



Optimizing environmental data services on federated Cloud and e-Infrastructures

Zhiming Zhao, Paul Martin, and Spiros Koulouzis

University of Amsterdam, Amsterdam, Netherlands (z.zhao@uva.nl)

The increasing volumes of data being produced, curated and made available by research infrastructures in the environmental science domain require services able to optimise the delivery and staging of data on behalf of researchers and other users of scientific data.

Specialised data services for managing the data lifecycle, for creating and delivering data products, and for customised data processing and analysis, all play a crucial role in how these research infrastructures serve their communities, and many of these activities are time-critical—needing to be carried out frequently within specific time windows.

We describe our experiences identifying the time-critical requirements of environmental scientists making use of computational research support environments. We present a microservice-based infrastructure optimisation suite, the Dynamic Real-time Infrastructure Planner, used for constructing virtual infrastructures for research applications on demand.

We demonstrate how our suite is used to optimise runtime service quality for use cases provided by the ENVRI research infrastructure communities.