EXPECT

DATA INFRASTRUCTURE SETUP

Deliverable D7.1

Authors:
Stephan Kindermann,
Daniel Fulla,
Pierre Antoine Bretonniere,
Marco Puccini,
Serena Lorenzini

Document information

Grant agreement	101137656
Project title	Underpinning infrastructure for the efficient and flexible analysis of large climate datasets
Project acronym	EXPECT
Project start date	1 April 2024
Related work package	WP7
Related task(s)	7.1, 7.2, 7.3
Lead organisation	DKRZ
Authors	Stephan Kindermann, Daniel Fulla, …
Submission date	30 September 2025
Dissemination level	PU (Public)

History

Date	Submited by	Reviewed by	Notes
17/07/2025	First draft, S. Kindermann, D. Fulla		First draft
29/08/2025	Stephan Kindermann, Daniel Fulla, Pierre Antoine Bretonniere, Marco Puccini, Serena Lorenzini	To be reviewed by Expect Theme 4	With comments to be addressed
08/08/2025	Theme4	Christian Lessig, Bernd Funke	Version submitted to external reviewers

30/09/2024	Daniel Fulla et. al.	Christian Lessig, Bernd Funke	Final version with comments from reviewers addressed

Disclaimer: Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them.

Table of content

List of tables

List of figures

About

The climate system is changing rapidly and some regions have seen increases in extremes beyond what is expected from climate model simulations. To support targeted climate adaptation strategies, EXPECT will enable trustworthy assessments and predictions of regional climate change including extremes by developing a prototype operational capability for integrated attribution and prediction of climate. This ambitious goal is closely aligned with the WCRP Lighthouse Activity on Explaining and Predicting Earth System Change.

EXPECT will identify and quantify the mechanisms by which physical processes govern regional climatic changes, including extremes, on inter-annual to multi-decadal time scales. It will do so by exploiting newly available climate simulations and Earth Observations (EOs), and by combining machine learning (ML) with physical methods.

The research will target fundamental knowledge gaps related to atmospheric circulation and land-atmosphere interactions, which represent major limitations in current climate predictions and projections, and in particular in understanding changes in European summer extremes.

To underpin the research, and benefitting the wider research community, EXPECT will develop tools to efficiently analyze a variety of large data sets in combination that are hosted in different repositories across institutions. This will facilitate the exploitation of recent investments into high-resolution climate models and E O data. EXPECT will further build data science capacity for the scientifically robust, efficient and reproducible analysis of the massive data assets, including novel ML approaches, and provide training for the climate science community and the next generation of researchers in particular. EXPECT will thus deliver significant scientific and technological advances for society and the climate science community that will last well beyond the project, in support of WCRP's strategic objectives.

Executive summary

This deliverable reports work done as part of work package 7 which establishes the foundational infrastructure for distributed data analytics within the EXPECT project. The goal is to enable future cross-institutional, scalable and FAIR data workflows that support diverse scientific goals of EXPECT. The objectives focus on integrating distributed compute and storage resources, enabling data ingestion from key climate datasets, and developing unified catalogues and tools for efficient, location-aware analysis.

Over the first 18 months, progress has been made in gathering partner requirements, tracking dataset usage, and integrating institutional infrastructures across major European centers such as DKRZ and CINECA. A first version of processing services is available for user testing and interactive computing platforms such as JupyterHub have been deployed to support user workflows. A first version of an integrated data catalogue system (integrating externally managed catalogs like ESGF) is in place.

This deliverable presents an overview of the status of the integration of infrastructure components enabling workflows that meet user needs and enable future cross-site interoperability. It also outlines planned enhancements to improve data staging, federated catalogues, and hybrid cloud-HPC processing models.

These developments contribute to a growing ecosystem of European climate data infrastructure, complementing and building upon initiatives such as ESGF, Copernicus, and Pangeo. The EXPECT project's distributed approach aligns with current EU priorities for FAIR data and scalable, collaborative research environments.

The technical solutions proposed in this deliverable are based on the requirements expressed by the partners of the consortium in the context of task 7.1 whose conclusions were presented in milestone M10 "Collecting requirements for shared data infrastructure and distributed processing (architecture, services)" delivered at M12.

Requirements on infrastructure from partners' feedback

In M10, available on the project wiki[1], the users were asked about the datasets they were using and planning to use during the project, considering their volume, where they were planning to access the data and with which tools, programming languages, and protocols. The main conclusions and how they have been taken into account in the design of the infrastructure are described in this document.
The principal datasets that people are planning to use come from ESGF (CMIP6, LSFMIP), the Copernicus Datastore Store (CDS) or MARS (ERA5 or mars seasonal forecasts) and some other external High Resolution Climate models data (Climate Digital Twin DestinE, EERIE, NextGEMS). CMIP6 and ERA5, which are available in Jasmin[2] and DKRZ are the most requested datasets.
Users either download the data from the CDS[3] and ESGF (~60%) and MARS[4] (~25%) to their local machines, or take advantage of the prefetched data from institutional repositories such as CEDA, DKRZ, BSC. The data access (bandwidth and complexity) is the main driver to choose where and how to run the analysis. This clearly indicates the usefulness of proposing the CINECA computing platform in addition to the BADC and DKRZ ones already mentioned. The currently prevalent approach of downloading the data to local disk for analysis also shows the importance and future benefits in terms of performance of the infrastructure described below. Through its online processing capacities, it is expected to significantly increase performance of the analysis, avoiding the overhead of the download.

To support effective collaboration between infrastructure partners (BSC, DKRZ, UREAD, CINECA), the project organizes regular coordination meetings both within and across infrastructure work packages, as well as broader project meetings such as the Theme 4 kickoff, Executive Board meetings, and the EXPECT General Assembly.

Bilateral meetings between specific partners, for example, DKRZ and CINECA, are organized whenever closer alignment is required. Collaborative development can further be supported through shared code repositories hosted on the DKRZ and BSCGitLab platforms, ensuring transparent version control and facilitating streamlined contributions across all teams.

[1] https://earth.bsc.es/expect/lib/exe/fetch.php?media=wiki:content:expect_user_survey-m10.pdf

[2] High-performance data analysis and storage platform developed and operated by the UK Centre for Environmental Data Analysis (CEDA)

[3] Climate Data Store

[4] Meteorological Archival and Retrieval System

EXPECT Infrastructure

An overview of the distributed infrastructure supporting climate analysis workflows in EXPECT is provided in Figure 1. The infrastructure is initially built around the three core data centers DKRZ, BSC and CINECA. DKRZ and CINECA provide storage and compute resources together with a range of associated data and compute services. These will be progressively integrated and extended as part of EXPECT. BSC complements these resources by hosting additional internal data servers. It also acts as a bridge to the evolving DestinE data infrastructure. The data storage, transport and compute services available at these sites which need to be integrated to enable cross-institutional workflows are summarized in the following subsections.

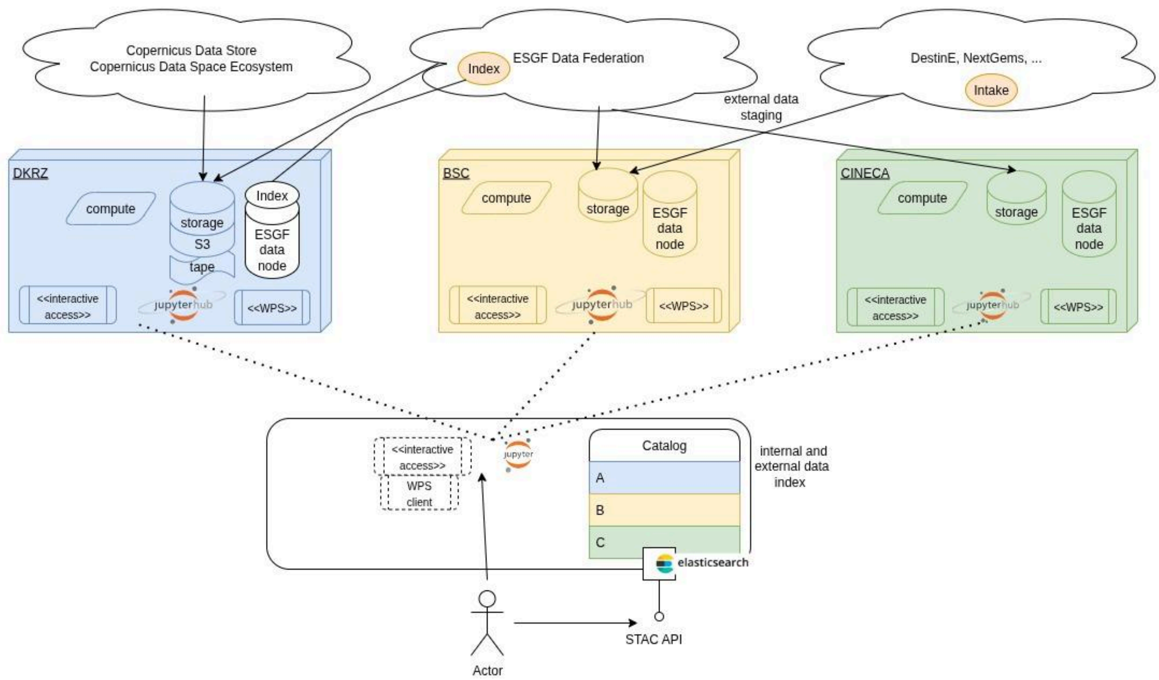

Figure 1 EXPECT Infrastructure Overview illustrates the distributed infrastructure connecting DKRZ, CINECA, and BSC, highlighting their complementary storage, compute, and data services.

To exploit DKRZ and CINECA resources as part of EXPECT workflows, scientists need to register at the sites and get access rights to resources and services:

To get an account at DKRZ and join the EXPECT project (project code: bk1444), users should first obtain a DKRZ user account by registering through the official portal. Once the account is approved, they should log in to the DKRZ user portal (LUV) and submit a request to join the EXPECT project via the project application interface.

At CINECA, users should begin by registering in the User Database (UserDB). After completing the UserDB registration and receiving approval from CINECA, users may have access to the EXPECT project systems hosted there.

With an account at either of the computer centers, users can access data storage, transfer data between sites, and use data processing capabilities. Work towards integration into a federated AAI[5] infrastructure (e.g. based on the EGI AAI[6] and aligned with EOSC AAI architecture efforts[7]) is not part of EXPECT, yet efforts are underway as part of other projects which involve EXPECT partners (especially DKRZ and BADC[8]/UKRI[9] e.g. in RI-SCALE[10]).

BSC provides additional infrastructure for standardized storage on disk and tapes, available exclusively to internal users only. This infrastructure supports datasets such as CMIP6, CDS seasonal forecasts, and ERA5. Additionally a dedicated workflow has been implemented to analyze DestinE data stored on the MareNostrum5 data bridge.

Data Storage and Transfer

Climate data analysis activities involve high volume data collections spread over distributed and different types of storage systems. On the one hand there is a need for fast, high-throughput storage access (e.g. based on disk or ssd hardware and associated high performance file systems) which are associated with (e.g. HPC) compute resources. On the other hand there is the need to temporarily share, transfer and store large amounts of data, which require automatable data transfer steps and associated data sharing systems e.g. based on cloud storage. Additionally there is the need to interact with long term archived data and thus access long term preserved data residing on e.g. tape-based storage systems.

[5] Authentication and Authorization Infrastructure

[6] EGI AAI: https://www.egi.eu/service/check-in-internal/

[7] EOSC AAI architecture task force: https://eosc.eu/advisory-groups/aai-architecture/

[8] British Atmospheric Data Centre

[9] UK Research and Innovation

[10] https://www.riscale.eu/

This section provides an overview of the EXPECT data storage and transfer infrastructure as well as first steps towards harmonizing and integrating data access services. A summary of problems encountered and open issues is given in section “data analysis workflows”.

Data Storage and transfer at DKRZ

DKRZ provides a comprehensive storage infrastructure[11] designed to support all stages of data-intensive research workflows. High-performance disk storage enables fast, frequent access to active datasets during simulations and analysis, while tape-based archival systems ensure long-term data retention and integrity. S3-compatible object storage services (swift- as well as minio-based) support cloud-native data sharing and collaboration. For long-term preservation and publication, platforms like the DOKU Archive and the World Data Center for Climate (WDCC) ensure traceability, accessibility, and compliance with FAIR principles. Extensive documentation is available, supporting the whole data management work cycle[12].

EXPECT data

At the core of DKRZ’s infrastructure lies a Lustre[13]-based high-bandwidth file system hosting core climate data collections as part of the Data Pool. Researchers can use the Data Pool as a shared community resource providing direct access to frequently requested climate data collections to ensure efficient workflows. All datasets include a README file detailing metadata, licensing, citation, and contact information. In collaboration with the respective external data providers, the DKRZ data management team ensures EXPECT data requirements are included. Specific collaborations include:

- ESGF: DKRZ hosts a large replica pool of CORDEX, CMIP5 and CMIP6 data. Current work is centered around preparations for hosting CMIP7 data based on a new ESGF2.0 infrastructure integration. To accomplish this, less frequently used data, e.g. CMIP5 and CMIP6 collections, will need to be removed from disk. EXPECT data requirements will influence this data removal process.

---

[11] DKRZ Levante HPC System: https://docs.dkrz.de/doc/levante/index.html
[12] DKRZ data services: https://docs.dkrz.de/doc/dataservices/index.html
[13] https://www.lustre.org/

- Copernicus: Climate data projection data (from CMIP5/6 and CORDEX) are provided to Copernicus via the DKRZ WPS interface (see section Data Processing at DKRZ) and regularly updated.
- In close collaboration with the EERIE project, DKRZ hosts EERIE data collections[14] enabling fine-grained (data analysis ready) data access. A dedicated data catalog has been developed based on STAC, which is integrated into the EXPECT data catalog system. Additionally, specific data access services for EERIE are provided, which support fine grained data access patterns.
- In close collaboration with the Warmworld[15] and NextGems[16] projects, DKRZ hosts high volume high resolution model data outputs.
- In close collaboration with DWD[17] and ECMWF, core ERA5 data collections are replicated and made available at DKRZ[18].

Smaller subsets of these data collections were also made available on S3 storage to support data sharing and for testing future data analysis workflows that more tightly integrate future (large) S3 storage capabilities at DKRZ (to be deployed during 2026).

To support future distributed data analysis workflows involving data transfers to DKRZ storage, data transfer tools and their integration with the different storage backends were reviewed and tested to identify the preferred tools for future work in EXPECT.

Besides operational tools at DKRZ[19], this also included Globus[20]. Globus is used, for example, in the ESGF data federation to enable parallel and high-volume data replication. Another tool is rclone[21], which was recently used successfully during a global high-resolution data hackathon[22]. Rclone is particularly well suited for syncing files on cloud storage.

[14] https://easy.gems.dkrz.de/simulations/EERIE/eerie_data-access_dkrz.html
[15] https://docs.dkrz.de/doc/dataservices/finding_and_accessing_data
[16] https://nextgems-h2020.eu/
[17] Deutscher Wetterdienst (German Meteorological Service)
[18] https://docs.dkrz.de/doc/dataservices/finding_and_accessing_data/era_data/index.html
[19] DKRZ data transfer tools: https://docs.dkrz.de/doc/levante/data-transfer/index.html
[20] Data Transfer with Globus: https://www.globus.org/data-transfer
[21] RCLONE: https://rclone.org/
[22] The World Climate Research Programme Global KM-scale Hackathon: https://www.wcrp-esmo.org/activities/wcrp-global-km-scale-hackathon-2025

Besides operational tools at DKRZ this included Globus, which is used e.g. in the ESGF data federation for parallel and high volume data replication and rclone[23] which was recently used successfully as part of a global high resolution data hackathon and is specifically suited for syncing files on cloud storage. The result of this evaluation are:

- The best option for data transfer is highly dependent on the use case (e.g. the storage locations of the involved data sets). We will concentrate essentially on three major options: uftp(and rsync) for conventional disk to disk transfers, Rclone for transfers involving cloud storage, and globus for (automatic) high volume data transfers.
- Rclone usage at DKRZ is currently hindered by ongoing work to define and deploy the future cloud infrastructure at DKRZ. The current infrastructure is outdated and faces sustainability problems and will be superseded by a new system to be deployed during 2026. Thus, real inter-institutional data transfer tests with CINECA only make sense later in 2026, also because of the cloud infrastructure deployment status at CINECA (see below).
- For all data transfer options, the independent user authentication and authorization (and thus the integration into the AAI[24] infrastructure) across HPC centers is currently hindering the development of more automatic data transfer workflows. When using lower-level tools like uftp[25], rsync and Rclone, users need to handle this by manual configuration of their institutional credentials. In the case of globus the integration into an integrated AAI system (e.g. provided by Globus or by integration with EGI[26]) is partially operational at DKRZ. Work is ongoing as part of the RI-SCALE project to integrate with the EGI AAI, yet this is currently limited to ESGF services at DKRZ and the roadmap for a full DKRZ AAI integration with Globus is open.
- Better data transfer integration of the DKRZ tape system[27] is currently hindered by the replacement of the current system with a new one, which will be deployed starting 2026. The EXPECT data sharing requirements, which especially include efficient support for interfacing cloud storage and the support for automatable high volume transfers (e.g. by globus support) were included into the requirements list for the new system.

[23] RCLONE: https://rclone.org/
[24] Authentication and Authorization Infrastructure
[25] UDP-based File Transfer Protocol
[26] European Grid Infrastructure, known as EGI Federation
[27] DKRZ HSM (tape archive): https://docs.dkrz.de/doc/datastorage/hsm/index.html

Data storage and transfer at CINECA

Initially, the most important data collection hosted at CINECA originates from OptimESM[28], which is a five-year project that aims to provide policymakers and society with crucial information about feasible mitigation pathways to achieve the Paris Agreement's goal of limiting global warming to 1.5°-2°C. The project also studies the consequences of temporarily or permanently overshooting these targets.

To achieve this, OptimESM brings together scientists to develop and improve a new generation of high-resolution Earth System Models (ESMs). These models are then used to produce new future climate simulations based on a range of new emission and land-use scenarios. The project will progress in five key areas, including delivering new knowledge on potential abrupt climate changes and their regional consequences, particularly in Europe and the polar regions.

EXPECT data

The EXPECT project will leverage data produced by OptimESM, which is published through the Earth System Grid Federation (ESGF) system. The OptimESM data is specifically hosted and managed on a dedicated ESGF data node[292] at CINECA. This direct integration is crucial as it enables CINECA to implement all the data processing and utilization services that are central to the EXPECT project's objectives, ensuring seamless access and efficient handling of the required information.

The ESGF community is currently undergoing a significant transition to a cloud-native architecture. This migration, referred to as the Next-Generation ESGF (ESGF-NG) infrastructure, involves moving away from legacy systems to embrace new technologies that ensure long-term sustainability and improved interoperability. As part of this transition, CINECA has proactively implemented a robust and future-proof solution for its data node. This implementation is based on Kubernetes orchestration, and it leverages S3-compatible object storage for efficient and scalable data archiving. This advanced setup ensures that the OptimESM data is managed on a modern, resilient, and high-performance infrastructure.

At CINECA, OptimESM data integration into the Earth System Grid Federation (ESGF) is supported by a modern Kubernetes[30]-orchestrated infrastructure featuring S3-compatible object storage and high-performance Lustre file systems. This hybrid

[28] https://optimesm-he.eu/
[29] https://esgf-optimesm.hpc.cineca.it
[30] https://kubernetes.io/ open source system for automating deployment, scaling, and management of containerized applications.

architecture enables seamless, high-speed access to distributed datasets and supports ESGF's ongoing transition to a cloud-native ESGF-NG platform. CINECA actively contributes to ESGF's sustainability and evolution through code contributions and best practices, positioning EXPECT as both a user and an enhancer of this global climate data federation.

Given the ongoing migration to ESGF-NG, the digital ecosystem is still in an active phase of development, and some components may still be subject to bugs or malfunctions. In this context, CINECA is an active participant in the community, contributing to the platform's stability and improvement. While the contribution is modest, it is nonetheless valuable. For example, CINECA has provided code contributions to key software components, as seen in the following pull requests:

- Add some specific configurations for ESGF data nodes which are S3 based deployments and to ensure Kubernetes Horizontal Pod Autoscaling (HPA): https://github.com/ESGF/esgf-docker/pull/258/files
- A small fix in the esgf-prepare tool, useful to prepare data to being published within an ESGF data node: https://github.com/ESGF/esgf-prepare/pull/84#issuecomment-3112438400

The following details the infrastructure and data processing experiments that were carried out up to M18.

- OptimESM outputs (dedicated ESGF data node)

- https://esgf-optimesm.hpc.cineca.it/thredds

Cloud storage

The storage component at CINECA consists of three parts: the ESGF node, the Kubernetes (K8s) cluster deployment and object storage. The ESGF data node serves as the CINECA node for EXPECT data. The K8s cluster was deployed using custom nodes on CINECA's on-premises ADA cloud.

Deploying the K8s cluster using a lightweight Kubernetes distribution (RKE2[31]) typically involves setting up 3 manager nodes (the control plane) to ensure high availability, as well as 6 worker nodes to handle workloads. Rancher is a powerful Kubernetes management platform that simplifies the deployment process through its intuitive user interface (UI) or command-line interface (CLI) tools.

[31] https://docs.rke2.io/

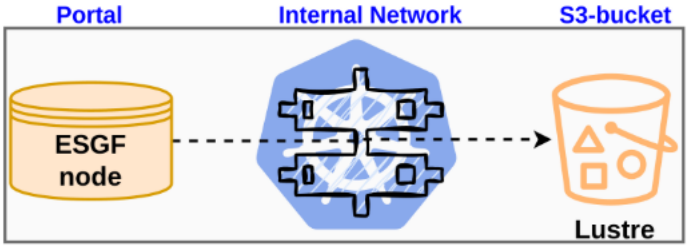

Figure 2: Cloud data storage setup

The data to be processed is stored in S3 object storage, which offers key advantages such as massive scalability, high durability and cost efficiency. The data is stored alongside rich metadata in a flat structure, making it ideal for unstructured data.

At Cineca, the S3 object store is integrated with the Lustre file system, which is a high-performance, distributed, parallel file system that is commonly used in high-performance computing (HPC) environments. It is optimised for fast, low-latency access to large datasets and supports POSIX-compliant file operations. Therefore, S3 with Lustre offers a powerful hybrid storage solution, combining the high performance of Lustre with the scalability and durability of S3.

A key advantage is that Lustre can be linked directly to an S3 bucket, enabling the high-speed processing of data stored in S3 with minimal manual data transfer. This setup provides fast, POSIX-compliant access to S3 objects as though they were local files. This is particularly advantageous for HPC, machine learning and big data analytics. As depicted in Figure 2, data transfer between the ESGF node and the S3 bucket takes place through a highly secure internal network.

It is worth mentioning that the above setup will be implemented and validated on the new Cineca GAIA cloud, in the future (post M20), which is likely to go into production in the coming months. The cloud will be integrated with cutting-edge GARDA storage technology based on VAST Data.

VAST is a modern data storage platform designed for massive unstructured data workloads and is a high-performance, all-flash system. Combining the speed of NVMe[32] with the scalability and efficiency of object storage, it supports file, block, and object access in one system. VAST is ideal for AI, analytics, and HPC applications, offering fast, scalable, cost-effective storage that avoids the complexity of traditional tiered systems.

[32] Non-Volatile Memory Express

Overview of data transfer options

To govern data transfer between devices over a network, one or more standardised protocols needs to be used. These protocols specify standardised rules and formats for data transfer and vary in terms of speed, security and complexity. The most suitable protocol will be chosen depending on the specific use within the project.

To move data to/from HPC systems, Cineca offers dedicated data transfer services that broadly fall into two categories: data movers and GridFTP. Data movers are dedicated, containerised nodes without interactive access that support only a limited set of commands (scp, rsync, sftp, wget, curl, rclone, AWS S3 and S3). GridFTP is also available on these nodes, but it can only be used via the globus-url-copy client, which must be run from the user's local machine.

Official documentation details the data transfer services (https://docs.hpc.cineca.it/hpc/hpc_data_storage.html).

Data Processing

Data processing within the EXPECT project is progressing through implementations at DKRZ and CINECA, where dedicated platforms are being developed to support scalable, efficient, and reproducible workflows for climate data analysis.

At DKRZ, the Levante HPC system is configured to enable both interactive and batch processing, using tools such as JupyterHub, Slurm, and the Rook Web Processing Service.

CINECA has deployed a cloud-native infrastructure based on Kubernetes, JupyterHub, and Dask, allowing for flexible, distributed data processing. These platforms can already support EXPECT use cases, and further work is ongoing to expand functionality and improve interoperability.

Data Processing at DKRZ

DKRZ offers HPC resources tailored for climate modeling and data-intensive workflows, with CPU-only, GPU-enabled, and interactive nodes. Storage relies on a high-performance Lustre-based file system for fast I/O. Typical workflows include data preparation (using Python, CDO[33], NCO[34]), batch simulations via Slurm, and postprocessing with data reduction and visualization. Automation through workflow managers and scripting enables efficient, reproducible research. Batch jobs run on Slurm, where users submit scripts specifying resources and commands. This non-interactive mode suits large, parallel simulations using MPI[35] or multi-threading.

Interactive data processing (jupyterhub)

DKRZ's JupyterHub provides an interactive, web-based environment on the Levante HPC system, enabling scalable, data-near analysis of large climate datasets like CMIP6 and ERA5. Users authenticate with DKRZ credentials, select resource profiles via Slurm, and launch sessions supporting Python, R, and Julia. Containerized environments and custom kernels allow reproducible workflows and advanced analyses.

For EXPECT, DKRZ's Jupyter notebooks repository offers hands-on tutorials and use cases for model data analysis. Notebooks demonstrate optimal HPC usage, processing tools, and data visualization. Users can clone the repository directly in JupyterHub, run demo notebooks, or create tailored kernels via conda/mamba for advanced workflows.

[33] Climate Data Operators
[34] NetCDF Operators
[35] Message Passing Interface

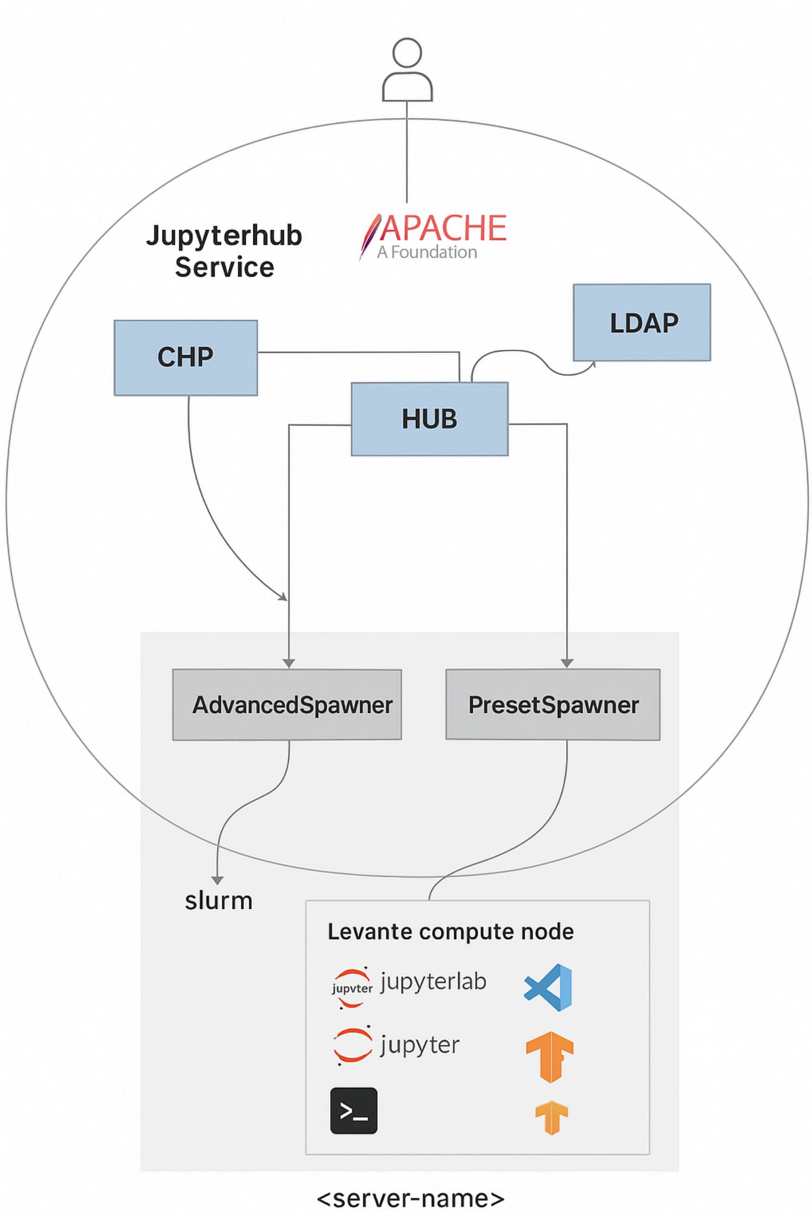

Fig 3 JupyterHub at DKRZ. The figure shows how DKRZ's JupyterHub service connects researchers to the Levante HPC system. Users log in through Apache with LDAP authentication, after which the JupyterHub hub manages sessions and routes requests via CHP. Resource allocation is handled by the DKRZSpawner, which offers both advanced and preset options and submits jobs to Slurm. Once scheduled, sessions run on Levante compute nodes, providing interactive environments.

Web Processing Service

Rook is a web-based climate data processing service developed under the ROOCS[36] project by DKRZ (Germany), CEDA/STFC (UK) and Ouranos (Canada). It offers standardized access to climate data processing via a Web Processing Service (WPS) API, enabling distributed workflows across diverse storage systems. Built on GeoPython tools, including pywps[37] (OGC WPS implementation[38]), clisops[39] (xarray-based climate data operations), rooki (a Python client for Jupyter integration), and daops[40] (data operation support), Rook facilitates processing close to the data, reducing large transfers and supporting FAIR principles.

Within EXPECT, Rook enables distributed, standards-based access to climate datasets and preprocessing services and integrates with interactive environments like Jupyter notebooks. Deployments include a permanent joint DKRZ & IPSL instance, a standalone DKRZ node integrated with ESGF and The ENES Climate4Impact[41] portal, and a CINECA installation optimized for S3-based HPC workflows. These run on virtual machines accessing local or cloud storage, enabling scalable and flexible processing. For example, the Copernicus Climate Data Store (CDS) uses a Rook adapter on a DKRZ VM to access CMIP6 and CORDEX datasets, complementing Copernicus' observational focus by providing model data processing with tools like regridding to harmonize spatial resolutions. Rook integrates with federated catalogs such as ESGF, Copernicus CDS, STAC, and Pangeo standards. However, its WPS interface currently lacks full compatibility with ESGF S3 storage instances. Regridding via clisops addresses challenges in grid inconsistencies and interpolation practices, offering a reproducible, programmable interface despite some community-level issues.

For hands-on use, the public WPS endpoint is rook.dkrz.de, with demo notebooks available at the rooki GitHub repository. Notably, the CMIP6 subsetting notebook

[36] Remote Operations On Climate Simulations
[37] PyWPS Web Processing Service implementation: https://pywps.org/
[38] Open Geospatial Consortium's Web Processing Service specification
[39] Climate Simulation Operations
[40] https://pywps.readthedocs.io/en/latest/
[41] https://www.climate4impact.eu/c4i-frontend/

supports time and bounding-box queries, and the time-components notebook demonstrates advanced subsetting and workflow examples.

Data processing at CINECA

The data processing setup involved deploying an easy-to-use, multi-user platform for processing data. The following details are from the testing and validation performed up to M18.

Overview of data processing components

The platform consists of several components deployed over K8s, including a user authentication mechanism, a JupyterHub interface and a dask gateway, as depicted in the following figure. Access to processing tools is currently restricted to authorised users (currently only internal staff) to ensure full access control and to avoid security breaches.

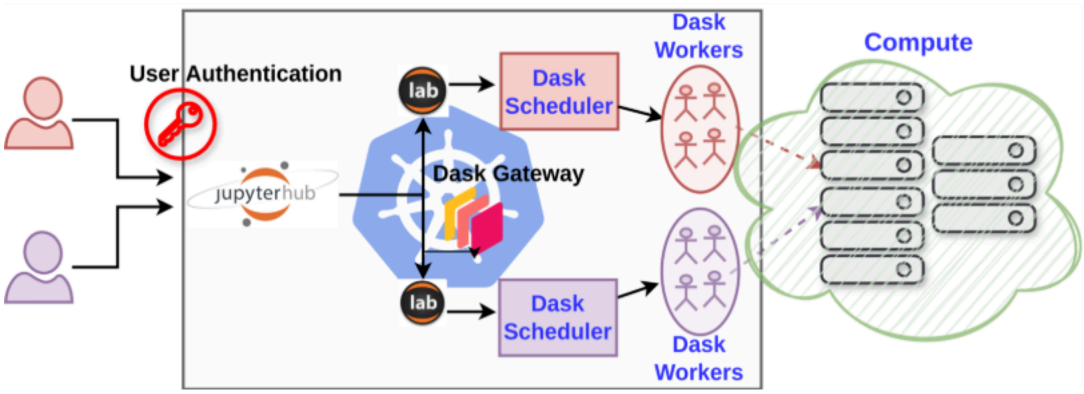

Figure 4: Multi-user Dask gateway implementation

The authentication has been enabled using X container/tool. The JupyterHub (*vide infra*) has been deployed with X container/Helm chart that allows multiple users to run their Jupyter notebooks directly on cloud infrastructure where the data is stored, enabling data-proximate computation. The third component of the data processing setup is the deployment of Dask-Gateway.

This is a powerful tool for managing and scaling Dask clusters. It does this in a secure, multi-tenant environment. This makes it ideal for distributed data processing in shared infrastructure like Kubernetes or HPC systems. It enables users to launch and connect to Dask clusters dynamically on demand while giving administrators control over resource limits and authentication. By separating the user interface from back-end cluster management, Dask-Gateway simplifies complex data workflows and enables efficient parallel computing across large datasets.

The deployment has been tested to address the several issues that includes; ensuring that the JupyterHub service exposes the cluster service through correct API endpoints, ensuring the compatibility of Dask and Dask-Gateway versions by avoiding outdated default images, enabling resource cleanup when the user shuts down the Dask cluster or closes the notebook server, enforcing user isolation via different namespaces to prevent two users from accessing the same Dask cluster, and examining the secure execution of Dask workflows.

Interactive data processing (jupyterhub)

JupyterHub is a multi-user orchestration layer that can be used to spawn and manage Jupyter notebook servers on shared infrastructure. Typically backed by Docker, Kubernetes, or other spawner backends, it provides centralised authentication, user session isolation, and resource management. It also supports integration with OAuth, LDAP, and other identity providers, making it suitable for secure environments. When deployed with Kubernetes, it enables scalable, containerised notebook environments for each user with configurable resource quotas and persistent storage. It is often the de facto choice for research platforms providing interactive computing environments with reproducible workflows.

JupyterHub can be effectively combined with Dask to provide computing resources and real-time performance monitoring via Dask diagnostics dashboards. When users run distributed Dask computations in their Jupyter notebooks within JupyterHub, Dask automatically launches a diagnostics dashboard providing live insights into task scheduling, memory usage and workload distribution. The following figure shows a screenshot of the typical diagnostics dashboard.

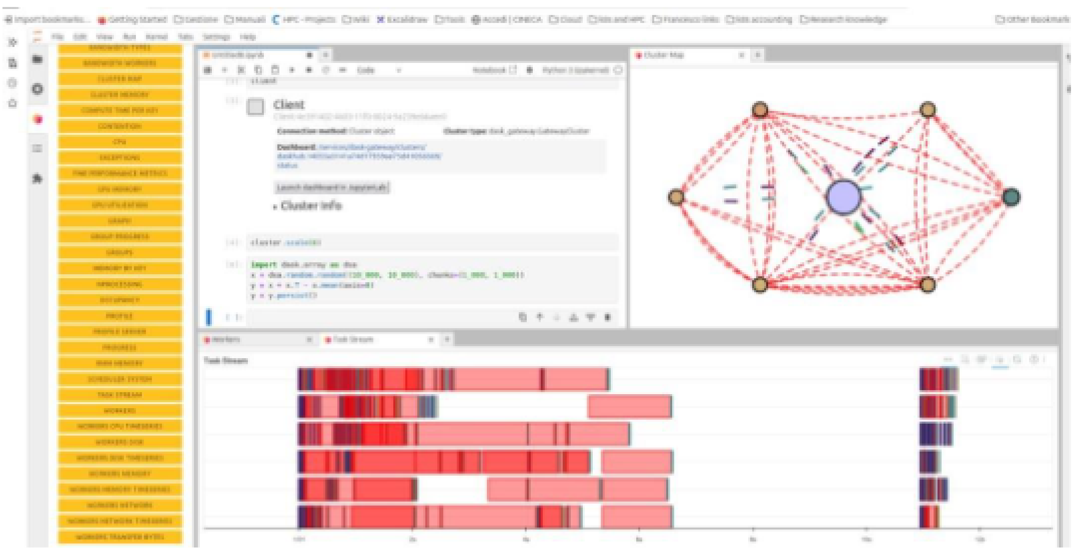

Figure 5: Dask data processing and diagnostics

Each user can then monitor their own computations via a link to the Dask dashboard. This is typically hosted on a dedicated port and can be securely proxied through JupyterHub. This integration is particularly beneficial in data-intensive environments, where performance visibility and scalability are essential. The implementation was tested to ensure the proper functioning of the dashboard, the isolation of individual user dashboards, and the possibility of exposing the dashboards to separate windows for visualisation convenience.

Web Processing service

The Rook-WPS service is part of a larger suite of tools known as ROOCS (Remote Operations on Climate Simulations). ROOCS implements a data-aware paradigm, which means it enhances data storage systems with the capability to analyze, process, and aggregate data directly where it resides. This approach allows for efficient processing of ESGF data by bringing the computation to the data, rather than moving large datasets.

The Rook-WPS service requires access to the same publicly exposed data made available through ESGF. Since the service is containerized, it can be hosted on a simple server or a virtual machine, as long as it can mount the volume where the data is stored.

This requirement, however, created a significant challenge for CINECA. The organization has updated its ESGF software stack to a new Kubernetes-based version that primarily uses S3 storage. This led to an extensive period of configuration and testing, as CINECA attempted to either use Rook directly with the S3 data or find a suitable workaround.

The only viable solution at present is to have an ESGF dataset replicated on both S3 storage (for publication via the new services) and a standard file system that can be mounted on the virtual machine hosting the Rook-WPS service.

Coincidentally, the OptimESM project data, due to its own specific project requirements, already possesses this dual-storage characteristic. Therefore, the next steps will focus on the configuration and deployment of the Rook-WPS service using this particular dataset.

Expect Catalog (of Catalogs)

Data collections which are used in EXPECT data analysis workflows are cataloged in different heterogeneous systems. However, currently a strong community effort focuses on harmonizing cataloging approaches based on Spatio-Temporal Asset Catalogs[42] (STAC) and the evolving associated STAC tooling ecosystem[43]. Thus, the newly developed Expect Catalog is based on STAC. It integrates external STAC catalogs but also deploys a dynamic STAC backend for additional data accessible at EXPECT partner sites. The basic structure is illustrated in Figure 6.

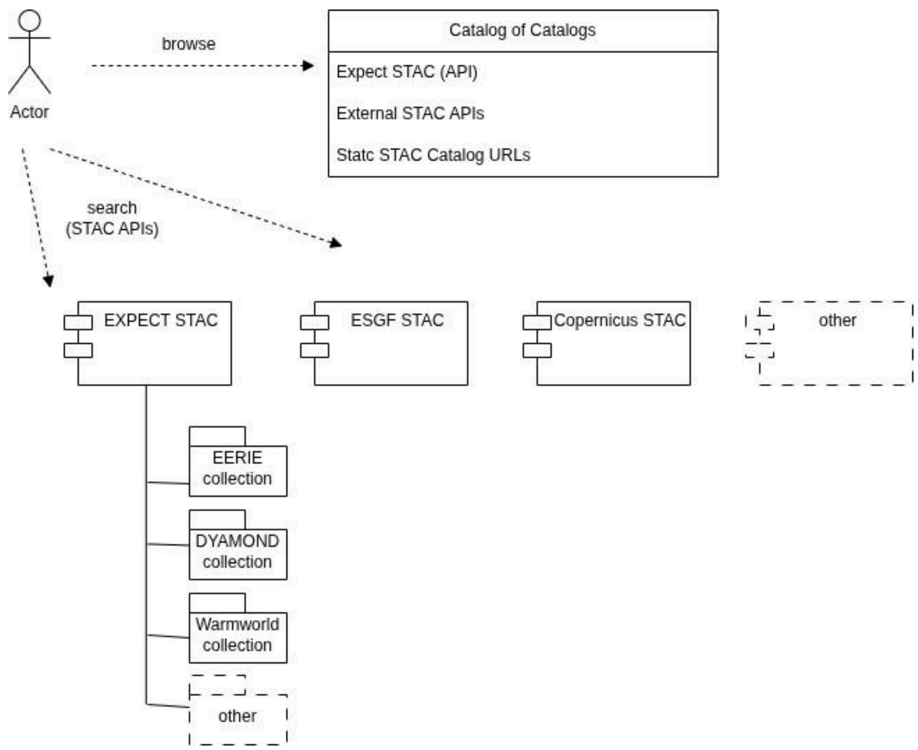

Fig 6. The EXPECT STAC catalog of catalogs

To be able to use tools from the rapidly evolving STAC tool ecosystem, the EXPECT catalog of catalogs can be browsed (with the the stac-browser tool[44]) based on a

[42] STAC Spatio Temporal Asset Catalogs https://stacspec.org/en
[43] STAC tools & resources: https://stacspec.org/en/about/tools-resources/
[44] Stac-browser: https://github.com/radiantearth/stac-browser

centrally managed json description[45]. This stac-browser view is available at https://discover.dkrz.de. This catalog of catalogs currently includes the following sub-catalogs:

- DKRZ / EXPECT STAC catalog: hosted at DKRZ, exposing climate model data from the following projects/efforts:
    - EERIE data collections
    - Warmworld data collections
- External STAC catalogs
    - ESGF (east) STAC catalog
    - Copernicus catalog
    - The DestinE data lake catalog

The EXPECT catalog of catalogs will be adapted, extended and refined during the project lifetime:

- The characterization of catalogs in the catalog of catalogs is currently based on metadata from the registry of research data repositories (re3data.org) effort. This characterization is quite generic and will be refined to include specific metadata useful to exploit climate model data repositories (e.g. providing information on the specific controlled vocabularies used to support data search).
- As mentioned before, many parallel efforts are ongoing to catalog climate data collections based on STAC and provide STAC APIs and catalogs. The EXPECT catalog of catalogs will follow these developments and associated ongoing harmonization efforts. For example, the STAC catalog that is currently integrated to provide access to data collections from the global ESGF data federation (e.g. CMIP6 and CORDEX) is based on the STAC catalog hosted at CEDA UK. It will be superseded by a STAC catalog which is fully synchronized with the ESGF (west) STAC catalog hosted in the US and will go operational to support CMIP7 later this year.
- There are also currently many early efforts to enable an automatic generation of static (as well as dynamic) STAC catalogs based on (e.g. temporary) data collections that are hosted on disk (and cloud) resources, associated metadata conventions, and controlled vocabularies. Different tooling support is currently under development, community wide (e.g. for

[45] EXPECT catalog of catalogs: https://tinyurl.com/ExpectCatalog

ESGF data publication) as well as at institutional level (DKRZ disk and tape data STAC catalog generators). Currently it is too early to decide on specific tool suites which can be proposed for e.g. EXPECT data providers wanting to register data in the EXPECT data catalog. Such data registration would support climate data sharing between climate data centers (and research institutes).

- To be able to exploit the EXPECT STAC catalog to support distributed climate analytics workflows as planned in workpackage 9, specific low level data access information will be included in the EXPECT catalog. This is necessary for users to have information on how to best access the information based e.g. on xarray and dask, as well as the low-level technical details necessary to configure client side tooling. For example, the exploitation of data collections based on STAC catalog entries is currently not as user friendly as those based on static intake catalogs[46], available for some data collections. The current DKRZ STAC catalog thus also includes information on data exploitation based on intake client side tooling. (see e.g. DKRZ static intake-esm data catalog[47])

Data analysis workflows

A key objective of the work in WP7 was to build the infrastructure foundation on which future work in WP9 can build upon to support distributed data analysis workflows involving high volume climate data collections in EXPECT. In Figure 7, a simple representative workflow example is provided, illustrating the necessary integration of the different infrastructure components that are summarized in the previous sections:

- Users interact with a centralized catalog to browse and search for the available data sets they need for the planned data analysis activity . This STAC catalog provides an overview of the data accessible at EXPECT HPC sites as well as via external data providers (e.g. Copernicus and ESGF). This is illustrated by the green lines in Figure 7.

[46] Intake: Taking the pain out of data access and distribution: https://intake.readthedocs.io/en/latest/index.html

[47] https://discover.dkrz.de/external/stac-dev-a722d9.gitlab-pages.dkrz.de/static/DKRZ-static-intake-esm.json?.language=de

● Users co-locate data at a centralized compute provider to be able to efficiently analyse the high-volume data sets. This involves data transfer from other EXPECT sites or external data providers (see red lines in Figure 7).
● Users trigger data analysis scripts at HPC sites. Either predefined scripts exposed as web processing services (WPS), or batch jobs running on the HPC and interactive analysis sessions e.g. based on Jupyter-hub exploiting parallel execution frameworks like such as dask (blue lines in Figure 7). As part of future work in WP9, broader support for parallel execution is envisioned by automatically exploiting distributed HPC centers, making the dask execution engines aware of data storage locations with the help of the STAC catalog.

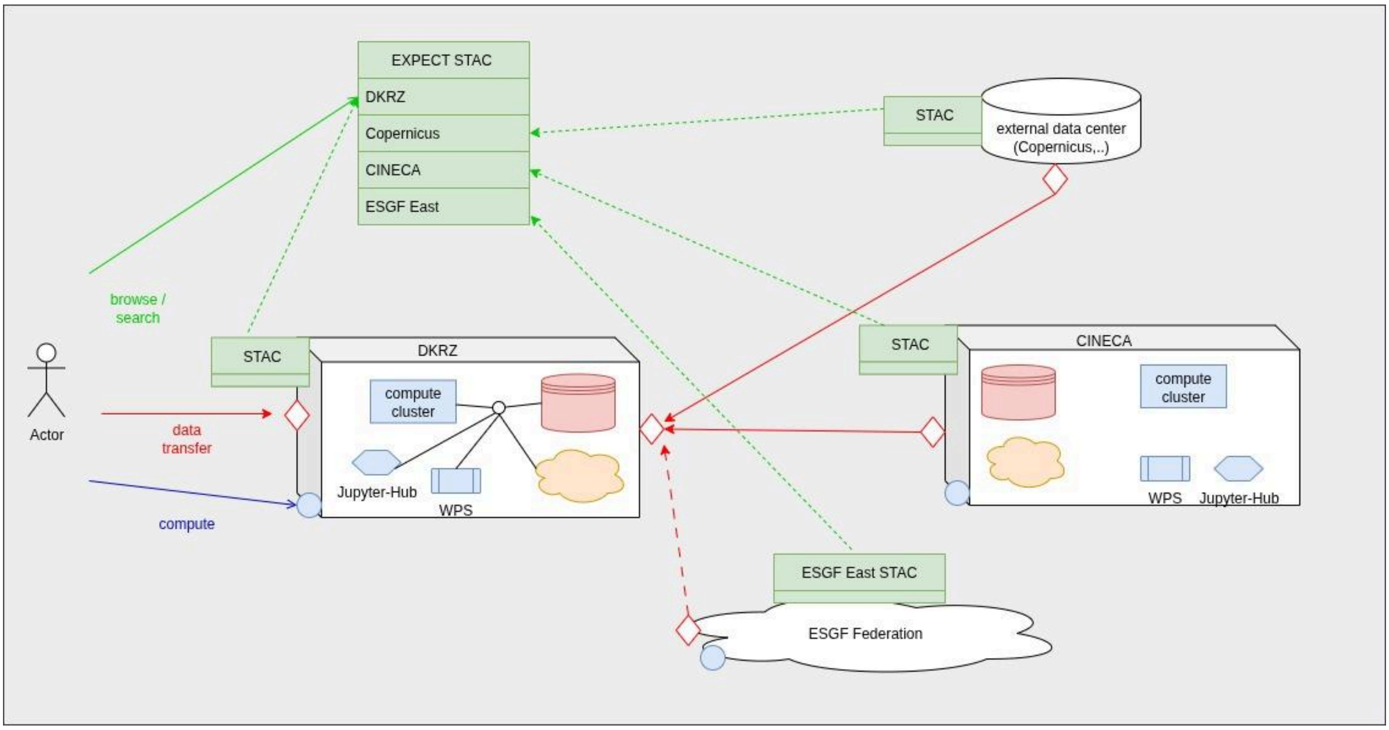

Figure 7. Distributed data analysis support

In the following a short summary of the status of work with respect to integration work is given:

- A first operational EXPECT catalog system has been established which can be used for the work starting in work package 9 in support of distributed workflow. This system consists of a catalog of catalogs and a STAC backend exposing a STAC API. Most of the data requirements (see section "Requirements on infrastructure") are addressed. There is a very dynamic evolution of STAC approaches currently at institutions and organizations worldwide, thus many adaptations are foreseen over the next year. Also, many aspects of the cataloging approach have to be refined, including usage of and agreement on specific STAC metadata profiles, harmonization of approaches towards data access and data aggregation. The STAC catalog of catalog support is currently limited. Catalog of catalogs metadata is currently not searchable via the STAC API as nested catalogs are at the moment not fully supported by STAC backend implementations[48]. Federated collection discovery[49] is currently explored e.g. in the NASA MAAP project[50]. Results will influence the EXPECT collection search approach.

- With respect to data transfer, different approaches have been reviewed but progress is currently hindered by the different storage systems (cloud, disk, tape) and their access policies at centers (mostly motivated by security considerations). Also DKRZ and CINECA are also currently in a phase of transitioning to new systems or system configurations. Thus, the most stable options for data transfer for now were selected for the next steps.
- Efficient compute capabilities/services integration with storage backends has severe security implications (open service API vs. protected internal HPC storage systems) which hindered the cross institutional use of e.g. web processing services by now. Future work will establish dedicated VPN channels between sites to address this.
- Access to interactive processing environments both at CINECA and DKRZ has been established. This allows direct access to data collections and supports parallel processing (e.g. based on DASK), which will be explored as part of work package 9.

Summary

During the first 18 months, progress has been made toward establishing a foundational infrastructure for distributed, FAIR-aligned climate data analytics. By leveraging and extending existing institutional capabilities at DKRZ and CINECA, the groundwork has been laid for scalable, interoperable workflows that address the growing volume and complexity of Earth system data.

A central pillar of this infrastructure is the development of a unified STAC catalog (https://discover.dkrz.de), which integrates a wide range of climate datasets, including CMIP6, EERIE, NextGEMS, into a single searchable framework. These catalogs support real-time API queries and enable metadata-driven workflows via

[48] Nested catalogs in stac-fastapi-pgstac · Issue #79 · stac-utils/stac-fastapi-pgstac
[49] Federated-collection-discovery: https://github.com/developmentseed/federated-collection-discovery
[50] https://maap-project.org/

tools like pystac-client and intake-stac. While STAC is inherently dynamic, static catalogs can also be produced on request for environments that require them.

EXPECT's work is tightly connected to the broader European climate research infrastructure landscape, particularly ENES-RI (European Network for Earth System Modelling Research Infrastructure), which coordinates efforts to develop, share, and exploit Earth system model data across Europe. EXPECT complements ENES-RI by focusing on federated data discovery, harmonization, and processing layers, thereby enhancing interoperability and user access across distributed resources.

The project also builds on the conceptual and technological framework proposed by FUTURA, a proposal for an EC-funded project aimed at enabling future climate data infrastructures grounded in cloud-native and FAIR principles. EXPECT adopts many FUTURA design patterns, especially regarding metadata standards, cloud storage integration, and API-first catalog access, contributing practical implementations and valuable feedback to this evolving ecosystem.

Complementing the storage layer, interactive, container-based data processing environments have been deployed on Kubernetes. This includes a secure, multi-user platform featuring JupyterHub and Dask-Gateway, enabling scalable parallel computation with dynamic resource allocation, session isolation, and real-time diagnostics dashboards. These tools empower users to perform interactive analysis directly where the data resides, reducing transfer overhead and improving performance.

In parallel, initial work to extend Rook's Web Processing Service (WPS) to support remote storage backends such as S3 and VAST multiprotocol systems has begun. Currently limited to local file access, future Rook versions will enable streaming data processing from remote storage and seamless integration with STAC Items, making Rook a flexible tool for cloud-native, metadata-driven workflows within EXPECT.

Future Work

While the infrastructure foundation is proposed, several key developments are planned for the next phase of the EXPECT project,

The full integration of dynamic STAC catalogs across participating sites will be pursued, including the indexing of additional datasets, improved support for real-time metadata updates and automatic catalog validation. For this EXPECT partners from DKRZ will also participate in upcoming STAC community hackathons[51]

[51] STAC Community Sprint, October 2025, Rome https://github.com/radiantearth/community-sprints/tree/main/14102025-esrin-rome-italy

especially targeting aspects of high-resolution climate data cataloging approaches (e.g. of ZARR data).

Development and demonstration of STAC-native workflow tools will also be a key focus. These tools will directly consume STAC metadata to initiate processing workflows, increasing reproducibility and automation. Integration with Rook and Dask workflows will be prioritized to support seamless, metadata-driven analysis pipelines.

Interactive data processing environments, such as JupyterHub and Dask-Gateway, will be expanded to support a broader user base. Enhancements will include implementation of OAuth or LDAP-based authentication, namespace-level isolation, and persistent user storage to ensure secure and scalable access to computing resources.

Ongoing alignment with the broader climate data infrastructure ecosystem remains essential. Active contributions to ESGF-NG and related open-source tools will continue, ensuring compatibility with evolving metadata standards, accessibility practices, and sustainability goals. Deeper collaboration with ENES-RI and the FUTURA initiative will be pursued to promote convergence toward a unified, cloud-ready European Earth system data ecosystem.