Polytope: Datacube Access for ECMWF's Weather Data

James Hawkes, Nicolau Manubens, Emanuele Danovaro, Stephan Siemen, Baudouin Raoult, Tiago Quintino

ECMWF, Reading, United Kingdom

James.Hawkes@ecmwf.int

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825532 and 824115.

Challenge

Every day, ECMWF produces approximately 120TiB of raw weather from high-resolution and ensemble forecasts. This data is used by our product generation system (PGEN) to compute derived data and push this data to our member states and other users. The raw data is also stored in the world's largest meteorological archive (MARS), currently holding over 300 PiB of primary data -- which is served around the world on demand.

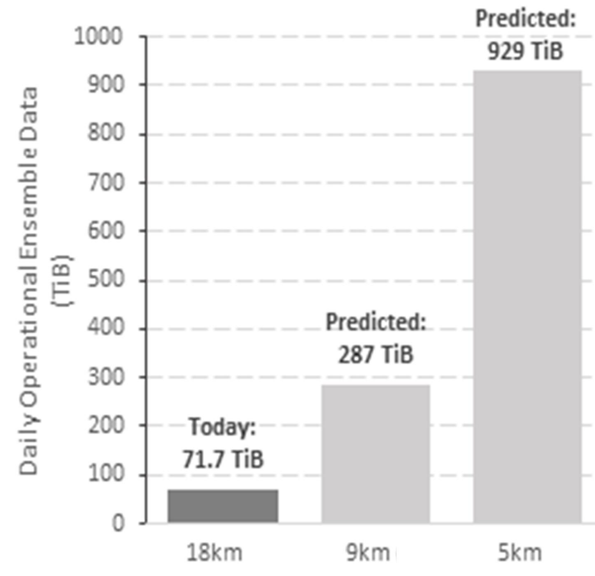

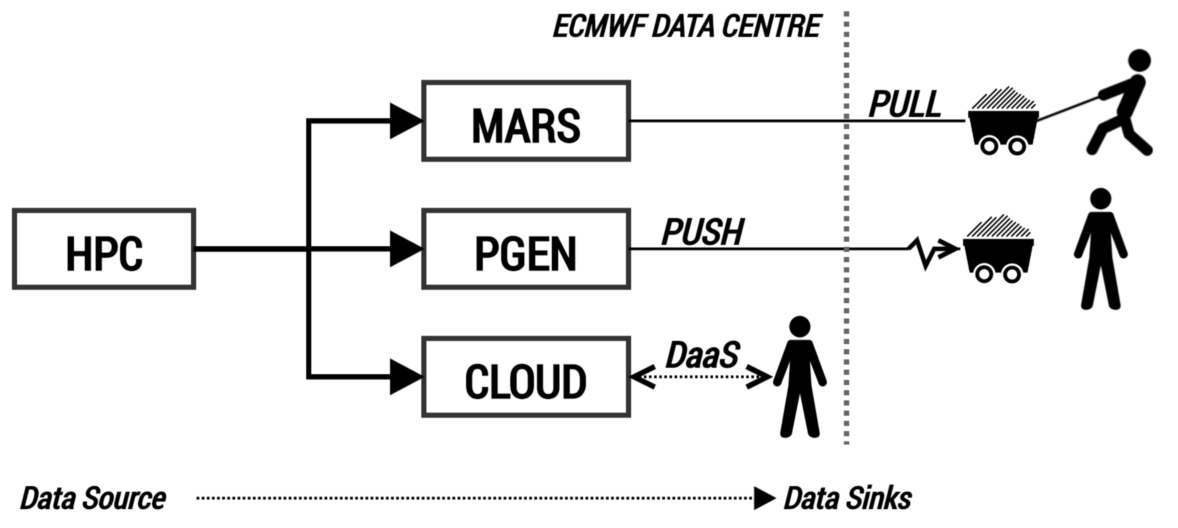

As the resolution of ECMWF's forecasts increase over the next few years, the amount of daily raw data produced will increase into the petabytes, and it will become difficult to move this data out of the HPC. To solve this problem, ECMWF is pursuing a convergence of HPC and Cloud, by building a cloud platform (the European Weather Cloud) in-situ with our HPC, and providing the cloud with "data-as-a-service" (DaaS). However, bringing the users close to the data is just one part of the puzzle. The bigger question is **how can users meaningfully interact with petabytes of data?**

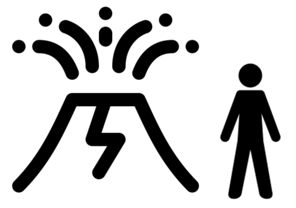

User A

Wants to retrieve wind data in a particular region, on demand, to quickly predict particle dispersion following an ecological disaster. They want to get the data quickly and the area of interest is dynamic.

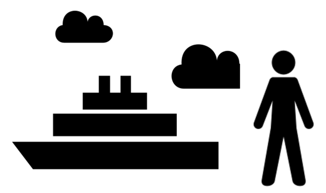

User B

Wants to retrieve the weather forecast along a spatio-temporal path, in order to optimize a ship's route across the ocean. They do not want to download entire regions or 4-D bounding boxes of data, they just want the few bytes they are interested in, in a simple format for quick data analysis and decision making.

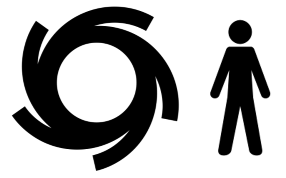

User C

Wants to perform machine learning on tropical cyclone data, and wants to retrieve weather data following multiple cyclone tracks. This user does not want to process or store large volumes of intermediate data in their already-expensive ML pipeline.

The current methods of data retrieval don't suit these users, because they would have to request large chunks of data and perform their own client-side data extraction. If these users could specify exactly what data they required, it would:

- **Reduce the post-processing burden on the user, and improve ease-of-access;**
- **Reduce the amount of data transmitted to, stored and processed by the user, improving efficiency.**

Concept

The raw meteorological datasets we wish to serve can be thought of as an n-dimensional array, or "datacube", with axes for forecast date, forecast timestep, parameter, vertical level, latitude, and longitude. Extracting data from these datacubes by defining extents along each orthogonal axis (the bounding box approach) is simple, but results in lots of excess data – especially in higher dimensions.

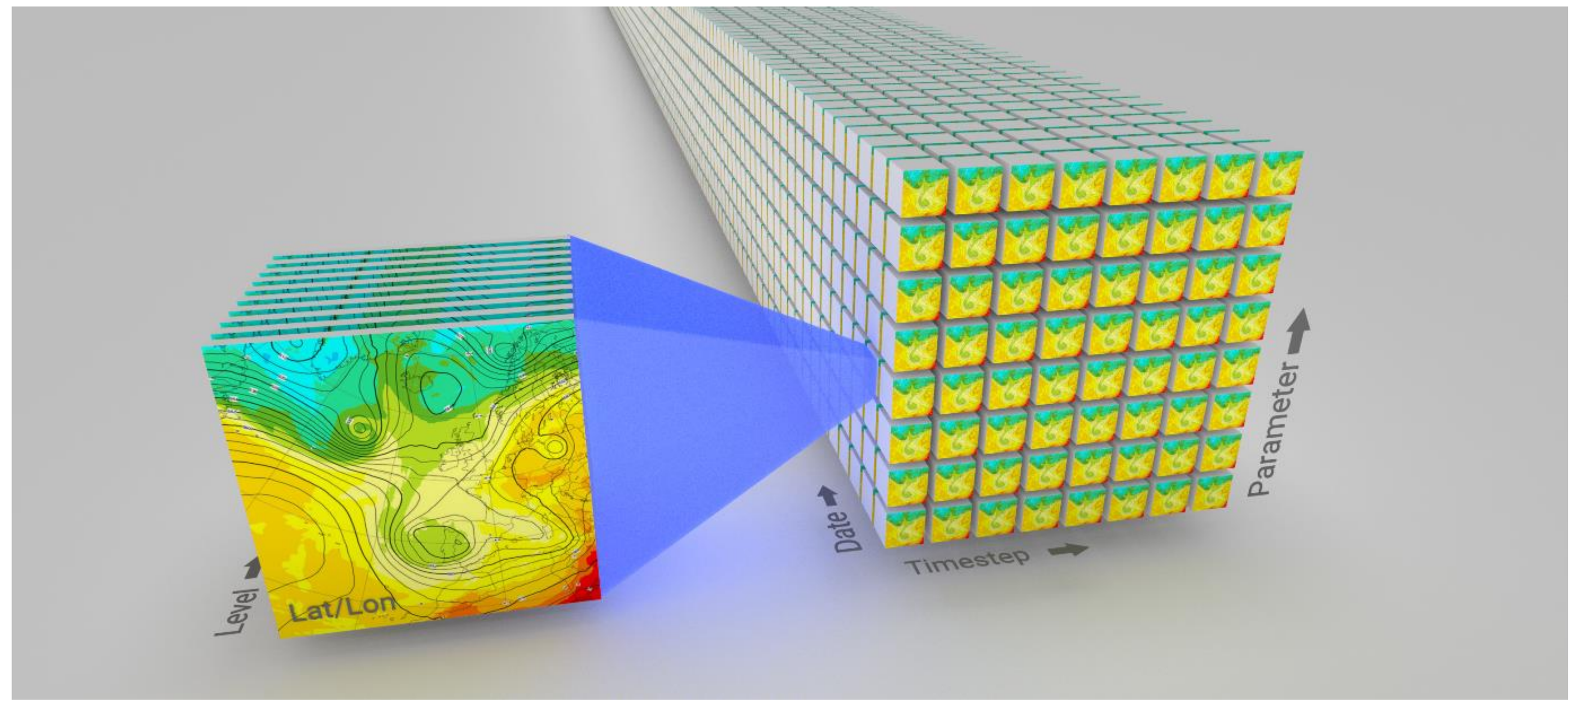

Conceptually, it is possible to define a subset of this n-dimensional space using a Polytope, an n-dimensional polygon stencil which cuts the datacube non-orthogonally. At ECMWF, we are extending our meteorological object store – the fields database (FDB) – to make it able to understand polytope-style requests for data extraction.

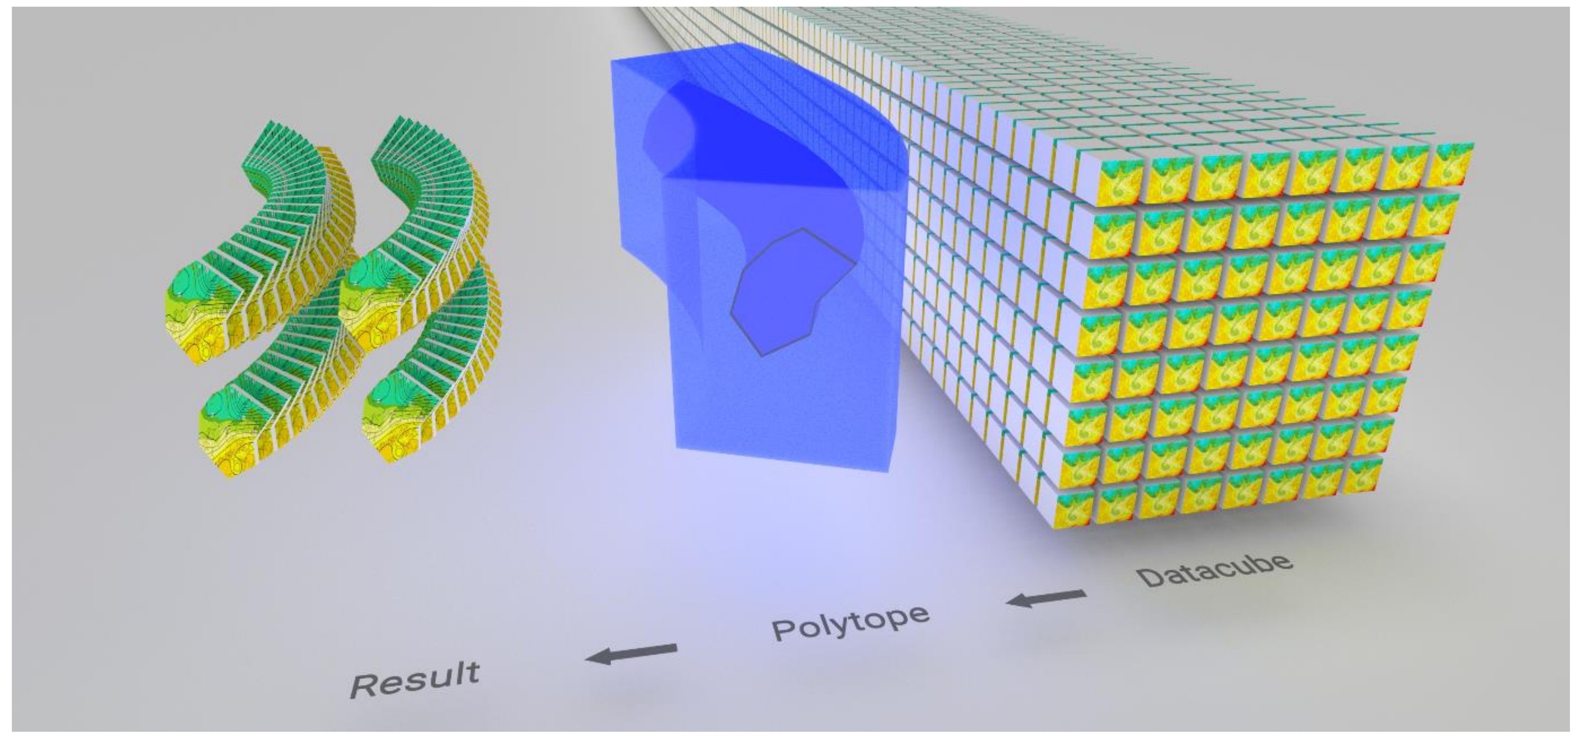

In parallel to these improvements, we are also designing a high-level interface, called 'Polytope', which allows users to easily create polytope-style requests. Polytope aims to return the data in an accessible format, allowing for quick and efficient integration with cloud-based workflows. It will offer helpful interfaces for common types of data extraction, such as trajectories, corridors, vertical profiles, areas and boundary frames; as well as providing a means to create your own polytopes.

Design

Polytope is a distributed RESTful service providing a simple API for data retrieval from a variety of data sources. When connected to an appropriate datasource, such as ECMWF's FDB, it will expose a rich toolkit of polytope-style data extraction APIs. Polytope is composed of a number of micro-services, which can be deployed anywhere using Kubernetes, Docker-Swarm, or on an ad-hoc basis.

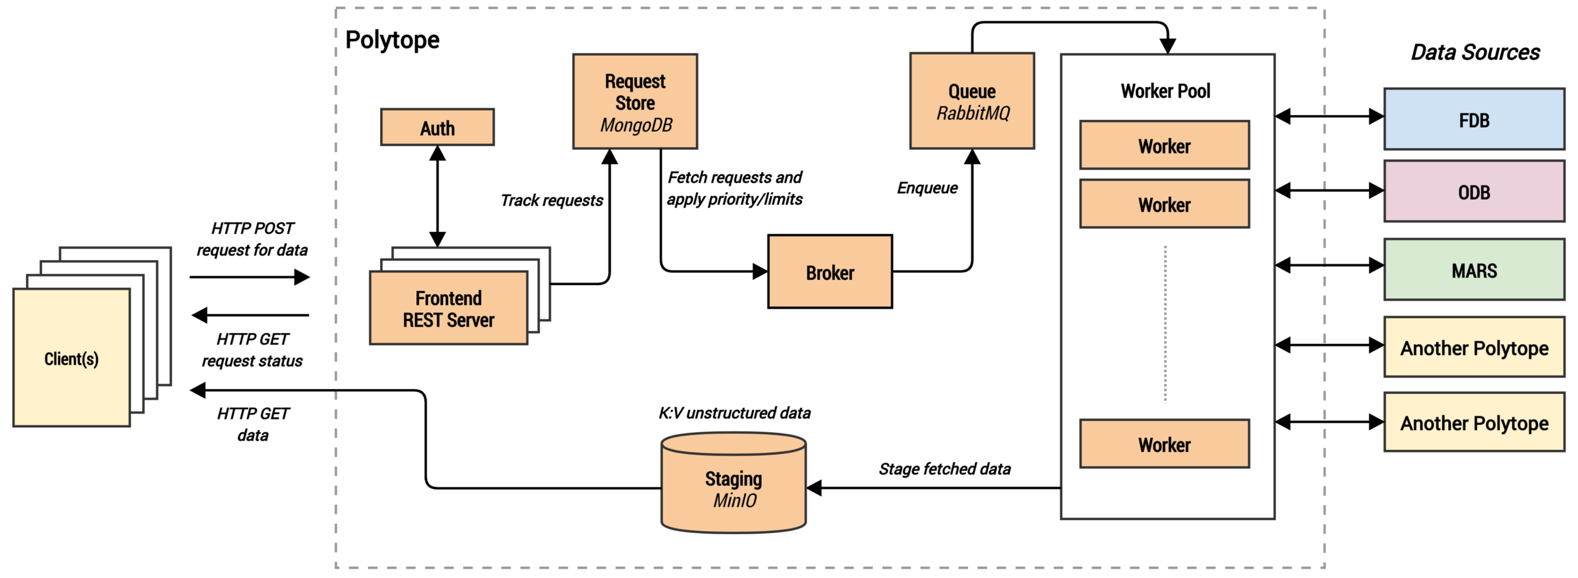

Polytope can be used client-less, simply using the HTTP REST interface, or more accessibly using the Polytope python client.

```
client = polytope.api.Client(user_email='abc', user_key='xyz')

request = {
    'stream' : 'oper',
    'levtype' : 'sfc',
    'param' : '165.128/166.128/167.128',
    'step' : '0',
    'date' : '2015-01-18/to/2015-01-19',
    'type' : 'fc',
    'class' : 'ei'
}

files = client.retrieve('ecmwf-mars', request)
arr   = xarray.open_dataset(files[0], engine='cfgrib')
```

Polytope is used in the HiDALGO project to serve data from ECMWF's HPC to multiple HPC and HPDA sites around Europe, including HLRS (Germany) and PSNC (Poland). HiDALGO provides modelling-as-a-service for crucial decision-making workflows, tackling human migration modelling, social network modelling and urban air pollution modelling – as well as pandemic modelling for the COVID-19 pandemic. Polytope is used in HiDALGO to provide a dynamic, on-demand data source for meteorological data.

It is also possible to connect multiple Polytopes together, creating a seamless federated data service between multiple HPC and Cloud centres. In this arrangement, data extraction can always be done in-situ with the corresponding data, ensuring efficient data retrieval and exchange between multiple sites. In the LEXIS project, Polytope is deployed at ECMWF (UK), LRZ (Germany) and IT4I (Czechia), providing a data exchange layer for complex weather and climate workflows spanning multiple computing centres (3 HPCs and 2 Cloud platforms).