



Impact Report

Advanced Computing for EOOSC

Jan 2021- June 2023

Foreword



EGI-ACE was EGI's latest flagship project between 2021-2023. We launched this endeavour in turbulent times – during physical meeting restrictions due to COVID-19 and later during limitations on open science and supply chains due to the war in Ukraine. And still, EGI-ACE not only succeeded in its activities and achieved its objectives but became one of the shining stars of the European Open Science Cloud project landscape.

We reached unprecedented scales of compute delivery across diverse sectors and scientific disciplines. We empowered a large ecosystem of thematic services and data spaces made for open science. We have federated providers and user supporters from over 30 countries. Our project not only brought EOSC to the next level but also expanded the EGI Federation and EGI Community with new members, new services, new partners and new users.

As project technical coordinator, I am delighted to bring you this brochure on behalf of the 33 consortium members, 24 third parties and more than 104,000 users we served in the past 30 months. The next pages provide a comprehensive overview of the achievements and impact of EGI-ACE on Open Science, Research, Collaboration, Innovation and Skills.

Enjoy reading!




Gergely Sipos

EGI-ACE Technical Coordinator

Head of Services, Solutions and Support at EGI Foundation

Table of Contents




03

Introduction


04

Key results



05

EGI-ACE services



11

EGI-ACE impact



11

Impact on Open
Science

13

Impact on
Research

15

Impact on
Collaboration

18

Impact on
Infrastructure
Innovation

22

Impact on
Skills and
Expertise

24

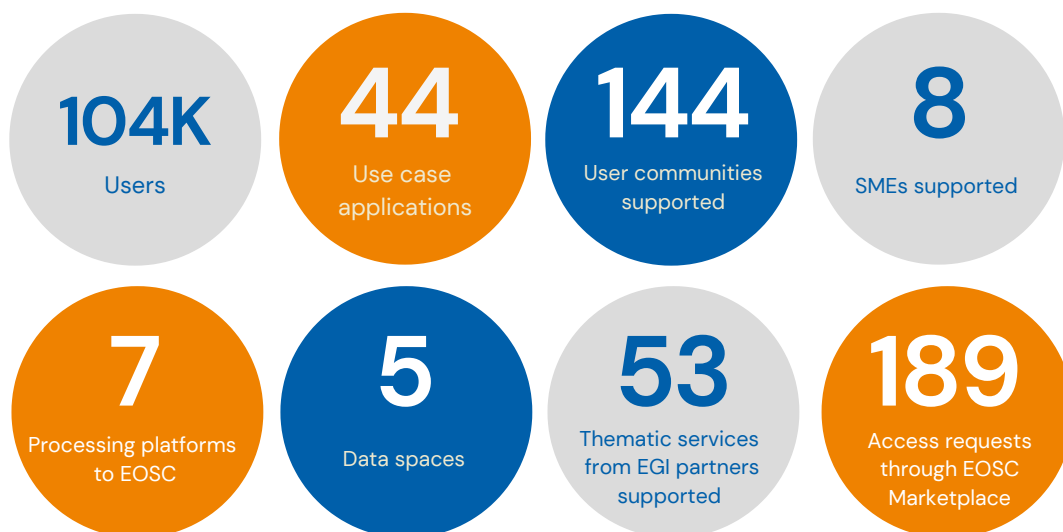
User
success stories

Introduction

EGI-ACE is a 30-month project coordinated by the EGI Foundation with a mission to **empower researchers from all disciplines to collaborate in data- and compute-intensive research through free-at-point-of-use services.**

The project consortium has built on the expertise and assets of the EGI federation members, key research communities, data providers and collaborating initiatives.

EGI-ACE has implemented the **Compute Platform of the European Open Science Cloud** and contributed to **EOSC Data Commons** with integrated computing, platforms, data spaces and tools aligned with major European cloud federation projects and HPC initiatives.



Key results

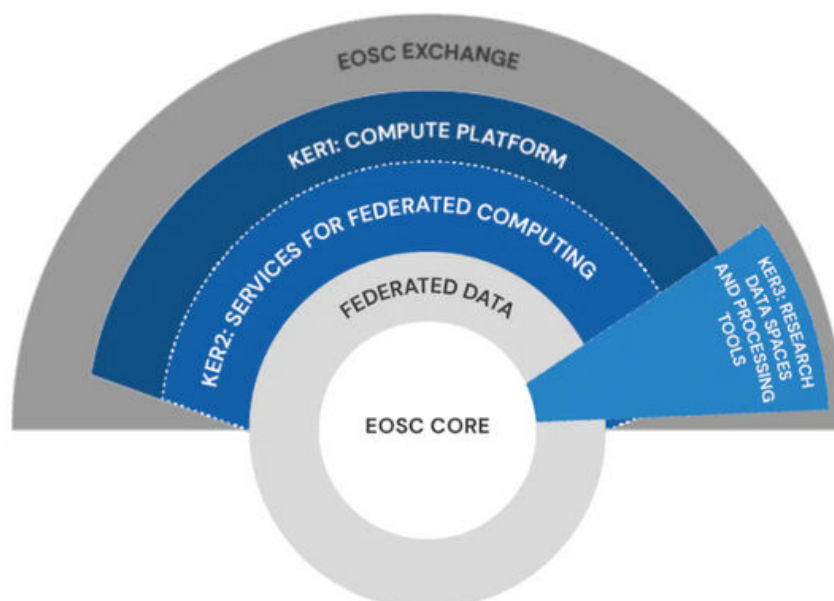
Over the two and a half years, the project presented tangible instances and stories of success, showcasing the capabilities and potential of federated resources within EOSC to drive science and support innovation.

The project concretely contributed to the design and implementation of EOSC through its three **Key Exploitable Results (KERs)**.

KER 1 EOSC Compute Platform services


KER 2 Services enabling federated computing in EOSC

KER 3 Data spaces and processing tools for EOSC



The EGI-ACE Key Exploitable Results (KERs in blue) overlaid on the EOSC Architecture

EGI-ACE services



KER 1

EOSC Compute Platform services

The EOSC Compute Platform is a free-at-the-point-of-use, distributed computing environment. It is built on a hybrid infrastructure composed of cloud computing resources, High-throughput computing (HTC) sites and High-Performance Computing (HPC) centres. It empowers users with higher-level services to ease the setup and operation of complex workflows, applications, containers, virtual research environments and data spaces on top of the hybrid infrastructure.

The Platform supports diverse data processing and analysis use cases. Thanks to EC and national funds, it provides free at-the-point-of-use services with user support and training for research infrastructures, communities, projects and the long tail of science. E-infrastructure providers joining the EOSC Compute Platform can benefit from the simplified integration with EOSC, streamlined user access handling and scalable resource allocation mechanisms, and various financial incentives.

Federated resource provider services

- [EGI Cloud Compute](#) Run virtual machines on-demand with complete control over computing resources
- [EGI Cloud Container Compute](#) Run Docker containers in a lightweight virtualised environment
- [EGI Online Storage](#) Store, share and access your files and their metadata on a global scale
- [SURF EGI High-Throughput Compute](#) Execute thousands of computational tasks to analyse large datasets

Compute and data federation services

- [EGI Check-in](#) Authenticate and Authorise at EGI and EOSC services in a uniform and easy way
- [IISAS Dynamic DNS Service](#) Register meaningful and memorable DNhostnameses for applications running on the EGI infrastructure
- [EGI DataHub](#) Access key scientific datasets in a scalable way
- [EGI Data Transfer](#) Transfer large sets of data from one place to another
- [STFC Rucio Data Management Service](#) Scientific Data Management
- [STFC CVMFS Content Distribution Service](#) Deploy software on the worldwide distributed computing infrastructure used to run data processing applications
- [OpenRDM EU](#) Use a complete solution to managing your research data, from the bench to the publication, all across Europe.

Platform services

- [EGI Workload Manager](#) Manage computing workloads in an efficient way
- [EGI Notebook](#) Create interactive documents with live code, visualisations and text
- [EGI Replay](#) Reproduce and share research on a notebooks-based platform
- [Elastic Cloud Compute Cluster \(EC3\)](#) Deploy a virtual cluster on top of IaaS infrastructures with a few clicks
- [Infrastructure Manager](#) Deploy complex and customized virtual infrastructures on multiple back-ends
- [Dynamic On Demand Analysis Service \(DODAS Portal\)](#) Use our multi-user on demand solution for both interactive (Jupyter based) and distributed data analytics
- [PaaS Orchestrator](#) TOSCA-based deployment orchestration service on multiple IaaS
- [DEEP training facility](#) Distributed training facility for Machine Learning, Artificial Intelligence and Deep Learning models

Value for Users

- Having a broad and growing spectrum of types of compute resources behind a single interface (EGI-ACE Open Call) and single allocation process.
- Streamlined support and communication through the use case shepherds and with access to experts for consultancy to select the most suitable types and combination of compute and related platform/data services to realise a use case.
- Build partnership with national providers for long-term usage of their compute resources.
- Free-at-point-of-use access to national and relevant international capacity.

- The distributed setup allows for computation to happen where the data is removing the need to move large amounts of data and any privacy concerns.
- Reusing software across providers with the use of AppDB removes the hurdle of installing software at every provider.

Value For Providers

- Providers who join the EOSC Compute Platform benefit from the simplified integration with EOSC
- Co-funding by the EC to deliver relevant services across national borders.

Value for EOSC

- The Platform allows the efficient use of European Commission and national funds by integrating cross-border with national access mechanisms, maximising the return of investment for all stakeholders.
- Computing infrastructure is wholly integrated with the EOSC ecosystem making EOSC an end-2-end destination.



KER 2

Services enabling federated computing in EOSC

EGI-ACE project delivered various services that ensure the efficient management of the EOSC Compute Platform as a federated environment. These enabling services include technical elements (Check-in, Configuration Database, Monitoring service, Usage Accounting system, and Helpdesk), as well as non-technical elements and coordination activities, such as an IT Management System, service security oversight, incident response team, and resource allocation team.

For providers, these services ensure a simplified and seamless integration into the Compute Platform and, therefore, with EOSC. For users, this provides a scalable resource allocation approach that considers local and European policies and the needs of national and international research communities. The project also developed a knowledge catalogue with best practices and training materials to help providers of the EOSC Compute Platform improve energy efficiency at data centres.

Value For Users

- A simplified landscape of protocols and tools to get information about and to interact with compute services in EOSC
- A scalable resource allocation approach that considers local and European policies and the needs of national and international research communities

Value for Providers

- The result enables standardised 'access to market' mechanism to service providers with built-in mechanisms to ensure fairness across the group of providers.
- Simplified order handling and customer relationship management (because of the first-line CRM team that pre-analyses and brokers orders to best-fitting providers because of the shepherd who liaises with the use cases).
- Being part of a computing community that advises each other on topics of shared interest (e.g. containers, green computing).
- Working with like-minded compute centres on harmonised policies, protocols, and approaches for service operation, architecture and funding.

See the services for federated computing



KER 3

Data spaces and processing tools for EOSC

Leveraging the EOSC Compute Platform, EGI-ACE has set up a thriving collection of thematic services in EOSC (data spaces and processing tools). The services are built by scientific communities, research infrastructures and projects by integrating scientific datasets and applications with the EOSC Compute platform. The project consortium includes Data Spaces and Processing Tools from Health and Medicine, Climate Research, Energy and Physical Sciences, Environmental Sciences, and Social Sciences and Humanities.

The project supported the setup of additional thematic services via its user engagement programme. The services are registered in EOSC Portal and are open for any researcher to perform scalable scientific data analysis. The services facilitate research and development activities, reduce development time for digital science, and contribute to the cross-fertilisation of knowledge across disciplines.

KER 3 includes 53 Thematic services, out of which 17 are provided by the EGI-ACE consortium and 36 services are provided by external stakeholders relying on EGI-ACE for infrastructure and platform capabilities.

Health and medicine

- [UseGalaxy.eu](#) Open, reproducible, web-based platform for data intensive research
- [Virtual Imaging Platform](#) European Galaxy Server – Open, reproducible, web-based platform for data intensive research
- [WeNMR – HADDOCK](#) Integrative modelling of biomolecular complexes
- [WeNMR – DisVis](#) Visualisation of interaction space between two molecules
- [WeNMR – PowerFit](#) Fit an atomic model with a electron density map
- [WeNMR – SpotOn](#) Identify Hot-Spots at protein-protein interfaces
- [WeNMR – AMPS-NMR](#) Web portal for the refinement of Nuclear Magnetic Resonance (NMR) structures of macromolecules

Climate research

- [OpenCoastS](#) On-demand operational coastal forecast service
- [ENES Data Space](#) Data science environment for climate data analysis on top of the EOSC Compute Platform

Environmental sciences

- [SeaDataNet WebOcean Data Analysis](#) Software package for the analysis, exploration and visualization of oceanographic and other environmental data with almost 100, 000 registrations
- [EMSO ERIC data services](#) Explore the Data Tools and download Data from EMSO Observatories
- [Disaster mitigation and agriculture COMCOT](#) – Tsunami wave propagation simulation portal

- GBIF Cloud Data Space Access to occurrence biodiversity data published by Portuguese providers, or by international providers for Portugal

Energy and physical sciences

- PROMINENCE Enabling HTC & HPC applications opportunistically across private, academic and public clouds
- LOFAR Science Processing Providing a service to generate science-ready LOFAR data, enabling discoveries in astronomy to happen faster and more easily

Social sciences and humanities

- OPERAS Metrics service Collect usage and alternative metrics for Open Access publications

Value for Scientific Communities

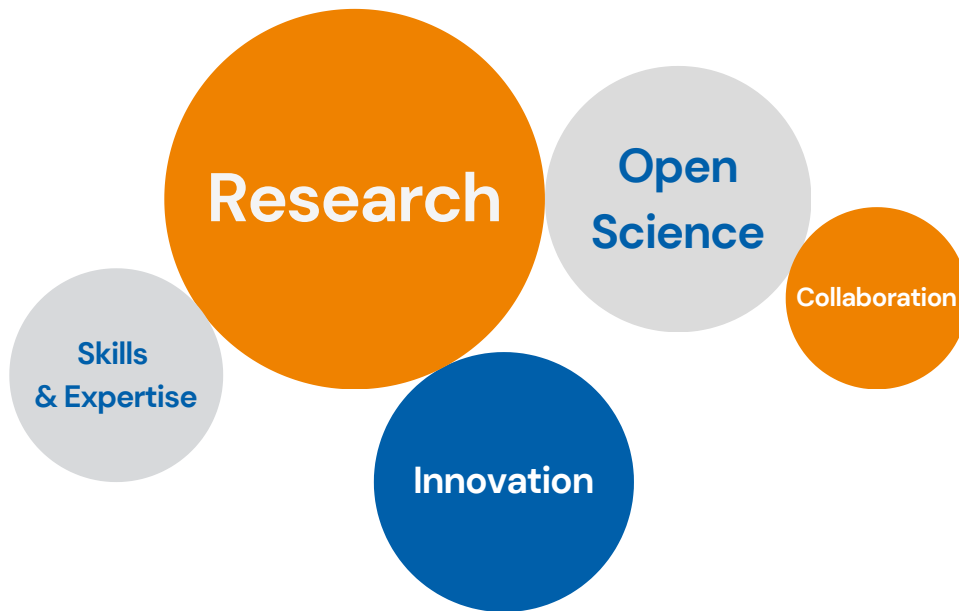
- Streamlined support and communication through the use case shepherds and with access to experts for consultancy to select the most suitable types and combination of compute and related platform/data services to realise a use case.
- Build partnership with national providers for long-term usage of their compute resources.
- Free-at-point-of-use access to national and relevant international capacity.
- Increase user base and reusability of their services
- Development of sustainability plans for operation beyond the lifetime of the project

Value for EOSC

- Contribution to the EOSC Data Commons through the setup and provisioning of 'Data Spaces'
- Contributions to multiple EOSC strategic objectives as defined by the PPP/SRIA.
- Development of guidelines and best practice approaches on how to set up and operate data spaces.



EGI-ACE Impact



1

Impact on Open Science

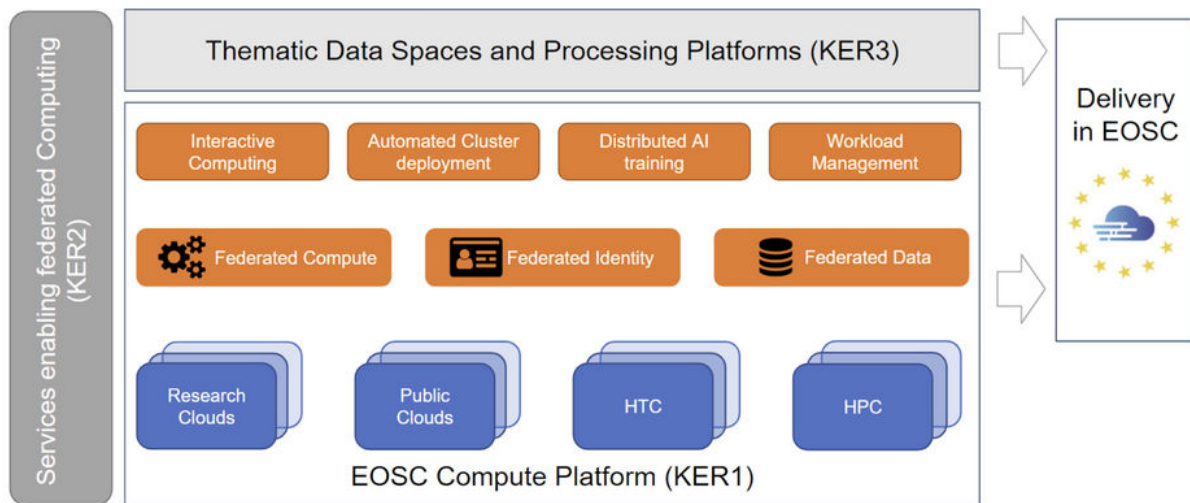
Scaled up EOSC

The overall objective of EGI-ACE was to “...deliver integrated computing, platforms, data spaces and tools as an integrated solution aligned with major European cloud federation projects and HPC initiatives.”. This objective was achieved by delivering **55 free-at-point-of-use compute services through the EOSC Portal**.

The services are grouped into the Compute Platform (19 services, forming KER1) and into Thematic Data Spaces and Processing Platforms (36 services,

forming KER3). The EOSC Compute Platform services provide a rich set of compute and big data management capabilities, offering solutions to interactive and batch data processing, Artificial Intelligence (model training and delivery), management of scalable clusters in clouds, and execution of massively parallel high throughput compute (HTC) applications. These functions are realised through a compute, data and authentication-authorisation federation layer, spanning over HTC, HPC, and cloud compute resources.

These 55 EGI-ACE services represent 9% of the total services (413) accessible in the EOSC Portal today.



Active EOSC

The EGI-ACE services are made available for users through the EOSC Portal and its Marketplace, as well as through an 'Open Call' which is a web form that is continuously available on the EGI-ACE website and attracts user communities to describe their access request alongside the usage case.

The EGI-ACE services attracted **189 access requests**, known as "orders", in the EOSC Marketplace and received **42 access requests** through the EGI-ACE Open Call. While the Marketplace orders were primarily coming from individual scientists or small research groups from academia and industry, the use case

submissions were dominantly from international research projects (typically Horizon 2020 and Horizon Europe) and from research infrastructures.

The **189 EOSC orders represent 35% of all the service access orders (538)** handled in EOSC during the 30 months of EGI-ACE. Moreover, 10 of the EGI-ACE services are among the 'most ordered top 20' services of EOSC, with the 'EGI Cloud Compute' and 'EGI Notebooks' services leading the overall list.

Through the different channels, the EGI-ACE services have been used by nearly 104,000 users during the 30-month project.

EGI-ACE service orders represented 35% of all EOSC service access orders.

10 EGI-ACE services in the 'most ordered top 20'.

Empowering Reproducible Open Science: A EGI-ACE Success Story

The EGI Notebooks and Replay services of EGI-ACE empower scientists and researchers to perform reproducible open science from a user-friendly user interface without having to deal with the complexities of the underlying infrastructure. EGI Notebooks delivers JupyterLab interactive data analysis environments where users can easily write notebook documents with text, code and visualisations that access datasets from the EGI infrastructure and from community repositories.

The computational analysis developed in EGI Notebooks can be shared for reproducibility via the EGI Replay service: users can publish their notebooks on open repositories (e.g. Zenodo, GitHub, or

GitLab) and generate shareable links to re-create the original computing environment and access the same datasets available on the EGI Replay service.

EGI-ACE supports instances of the Notebooks and Replay services for a long tail of science users (currently with more than 200 users registered) and for eight international communities and projects: Auger (astrophysics), EISCAT_3D (Atmospheric physics), PaNOSC (Photon and Neutron sciences), Reliance (Data Cubes, Text mining), Biomed (Bioinformatics), CESSDA (Social sciences), LETHE (Life sciences), and C-SCALE (Earth Observation).

More than 20,000 sessions have been started on the services during EGI-ACE, and more than 200 different environments were re-created in Replay since May 2022, since the opening of the service.

2

Impact on Research

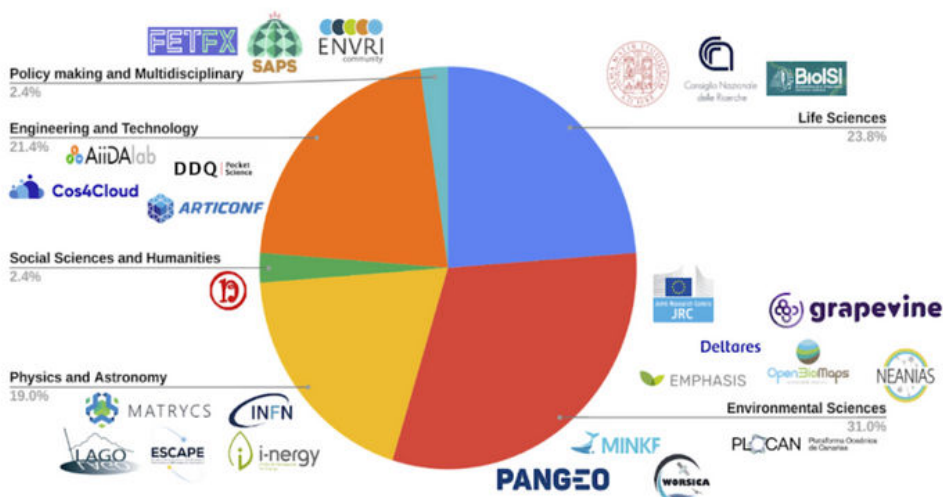
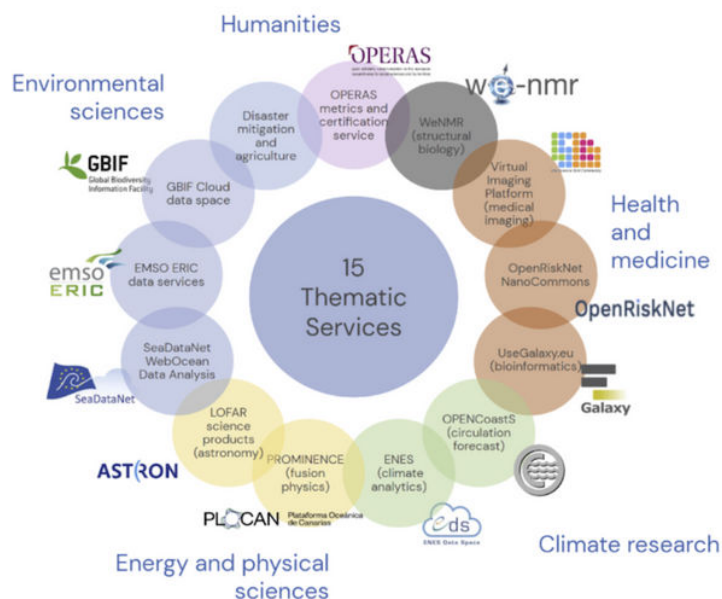
Fuelling Thematic Services

EGI-ACE worked intensively with scientific communities to set up and deliver Thematic Services in diverse disciplines to contribute to the implementation of the EU Data Strategy and to support the Green Deal, Health, Fundamental Research and Social Sciences and Humanities.

Fifteen scientific groups integrated application and scientific data with the EOSC Compute Platform, delivering 17 scalable data hosting and processing services for the EOSC Exchange layer.

The consortium collaborated with additional scientific groups through EOSC Marketplace and EGI-ACE Open Call, resulting in **36 new Thematic Services**. The overall EGI-ACE-related Thematic Service portfolio reached 53, with 35 registered in the EOSC Portal. These 53 services served nearly **104,000 users** over 30 months, resulting in **69 published papers** by providers and end users. In addition, EGI-ACE received **78 citations in scientific publications**.

17 Thematic Services from the consortium



36 Thematic Services from external communities

Empowering big Geoscience data analytics – PANGEOs success story

Pangeo is first and foremost a community promoting open, reproducible, and scalable Big Data geoscience. The community provides documentation, develops and maintains software, and deploys computing infrastructure to make scientific research and programming easier. Pangeo software

ecosystem involves open source tools such as xarray, iris, dask, jupyter, and many other packages. The Pangeo community in Europe sometimes referred to as “Pangeo@Europe”, aims to promote Pangeo in European time zones and increase the visibility of European contributors to the Pangeo ecosystem, independently from their institutions.

EGI-ACE helped Pangeo@Europe to have a shared deployment of its analytics tools in the cloud, where scientists and/or

technologists use and contribute to Pangeo software, can exchange know-how and provide feedback.

The EGI-ACE Pangeo deployment can speed up the learning process because geoscientists can learn with real-world examples how to access, analyse,

visualise and share data, making their research work FAIR and reusable. **During one year of operation the setup served over 100 scientific users from more than 10 countries, who consumed over 4 Million CPU-hour computing from EGI-ACE.**



Find more user success stories at the end of this report

3

Impact on Collaboration

Partnerships with peer-infrastructure outside Europe

During the proposal preparation phase, EGI-ACE started to setup 'Memoranda



of Understanding' (MOUs) with e-infrastructures outside Europe, based on a shared vision of supporting compute-intensive Open Science with federated compute infrastructures. Based on the MoUs, the project performed a joint set of actions with the infrastructures and achieved the following impact:

- **Computer Network Information Center (CNIC) of the Chinese Academy of Sciences (CAS):** established a federated cloud compute and storage testbed, federating data centre facilities from the EGI Federation and the CSTCloud infrastructure, using the EGI Federated Cloud approach. The EISCAT Research Infrastructure (Atmospheric Science) and the WeNMR community (Structural Biology) validated the setup.
- **Open Science Grid (OSG) in the United States:** EGI and OSG continued their long-standing collaboration under the EGI-ACE umbrella, focusing on increased computational power delivery for the science gateways of the WeNMR Structural Biology community. The consumption increased by ~30% from April 2020 due to the extra demand for COVID-related simulations. During the EGI-ACE project, the HADDOCK portal

of WeNMR was used by more than 30,000 researchers from over 130 countries.

- **Inter-University Institute for Data-Intensive Astronomy (IDIA), operator of ilifu, a node in the national data infrastructure of South Africa:** EGI and IDIA collaborated on adopting the EGI Federated Cloud approach at the ILIFU-UCT cluster, and to integrate radioastronomy-related computational workflows with the infrastructure. The goal of the experiment was to make South African workflows portable to European infrastructures, enabling

scientists in Europe to run those on infrastructures in a sustainable way on infrastructures that are local to them. The collaboration finished with successful cloud integration and application porting tests.



Partnership with ESFRI research infrastructures

Research infrastructures (RIs) represent one of the most important user segments for EGI. The EGI-ACE consortium includes RIs that reached mature status and joined the EGI Federation in 2021 (EMSO-ERIC representing the EMSO Members, CMCC representing the IS-ENES consortium, MARIS BV representing the SeaDataNet AISBL Members, and the EISCAT Scientific Association representing the EISCAT Council Members).

During its implementation phase, the project established collaboration with seven additional RIs on the ESFRI roadmap. Four of these partnerships focused on service delivery for the RI consortia, typically in the form of a cloud.




resources where the RI community host thematic portals, applications, and notebooks:

- METROFOOD-RI (Food and Nutrition Sciences)
- e-RIHS (Heritage Science)
- Instruct-ERIC (Structural Biology)
- EU-OPENSREEN-ERIC (Chemical Biology).

Within the additional three new partnerships with SLICES, EBRAINS and SoBigData, all belonging to ESFRI's 'Digital Infrastructures' segment, our joint work resulted in a shared work plan aimed at

reducing the environmental footprint of European compute infrastructures. The work plan is scheduled for implementation in Q4 of 2023, following the conclusion of EGI-ACE.



Partnership with industry

During EGI-ACE, the EGI community established a formal engagement structure towards industry: the EGI Digital Innovation Hub (EGI DIH). The DIH is a virtual space where companies and technical service providers meet to test solutions before investing. The EGI DIH is a new type of partnership where companies and technical service providers can collaborate, test solutions, and access advanced computing services to drive digitalisation and enhance productivity.

In the context of the DIH, the EGI-ACE project:

- Validated commercial cloud integration scenarios with T-Systems and CloudFerro. Both companies operate infrastructure-as-a-service clouds, and during the project, these were successfully integrated with the EGI Check-in authentication- authorisation service, establishing a foundation for federated access across academic and commercial clouds in EOSC.
- Established a business relationship with MathWorks, the creator of MATLAB, Simulink, and more than 90 add-on products that engineers and scientists apply to their exploration, design, and development efforts.

MathWorks and EGI partnered to offer a web-based version of MATLAB via EGI's Notebooks service. The setup enables MATLAB users to run MATLAB code on EGI infrastructure to analyse and visualize data, develop algorithms, and share research output with others via the Replay framework hosted by EGI.

- Attracted 15 SMEs as users to the EGI-ACE services with the Open Calls. Companies from multiple domains (Energy, Open Science, Agriculture, IT, Mobility, Logistics) and countries (Norway, Spain, Italy, Greece, Portugal, Germany, Finland, Malta, Ukraine) required support for using EGI services, dominantly EGI cloud compute and DataHub. The companies have been supported via a collaboration with the EOSC DIH.



Empowering Private Sector Innovation: Binare's Success Story



Binare Oy, a cybersecurity startup from Finland, partnered with EGI and obtained support from EGI DIH within the scope of EGI-ACE and EUHUBS4Data projects. Binare's goal with these projects was on addressing cybersecurity challenges in the Internet of Things (IoT) domain encompassing billions of devices and and a significant market value.

From 2021 to 2023, EGI provided extensive support to Binare, managing projects and offering computing & storage resources. With EGI's technical

expertise and network of computing centers across Europe, Binare could concentrate on their project's core challenges rather than on troubleshooting and resource management. As a result, Binare highly recommends EGI as a reliable partner for anyone involved in computing-based innovation.

Binare would like to express their gratitude to EGI for their exceptional support and looks forward to future collaborations with EGI and its partners to address current and future cybersecurity and IoT challenges.

4

Impact on Infrastructure Innovation

Cloud integration programme


The EGI Cloud federation, with 23 federated providers in 2021, was an important asset to the project. EGI-ACE has been running an ambitious extension programme to increase the capacity and geographical footprint of this federated infrastructure. The programme offered 3 integration scenarios:

- **Entry level:** joining the 'AAI federation'. This level was the target with the T-Systems and CloudFerro commercial providers and provides single sign-on across clouds and between EGI-ACE and EOSC services.
- **Medium level:** joining the application federation (through AppDB) and the data federation (through DataHub) besides the AAI federation. This level was targeted by national cloud providers of Moldova, Latvia, Armenia, and Georgia to enable the exchange of applications and scientific datasets between EOSC and those national clouds.
- **Full integration:** joining the monitoring, accounting and helpdesk systems of EGI besides the AAI, AppDB and DataHub integrations. This level was reached by OpenStack clouds from Greece (GRNET), Italy (INFN CNAF), Germany (GSI), Hungary (ELKH Cloud), Ireland (Walton Institute), China

(CSTCloud), and South Africa (ILIFU). Such level of integration makes it possible to involve these clouds in



international resource allocation and delivery scenarios with Service Level Agreements.



Open National Infrastructures – Resource allocation approach

EGI-ACE put in place an innovative resource allocation model that combined multiple funding streams to reimburse the cost of services:

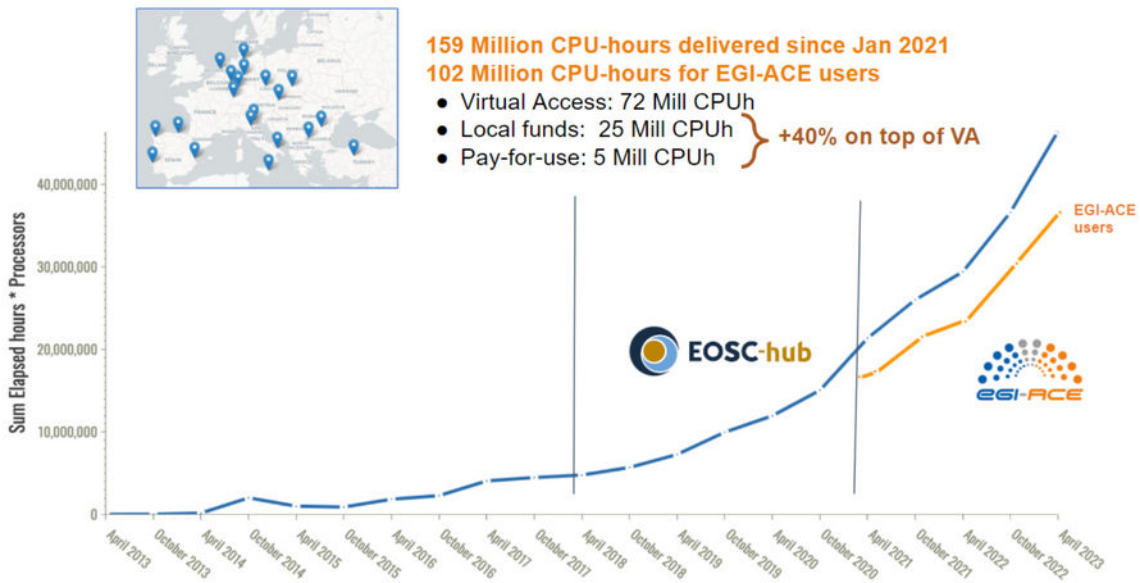
- 1. Virtual Access funding** went directly from the EC to 33 providers. This enabled access to any user community who met the baseline eligibility criteria. Approximately 5.5 Million EUR was distributed to 33 providers during 30 months in this way, enabling the delivery of **73 Million CPU-hours of computing capacity**.
- 2. National and institutional funding** from ministries and other types of funding agencies to providers to serve national or thematic user groups according to specific institutional/national missions and priorities. Over 100 providers participated in compute delivery this way, providing **24 Million CPU-hours compute capacity** in 30 months to the use cases. The intensive use of this funding instrument lowered the need for Virtual Access funds and significantly expanded the provider base and available capacities that the project could tap into.
- 3. Research community funding** brought by the users to providers to consume services on a “pay-for-use” basis. Such funding was available for

the users from their own EC-funded project grant, and EGI-ACE compute providers delivered 5 Million CPU-hours in this way to 9 communities* in 30 months.

To enforce the use of local resources and to minimise the stress on EC funding EGI-ACE prioritised the 2nd funding stream during compute allocations, allocating cloud(s) for new use cases based on their existing national/institutional commitment rather than available EC funding. This approach ultimately built stronger and more sustainable provider-user relationships that remain active even after the end of the project.

* BD4NRG, PITHIA-NRF, DECIDO, AI4PublicPolicy, StarWAI, LETHE, DIGITbrain, Exprivia ESA – Benchmark, PolicyCLOUD

The cloud providers delivered **102 million CPU-hour capacity** in total for the EGI-ACE users and user communities. Thanks to the innovative resource allocation process, **EOSC users gained +40% capacity** on top of the Virtual Access funding. The consumption follows an exponential curve, predicting further growth for the next years.



Overall EGI Cloud uptake during EGI-ACE

Compute Continuum and Data Spaces

The infrastructure layer of the EOSC Compute platform initially built on compute cloud, compute container and High Throughput Compute facilities. During EGI-ACE, this layer was extended with High Performance Computing (HPC) systems. 4 HPC facilities have been brought into EGI CESGA (Spain), IICT-BAS (Bulgaria), LIP/INCD (Portugal), and TUBITAK (Turkey), and defined how the HPC extension to the EOSC compute platform can be achieved with

- EOSC-compliant, federated access management via the EOSC Portal
- Availability and reliability monitoring for the federated HPC sites
- Integrated usage accounting across HPC, cloud and HTC sites
- Access to distributed, federated data for jobs from the HPC facilities
- Portable containers across cloud, HTC and HPC systems.

Furthermore, the project also made improvements to the EGI Container Compute service, reaching a system with

fully managed container clusters for users, simplifying the setup and execution of applications in the EGI cloud federation even further.

The HPC compute continuum capabilities have been validated in the project by four scientific pilots:

- Climate research use case from ENES (CMCC, Italy)
- High Energy Physics simulations for the High Luminosity run of the Large Hadron Collider (CERN)
- Cross-platform fusion workflows (UK Atomic Energy Authority)
- Photon and neutron science use case from the ELI-NP Research Infrastructure (IFIN-HH, Romania).

Service providers:



Use case communities:



The project built 5 scientific Data Spaces on top of the Compute Platform. Data Spaces host and integrate both data and online applications into a single unit, enabling the scalable analysis of big datasets on a compute continuum.




Success story: ENES Data Space for climate simulations

ENES is a European network of 50+ partners collaborating to accelerate climate progress, Earth system modelling, and understanding. They support global/regional climate simulations, adhere to Earth System Grid Federation data standards, and play a crucial role in IPCC assessments and EU policy climate projections.

ENES partnered with EGI to set the 'ENES Data Space' in EOSC, which delivers an open, scalable and cloud-enabled data

science environment for climate data analysis on top of the EOSC Compute Platform. The ENES Data Space consists of a JupyterLab instance, including a large set of pre-installed Python libraries and a ready-to-use Ophidia HPDA framework instance for running data manipulation, analysis and visualization tasks on EGI Cloud and HPC resources. The ENES Data Space hosts (open) data from the ESGF federated data archive on compute clouds to support researchers in realistic climate model analysis experiments.

Read the whole story: <https://www.egi.eu/case-study/enes/>



5

Impact on Skills and Expertise

26 webinars
24 training events
with over
1000 participants

Building human capacity with shepherds, training and co-design

The project has organised support for the received use cases through technical experts, also called 'shepherds' and Competence Centres.

One shepherd was assigned for each received use case and acted as the main point of contact, with the overall responsibility for the user's success. Shepherds performed a technical analysis of the use case, identified the most suitable services and providers, and brought domain experts into the support as needed. All the supporters formed a 'Competence Centre' for the use case – a

cross-functional group with the right skills and resources to reach successful implementation. The project established a network of 22 shepherds from 9 institutes: CMCC – Italy, UPC – Spain, INFN – Italy, CYFRONET – Poland, Fraunhofer SCAI – Germany, CSIC – Spain, IN2P3 – France, CNRS, France, EGI Foundation – the Netherlands.

The project also delivered a training programme to enhance the knowledge and skills of our users, our consortium members and external service providers. The programme included 26 webinars and 24 extended training events (remote in 2021–2022, face-to-face and hybrid from September 2022). The trainings have been attended by over 1,000 participants in 30 months.

Success story: Empowering Ukrainian researchers

EGI-ACE organised an online training workshop for researchers from Ukraine in December 2022. The goals of the event were to:

- Provide guidance on how to get access to the European resources of the EGI compute federation, with focus access to the GPU resources to run AI models.

- Provide general consultancy and assistance on distributed computing and data management for Researchers and Innovators in Ukraine.

The online event was attended by 13 Ukrainian researchers from different scientific disciplines and also by representatives of the Bogolyubov Institute for Theoretical Physics, the Ukrainian EGI member.



Improved service maturity through FitSM

FitSM is a lightweight standard for IT service management. It brings order and traceability with simple, practical support. It provides a standard conceptual and process model setting out realistic requirements and pragmatic implementation.

EGI-ACE delivered FitSM training to increase the maturity of service management within the consortium and among IT service providers of EOSC in general. During COVID-19, the EGI Foundation successfully developed and established an online variant of the 'foundation' level FitSM course, the most popular course of the 3-level FitSM programme.

Since January 2021, the project has delivered 18 FitSM training sessions: 13 foundation and five advanced levels.

Eleven courses were held as 'in-house training', dedicated to the staff of a certain organisation/project/community, while 2 were 'open courses', being open for any attendee.

The 14 events were attended by over 200 participants who gained 167 FitSM Foundation Certificates and 46 FitSM Advanced Certificates.



The **14** events were attended by over **200** participants who gained **167** FitSM Foundation Certificates and **46** FitSM Advanced Certificates.

User success stories

Here are examples of remarkable success stories that have emerged from EGI-ACE Open Call for Use Cases.

To pursue the goal of broadening the demand for the EOSC across diverse sectors, EGI-ACE has set up an open call

system that has enabled various communities worldwide to participate and apply for sponsored, free-at-the-point-of-use services. These include access to excellent infrastructure and platform services, personalised user support, and comprehensive training programmes.



Environmental & Earth sciences

Cos4Cloud

Cos4Cloud is a Horizon 2020 project empowering citizen science with deep Machine Learning and video recognition. The federated and distributed resources of EGI-ACE supported the deployment of the MOBIS (Mobile Observation Integration Service) back end, unlocking new possibilities for collaborative environmental research.

eHoney

EGI-ACE supported the eHoney project in their groundbreaking research on climate change's impact on biodiversity. We provided essential services for processing and downloading large-scale genomic data, empowering their computational analysis of honey-derived environmental DNA.

EMPHASIS

EGI-ACE also supported the EMPHASIS infrastructure, empowering researchers in plant phenotyping across Europe. To name a few examples of our support, EGI-ACE provided crucial resources including online

storage, PU computing, the EGI notebook, and authentication via EGI Check-in.

GRAPEVINE

GRAPEVINE aims to adopt a set of models to train a predictive model based on Machine Learning techniques to improve the prevention and control of grape diseases in the wine industry. The allocation of our HPC resources enabled the project to continue its weather simulation tasks when its initial resources were exhausted.

PANGEO

PANGEO is a worldwide community-centric, open, collaborative and scalable big geoscience data analytic ecosystem where scientists, developers, and research software engineers can contribute to developing software and infrastructure and work on Big Data Geoscience research problems. EGI-ACE helped PANGEO to create a common platform for its users and onboard European researchers on this Pangeo EOSC infrastructure.

OBSEA

OBSEA is a ground-based cable observatory located in Spain, in a protected fishing area. The project relies on EGI Cloud Compute, Check-in and Online Storage to deploy its infrastructure's data acquisition and visualisation services.

EMSO ERIC

EMSO is the European multidisciplinary research infrastructure coordinating seafloor and water column observatories. EGI has contributed to operating different EMSO ERIC services, including the data portal, which received thousands of visits from different countries and continues to do so thanks to the support of EGI-ACE.

EISCAT_3D

The EISCAT_3D radar system is set to revolutionise Earth's upper atmosphere and space environment monitoring. The EISCAT Data Access Portal is being developed, allowing researchers to browse, analyse, and download data while ensuring long-term archiving. EGI-ACE services, including the

Workload Manager, Cloud Compute, Online Storage, Check-In, and Perun, have been instrumental in supporting the development of this portal and the underlying distributed compute infrastructure.

ENES

ENES is European Network for Earth System modelling that accelerates progress in climate, Earth system modelling and understanding. The ENES Data Space established by EGI-ACE delivers an open, scalable and cloud-enabled data science environment for climate data analysis on top of the EOSC Compute Platform.

RELIANCE

RELIANCE provides a suite of innovative and interconnected services integrated into the EOSC and used by the EOSC scientific communities to support thematic and multidisciplinary research in Earth Science. The project has leveraged and integrated various EGI services with ROHub, the RELIANCE platform for managing RO-crate Research Objects, including EGI Check-in, Notebooks, Replay and DataHub.



Health & Medicine

BioISI

PypKa is a Poisson-Boltzmann-based pKa predictor for proteins using 3D structures as input. Considering that the performance of the PypKa cloud service is scaling almost linearly concerning the number of CPU cores, the PypKacloud server was deployed on the cloud resources of the EGI cloud infrastructure, and has relied on several EGI-ACE services.

CNR Covid use case

This success story deals with antiviral and virucidal activities against SARS-CoV-2 and how researchers can improve the development of antiviral agents by gaining insight into the molecular mechanism of spike protein inactivation after UV-C exposure. EGI provided 'free at the point of use' computing resources and consultancy services to implement the analysis.

WeNMR

WeNMR is a global Virtual Research Community focused on structural biology and life sciences, supported by EGI. It offers researchers worldwide a platform to gain insight into biological macromolecules' structural and dynamic properties at atomic details, which is critical for understanding most cellular processes. Now available on the EOSC Marketplace, the WeNMR services have continuously evolved to simplify access, increase throughput, and ensure access to sufficient computational resources.

BIOMED

Biomed is an international Virtual Organisation that supports the Life Sciences community, particularly in medical image analysis, bioinformatics, and drug discovery.

Its flagship science gateway, the Virtual Imaging Platform (VIP), offers researchers access to multiple applications and vast storage and computing resources. From the EGI Infrastructure, 14 HTC and cloud resource providers are supporting the biomed community by allocating computing and storage resources.

NBIS

NBIS is a distributed national research infrastructure in Sweden offering bioinformatics support in various projects, depending on the user's needs. EGI provides NBIS with Cloud Computing and Online Storage services needed to develop portals for Life Science researchers worldwide.



Engineering

DataCloud

DataCloud is a Horizon 2020 project that develops methods for the entire lifecycle of Big Data pipelines on diverse infrastructures for efficient processing and monitoring. The project used EGI Cloud Compute and Online Storage to distribute the computational task to a scalable compute platform and to store intermediate results from the user jobs, and VMOps Dashboard and the Infrastructure Manager to configure their resources and perform tests. The EGI Check-In was used for authorised access and EGI Applications Database to configure and deploy underlying services.

I-Energy

The I-ENERGY project aims to develop and demonstrate innovative AI-as-a-Service Energy Analytics Applications and Digital Twins services through 9 pilots. With EGI's support, the I-ENERGY project successfully

developed 16 AI services, which have the potential to change the energy industry by improving efficiency and reducing costs.

MATRYCS

With buildings accounting for nearly 40% of EU energy consumption, the building sector should play a key role in effective climate policy. MATRYCS is an EU-funded project aiming to elevate building energy management and boost efficiency through improved data processing, analysis and aggregation. The three layers of the MATRYCS Reference Architecture take advantage of the capabilities provided by the EGI Federated Cloud Infrastructure. Solutions offered by EGI played a crucial role in advancing the development of the data-driven reference architecture for scalable big data management and analysis of energy-efficient smart buildings.



Physics

IceCube

IceCube is the world's premier facility to detect neutrinos with energies above approximately 10 GeV located at the geographic South Pole. IceCube is a remarkably versatile instrument addressing multiple disciplines, including astrophysics, particle physics, and geophysical sciences. With the recent upgrade of the IceCube detectors, the computing demand for running and testing IceCube piloting activities will increase significantly, which will cause an even greater demand for GPUs.

IceCube relies on the EGI computing and storage resources such as the EGI Cloud Compute and the EGI Online Storage to distribute the computations and to extend the initial pool of GPU resources for running and testing IceCube piloting activities. User registration and authentication mechanisms are connected to EGI Check-In and finally, IceCube uses EGI's technical support to access and use our accelerated computing services smoothly.



Social Sciences and Humanities

OPERAS

OPERAS is the European Research Infrastructure for the development of open scholarly communication in the social sciences and humanities. It provides the research community with the missing brick it needs to find, access, create, edit, disseminate and easily and efficiently validate social sciences and humanities outputs across Europe.

Besides Check-in and Cloud Compute, OPERAS uses the cloud-based EGI Online Storage to distribute the computational task to a scalable compute platform and to store intermediate results from the user jobs. Additional consultancy and technical support enable smooth integration of the our services.

Find all our user success stories online



Advanced Computing for EOSC

We want to extend our heartfelt appreciation to all the project partners, service providers, and users who participated in our calls for use cases. Your expertise, contributions, and active engagement have driven innovation and shaped the future of our services and EOSC. Together, we have created a thriving ecosystem that unlocks new possibilities and advances research. Thank you for being an integral part of our journey.

Scan this QR code to see all our project partners.



EGI-ACE has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101017567.

EGI-ACE

[www.egi.eu/project/egi-ace/
contact@egi.eu](http://www.egi.eu/project/egi-ace/contact@egi.eu)



egi-ace-po@mailman.egi.eu



twitter.com/egi_infra



[linkedin.com/company/egi-foundation](https://www.linkedin.com/company/egi-foundation)