Earth System Grid Federation Future Architecture, Copernicus, Cloud and ESA

ClimateData.ca Meeting, 23 November 2021

Philip Kershaw, Technical Manager
Centre for Environmental Data Analysis

Science and Technology Facilities Council
Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Earth System Grid Federation: a globally distributed data archive for climate data

31	170	24	12,870,551	722,400,940	25.24
Projects	Countries	Data nodes	published datasets	files downloaded	PB downloaded

ESGF Dashboard: http://esgf-ui.cmcc.it

ESGF Federation

Data usage

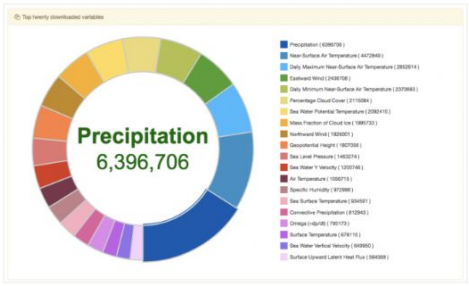

Data publication

UKRI Science and Technology Facilities Council / Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

ESGF – Application and Evolution

ESA Climate Change Initiative Open Data Portal

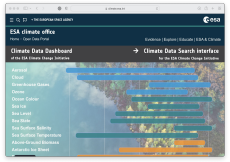

Public Cloud Public Dataset Programme(s)

ESGF Frontend

Data lakes – Government-sponsored

CMIP5 >> Earth System Grid Federation >> CMIP6

ESGF Future Architecture >> ESGF 2.0

Institutional-based hosting

CMIP/CORDEX for Copernicus Climate Data Store

UKRI Science and Technology Facilities Council / Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

ESA Climate Change Initiative Open Data Portal

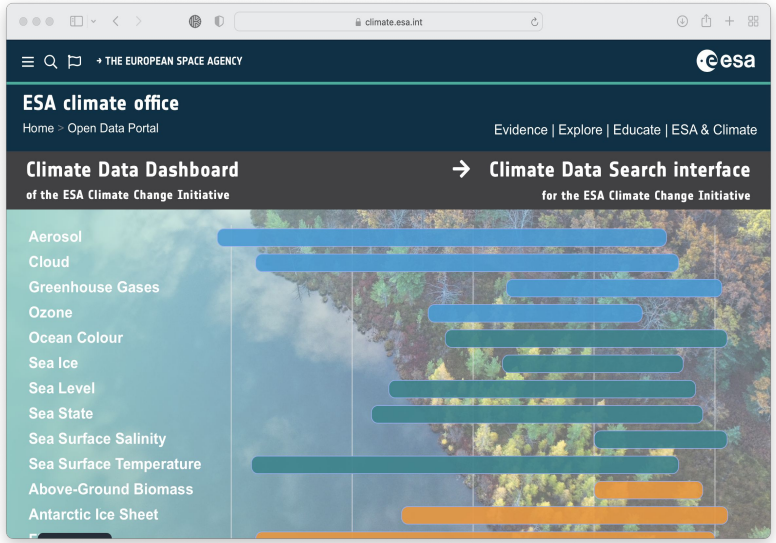

2 Phases:

1) Leveraged ESGF
    1) Quick win with search and download
    2) Bespoke search API incompatible with other community standards - OGC CSW
    3) THREDDS Data Server couldn't scale to our needs

2) Redesigned to address issues
    1) OpenSearch API replaced ESG Search
    2) Scalable data service with Kubernetes
    3) Zarr format cache of netCDF data on object store for performance

C3S 34[a-f] Projects for the CDS

- Architected a system for delivering resilient CMIP5 and CORDEX data access for the CDS by creative application of federated architecture for Earth System Grid Federation

ESGF: an international federation of nodes providing a network of access points to model data

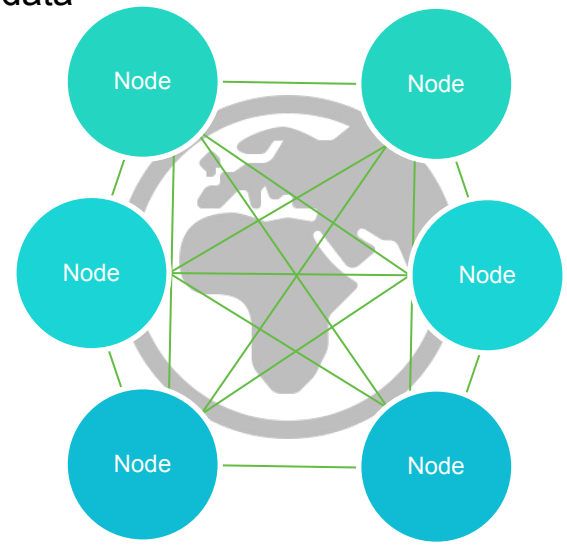

C3S 34a/b system: a single resilient point of access to data delivered through replication and redundancy

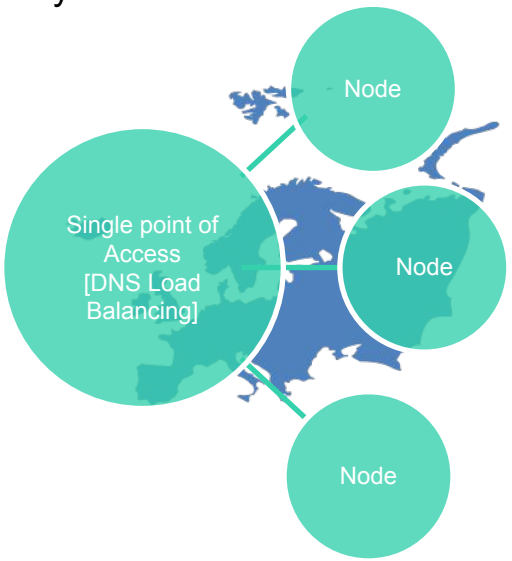

UKRI Science and Technology Facilities Council | Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

C3S Resilient CMIP and CORDEX Data Access

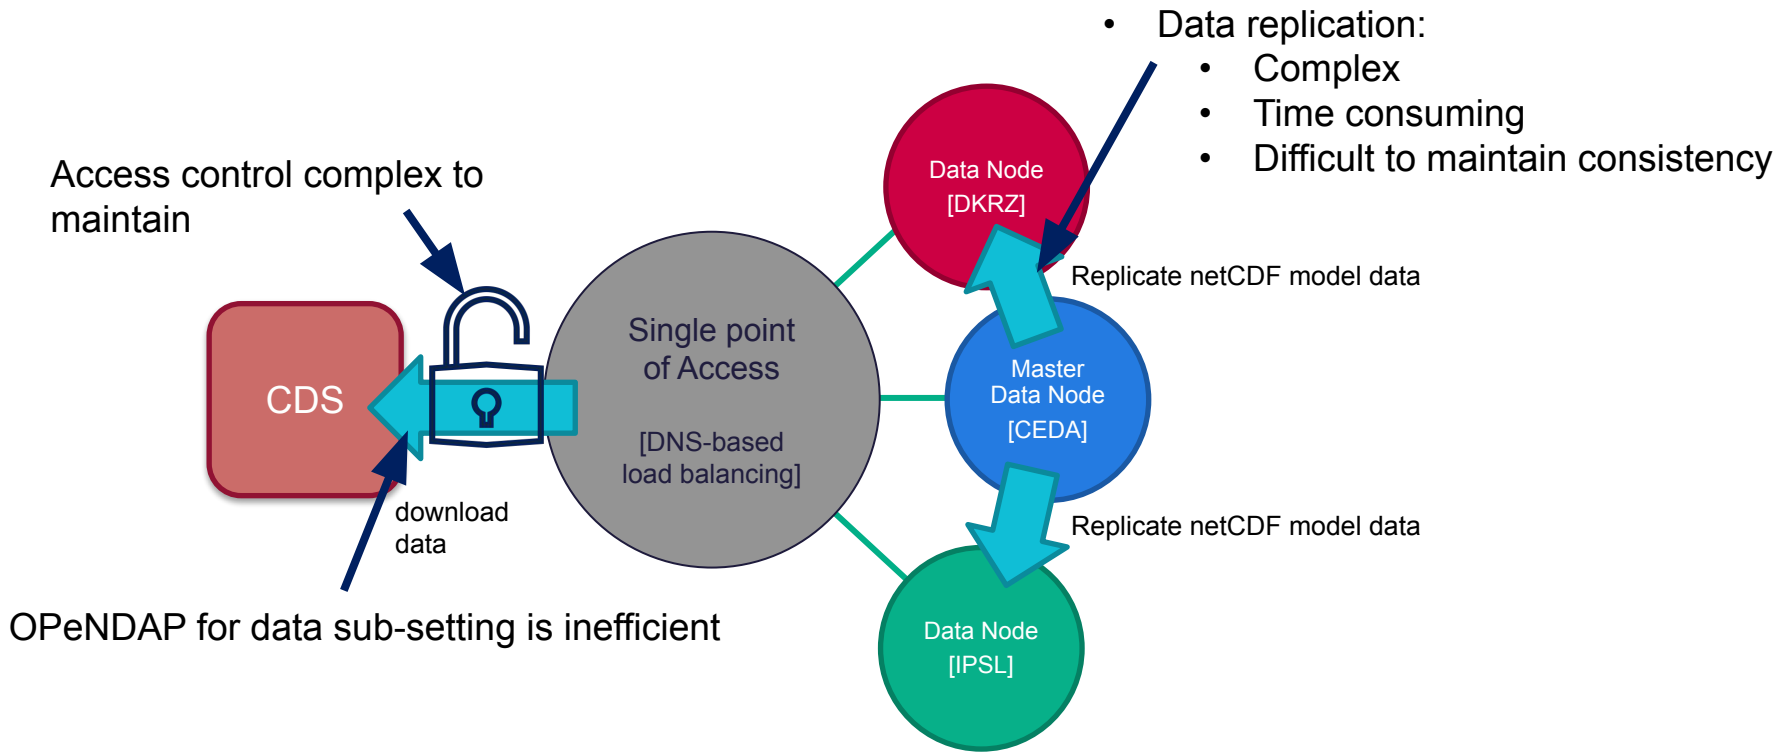

- Data replication:
  - Complex
  - Time consuming
  - Difficult to maintain consistency

Sub-setting Services for C3S 34e Project

Climate Change Service

climate.copernicus.eu

C3S 34e Project

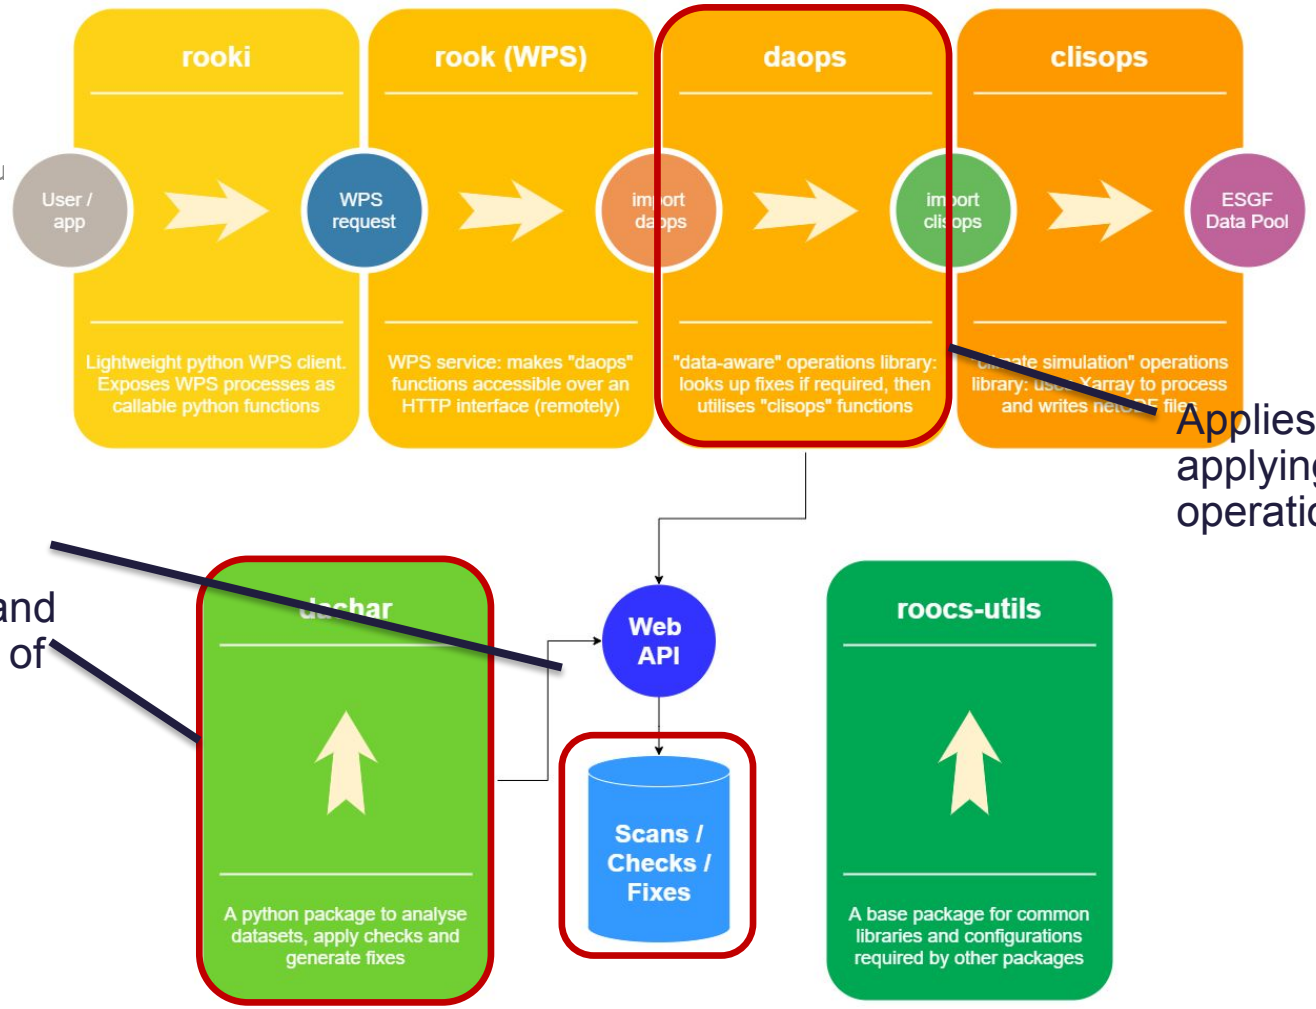

Applies fixes to datasets before applying subsetting/regridding operations

Analyse datasets and make an inventory of fixes

Credit: Ag Stephens, CEDA

UKRI Science and Technology Facilities Council / Natural Environment Research Council

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

https://nbviewer.org/github/roocs/rooki/blob/master/notebooks/demo/demo-rooki-subset-by-point.ipynb

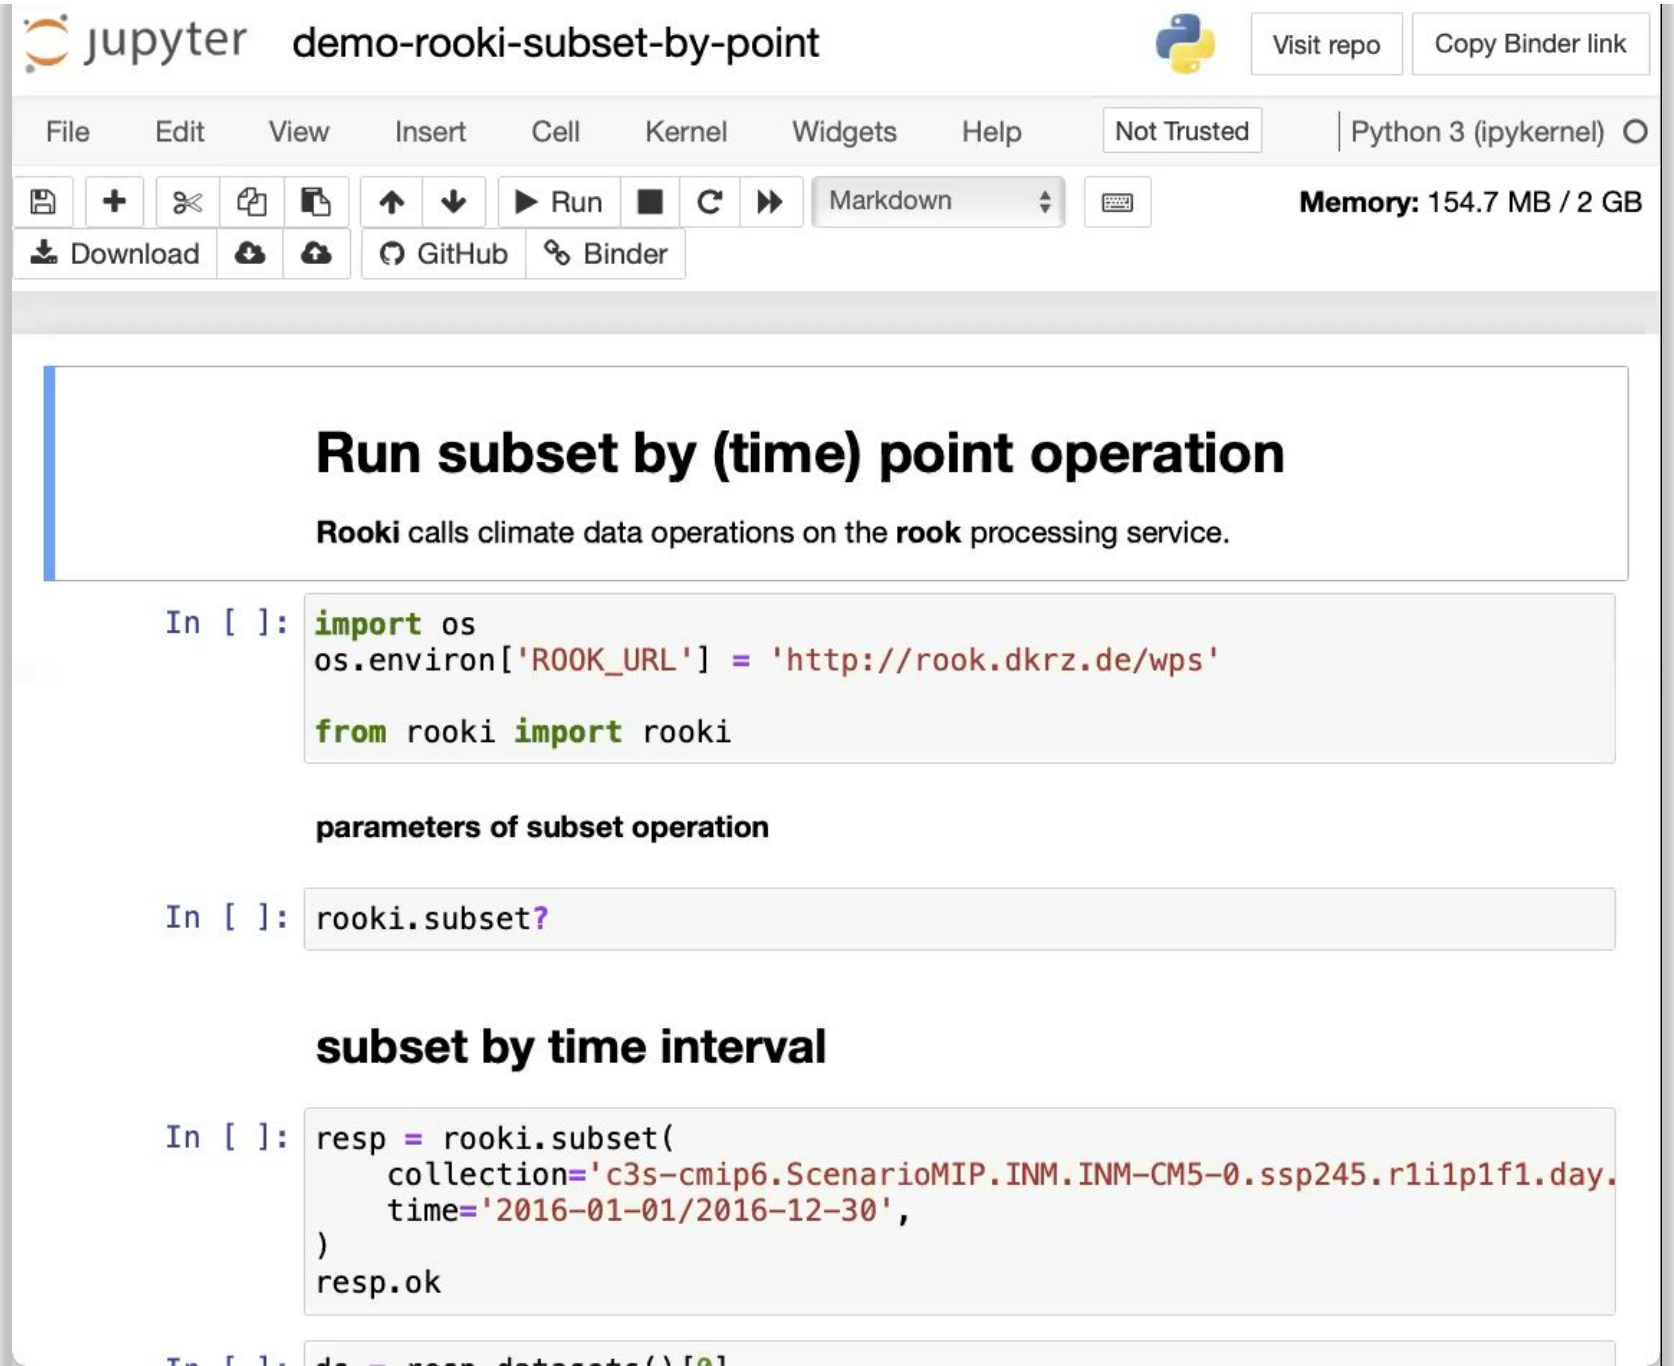

Science and Technology Facilities Council
Natural Environment Research Council

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

ESA Earth Observation Exploitation Platform Common Architecture (EOEPCA)

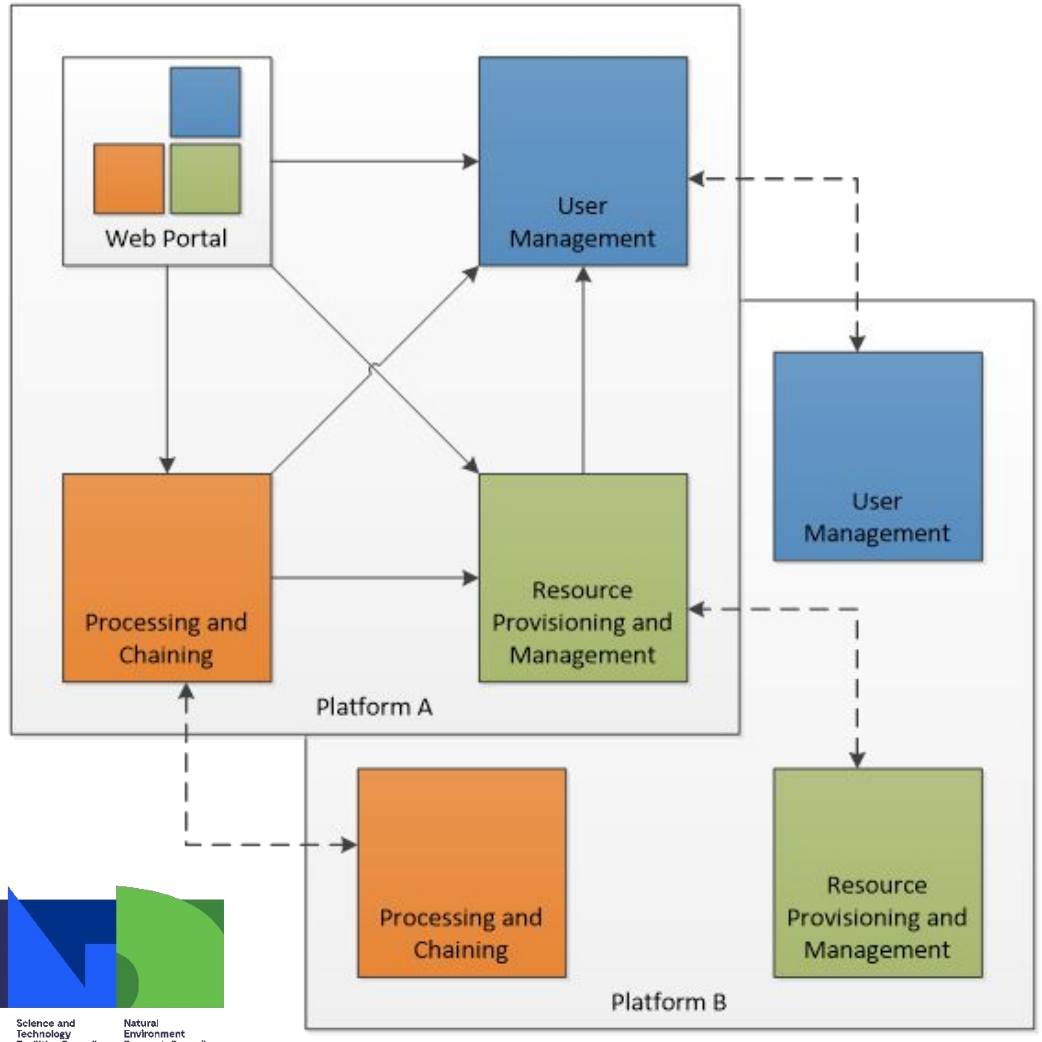

- Architectural blueprint for the federation of *platforms*
- Interlinked with the work on recent OGC testbeds
- CEDA involved with consultancy role
- Processing and chaining of particular interest
  - ADES and EMS
  - Ability to push customised shrink-wrapped processes to 3rd party WPS instances
- Innovations with ID management: UMA

ESGF – Application and Evolution

ESA Climate Change Initiative Open Data Portal

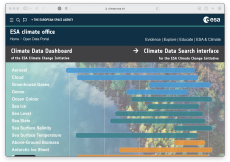

Public Cloud Public Dataset Programme(s)

ESGF Frontend

Data lakes – Government-sponsored

CMIP5 >> Earth System Grid Federation >> CMIP6

ESGF Future Architecture >> ESGF 2.0

Institutional-based hosting

CMIP/CORDEX for Copernicus Climate Data Store

UKRI Science and Technology Facilities Council / Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

ESGF Future Architecture

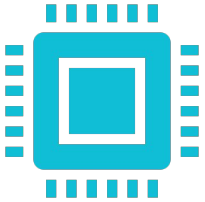

Platforms and systems administration

Modular, scalable architecture: Containers, Kubernetes

Embrace infrastructure-as-code approach

Search services

Modernise, centralise and simplify

Use community standards: STAC

ID Management and Access Entitlement

Modernise, centralise and simplify

Use industry standards: OpenID Connect / OAuth 2.0

Progress and Achievements

- Container and Container+Kubernetes installs available
- **Deployed on AWS (GFDL) and at CEDA**

- Major community engagement on use of STAC for ESM data
- Prototype developed by CEDA
- Integration tests
- CoG and MetaGrid futures??

- OpenID Connect / OAuth 2.0 done
- New Authorisation system with Open Policy Agent
- Authentication integrated with C4I in test

UKRI Science and Technology Facilities Council | Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

ESGF Future Architecture

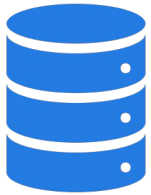

New modes for Data Access + Storage

Augment trad. file serving with object store

New models for aggregation and subsetting, retire OPeNDAP

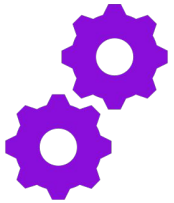

Compute Services

Important but no consensus for ESGF-wide standard offering yet

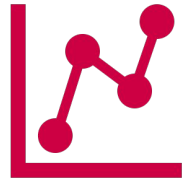

Metrics Collection

Leverage advances in industry with standard tooling to exploit - Prometheus and InfluxDb, Grafana

Progress and Achievements

- Factored out TDS
- Test CMIP6 data caches on object store at CEDA and DKRZ

- Ported C3S WPS Data Reduction Services for use in ENES CDI
- Used with Climate4Impact
- Reboot of Compute Working Team

- New Metrics system integrated with CMCC

UKRI Science and Technology Facilities Council / Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Future Architecture Node – Phase 1

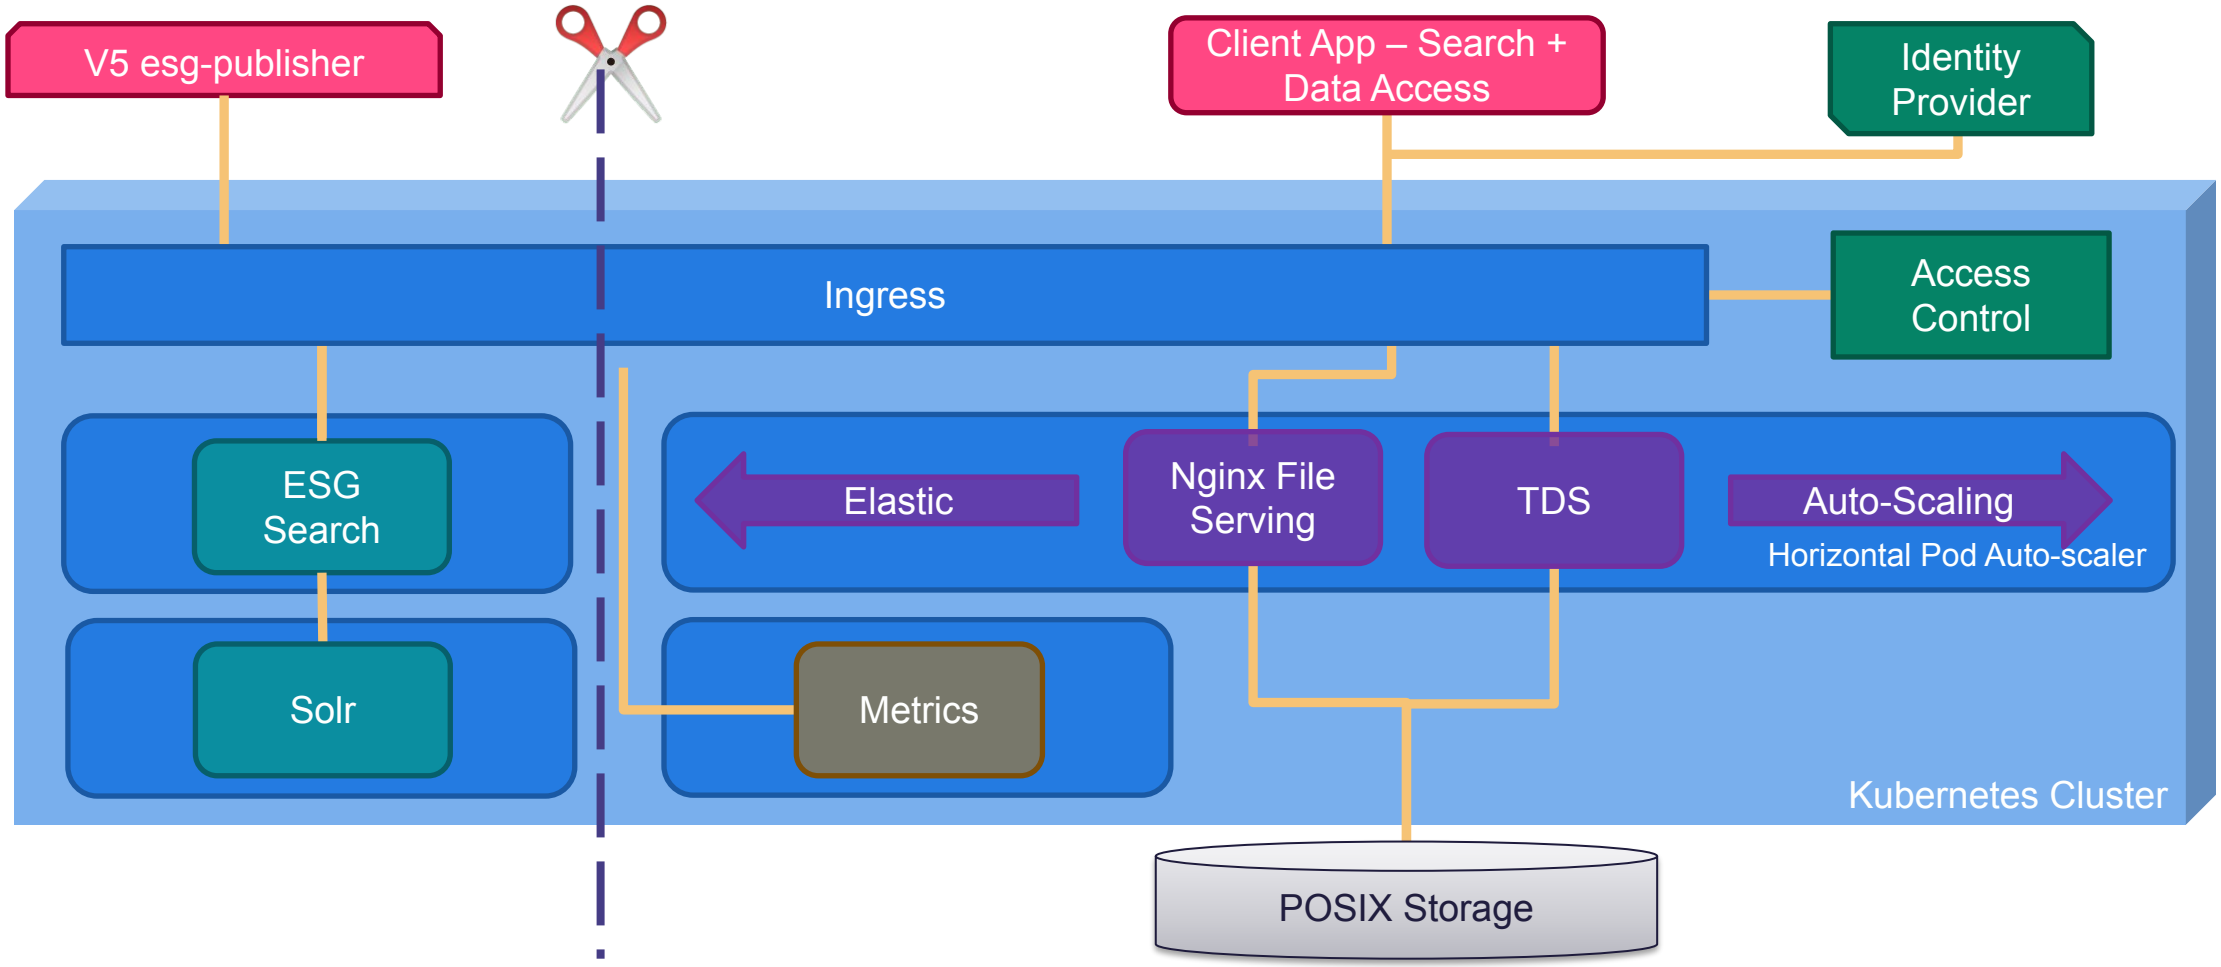

Future Architecture Node – Phase 2

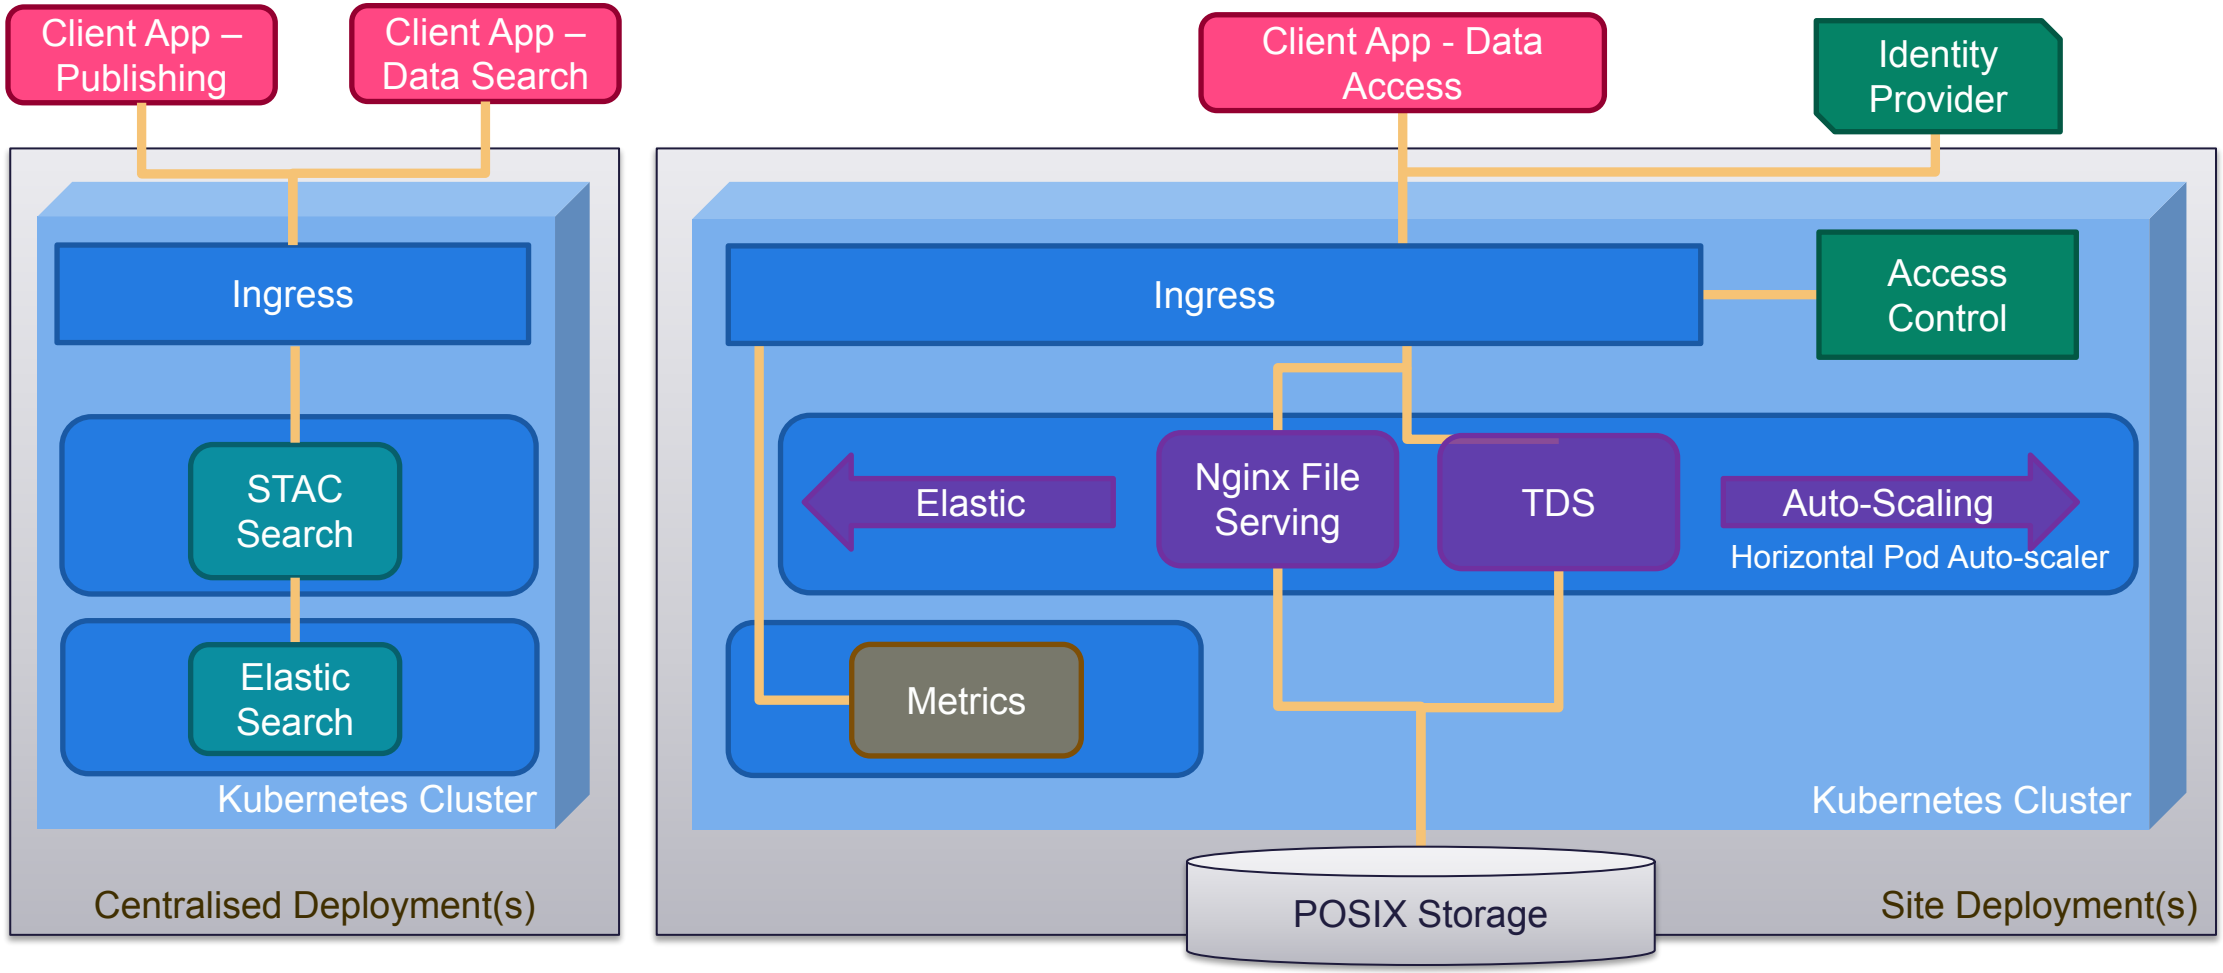

UKRI Science and Technology Facilities Council / Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

STAC API for ESGF

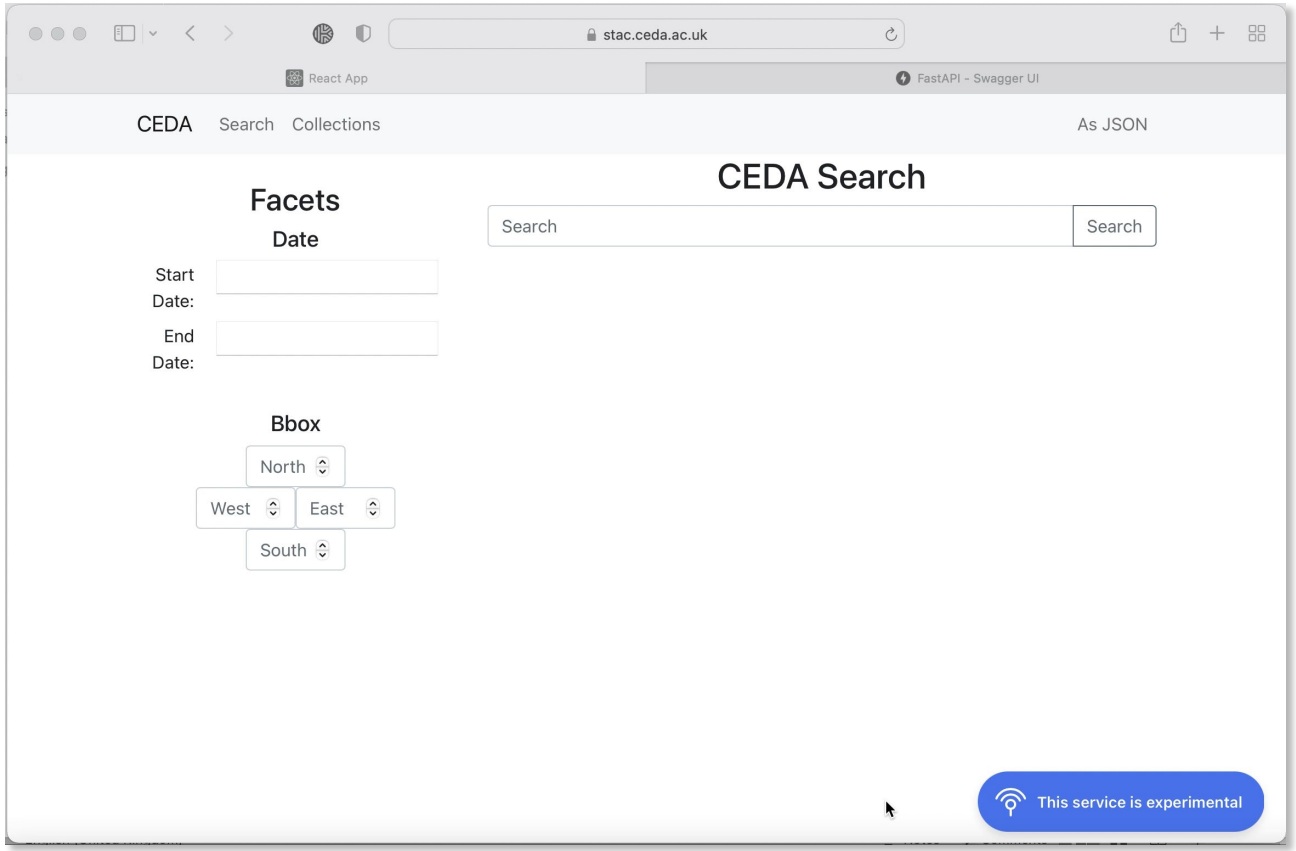

- Implementation of STAC API
- ElasticSearch backend
- Filter extensions to support faceted search
- Fully featured STAC equivalent API to ESG Search
- Simple frontend created to demonstrate its features

UKRI Science and Technology Facilities Council Natural Environment Research Council

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

IS-ENES3 - Data Analytics using Notebooks/icclim

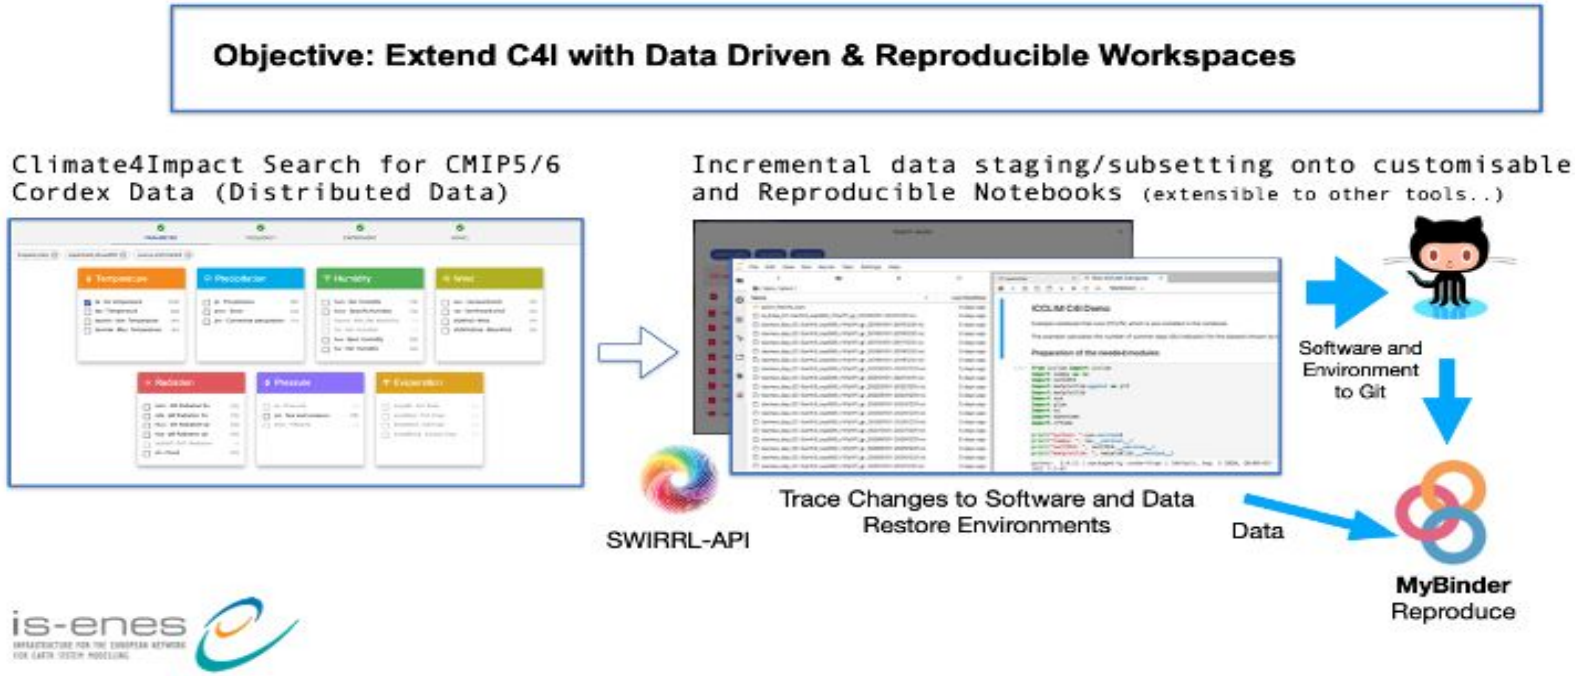

DestinE and Blueprint Architecture

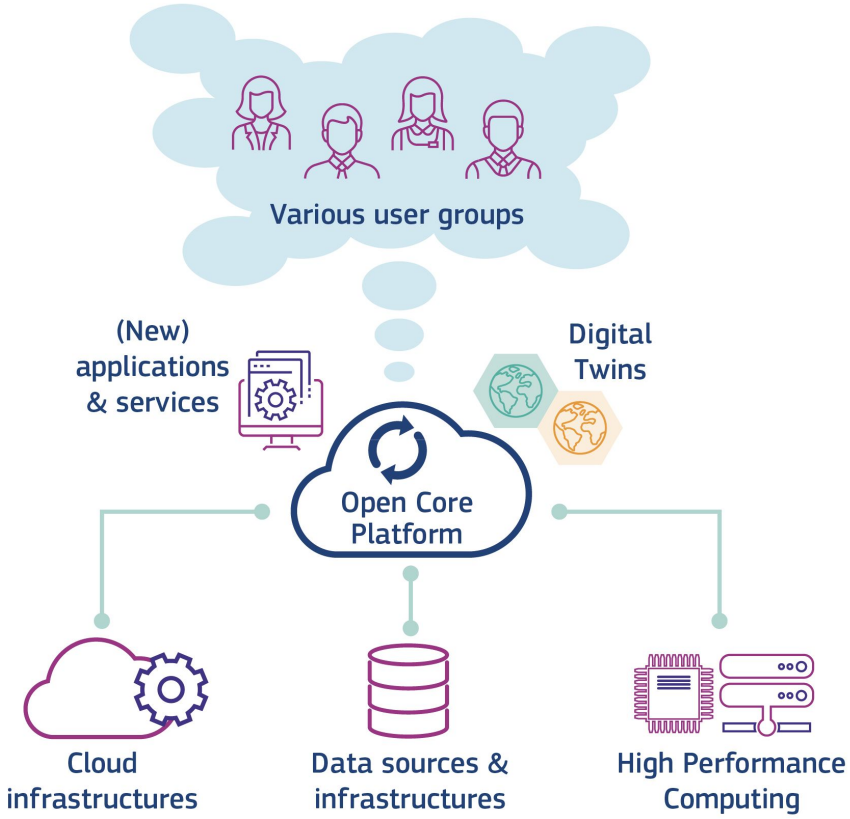

Destination Earth (DestinE) - major EU initiative:

- “to develop a very high precision digital model of the Earth (a ‘digital twin’) to monitor and predict environmental change and human impact to support sustainable development”

JASMIN

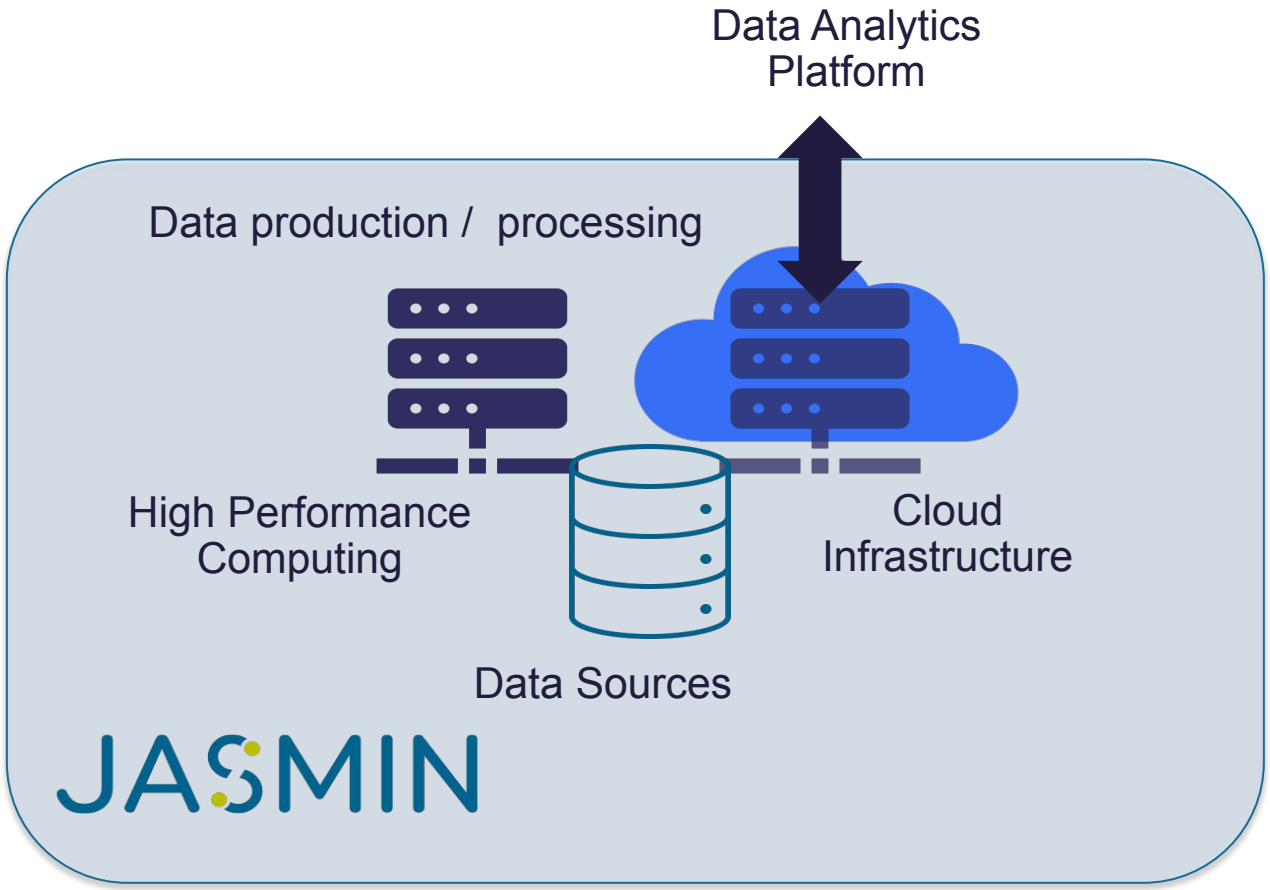

ESA Digital Twin Earth (DestinE) Precursor - land surface modelling and climate

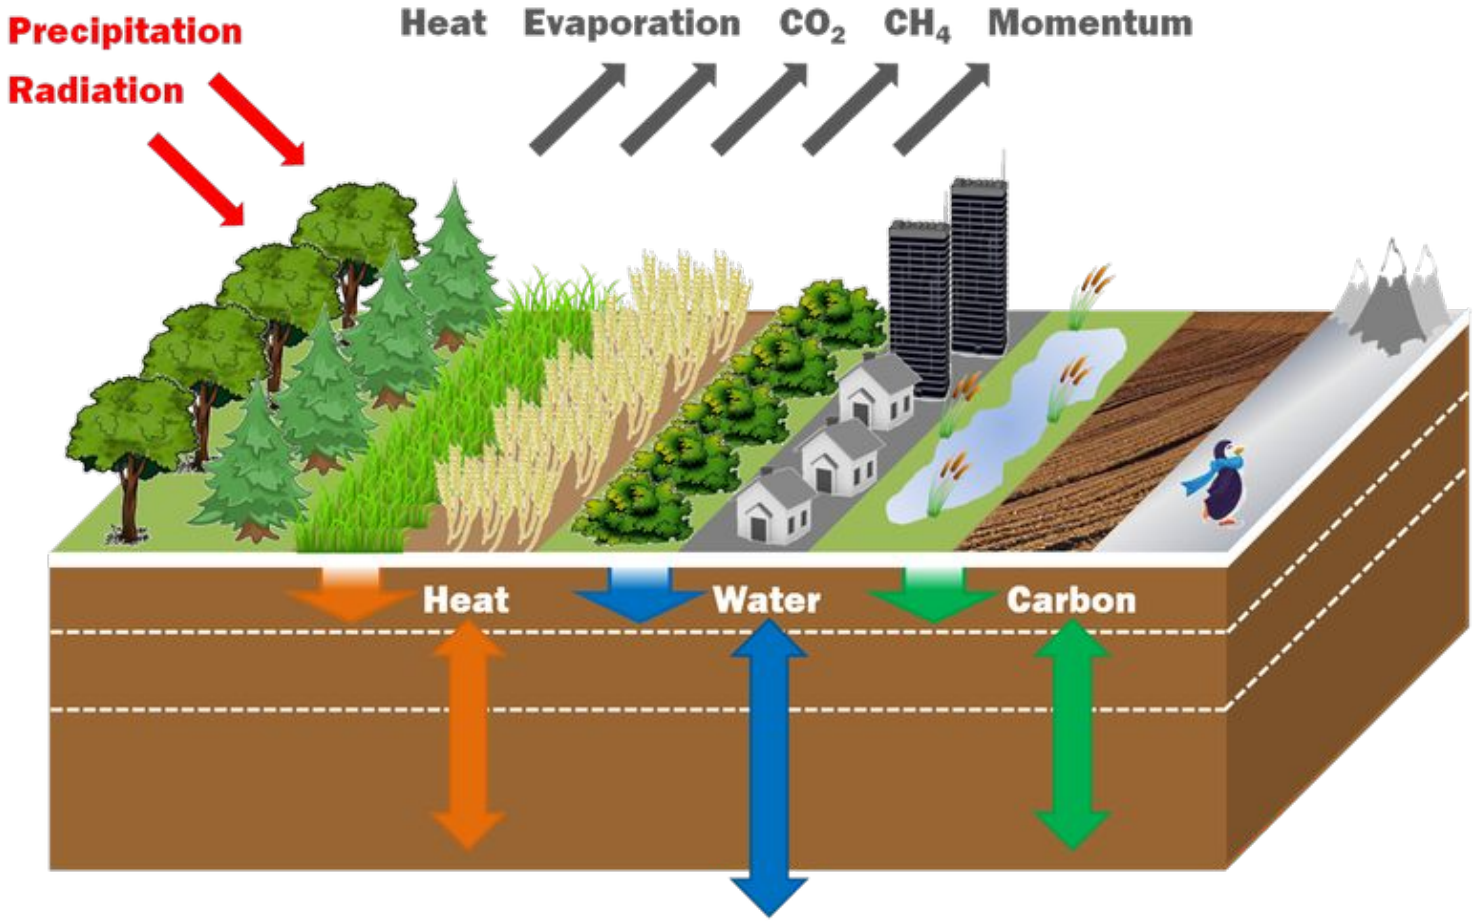

- Using JULES (Joint UK Land Environment Simulator)
  - the land surface component in the Met Office Unified Model
- Improvements with Data Assimilation
  - LaVEnDAR (The Land Variational Ensemble Data Assimilation fRamework)
  - Feed in satellite observations – SIF and SMAP data

What could be the future impact of climate change on the soil moisture?

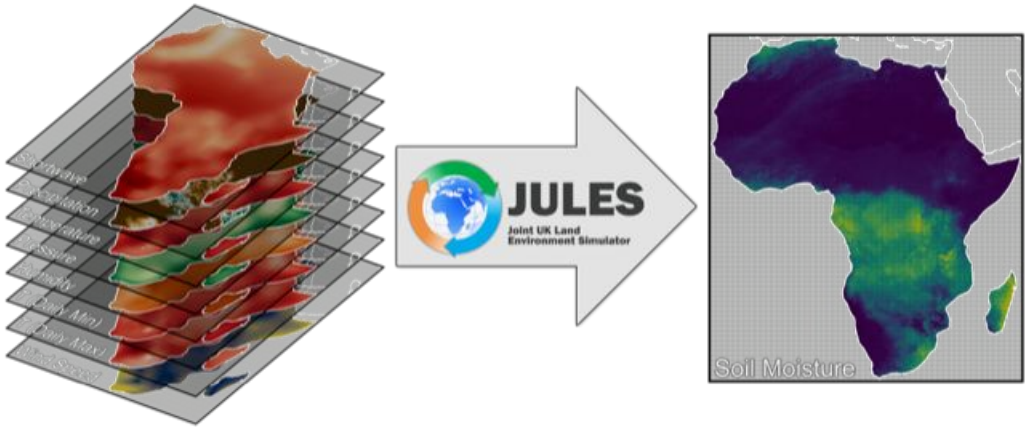

- JULES driven with climate projections from ISIMIP data (Inter-Sectoral Impact Model Intercomparison Project)

Make a surrogate AI model to JULES

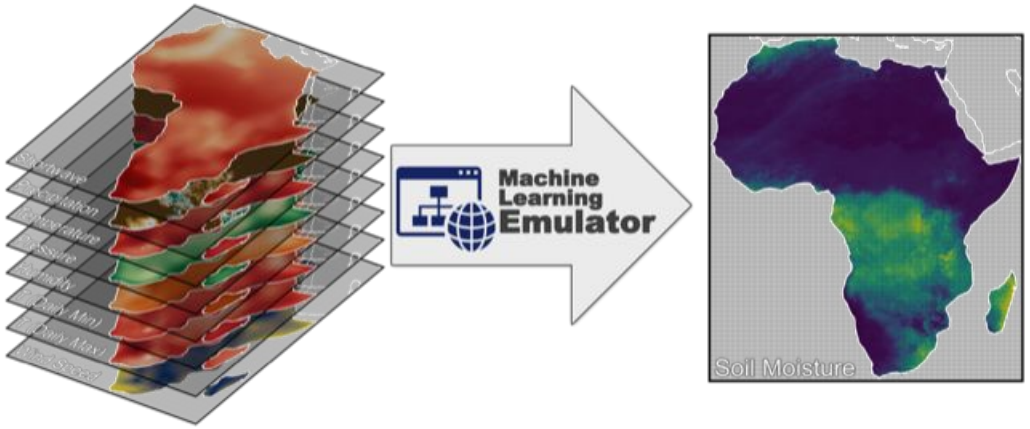

- Experimented with Machine Learning (ML) techniques
- Goal: a general-purpose algorithm -

  `time series of daily weather data  time series of soil moisture data`

- Successfully applied XGBoost (eXtreme Gradient Boosting) algorithm.
- trained on up to 1000 grid cells, representative of the various biomes in continental Africa
- Demonstrated to accurately emulate JULES output at other locations
- The credibility of the model is enhanced by its transparency and explainability

Digital Twin Precursor on JASMIN: HPC for data production, cloud for analysis

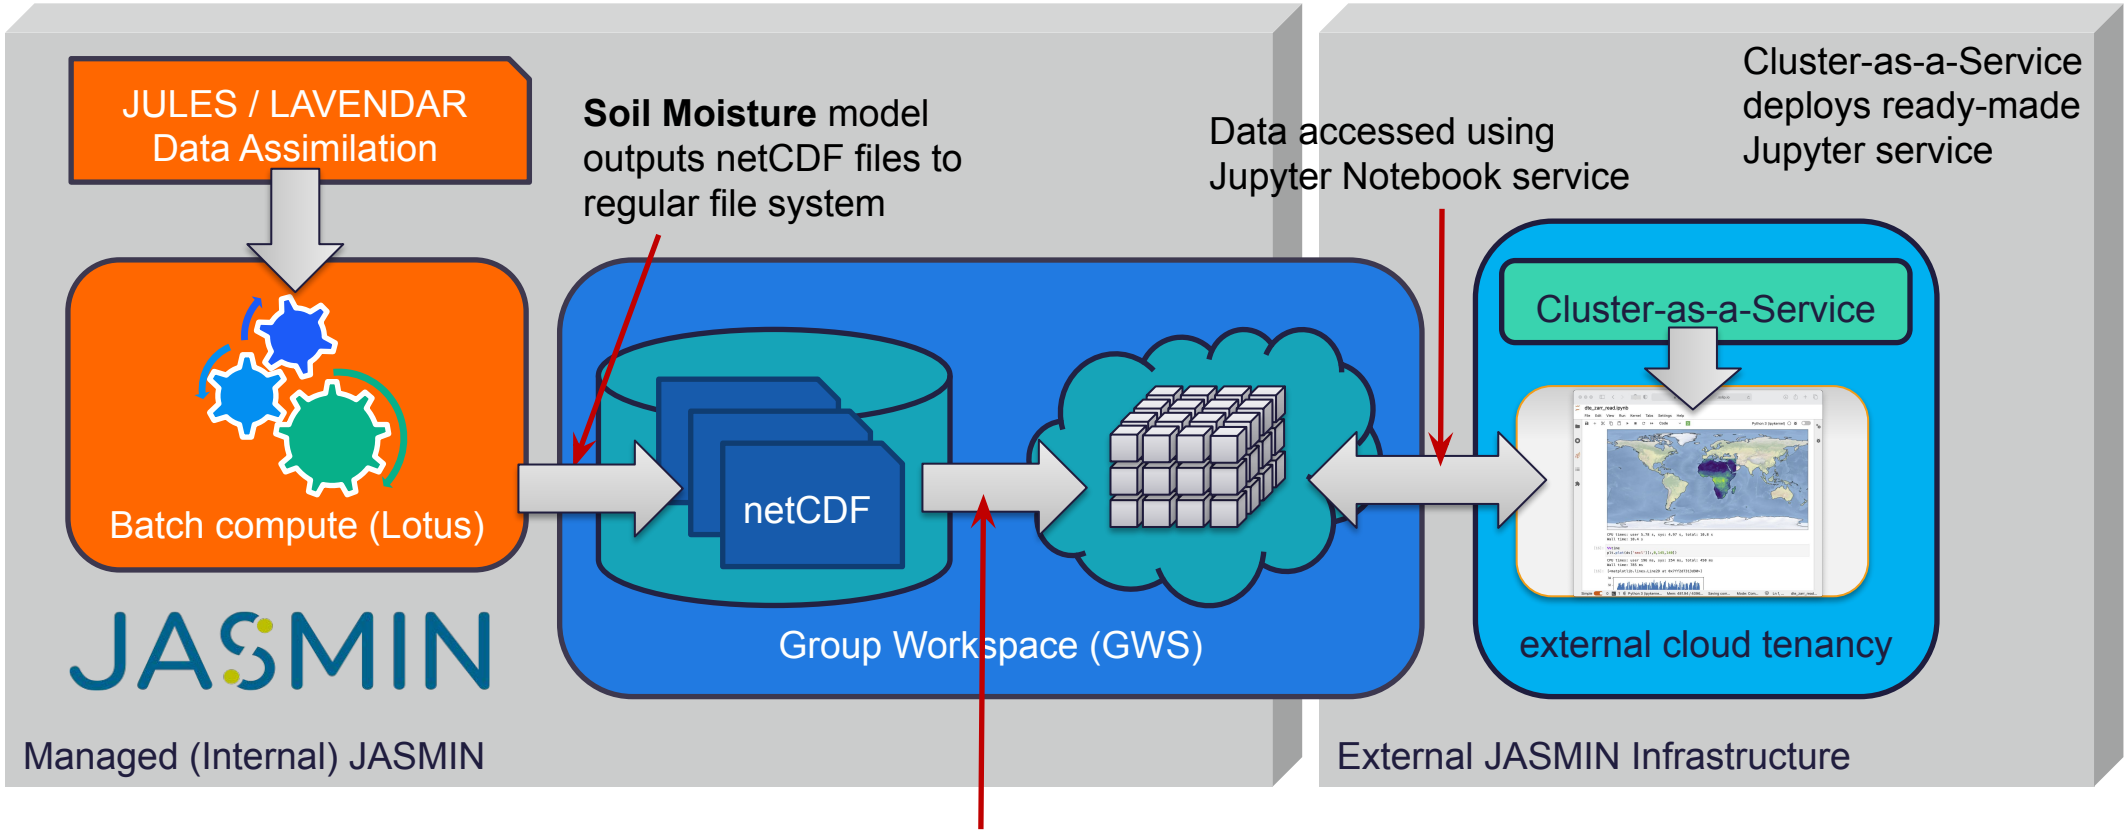

Arrangement of data and efficient access

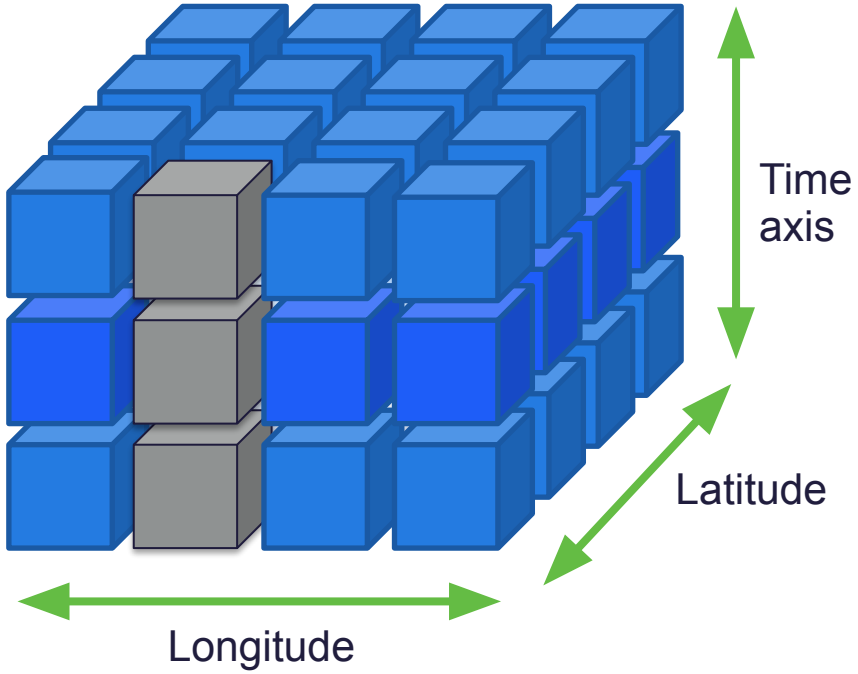

- Data output from models as netCDF format
- Data in files arranged in spatial dimensions one per time step
- But predominate access pattern for analysis of climate data in the project is time series query (grey blocks)

Object Store: Different storage strategies showed radically different performance

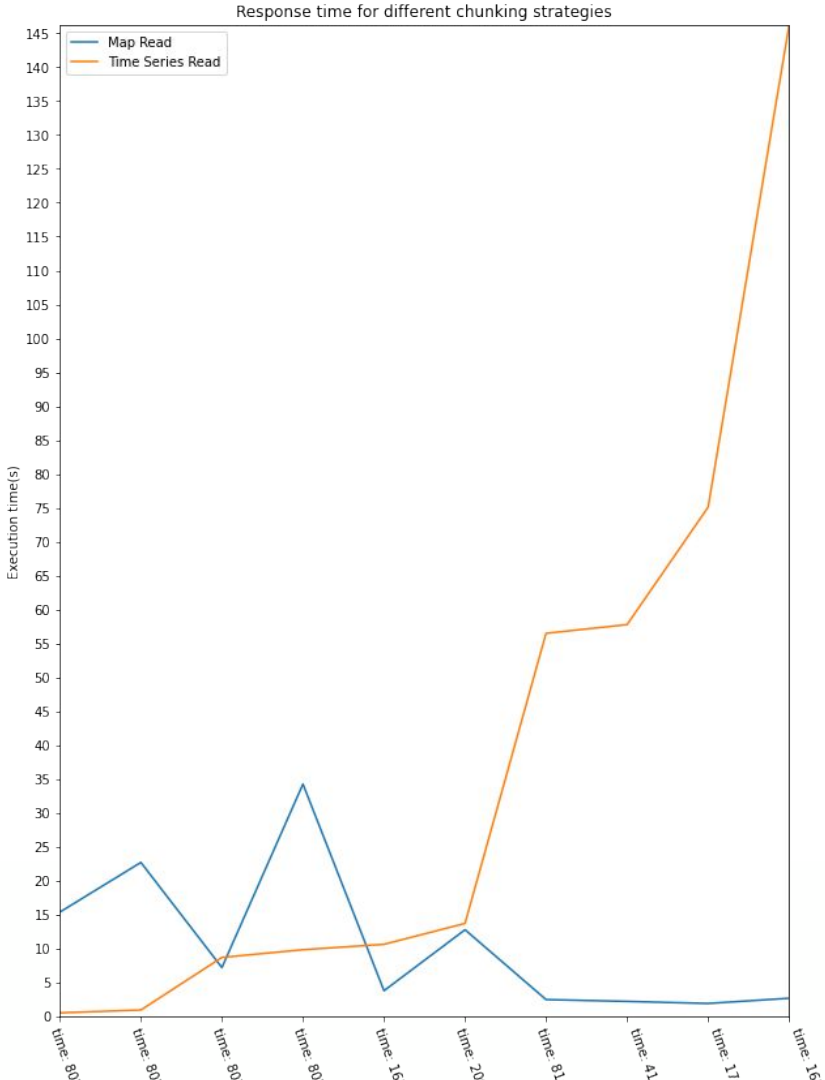

- We experimented with different storage chunking arrangements
- 20-year dataset of soil moisture

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Using Object Store for re-arrangement of data to suite our access patterns

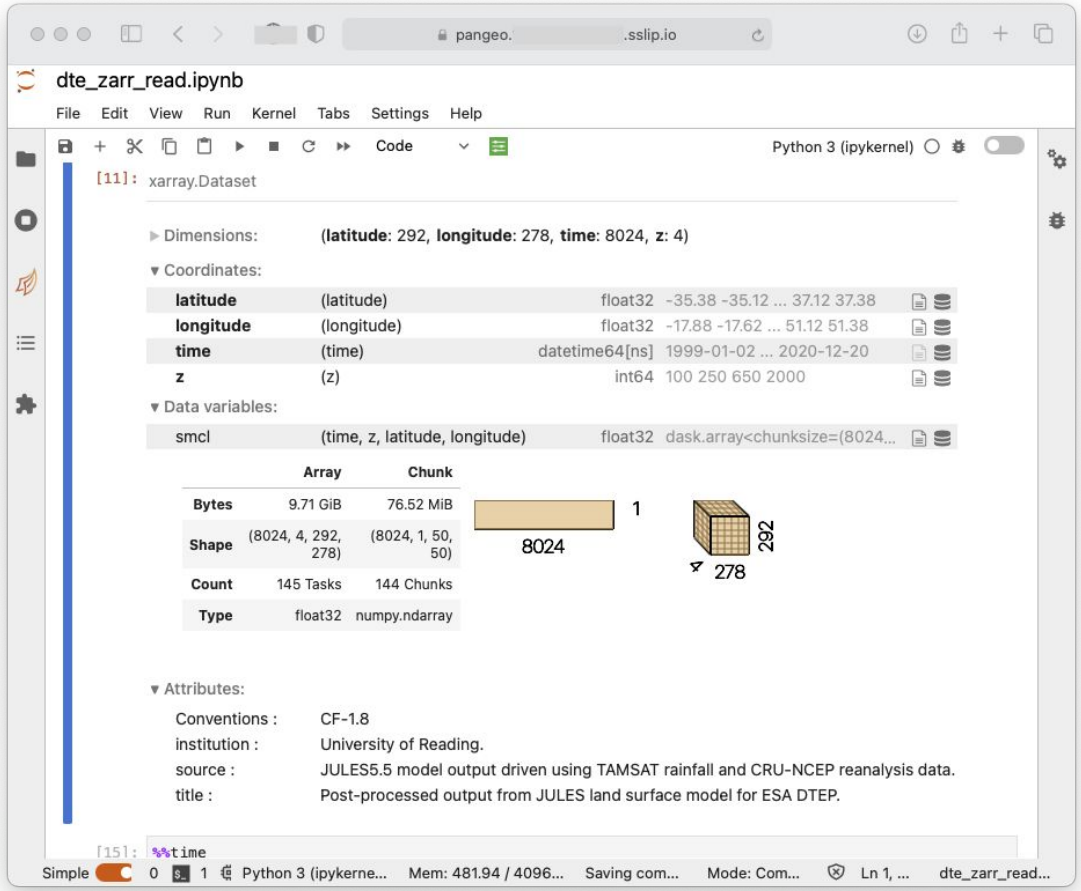

- Using zarr and xarray Python libraries to store and access the data
- Chunked data into a series of strips along the time axis

re for Environmental Analysis
ND TECHNOLOGY FACILITIES COUNCIL
ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Facilities Council
Research Council

Rechunking of data made possible interactive maps with long time series

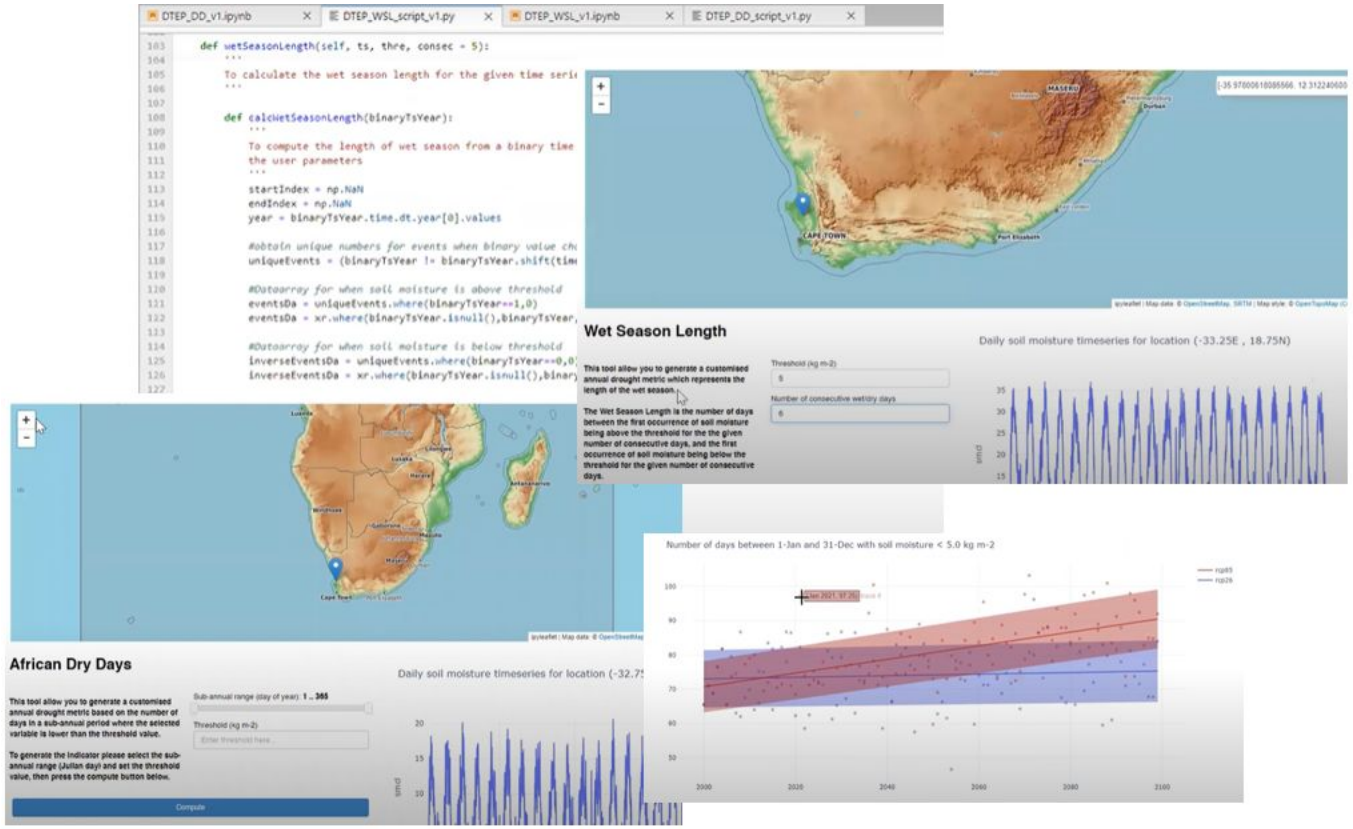

re for Environmental Analysis
ND TECHNOLOGY FACILITIES COUNCIL
ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Facilities Council
Research Council

Take home message: object store for analysis-ready cache specific to project needs

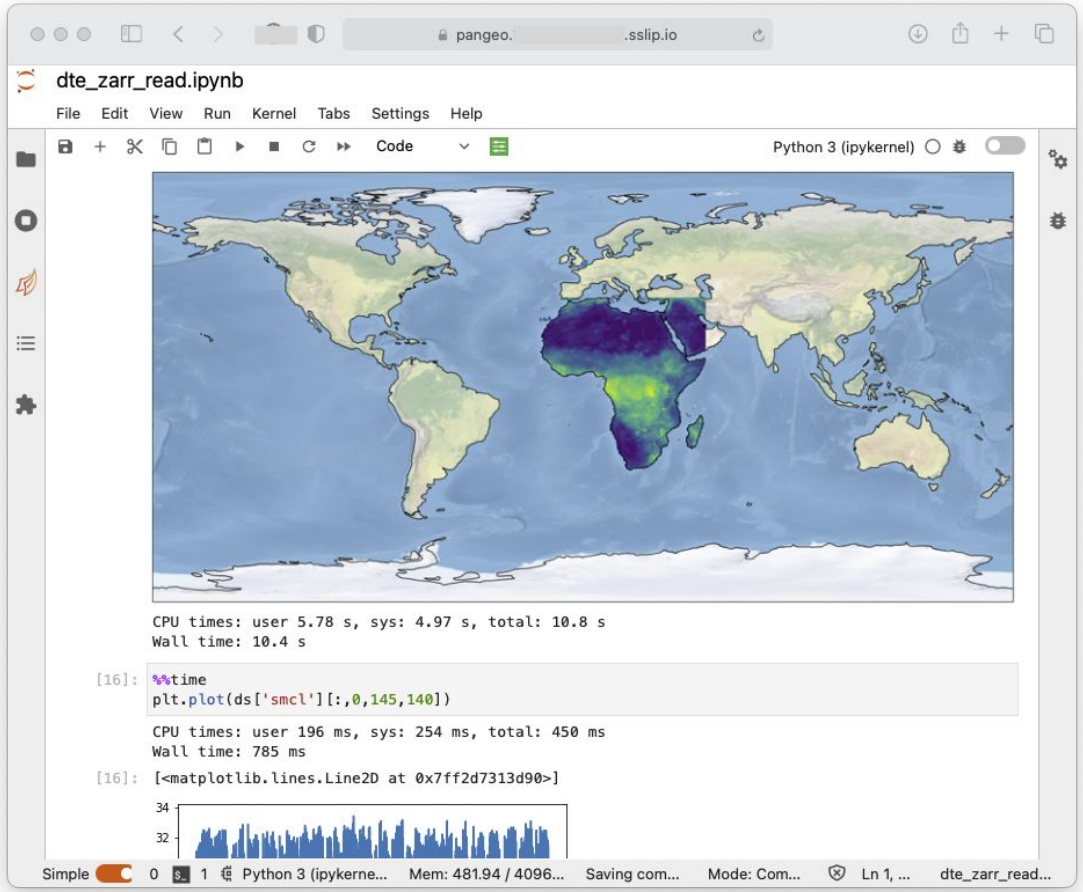

- Object store can be efficient for access on cloud
- It is essential to orient data storage to suit predominant access patterns
- Good news – re-writing data into different orientations was fast

re for Environmental Analysis
ND TECHNOLOGY FACILITIES COUNCIL
ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Facilities Council
Research Council

Futures

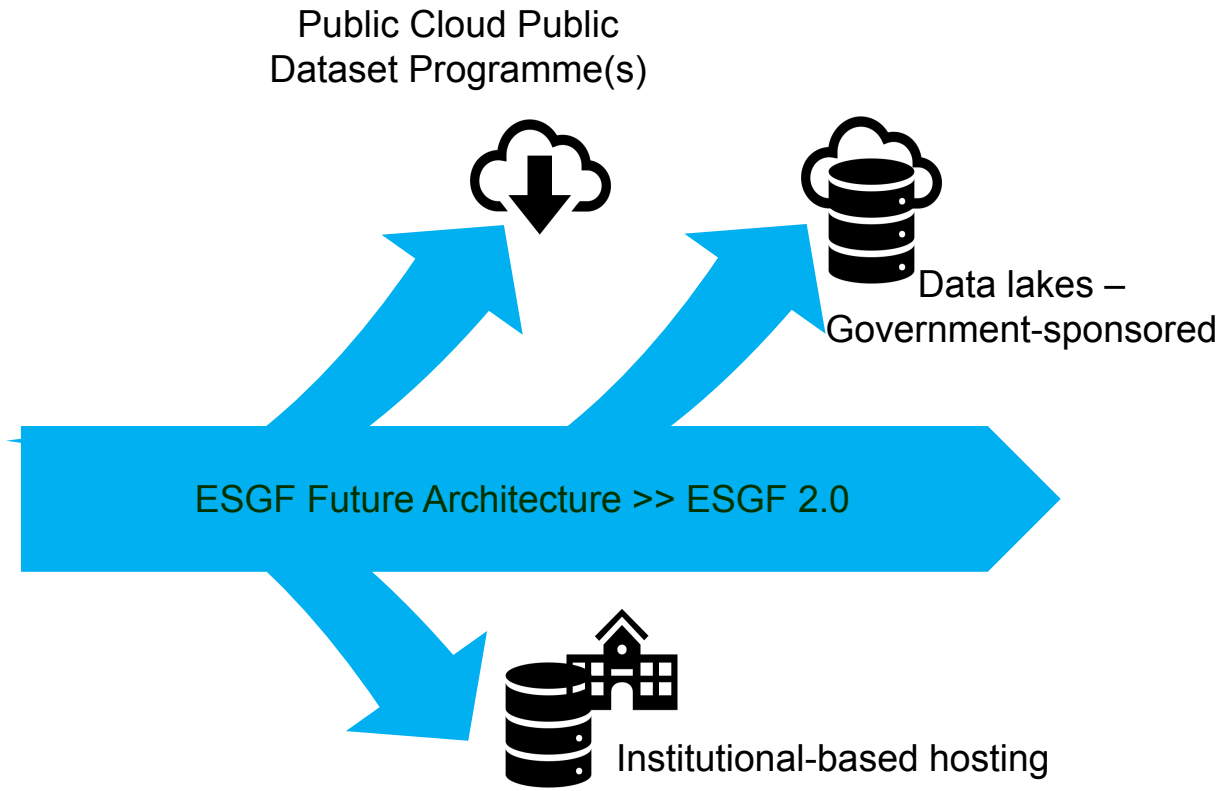
Public Cloud Public
Dataset Programme(s)
Data lakes –
Government-sponsored
ESGF Future Architecture >> ESGF 2.0
Institutional-based hosting

Acknowledgements + Further Info

ESGF Future Architecture Report: https://doi.org/10.5281/zenodo.3928222

IS-ENES3 website
https://is.enes.org/

@ISENES_RI
@cedanews
@PhilipJKershaw

Contact us at
is-enes@ipsl.fr

Subscribe to the
IS-ENES3 H2020
Youtube channel !

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement N°824084

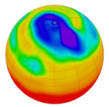

Centre for Environmental Data Analysis
SCIENCE AND TECHNOLOGY FACILITIES COUNCIL
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL