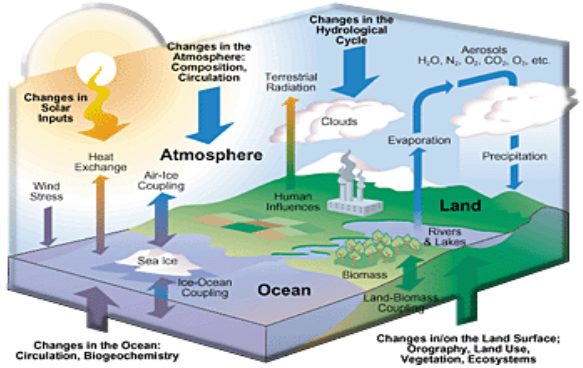

Traditional climate science domain

Model complexity

Model resolution

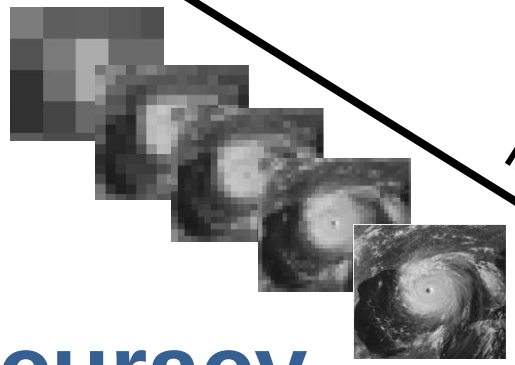

➔Accuracy

Today: we need high-resolution, 'Earth system' model ensembles to perform at all scales!

The data challenge

- No user can handle all our data in real-time
  - Much of ECMWF (Ensemble) forecast stays unused!
  - ECMWF always looks for new ways to give user access to more of its forecasts
  - Not made easier by domain specific formats & conventions

- Dissemination system
  - Sophisticated push system to disseminate 100TBs in real time across the world

- Web services
  - Develop & explore (GIS/OGC) web services to allow users to request data on-demand

The Key Challenge:

How do we improve user access to such volumes of data?

How can you access data today?

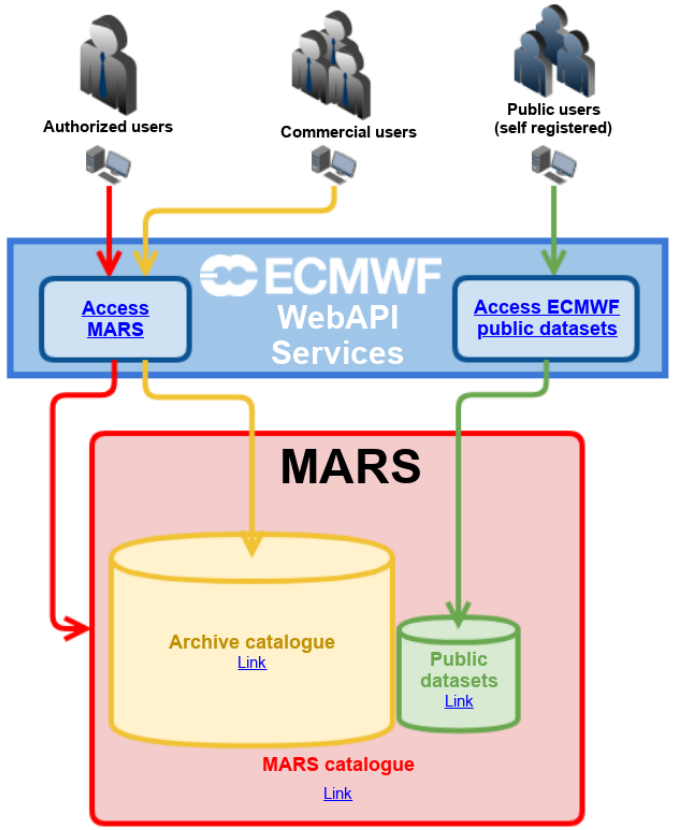


```
server.retrieve({
    "class": "er",
    "dataset": "era15",
    "date": "1991-11-01/to/1991-11-30",
    "levtype": "sfc",
    "param": "167.128",
    "repres": "gg",
    "step": "0",
    "stream": "oper",
    "time": "12:00:00",
    "type": "an",
    "target": "output",
})
```

- Try it out yourself:
  https://pypi.org/project/ecmwf-api-client/

  https://apps.ecmwf.int/datasets/

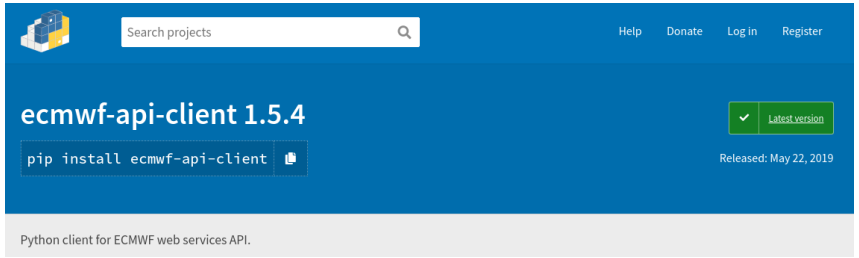

How can you access data today?

Copernicus Climate Data Store (CDS)

- New portal to find / download and work with Copernicus climate change data

- Many data sets too large for users to work locally
→ therefore it offers server side processing

- High-level descriptive Python interface
  - Allow non-domain users to build apps

- Try it out yourself:
  https://cds.climate.copernicus.eu

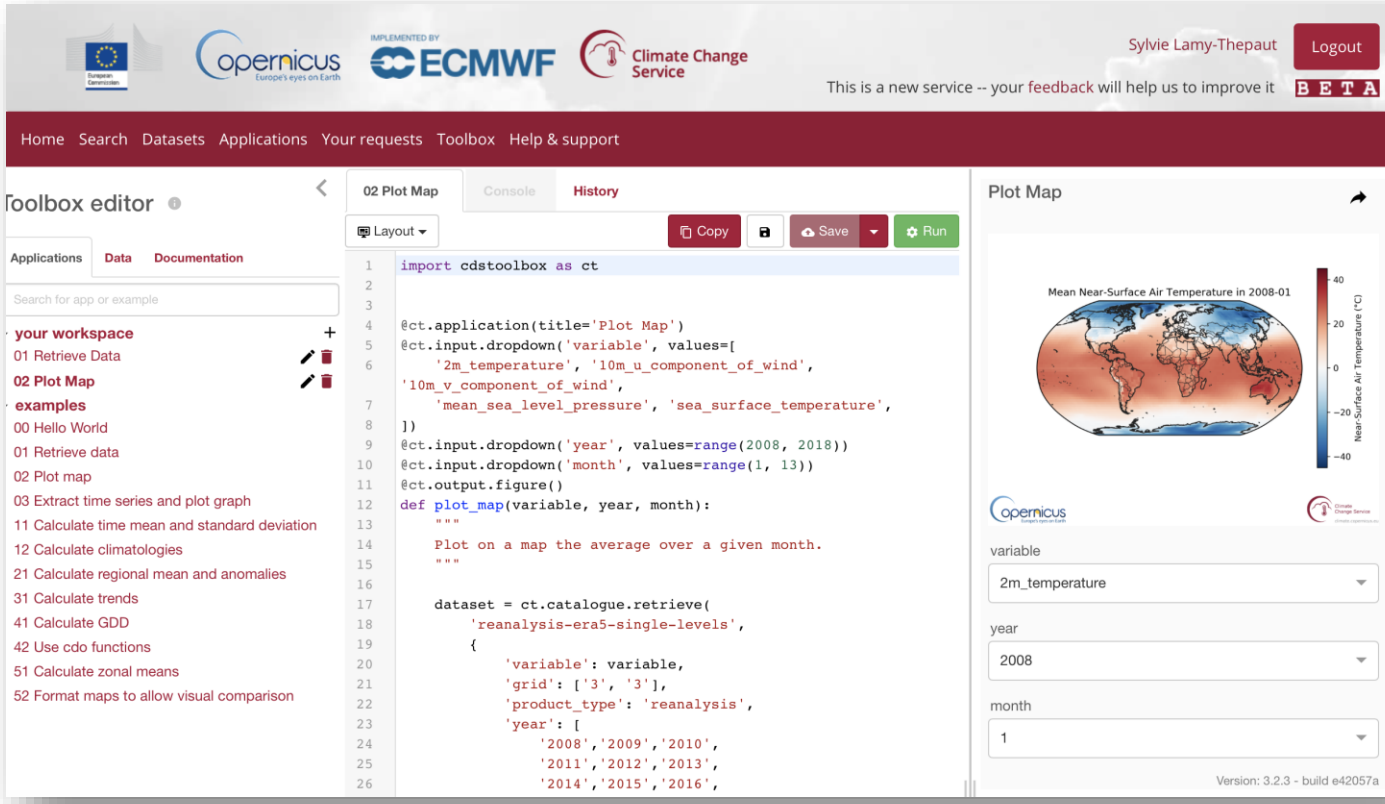

CDS toolbox

ECMWF Novel Data Flows

Producer

FDB

PFS

Cloud

Consumer

- Bring **users** to the data
- Use data while it is **hot**
- Access using **scientifically meaningful metadata**

Archive

HDD

Tape

MARS

Introduction to ECMWF

The Data Challenge

HiDALGO & ECMWF

HiDALGO:

HPC and Big Data Technologies for Global Systems – European project funded by the Horizon 2020 Framework Programme of the European Union carried out by 13 institutions from seven countries.

The Vision:

To advance technology to master global challenges

The Mission:

To develop novel methods, algorithms and software for HPC & HPDA to accurately model and simulate the related complex processes. To also enable coupling of pilots to external data sources (e.g. ECMWF).

Pilot Test Cases:

1. Migration pilot (Derek Groen, Brunel University, UK).
2. Air pollution pilot (Zoltán Horváth, SZE, Hungary)
3. Social networks pilot (Robert Elsässer, PLUS, Austria)

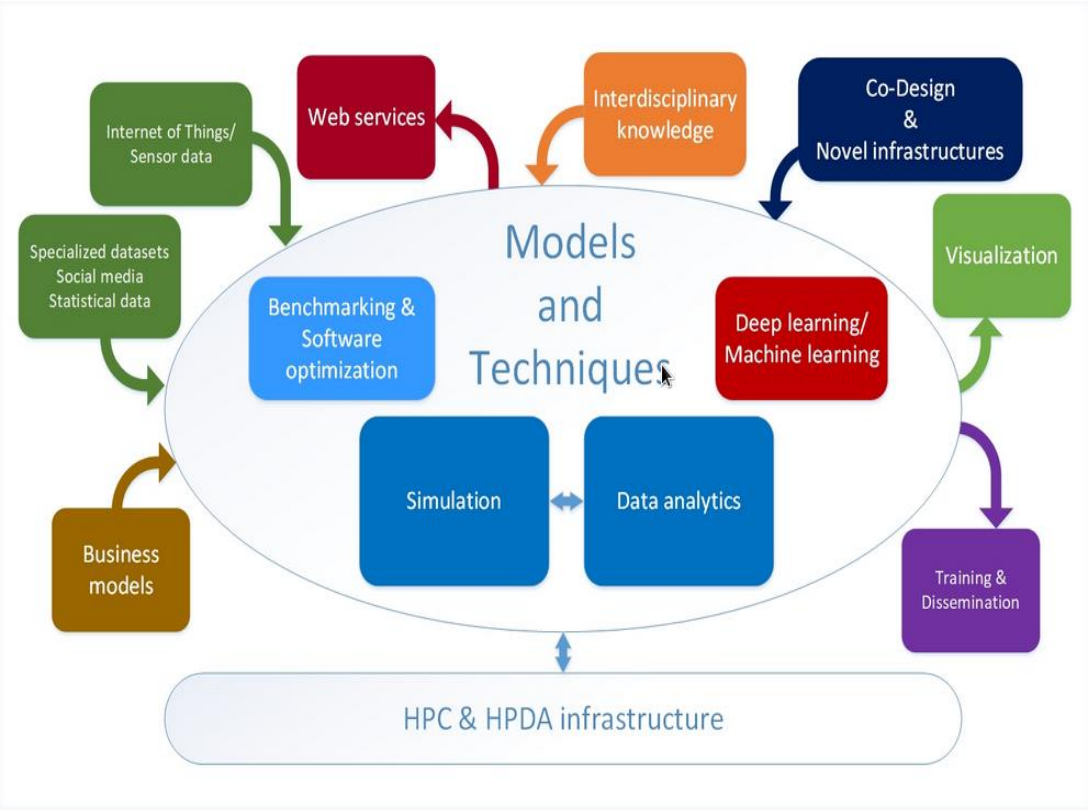

ECMWF's role

Enable coupling as a means to build a workflow

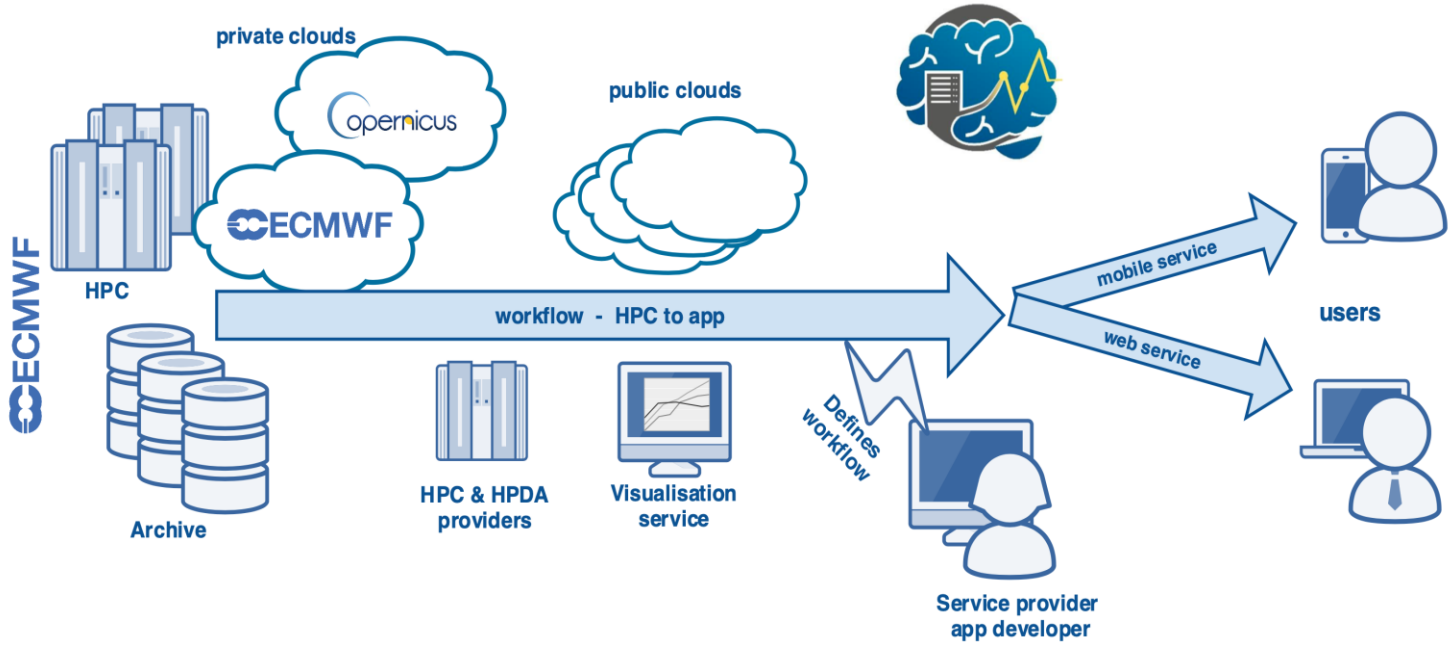

With a "closed" HPC system, ECMWF brings in valuable experience on how these systems can be integrated in larger workflows --> this can be a model for many similar HPC systems around Europe!

HiDALGO HPC & Cloud Facilities

HLRS

- Cray XC40 “Hazelhen”
- 7.4 PFLOPS
- 185,088 cores

PSNC

- Huawei CH121 “Eagle”
- 1.4 PFLOPS
- 32,984 cores

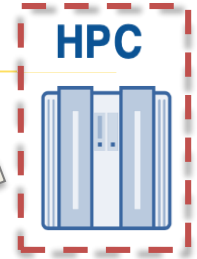

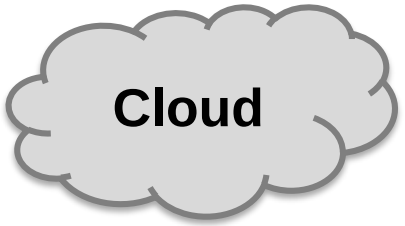

Data analytics

Deep learning/
Machine learning

Weather and Climate Data Coupling

Two step approach to coupling

Step 1: Static coupling (1st year of project - 2019)

- coupling with static reanalysis (climate) data for the purposes of pilot model calibration

Completed!

Step 2: Dynamic coupling (2nd year – 2020 onwards)

- coupling with forecast data via a REST API

Climate Data Store (CDS)

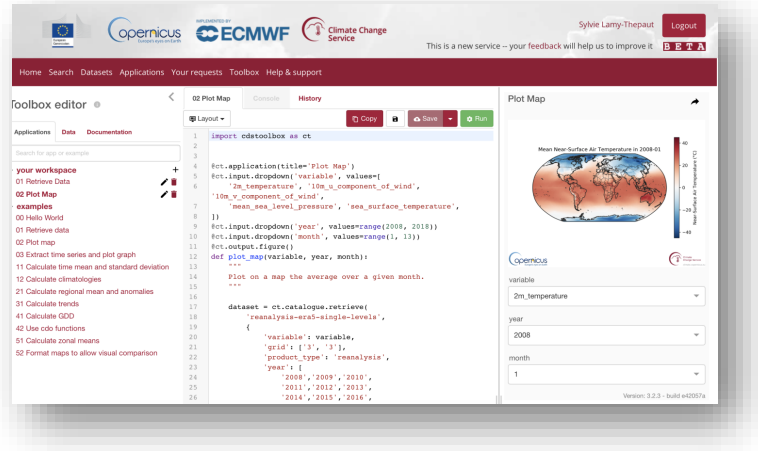

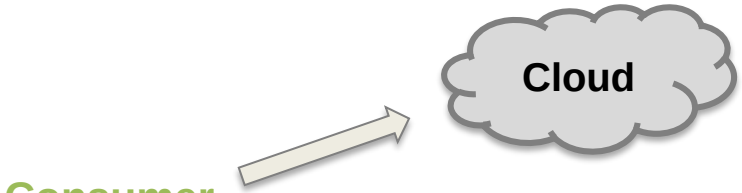

- Bring **users** to the data
- Use data while it is **hot**
- Access using **scientifically meaningful metadata**

Vision:

To enable users to build custom workflows utilizing ECMWF's weather forecast and climate data

Providing ECMWF data to the pilot applications

The main requirements:

1. Bring users to the data and avoid moving the data out of the data centre.
2. Provide users with computing resources collocated directly with data.
3. Align with data-centric approach of *"move the compute, not the data".*

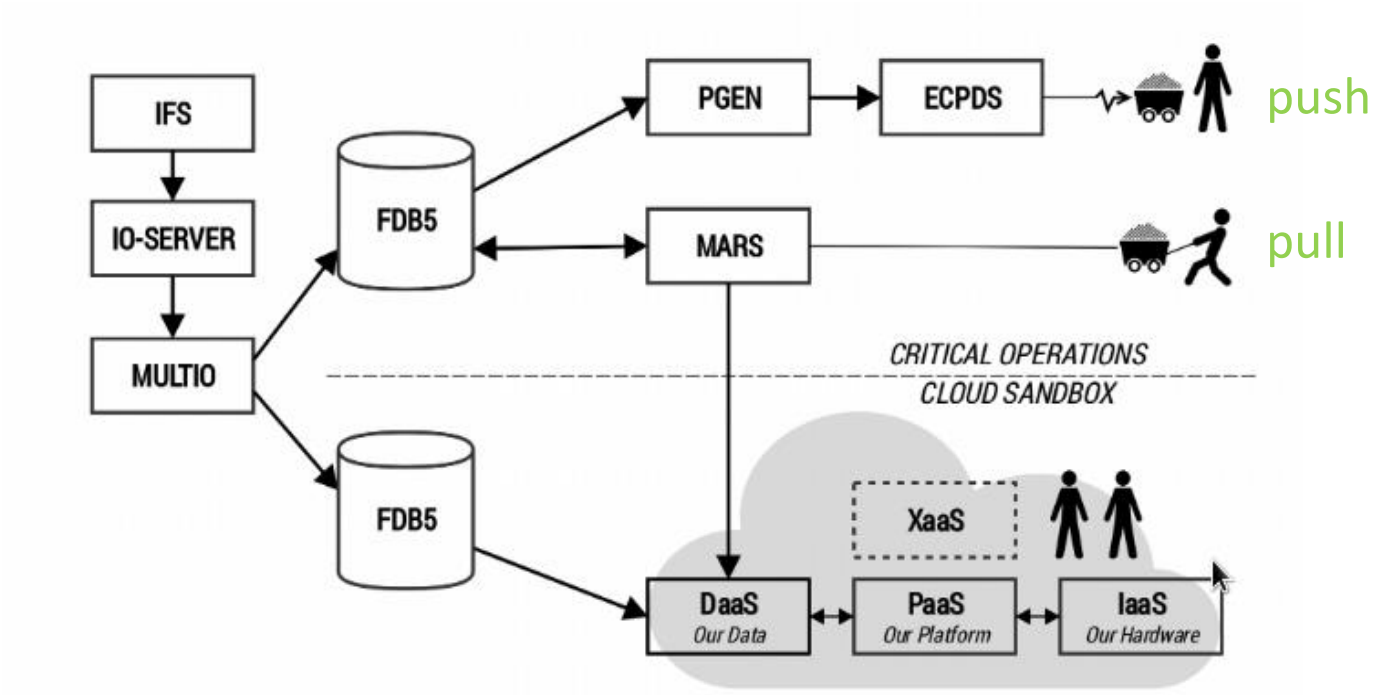

How to enable this:

1. Mechanism to pull/push data from ECMWF.
2. Mechanism to run custom post-processing at ECMWF.
3. Mechanism to explore what data and processing options ECMWF offers.

Cloud Data-as-a-Service: Polytope

Service designed for efficient provisioning of meteorological data to Cloud and HPC applications

- Under development at ECMWF
  - Deployed internally at ECMWF
  - Accessible externally
  - Beta-tested via European Weather Cloud
- Exposes a RESTful API
- A CLI and python API aid the users interacting with the Polytope API
- It interfaces MARS directly
- Will implement hyper-cube data access

Step 1: submit request

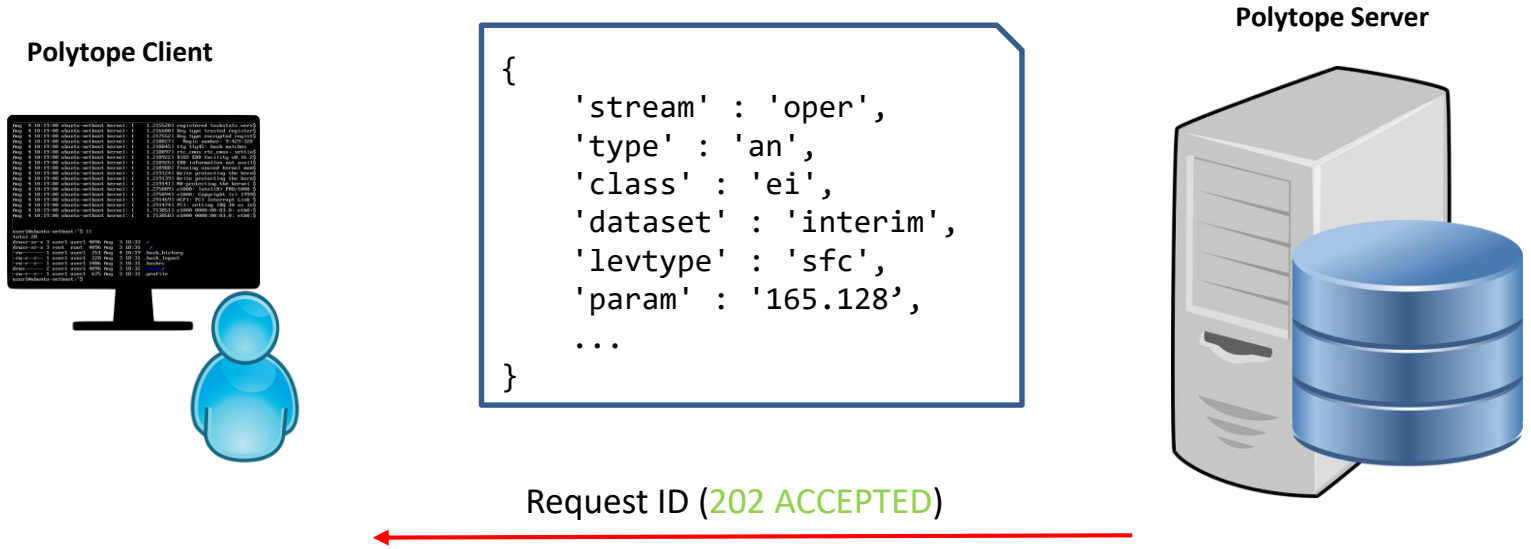

Cloud Data-as-a-Service: Polytope

Service designed for efficient provisioning of meteorological data to Cloud and HPC applications

- Under development at ECMWF
  - Deployed internally at ECMWF
  - Accessible externally
  - Beta-tested via European Weather Cloud
- Exposes a RESTful API
- A CLI and python API aid the users interacting with the Polytope API
- It interfaces MARS directly
- Will implement hyper-cube data access

Optional step: list requests

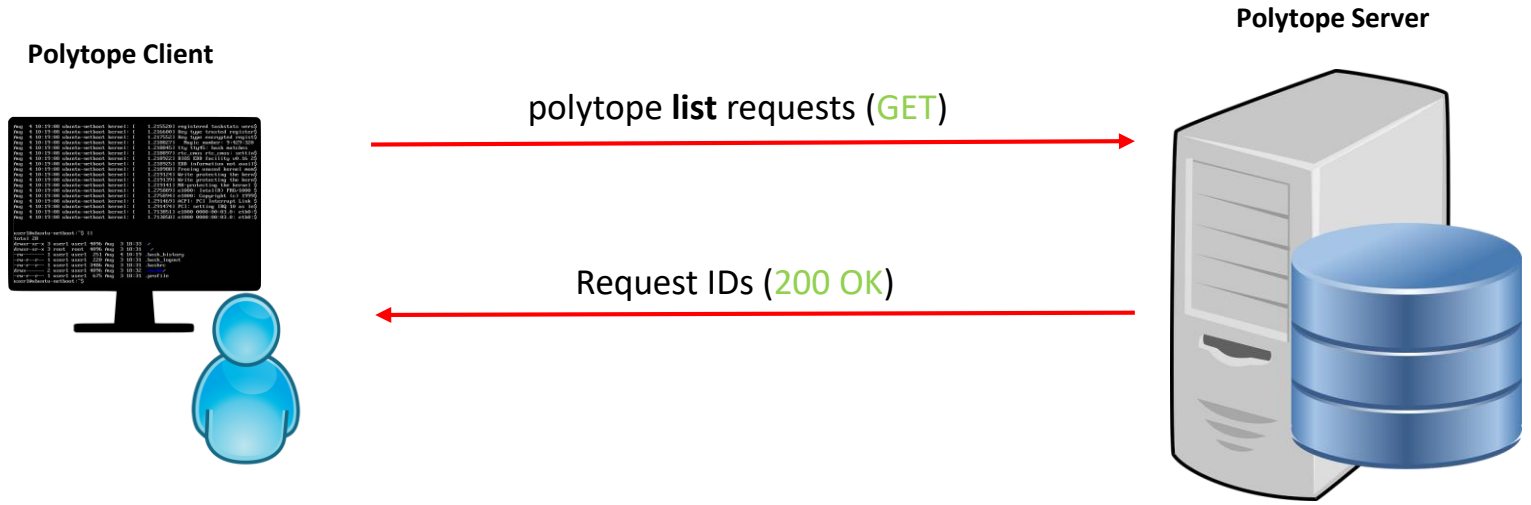

Cloud Data-as-a-Service: Polytope

Service designed for efficient provisioning of meteorological data to Cloud and HPC applications

- Under development at ECMWF
  - Deployed internally at ECMWF
  - Accessible externally
  - Beta-tested via European Weather Cloud
- Exposes a RESTful API
- A CLI and python API aid the users interacting with the Polytope API
- It interfaces MARS directly
- Will implement hyper-cube data access

Step 2: poll for data

Cloud Data-as-a-Service: Polytope

Service designed for efficient provisioning of meteorological data to Cloud and HPC applications

- Under development at ECMWF
  - Deployed internally at ECMWF
  - Accessible externally
  - Beta-tested via European Weather Cloud
- Exposes a RESTful API
- A CLI and python API aid the users interacting with the Polytope API
- It interfaces MARS directly
- Will implement hyper-cube data access

Step 3: delete completed request

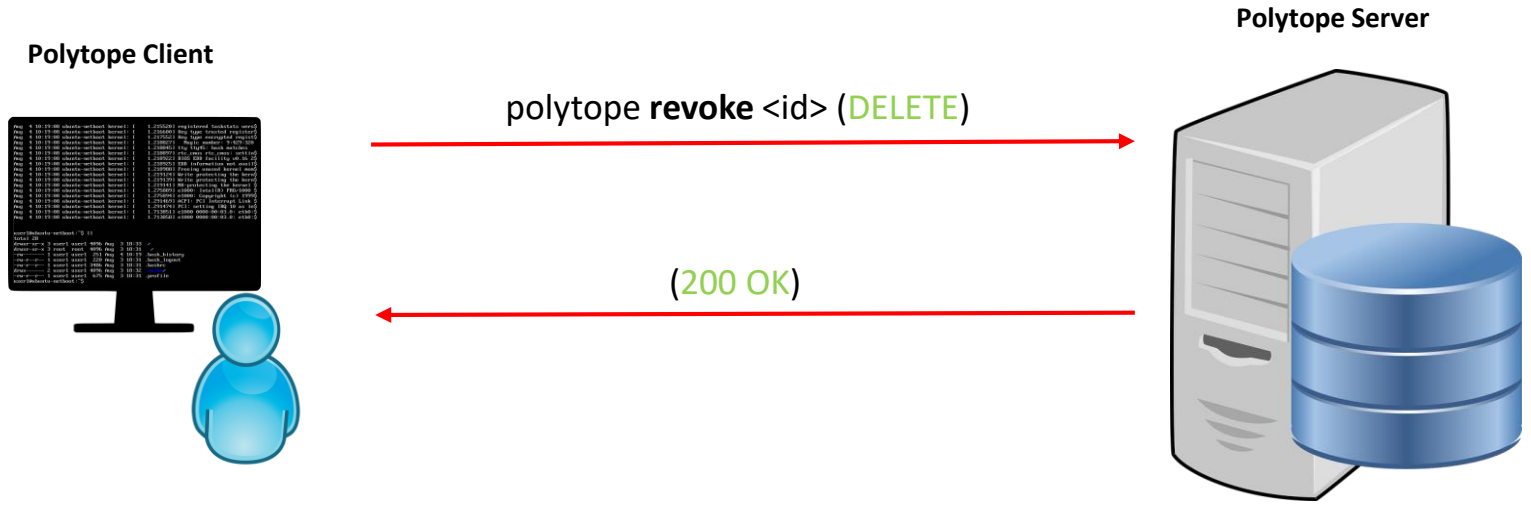

Cloud Data-as-a-Service: Polytope

The system has been developed as a set of independent services for scalability (elastic architecture, muti frontend, workers, etc.), ease of deployment (Kubernetes support), with a shallow software stack.

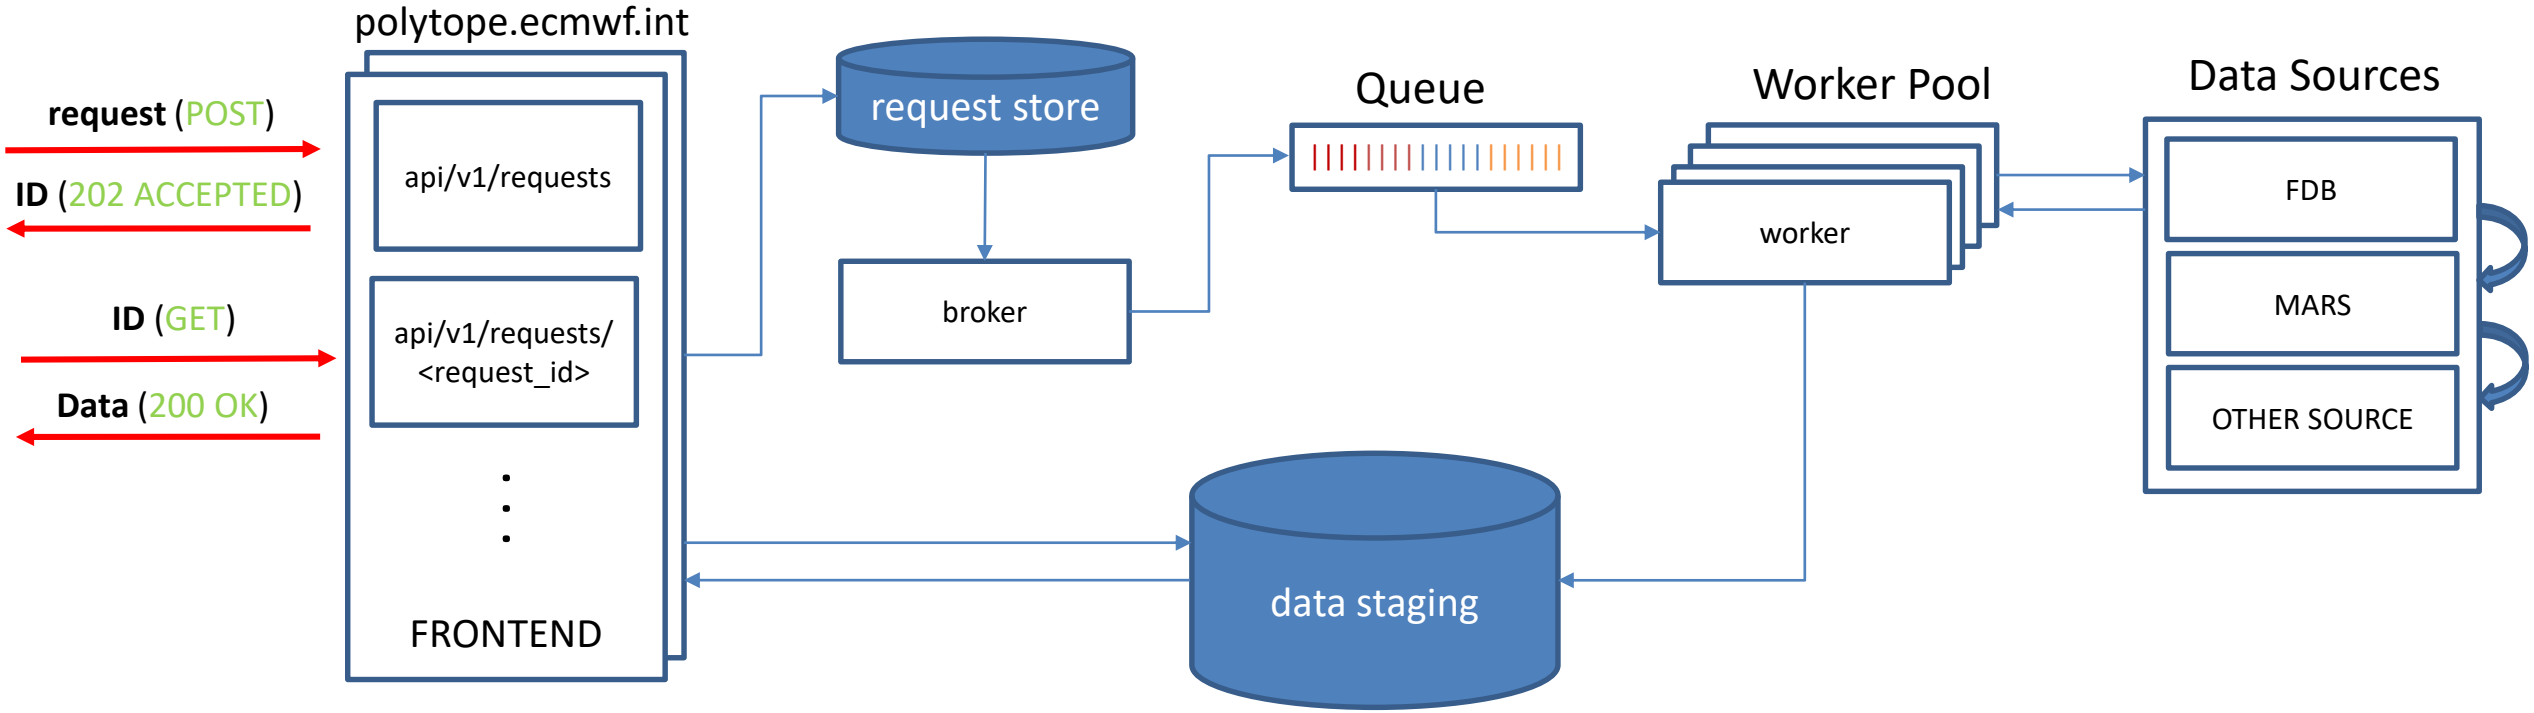

THANK YOU !

QUESTIONS ?

Acknowledgements

This work has been supported by the HiDALGO project and has been partly funded by the European Commission's ICT activity of the H2020 Programme under grant agreement number: 824115. This paper expresses the opinions of the authors and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this paper.