


Highly Scalable Services for High-Throughput Data and Notifications: Polytope and Aviso

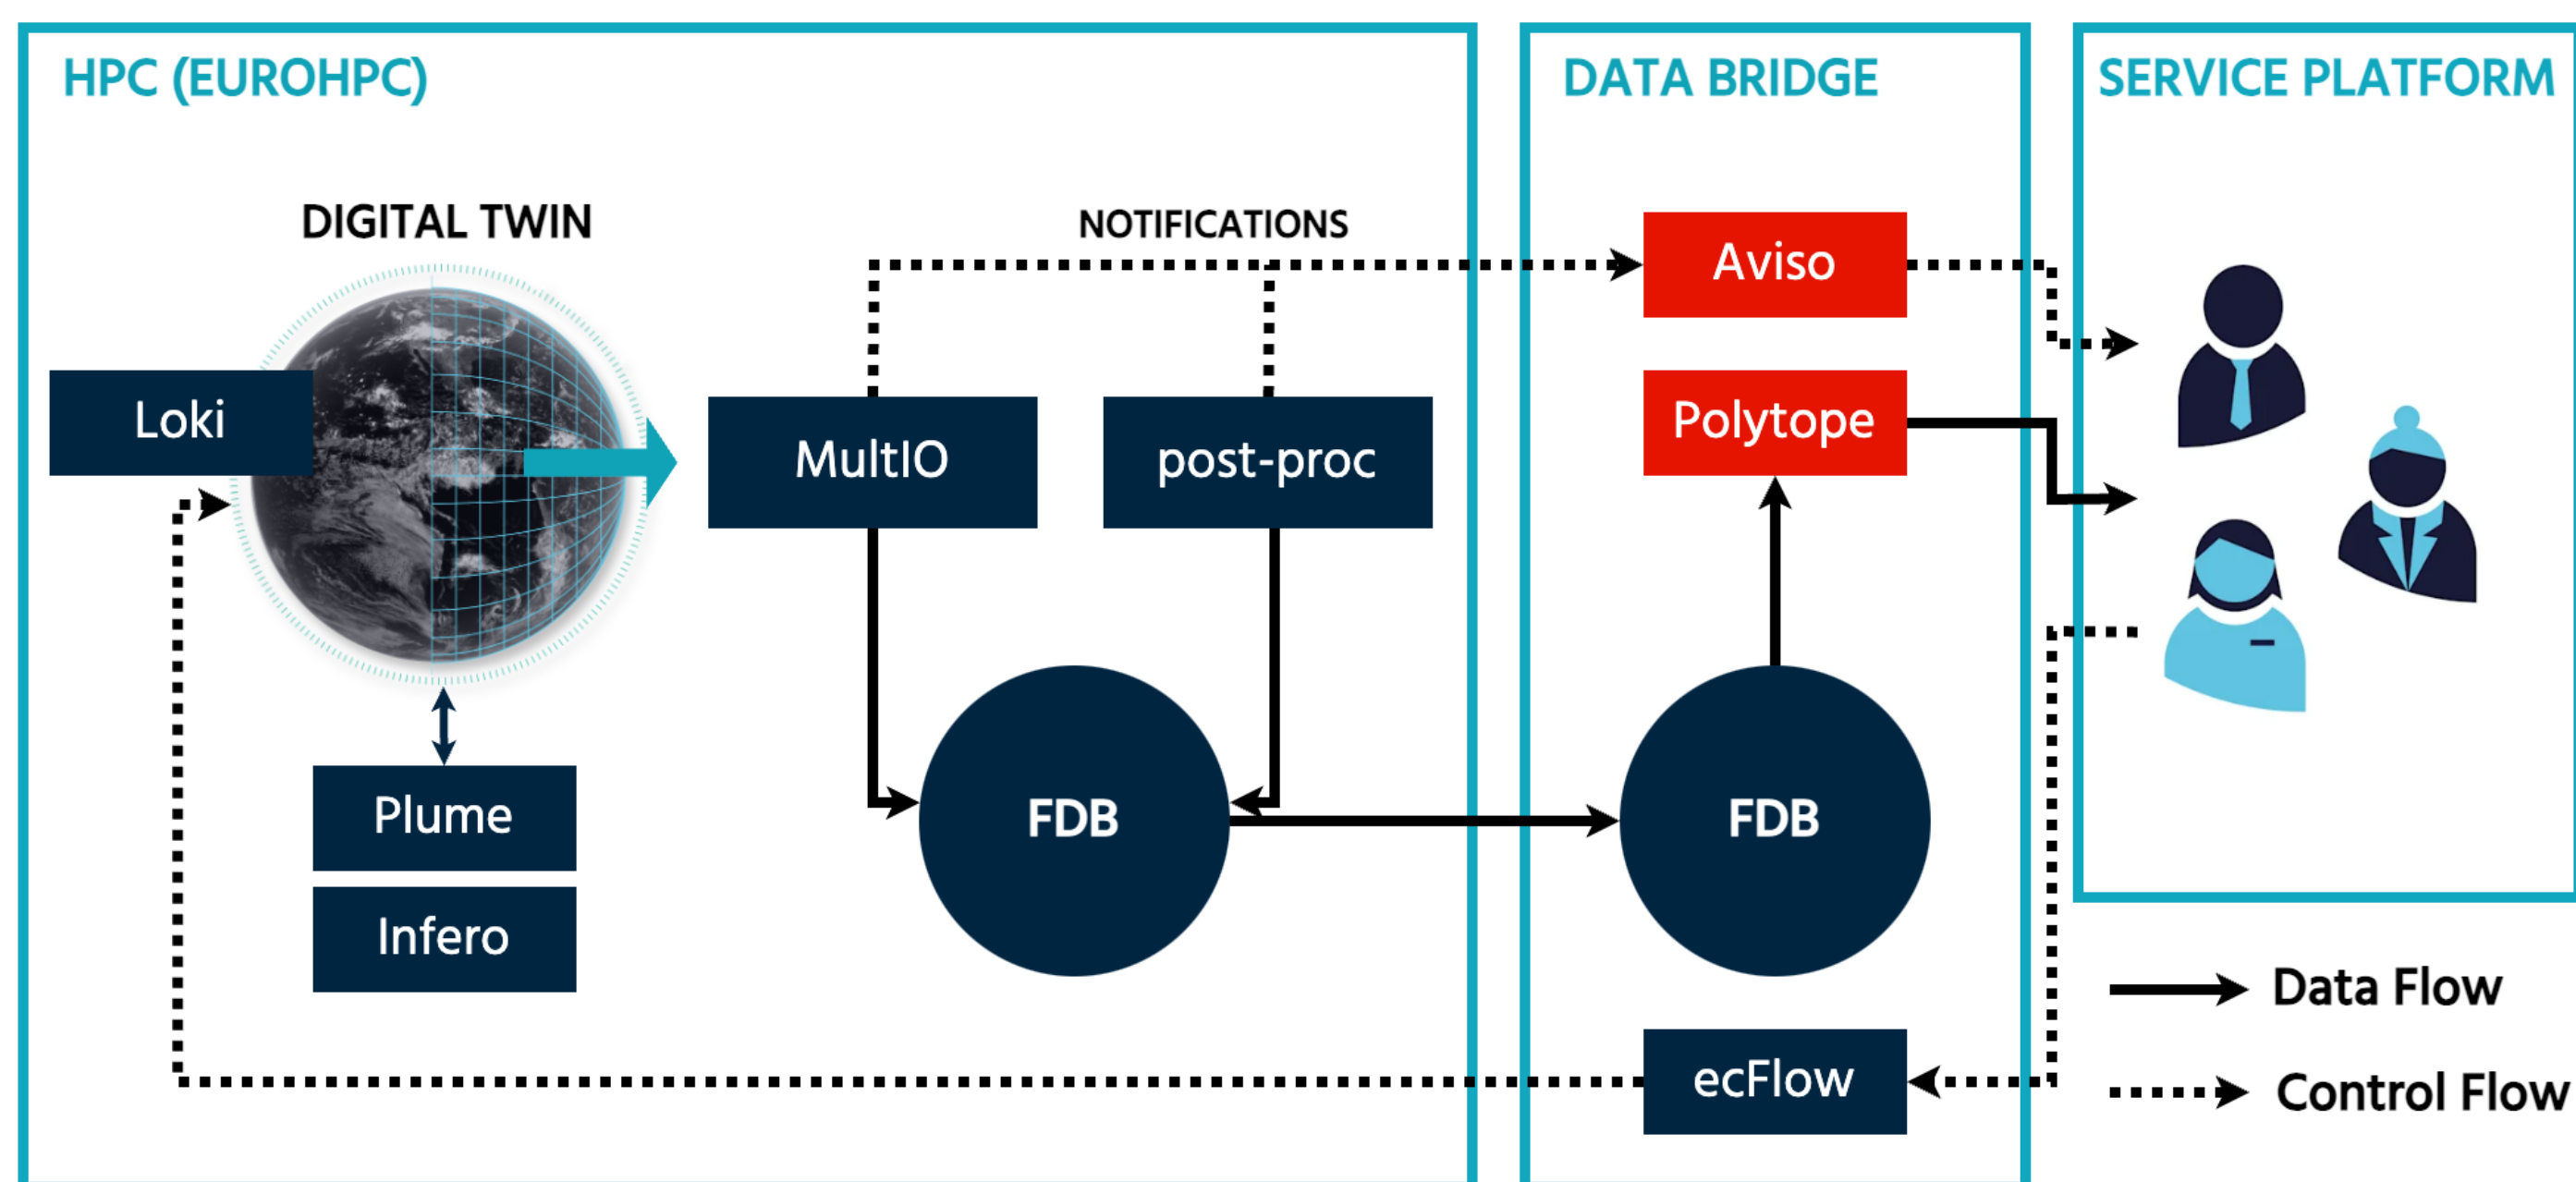


Samet Demir^{1*}, Peter Tsrunchev¹, James Hawkes¹, Tiago Quintino¹

(1) ECMWF; (*) samet.demir@ecmwf.int

1. Digital Twin Engine

At the heart of Destination Earth (DestinE), the **Digital Twin Engine (DTE)** plays a pivotal role in integrating diverse resources and components to create a seamless digital replica of our planet. The DTE ensures that all system components, from scientific data models to advanced digital technologies, operate smoothly and cohesively. By deploying sophisticated software and services across different layers of the system, the DTE leverages cutting-edge advances in supercomputing, machine learning, and European research to support DestinE's mission of creating a sustainable future for Europe.



High level visualization of the Digital Twin Engine

A high-level visualization of the **DTE** is presented in the image above, highlighting **Aviso** and **Polytope** in red, as they are the central focus of this poster.

2. Aviso

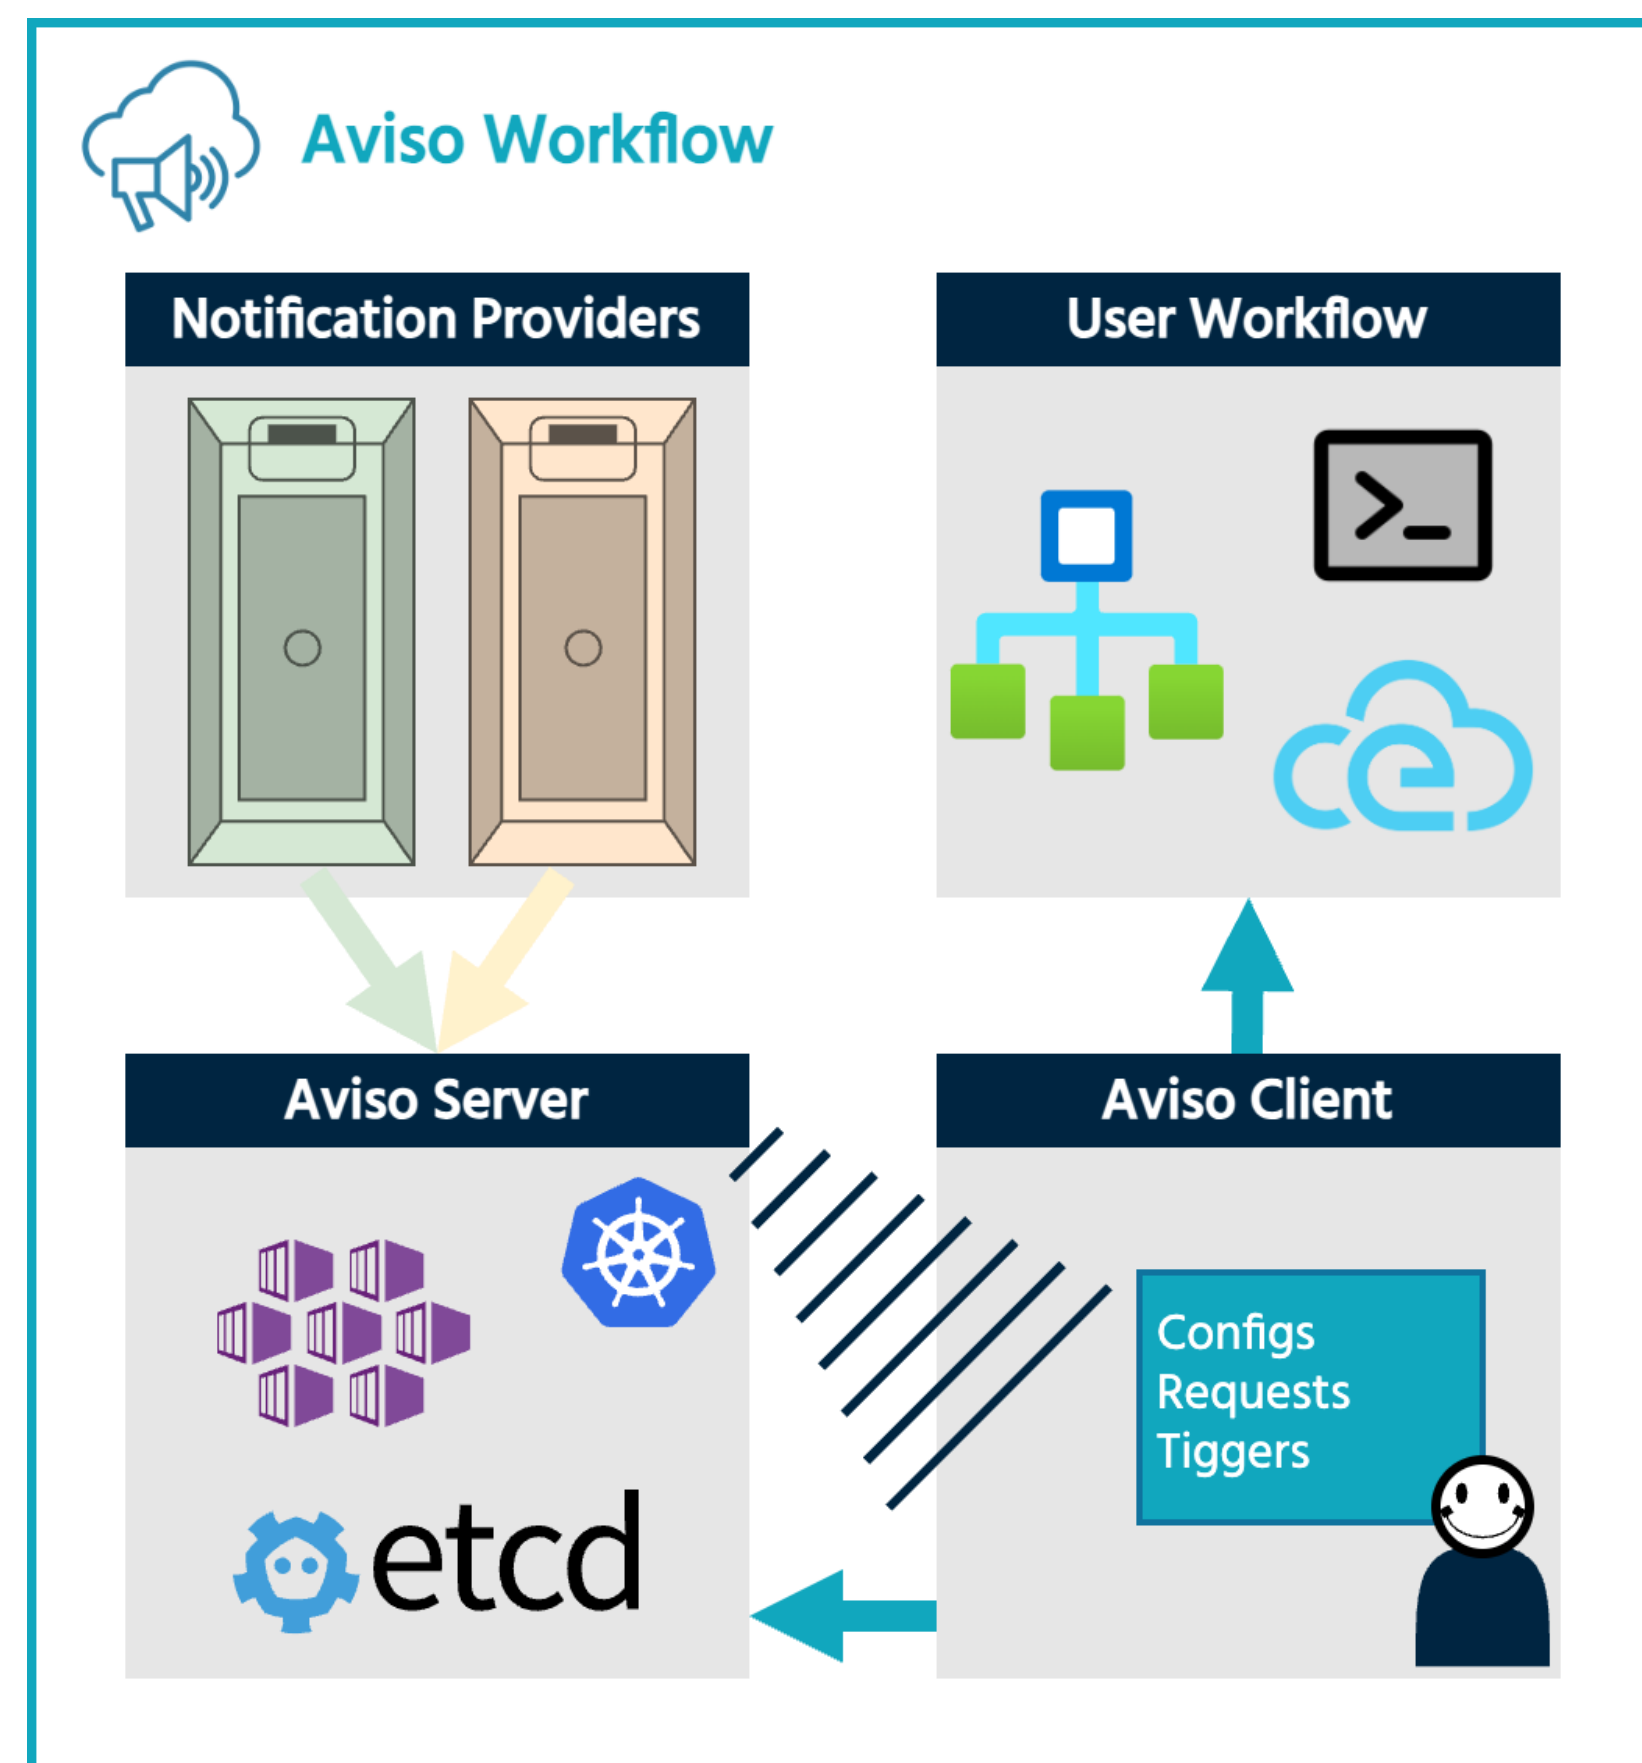


Diagram illustrating the Aviso workflow, showing the step-by-step process of how notifications are delivered to the user through each system component.

Aviso provides several key features:

- Event notifications to initiate automated workflows
- User-defined workflow triggers for seamless process automation
- Flexible "When <this> happens, Do <that>" logic
- Persistent event history with the ability to replay past notifications
- Operates independently of HPC or cloud environments
- Protocol agnostic
- Built for high reliability, ensuring performance in time-critical scenarios

While Aviso delivers lightweight messages, the payload is designed to inform users about specific events or data locations (such as a URL), rather than delivering large data files directly. The system's client-server architecture, based on a persistent key-value store, ensures scalability and robustness, enabling it to manage events across diverse HPC centres and cloud environments. The client side is a lightweight Python application that offers a Python API and a CLI interface for seamless integration into the users' workflows. Lastly, the notifications can be exchanged using CloudEvent messages, which enables the creation of workflows that extend across multiple data centres and cloud-based infrastructures. The general workflow of Aviso can be summarized as below:

1. Aviso client allows an End-User to subscribe to an event and to program a trigger
2. Aviso client polls Aviso server for changes to the defined event
3. A notification provider submits a notification to Aviso server
4. The subscriber is notified with a new event
5. The event triggers the user's workflow

3. Polytope

Polytope is an innovative data extraction service designed to improve access to vast, petabyte-scale data sets. As a crucial component of DTE, Polytope allows users to request precisely the data they need from complex, multidimensional datasets, using an n-dimensional polygon (polytope) to define their area of interest. This approach avoids unnecessary data extraction, improving both efficiency and accessibility.

As data from high-resolution simulations increases, Polytope addresses the challenges of efficiently accessing petabyte-scale datasets. It eliminates the need for users to manually sift through large, rectangular data blocks, enhancing the speed, accuracy, and agility of data retrieval. This makes it ideal for applications like real-time environmental monitoring, where rapid access to specific data is crucial.

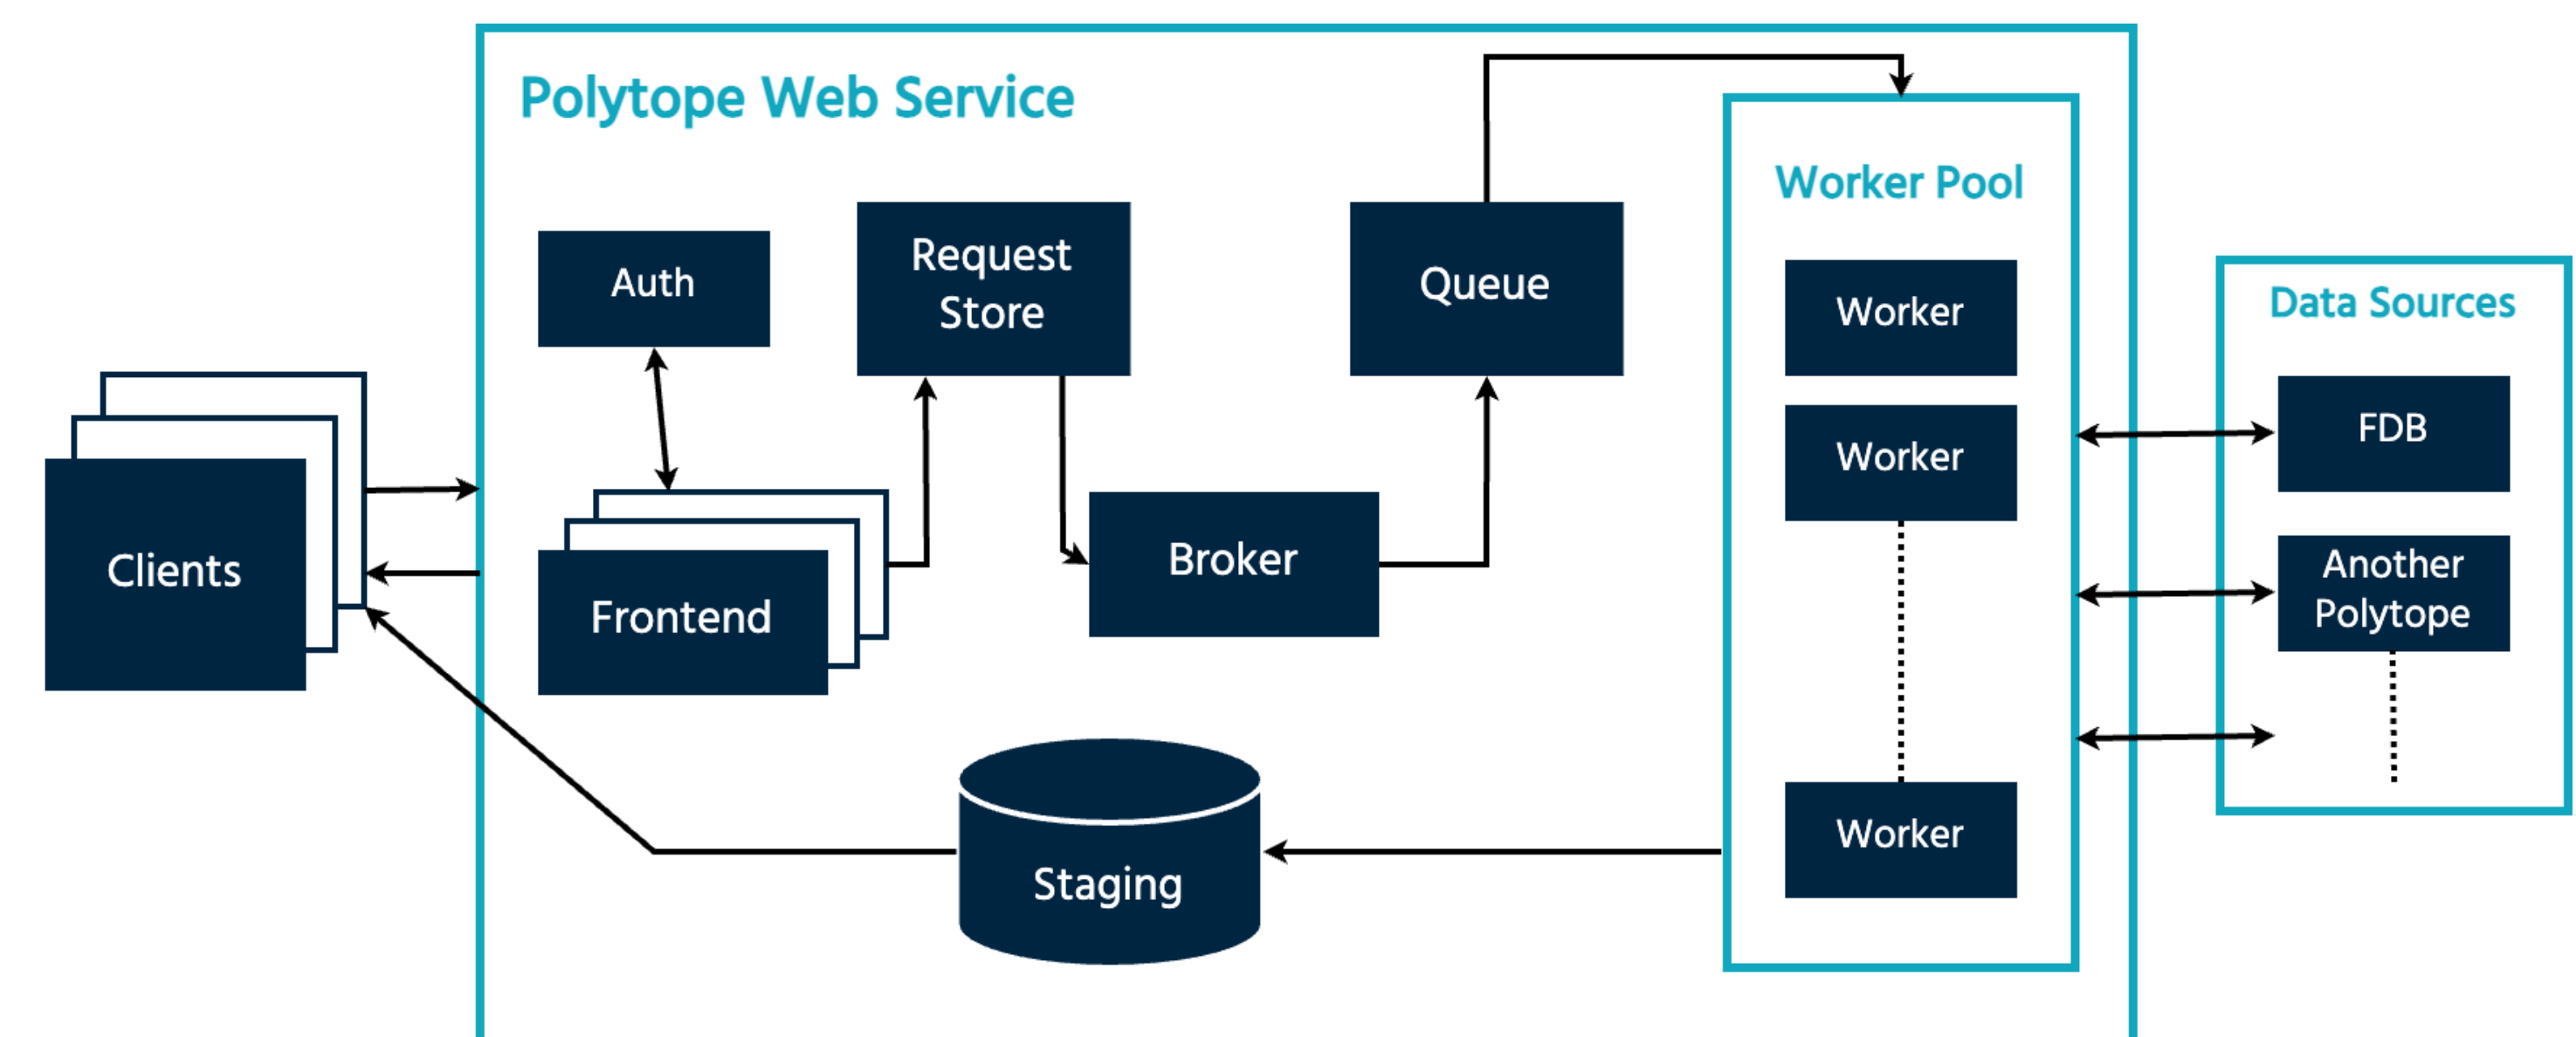


Diagram demonstrating the procedure of data delivery to the user via each system component.

Polytope integrates seamlessly with the Fields DataBase (**FDB**) and other DTE systems, providing a simple, RESTful API for data retrieval. Its distributed architecture, built on microservices, allows easy deployment across multiple environments, ensuring smooth data access wherever it is needed. By improving the efficiency of data extraction, Polytope reduces computational overhead, saving both time and energy.

4. Seamless Data Interaction

The combination of Aviso and Polytope within the Digital Twin Engine opens up interactivity for users across various industries. Together, they create an integrated system that allows users to extract precisely the data they need, in real time, as events unfold.

This synergy enables a level of responsiveness that was previously impossible. For example, in environmental monitoring, sectors like agriculture, energy, and disaster management can use Aviso to receive alerts about incoming storms or shifting weather patterns. At the same time, Polytope allows them to extract region-specific data, such as wind speeds or precipitation levels, precisely where and when it matters most. This tailored access to data ensures that decisions are not only timely but also informed by the most relevant and up-to-date information.

This integrated system not only enhances decision-making but also promotes sustainability across various sectors. For instance, offshore wind farm operators can receive real-time notifications through Aviso about approaching wind storms, enabling them to adjust turbine operations to safeguard equipment and optimize performance. Simultaneously, Polytope allows them to extract precise data on wind speeds and directions in their specific area, providing the actionable insights needed to make informed adjustments quickly.

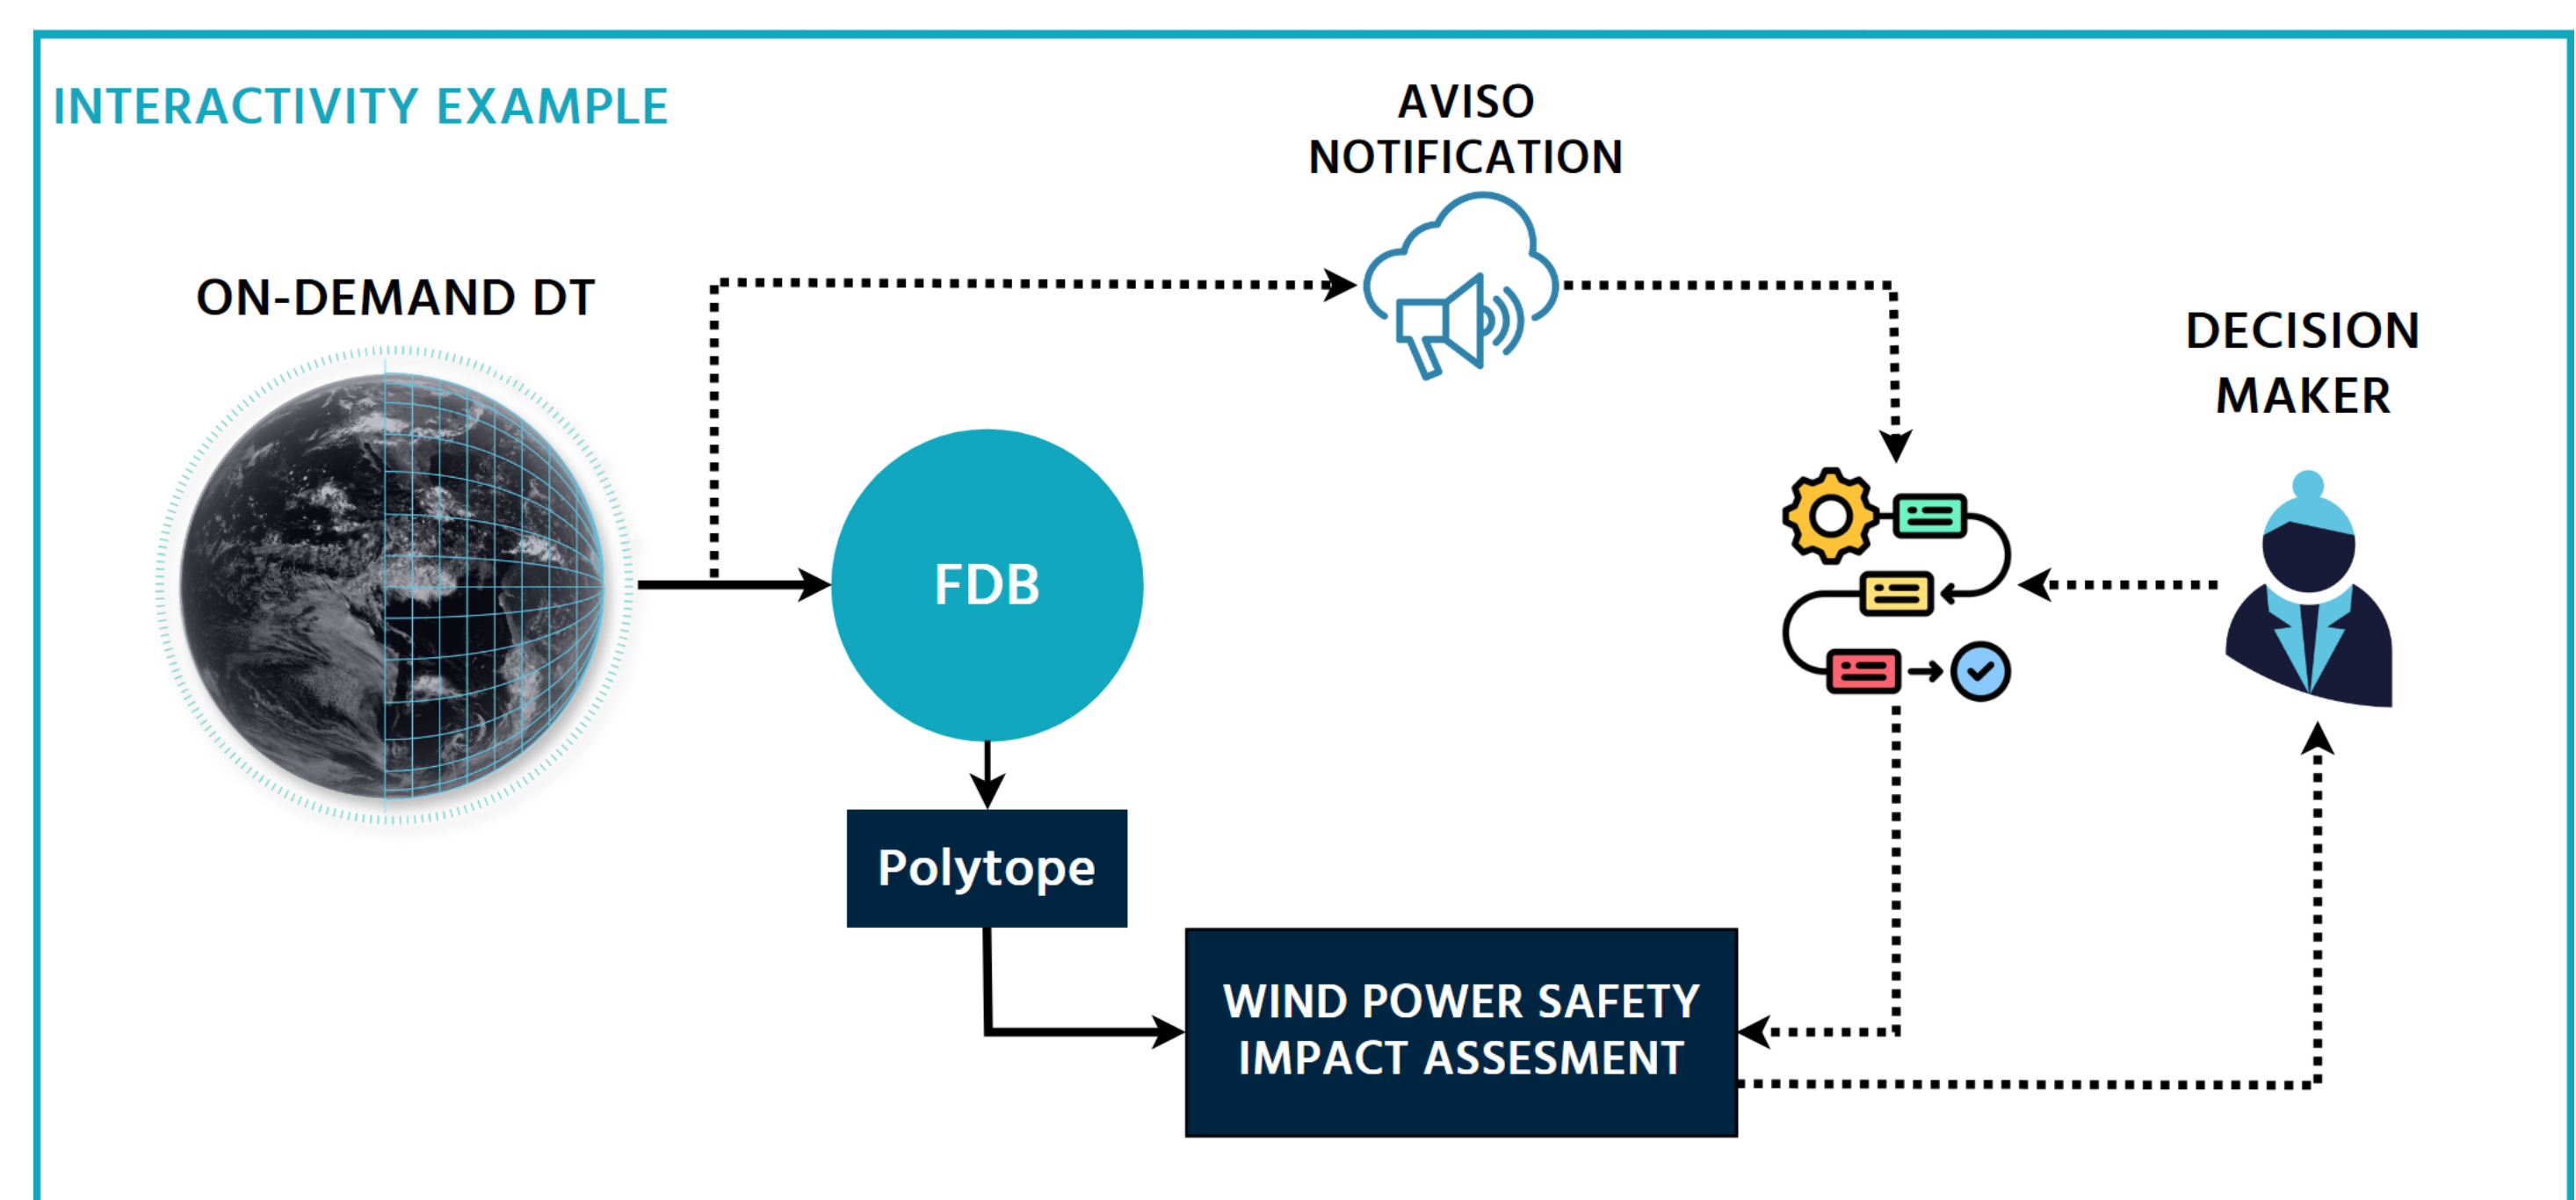


Diagram illustrating an example workflow that utilizes both Aviso and Polytope for wind power safety.



Funded by the European Union

Destination Earth

implemented by

