SPECTRUM

D5.2 Interoperable access policies: analysis and recommendations

Status: UNDER EC REVIEW

Dissemination Level: Public

Funded by
the European Union

Disclaimer: Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. SPECTRUM is funded by the European Union - Grant Agreement Number 101131550 - www.spectrumproject.eu

Abstract

Key Words	Infrastructures, research infrastructures, access policies, IT security, authentication, federation, user agreements

This deliverable documents the analysis of access policies of 20 European e-Infrastructures in the areas of High-Performance Computing, High-Throughput Computing, Cloud Computing, and Data, which are open for use by European scientists. It compares and contrasts them with the requirements of European research communities with a focus on the High-Energy Physics and Radio Astronomy fields and the use cases described in Spectrum Deliverable D5.1 (Representative use cases: analysis and alignment) and provides recommendations for the future evolution of the access policies to meet the current and future needs of European researchers.

Document Description			
D5.2 Interoperable access policies: analysis and recommendations			
Work Package Number 5			
Document Type	Report		
Document status	Under EC Review	**Version**	1.0
Dissemination Level	Public		
Copyright status	CC BY This material by Contributing Parties of the SPECTRUM Consortium is licensed under Creative Commons Attribution 4.0 International License.		
Lead Partner	FZJ		
Document link	https://documents.egi.eu/document/4072		
Digital Object Identifier	https://zenodo.org/records/15647609		
Author(s)	• Hans-Christian Hoppe (FZJ) • Luis Cifuentes (FZJ) • Xavier Salazar (EGI)		
Contributing Authors	• Jeff Wagg (CNRS) • Tommaso Boccali (INFN) • Kristen Lutz (SURF) • John Swinbank (ASTRON)		
Reviewers	• Raymond Oonk (SURF) • Fabio Affinito (CINECA)		
Moderated by	• Patricia Ruiz (EGI)		
Approved by	• Sergio Andreozzi (EGI) - on behalf of AMB		

Revision History			
Version	**Date**	**Description**	**Contributors**
V0.1	07/01/2025	First draft	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ)
V0.2	21/03/2025	Second Draft	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ) Xavier Salazar (EGI)
v0.3	31/03/2025	Internal Review	Raymond Oonk (SURF) Fabio Affinito (CINECA)
V0.4	23/05/2025	Rearranging content, adding appendices, Extended gap analysis and recommendations	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ) Xavier Salazar (EGI), Kristen Lutz (SURF) John Swinbank (ASTRON) Tommaso Boccali (INFN) Jeff Wagg (CNRS)
V0.5	06/06/2025	Internal Review	Raymond Oonk (SURF)
V0.6	11/06/2025	Version incorporating the suggestions of the internal Deliverable review.	Luis Cifuentes (FZJ) Hans-Christian (FZJ)
V0.7	13/06/2025	AMB Approval	Sergio Andreozzi (EGI)
V1.0	13/06/2025	**Final**	

Terminology / Acronyms

Terminology / Acronym	Definition
AAA	Authentication, Authorization, and Accounting
AAI	Authentication and Authorization Infrastructure
ACL	Access Control List
AI	Artificial Intelligence
AMB	Activity Management Board
AMD	Advanced Micro Devices
API	Application Programming Interface
ARM	Acorn RISC Machine
ATLAS	A Toroidal LHC ApparatuS
BLAS	Basic Linear Algebra Subprograms
BSC	Barcelona Supercomputing Center
CEA	Commissariat à l'énergie atomique et aux énergies alternatives
CEP	Central Processing
CERN	European Organization for Nuclear Research
CLI	Command-Line Interface
CMS	Compact Muon Solenoid
CoP	Community of Practice
CPUs	Central Processing Units
CSCS	Centro Svizzero di Calcolo Scientifico
DoA	Description of Action
DOI	Digital Object Identifier
EAB	External Advisory Board
EByte	1018 Bytes
EFlop/s	1018 Floating point operations per second
EFP	European Federation Platform as contracted by the EuroHPC JU
EGI	European Grid Infrastructure
EOSC	European Open Science Cloud

EPCC	Edinburgh Parallel Computing Centre
EPYC	Efficient Performance Yield Core (AMD Processor)
EPCC	Edinburgh Parallel Computing Centre
ERUM	German Acronym for “Research into the Universe and Matter”
EU	European Union
FLOP	Fast Fourier Transform
FZJ	Forschungszentrum Jülich
GA	General Assembly
Gbps	10^9 bits per second
GByte	109 Bytes
GCS	Gauss Centre for Supercomputing
GDPR	General Data Protection Regulation
GENCI	Grand Équipement National de Calcul Intensif
GFlop/s	109 Floating point operations per second
GPUs	Graphics Processing Units
GUI	Graphical User Interface
HDF5	Hierarchical Data Format version 5
HEP	High Energy Physics
HLRS	Höchstleistungsrechenzentrum Stuttgart
HPC	High Performance Computing
HTC	High-Throughput Computing
HTTPS	Hypertext Transfer Protocol Secure
HW	Hardware
I/O	Input/Output
ICSC	Centro Nazionale di Ricerca in HPC
ICTS	Singular Scientific and Technical Infrastructure
JSON	JavaScript Object Notation
JSC	Jülich Supercomputing Centre
JU	Joint Undertaking
KER	Key Exploitable Result

KIT	Karlsruhe Institute of Technology
KPI	Key Performance Indicator
LDAP	Lightweight Directory Access Protocol
LHC	Large Hadron Collider
LRZ	Leibniz-Rechenzentrum
LOFAR	Low Frequency Array
MAAC	Minho Advanced Computing Center
MFA	Multi-Factor Authentication
ML	Machine Learning
MPI	Message Passing Interface
MS	Milestone
NAS	Network Attached Storage
NFS	Network File System
NFDI	Nationale Forschungsdaten Infrastruktur
NHR	Nationales Hochleistungsrechnen Allianz
NIKHEF	National Institute for Subatomic Physics
NL	Netherlands
NVIDIA	NVIDIA Corporation
NVMe	Non-Volatile Memory Express
OpenMP	Open Multi-Processing
OS	Operating Systems
PByte	1015 Bytes
PFlop/s	1015 Floating point operations per second
PO	Project Objective
POSIX	Portable Operating System Interface
PUNCH4NFDI	Particles, Universe, Nuclei and Hadrons for the NFDI
QC	Quantum Computing
QoS	Quality of Service
RA	Radio Astronomy
RES	Red Española de Supercomputación

REST	Representational State Transfer
RI	Research Infrastructure
R/W/M	Read/Write/Modify
SCP	Secure Copy Protocol
SKA	Square Kilometer Array
SKA-IAM	SKA Identity and Access Manager
SKAO	Square Kilometer Array Observatory
SLA	Service Level Agreement
SME	Small and Medium-sized Enterprise
SRIDA	Strategic Research, Innovation and Deployment Agenda
SRCNet	SKA Regional Center Network
SSD	Solid-State Drive
SSH	Secure Shell
SSO	Single Sign-On
SW	Software
TByte	1012 Bytes
TFlop/s	1012 Floating point operations per second
UK	United Kingdom
VPN	Virtual Private Network
WAN	Wide Area Network
WG	Working Group
WLCG	Worldwide LHC Computing Grid
WP	Work Package
WPL	Work Package Leader

Table of Contents

List of Figures

List of Tables

Executive summary

This document contains the results of a detailed study on access policies in force or proposed for 20 European e-Infrastructures, which serve European research communities and data-intensive use cases with significant computing requirements. It defines an analysis template covering the different aspects of access policies, briefly describes the rationale for selecting the specific e-Infrastructures studied (which are either deployed and operational or planned to become so within the next 1.5 years), characterises these infrastructures, and documents the full analysis results in an Annex. Based on the study of 14 use cases in Spectrum Deliverable D5.1 (Representative use cases: analysis and alignment) and material collected via the Spectrum Community of Practice, the studied access policies are compared to research community requirements, gaps are identified and recommendations formulated for the future evolution and improvement of access policies to meet the scientific user needs.

1. Introduction

This document constitutes Deliverable D5.2 (Interoperable access policies: analysis and recommendations) of the SPECTRUM project. It contains an analysis of access policies and federation methods for advanced compute and data resources in European e-Infrastructures and compares these policies and methods with the needs of European research communities mainly in the high-energy physics (HEP) and radio-astronomy (RA) areas, which are detailed in the companion SPECTRUM Deliverable D5.1 (Representative use cases: analysis and alignment). Based on this "gap analysis", this document provides recommendations for improving the interaction between scientific end users and their applications/workflows and compute and data infrastructures, emphasising ease-of-use for scientists, interoperability, and security. These recommendations cover mid-term (up to three years) and long-term (up to seven years) periods. They are aligned with the recommendations provided in the companion SPECTRUM Deliverable D5.3 (Landscape of RIs: Technologies, Services, and Gaps).

A detailed study covering *all* planned or operational e-Infrastructures in Europe was not feasible within the time and effort provided by Spectrum – a selection amongst the plethora of such infrastructures had to be made; the guiding principles were (i) coverage of the key infrastructures for the HEP and RA scientific areas, (ii) focus on significant size/user-base infrastructures in these fields and European science at large, and (iii) representativity for the larger set of e-Infrastructures for high-performance, high-throughput or large-scale AI computation or large-scale data used by European open science. For the latter, suggestions from the SPECTRUM community of practice (CoP) were also considered.

Likewise, the number of scientific use cases in Europe which rely on compute/data e-Infrastructures (or plan to do so) is enormous; for our analysis, the 14 use cases discussed in the SPECTRUM Deliverable D5.1 have taken precedence, with additional requirements added from the results of the Spectrum Community of practice data gathering activities.

The analysis of each e-Infrastructure follows a common template presented in detail in Section 2. Section 3 presents the list of e-Infrastructures studied, along with an analysis of each e-Infrastructure according to the template detailed in Annex 1. Section 4 presents a concise summary of that analysis, references the list of scientific use cases covered by D5.1, and discusses their requirements in comparison to the access policies and methods available or planned to become available in the short term (within a year). The most important outcome of this section is the list of identified gaps, and based on these, Section 5 presents recommendations for the extension or evolution of e-Infrastructures to better address the needs of scientific user communities and use cases.

2. Access Policies Analysis Template

This section outlines the template for analyzing access policies and the methods to be followed in Annex 1 (the actual analysis), Section 4 (summary of the analysis, use case requirements, and gaps), and Section 5 (recommendations for extension/evolution). The content is derived from information made available online through Web presences, papers, and service portals, as well as direct communication with e-infrastructure providers.

2.1. Obtaining Access

- Targeted users – who can obtain access
  - Including nationality/location, affiliation, occupation
- Process – how can access be obtained
  - Including writing of access requests/proposals, approval process/criteria, and time to approval
- Reporting requirements (for the end-users)
  - Including frequency and required contents of reports by infrastructure end-users

2.2. Access Tracks and Modalities

- Access tracks and variants
  - Reflecting different capabilities, quotas, and timescales
- Access modalities
  - Including batch vs. interactive access, login/ssh/scp vs. application programming interfaces (APIs), service interfaces, or (domain-specific) Web portals

2.3. Summary of Accessible Resources

This section is distilled from the more detailed description in SPECTRUM Deliverable D5.3 (Landscape of RIs: technologies, services, gaps).

- Compute resources
  - Number of nodes, CPU/GPU platforms, node performance
- Data resources
  - Storage abstraction (file system, object store, …), capacity
- Data transfer
  - Data transfer methods
  - Indication of possible data transfer bandwidths,
  - Accessible resources – kind and volume (summary)

2.4. Access Management and Security

- Identity management, authentication, authorisation, accounting
  - How do end-users specify & substantiate their identity?
  - What authentication and authorization methods are used by the e-Infrastructure?
  - What is the model for accounting resource use (per user, per project, …), and how are quotas/budgets handled?
- Security methods and processes
  - Such as the need for encryption in transfer and storage, as well as controlled access to the internet, etc.

2.5. Rules and Assurances

- End-user policies & rules
- Rules and support for fair use, security, and data protection
  - Including the quality of service guarantees offered

- Available end-user support
  - By the e-Infrastructure or by centers/nodes of the infrastructure

2.6. Monitoring, Evaluation, and Evolution

- KPIs monitored by the e-Infrastructure
  - To assess and ensure infrastructure performance
- Means of engaging with end-users
  - Proactive information (in case of problems) and other end-user engagements outside of regular end-user support
- Means of evaluating, improving, and/or evolving the e-Infrastructure
  - Based on recorded data, user feedback, or surveys, etc.

3. Selection of e-Infrastructures Studied

The SPECTRUM DoA specifies a minimum count of 15 e-Infrastructures to be studied in tasks T5.2 and T5.3. The total number of such infrastructures used by European scientists across the fields is much larger. Even if one restricts the analysis to the HEP and RA fields, a selection would have to be made, as such infrastructures exist at different levels, namely transnational/European, national, or regional.

In addition, the services provided by e-Infrastructures roughly fall into several categories: HPC/AI[1] computing (designed to run closely-coupled applications/workflows with a high degree of parallelism), high-throughput computing/HTC[2] which targets running a large number of tasks at the same time, which themselves have only a low degree of parallelism, and data infrastructures (which provide access to input data for scientific use and storage space for results data). A fourth category, which we label as "Cloud-like" infrastructures, is differentiated by the prevalent use of cloud-native, service-based interfaces (pioneered by cloud service providers) and the capability to transparently utilize resources at multiple geographic locations, potentially provided by several third parties.

Three key criteria for selection by T5.2 and T5.3 were that the respective e-Infrastructures (i) are designed for and in active use by European scientists (or such use is firmly planned within a year of writing this Deliverable), (ii) are operated by European organisations and hosted in Europe, and (iii) provide the required amount of information about access policies and technical capabilities and capacity.

This subsection presents the list of e-Infrastructures investigated within the four categories mentioned above. Unless stated otherwise, both T5.2 and T5.3 did analyse this list from their respective, orthogonal points of view (access policies vs. compute/data capabilities/services offered).

3.1. HPC-oriented e-Infrastructures

HPC-oriented infrastructures provide large compute capabilities, are designed for executing tightly coupled applications with high degree of parallelism (up to 1000s/10000s of server nodes with single/dual CPUs or 1-8 GPUs each) or workflows containing such with high efficiency, and thus rely on supercomputers of Peta- (10^{12} floating-point operations per second or Flop/s) to Exascale (10^{18} Flop/s) aggregated performance. Such systems rely on inter-node fabrics of the highest possible performance (both in terms of latency and bandwidth) to enable scaling highly parallel applications up to a desired performance.

- **EuroHPC Joint Undertaking/EuroHPC JU HPC infrastructure**
  - At the time of writing, the EuroHPC JU HPC infrastructure consisted of eight deployed and operational hosting sites, with four more planned to become operational in 2025 and 2026.

[1] This specifically refers to AI training, which requires very large degrees of parallelism and high-performance communication between nodes, similar to typical HPC workloads.

[2] "Batched" AI inference has similar workload characteristics, with inference data being passed through a potentially large and dynamic set of inference tasks running on a single GPU (or a part of it) or on a CPU. It often requires highly efficient support of small data types, though.

  - The EuroHPC JU distinguishes between mid-range systems (up to a few PetaFlops/s), Petascale systems (up to approximately 100 PetaFlops/s), and pre-Exascale/Exascale systems (above 100 PetaFlops/s).
  - The EuroHPC JU co-funds the costs required for system purchase, installation, and operation, typically covering 50% of the costs. The remaining costs are borne by the hosting sites and/or national or regional funding authorities. The capacity funded by such sources is available through other, mostly national organisations, for instance through GCS in Germany.
  - EuroHPC JU defines a core set of access policies and rules, which are discussed in detail in section 6.1.
  - This infrastructure offers five different access modalities (differentiated by the purpose of access and scale of resources provided): benchmark access, development access, regular access, extreme-scale access, and AI & data-intensive applications access.
  - T5.2 (and this Deliverable) studied these five access modalities, whereas T5.3 (and D5.3) investigated the set of current or short-term planned HPC/AI sites and their systems.
  - Until mid-March 2025, the EuroHPC JU awarded a total of thirteen “AI factories” to European consortia - these combine new AI-optimised processing capabilities with the provision of high-level services for AI end-users in science and industry; due to the timing, D5.2 and D5.3 can only mention this effort, yet not provide any details on access policies or specific technology and services provided.
  - Additionally, in Q4/2024, the EuroHPC JU awarded a contract for deploying a “federation platform” across their HPC/AI systems. Technical information has started to become available at the beginning of 2025; therefore, the analysis of this development in terms of access policies is of a preliminary nature.
  - Finally, the EuroHPC JU has also decided to fund eight deployments of Quantum Computing systems to European consortia - while initial information on the technical design of these systems is known, and some systems are nearing actual deployment, specific information on access policies is not available. These systems are co-located and integrated with HPC systems, which in turn are available under the five aforementioned access modalities; however, their specific access rules are not known at the time of writing.
  - In April 2025, the EuroHPC JU issued a call for expression of interest on ideas for establishing AI GigaFactories in the EU. These will be large-scale AI compute & data facilities primarily designed for developing, training and deploying large AI models and applications. The AI GigaFactories will achieve a size of approx. 100000 advanced AI processors, which is 4x the size of the AI factories. Their objective is to further strengthen EU research, commerce and industry, facilitate the creation of new AI solutions, and ultimately realise the vision of Europe as an “AI continent”. The required investments will be massive (3000-4000 MEUR total cost of ownership), and partnership with the private sector is planned.

    The call for proposed GigaFactories was not yet published at the time of writing; it was suggested that the EuroHPC JU will handle the evaluation and selection procedure.

- Gauss Center for Supercomputing/GCS in Germany
  - GCS is an association of the three largest German supercomputer centres (Höchstleistungsrechenzentrum Stuttgart/HLRS, Jülich Supercomputing Centre/JSC, and Leibniz Rechenzentrum Munich/LRZ.
  - GCS provides access to German scientists for a collection of their HPC systems, including parts of EuroHPC JU systems, across scientific disciplines.
  - GCS will provide access to its own contingent of the Jupiter Exascale-supercomputer at JSC; Jupiter's initial availability is planned for H2/2025.
- Nationales Hochleistungsrechnen/NHR Alliance in Germany
  - The NHR alliance is a German association of twelve “tier-1” HPC centres (Technical Universities of Aachen, Darmstadt, Dresden and Kaiserslautern, Universities of Berlin, Frankfurt/Main, Göttingen, Mainz, Nuremberg/Erlangen, Paderborn and Saarland, and the Karlsruhe Institute of Technology).
  - NHR provides access to mid-range HPC systems to German and European scientists, with systems designed to support a subset of scientific disciplines.

- Red Española de Supercomputación/RES in Spain[3]
    - The RES is a Singular Scientific and Technical Infrastructure (ICTS) distributed throughout Spain.
    - Currently, the RES consists of 14 nodes located in various research centers and universities in Spain, which offer their computing and data exploitation services through competitive calls issued by the RES.
- Grand Equipement National de Calcul Intensif/GENCI in France
    - GENCI operates three national-scale supercomputers (currently the Adastra, Jean Zay, and Joliot Curie systems).
    - GENCI, with partners in the Jules Verne consortium, will provide access to the second European Exascale-supercomputer Alice Recoque starting from 2026, in addition to the EuroHPC JU.
- Centro Svizzero di Calcolo Scientifico/CSCS in Switzerland
    - CSCS operates the Swiss National supercomputing resources (currently the Alps systems) and offers access to scientists.
    - CSCS also provides key computing services to the Swiss weather service and other research institutions in Switzerland.
- Edinburgh Parallel Computing Centre/EPCC in the UK
    - EPCC operates the largest UK HPC system (currently Archer 2) and makes it available for scientific and industrial use across disciplines.
- Centro Nazionale di Ricerca in HPC, Big Data e Quantum Computing/ICSC in Italy
    - This Foundation, born thanks to the “Next Generation EU” funding, aggregates Italian supercomputing facilities (including CINECA Tier-0 resources and INFN Tier-1), as well as the network and storage facilities.
    - ICSC allocates HPC and Cloud resources to Italian researchers and SMEs.

3.2. HTC-oriented e-Infrastructures

HTC-oriented infrastructures provide the capability to run a large number of independent application instances in parallel, each of which has at most a moderate degree of parallelism (each instance running on O(10) cores, usually on a single node, and ensembles or workflows consisting of up to O(1000) instances per computing site, with little or no communication between the instances). The focus is on the aggregated throughput in terms of the number of application instances run rather than on providing the highest aggregated performance for a highly-parallel, tightly-connected application. This leads to a different system architecture and configuration, as inter-node communication is less important than for HPC systems. Typical HTC systems focus more on single-core or single-node performance and utilize scaling out to high numbers of nodes to increase application throughput in a linear way. An important element is the provision of local “scratch” storage space (usually per node) to avoid overloading shared file systems; data is staged into and out of these scratch spaces at the beginning and end of jobs. Conceptually, HTC systems are close to typical Cloud systems.

- Worldwide LHC Computing Grid (WLCG) with HTC-oriented centres across the world and CERN in Switzerland as a hub
    - WLCG provides global digital resources for the storage, distribution, and analysis of the data generated by the Large Hadron Collider (LHC).
    - WLCG geographically distributes its workload over 170 compute centers worldwide, which in 2025 comprises about 1.4M CPU cores and 1.5 ExaBytes of storage.
    - The focus in this subsection is on the compute side of WLCG, which handles data processing and analysis tasks for the global high-energy physics community.
- National Institute for Subatomic Physics (NIKHEF) in the Netherlands
    - NIKHEF, together with SURF, is a representative example of a (geographically distributed) Tier-1 WLCG centre.
    - The Tier-1 facility, via the Dutch national call for Computing Time, also provides resources to other scientific domains (e.g., Radio Astronomy, Engineering, Life Sciences, and Climate).

[3] https://www.res.es/en/about-res/nodes.

- EGI HTC-Oriented e-Infrastructure
    - EGI federates HTC-oriented computing services for different scientific user communities, acting as a middle layer between scientists and the actual providers of compute and data resources.
- Square Kilometer Array (SKA) regional centers across Europe
    - The SKA regional center network (SRCNet) nodes will enable HTC-style data manipulation and analysis of radio-astronomy data.
    - The SRCNet nodes are in the final planning stages, with the first nodes scheduled for deployment in 2025.
- Low Frequency Array (LOFAR) Central Processing (CEP) infrastructure in the Netherlands at the University of Groningen
    - This provides the central signal processing facility for the LOFAR radio-interferometry instrument.
    - This signal processing facility is supported by a distributed and federated data lake, known as the LOFAR Long Term Archive (LTA), with sites in Poland (PSNC), Germany (FZJ), and the Netherlands (SURF). Additionally, higher-level processing of LTA products is also done at the compute infrastructure near/at these data sites.

3.3. Data-oriented e-Infrastructures

Data-oriented infrastructures provide mid/long-term storage for significant amounts of data and support the sharing of such data; they can make data from scientific instruments or observations accessible to scientists for further processing, or enable these to store results of computation or scientific analysis and make these available to the larger scientific community.

- Worldwide LHC Computing Grid (WLCG) with data-oriented centres across the world and CERN in Switzerland as a hub
    - WLCG provides federated, global digital resources for the storage, distribution, and analysis of the data generated by the Large Hadron Collider (LHC).
    - The focus in this subsection is on the data side of WLCG, which stores data (ca. 1.5 EByte in 2025) recorded by the four LHC experiments and makes it available for further processing or analysis by the worldwide high-energy physics community.

- SKA Regional Center Network (SRCNet) with federated, data-oriented centres across the world. A centralized hub for the SRCNet is not foreseen at this time.
    - The data nodes that make up the SRCNet will receive, preserve, and disseminate the products created by the SKA Science Data Processors located near the radio-telescope sites in South Africa and Australia.
    - The SRCNet is being built up in its first prototype version v0.1, to be tested in 2025. In this Deliverable, we focus on the data part of the SRCNet.

- EBRAINS Neuroscience e-Infrastructure
    - EBRAINS provides access to neuroscience data, computational models, and software tools for researchers, clinicians, scientists, and students.
    - It is based on a two-tier organisational structure with a central hub located in Brussels, Belgium, and currently eight national nodes (France, Germany, Greece, Italy, the Netherlands, Norway, Poland, Sweden).
    - EBRAINS is the result of a EU-funded project, which in itself is based on results from the "Human Brain" lighthouse initiative.

- Low Frequency Array/LOFAR long-term data archive in Germany, the Netherlands, and Poland
    - The LOFAR long-term archive (LTA) provides access to the complete set of radio-astronomy data produced by the LOFAR instrument.

- ErUM[4] Data Hub in Germany
    - This infrastructure is a central networking and transfer node for data related to exploring the universe and matter.
    - It supports eight communities from the field of physics (nuclear particle, ionizing radiation) and astronomy/astrophysics (observatories, astroparticles).

- PUNCH4NFDI in Germany
    - Supports particle, astrophysics, astroparticle, hadron, and nuclear physics communities.
    - Uses the NFDI[5] data infrastructure architecture.

- Copernicus data spaces
    - Copernicus makes observation data from a global network of Earth observation satellites (including the Sentinel missions) available to scientists and the general public.

3.4. Cloud-oriented e-Infrastructures

In contrast to the three categories of infrastructures described above, Cloud-oriented e-Infrastructures often provide a combination of compute (HTC and increasingly HPC) and data services. Their distinctive features are (i) the provision of general-purpose service-based interfaces, initially introduced and popularized by commercial Cloud vendors, and (ii) the use of resources that are geographically distributed and/or owned and controlled by different entities. Increasingly, HPC, HTC, and data e-Infrastructures are moving towards (also) providing Cloud-like services to extend their user base and improve ease-of-use.

For SPECTRUM D5.2 and D5.3, we decided to focus on non-commercial Cloud e-infrastructures and did not include commercial Cloud providers. While these are being used directly in specific fields of science and research, the HEP and RA sectors dominantly rely on non-commercial, “public” infrastructures; sometimes, such infrastructures (such as the EGI Federation) do themselves use commercial Clouds as back-ends, or involve commercial players for implementation, installation or operation of parts of the infrastructure..

- EGI Federation
    - This infrastructure provides data, Cloud, and HTC services to European scientific communities as a “one-stop shop”.
    - It offers a Cloud interface (Federated Cloud) that federates a set of diverse Cloud resource providers and offers these in a uniform manner.
- SURF Data Processing (Grid and Spider infrastructures) in the Netherlands
    - SURF Grid offers the HTC Grid platform services on top of (in-house) OpenStack cloud. The Grid service participates in the NL Tier-1 site for WLCG, but also caters to other scientific domains. It is optimised for solving large-scale, data-intensive computational problems and offers efficient storage of large amounts of data.
    - SURF Spider offers a platform mixing elements of traditional HPC and HTC. It is deployed on top of an in-house OpenStack cloud and can also easily be cloned for custom use. Spider is optimised for high-throughput computing and offers scalable processing of large datasets.
- European Open Science Cloud/EOSC Federation
    - This infrastructure follows a two-tier architecture, comprising a central EOSC EU node procured by the European Commission (EC) and a number of EOSC pilot nodes located in different countries and serving various thematic research communities. The deployment of these pilot nodes started in March 2025 and is currently in the build-up phase.
    - The EOSC Federation comprises multiple EOSC Nodes that collaborate to share and manage scientific data, knowledge, and resources across various scientific disciplines and geographical areas.
- Simpl data federation platform
    - Simpl is a federated platform providing unified, safe, and secure data access and interoperability among European data spaces
    - It is being implemented under a commercial contract for the European Commission, and is positioned as the central middleware to federate the diverse set of existing and future common European data spaces

[4] ErUM is an acronym of “Erforschung von Universum und Materie” (research into the universe and matter).
[5] NFDI is an acronym of “Nationale Forschungsdateninfrastruktur” (national research data infrastructure).

4. Access Policy Summary, Requirements, and Gaps

This section begins with a concise summary of the detailed access policy analysis results in subsection 4.1; the details can be found in Annex 1. Following this, it recapitulates the use case requirements as identified by Spectrum Deliverable D5.1 (Representative use cases: analysis and alignment), which are relevant to the access policy topic in subsection 4.2. Priority is given to cross-cutting requirements shared by several of the studied use cases. Finally, Subsection 4.3 compares the access policies offered by the e-Infrastructures studied with the requirements, and the gaps found are discussed.

4.1. Access Policy Analysis Summary

4.1.1. Obtaining Access

From the analysis of access policies in section 6, two distinct methods of handling access requests and providing access emerge:

1. Allocating resources according to the result of reviewing a specific access proposal according to the rules of the resource provider or of the organisation responsible for a resource call-for-proposal, which focus on scientific merit and novelty, and often require a peer review. Access is granted for a specified period (usually one year), and extension proposals are typically supported by most providers. Otherwise, the mechanism is designed to disallow repetitive access proposals, and in some cases, it also restricts principal investigators (PIs) from submitting multiple different access proposals.
2. Based on the agreement between a research community and an e-Infrastructure, and in some cases involving a third party as a broker, access is provided to scientists based on their demonstrated membership in that research community. The allocation period can be an arbitrary amount of time, and entitlement lapses when the scientist leaves the research community (or "virtual organisation").

In the case of 2, agreements can include pledges of funding for the e-Infrastructure and SLAs defining the services guaranteed. In case 1, services are provided on a best-effort basis.

Some federation efforts (like the EuroHPC JU EFP) will provide the (mainly) software platform to facilitate access to all resources in the e-Infrastructure using the same identity and interfaces, yet do not handle allocations that are usable across all these resources.

A significant part of the e-Infrastructures covered (such as the HPC infrastructures) do require submission of reports detailing the work done and results achieved during use of the granted access privileges, and/or require the right to publish information on the use of their resources, derived from the usage reports or from dissemination material requested from their end-users.

4.1.2. Access Tracks and Modalities

Traditionally, HPC and HTC centers focus on batch processing, governed by a scheduling/orchestration system like Slurm. Here, end-users submit job scripts (containing the amount and type of compute resources required), which are run at a time determined by the scheduling system in unattended mode. This approach enables the efficient use of available HPC resources and is well-suited to large-scale parallel jobs. Interactive access is possible in two ways: via special head or login nodes (potentially competing with other users) or sets of (partial) nodes allocated via the batch system (with exclusive access).

4.1.3. Access Management and Security

Common mechanisms used by the studied e-Infrastructures include using password-protected certificates/keys (such as SSL key pairs) provided by the end-user, or relying on tokens or time-limited keys provided by a central single-sign-on service (which in turn uses userid/password or certificate/key authentication). Due to security concerns, centers have started to enforce multi-factor authentication (MFA), which involves an "interactive" authentication step involving a resource (such as a mobile phone) in the possession of the end user and the end user him/herself.

To avoid end-users having to go through MFA (or providing the password to a locked key/certificate) each time, systems accept repeated access from the same identity for a certain validity period after an initial, MFA-based authentication, or they provide an access token to be used for repeated accesses with a time-limited validity. The length of the validity period varies, yet is customarily limited to several hours, with CSCS allowing token-based access for a full day.

For the widely adopted authentication method using SSH key pairs, end-users can avoid repeatedly entering the private key password, for instance, by keeping an initial connection open for a full session or by automating the password entry using Linux mechanisms. Unfortunately, this creates security vulnerabilities, particularly for the second example, which would require the storage of cleartext passwords on file or in memory.

Source address filtering has been adopted by several HPC centers (EPCC and JSC being examples); in effect, end users have to provide a list of "known good" IP addresses from which legitimate access for them will occur.

4.1.4. Rules and Assurances

All e-Infrastructures studied require their end users to comply to acceptable use policies; while they differ in detail across the infrastructures, they contain a core of rules banning malicious behaviour (which would impact the operation and/or other end users), oversubscription of scarce resources (such as access or login nodes, or flooding a batch system with jobs), and use for commercial purposes (unless explicitly approved). End users are also generally prohibited from utilising identities and/or access privileges from other persons or projects.

Use or generation of personal information[6] and in particular special category data[7] does impose potentially significant obligations[8] and requires state-of-the-art technical and operational protection measures. These can impact both end-users and infrastructure operators. Most of the e-Infrastructures, therefore, rule out the use and creation of such data, or require specific data protection agreements to be signed and executed.

As mentioned in Section 4.1.1, it is customary for e-Infrastructures to request information about the use of their resources for publication purposes.

The e-Infrastructures studied in general give "best effort" assurances regarding availability and operation of their resources, and will not accept liability for direct or consequential damages caused, for instance, by non-availability of services, technical or operational faults, or activities of their personnel. Should end users require SLAs with hard guarantees, these will need to be negotiated - one example is the CSCS HPC e-Infrastructure, which also hosts the MeteoSwiss[9] weather predictions.

4.1.5. Monitoring, Evaluation, and Evolution

System and center operators all monitor and evaluate the performance of their local infrastructure in site-specific ways. This includes uptime, incidence of faults, utilization of compute, storage, and network resources, as well as usage statistics derived from end-user and project accounting. This data informs the operation of resources/sites (such as the scheduling of maintenance or the incremental addition of resources), and it forms the basis for justifying public funding received from funding authorities or projects/communities. The key observation is that such measures are handled in a site-specific (and often non-public) way for most of the studied e-Infrastructures.

[6] As defined by the GDPR regulation.

[7] Which includes health and medical data.

[8] Including, but not limited to strict confidentiality, prompt notification in case of data leaks, right of natural data owners to request data deletion.

[9] MeteoSwiss is the Federal Office of Meteorology and Climatology of Switzerland (https://www.meteoswiss.admin.ch/).

At the long-term/strategic level, an ongoing evaluation of KPIs relevant to the "end customers" and funding contributors at the infrastructure level is an important factor for improving the quality of service to end users and/or funding organisations; combined with predictions from the end user communities about future capacity and capability needs, these inform planning of e-Infrastructure extensions and refreshes, including purchases of large-scale compute and storage resources. Most of the studied e-Infrastructures do not make the evaluation processes, data, or the rationale for purchasing decisions public.

Exceptions include the EuroHPC JU HPC infrastructure and the WLCG; the former has publicly defined high-level KPIs and metrics, yet keeps the process used for evaluation of these and reaching decisions on adaptations/extension of their infrastructure (which involves the EuroHPC JU governing board with representatives from the governments providing funding) private and confidential. While the EuroHPC JU reaches out to its end-user community, the degree of influence it has can therefore not be conclusively established.

WLCG, on the other hand, provides data on infrastructure usage and KPIs in an open manner, and material on future requirements posed by HEP experiments and end-user communities is also publicly available.

4.2. Use Case Requirements Summary

4.2.1. Obtaining Access

Large experiments like LHC or SKA are operated for a long time, plan their evolution many years in advance, and depend on long-term commitments for compute/data resources. Obtaining and managing resource allocations over several years (potentially spanning the lifetime of a compute resource) is a critical factor for uninterrupted access by researchers, enabling the amortization of work on site-specific adaptations/optimizations. Resource providers can also benefit from the longer-term planning of experiments and map out the evolution of their compute/data resources according to future experiment needs. Besides HEP and RA, many other scientific fields have the same characteristics (optical telescopes, earth observation, neutron sources, advanced microscopy, etc.).

Federated e-Infrastructures promise the end user to be able to use a wide selection of available resources (either by manual selection of specific resources by the end user or automatic selection by the infrastructure based on availability and suitability for the specific service/activity requested by the end user) using the same identity, methods and interfaces. Such federation services are available in important European e-Infrastructures (e.g., WLCG), or will become available in the short-term (like for the EuroHPC JU HPC and AI infrastructure).

The next level of federation is the provision of resource allocations which are valid across the various nodes/sites of an e-Infrastructure – this would enable end-users to select the specific resources to be accessed at will, or enable automatic selection/brokering (to cope with spikes in demand (by using additional resources) or with resource outages). From a service provider perspective, such allocation across sites/nodes can be seen as problematic, for instance since it complicates resource capacity planning and limits control of the circle of users on their resources.

4.2.2. Access Tracks and Modalities

Scientific workflows in HEP and RA require support both for batch processing and interactive use of computing resources. Batch processing runs applications and workflow steps in unattended mode, and it is appropriate for simulation and data processing tasks running on a significant number of nodes with potentially significant runtime; interactive access is important for scientists to explore/analyse data, visualize data or results, and steer computation or modify models in near-real time.

Interactive platforms, such as Jupyter Notebooks and integrated data analysis environments, must be able to run efficiently in HPC and HTC centers and access data in high-performance storage tiers efficiently.

Domain scientists profit from having high-level interfaces to tasks that they commonly run on compute and data resources; these can take the form of services or APIs, potentially with added, Simple to use frontends (portals) which automatically query the user for all required input parameters and data locations and potentially perform consistency checks before starting a job or task.

The WLCG e-Infrastructure demonstrates the benefits of providing domain-specific interfaces and workflow systems (for the LHC experiments and the analysis of their data and simulations using it). Other infrastructures offer service- or API-based interfaces for general use (for example, CSCS).

4.2.3. Access Management and Security

The use cases do not express a general preference for any particular method of specifying and substantiating a user's identity, nor for other authentication/accounting data, where applicable. It is required that access to compute and data resources uses the same identity, and that this identity is valid across all sites of a federated e-Infrastructure.

Authentication and authorization mechanisms must support the unattended execution (except for the original launch operation) of workflows and their steps across the targeted e-Infrastructure. Such workflows can be complex and involve long-running steps, requiring authorization from a system/infrastructure to remain valid until the workflow execution is completed, without the need for human intervention.

Data security is important for ensuring long-term retention, availability, and protection against unauthorized modification. The use cases considered in Spectrum require protection against unauthorized access only to support the scientific process by enforcing a "blackout period", which gives teams that collected/generated the data the time to publish first. Other than that, guarding data confidentiality is not a requirement, since no personal data is stored/processed, and the open scientific data is intended to be publicly accessible.

4.2.4. Rules and Assurances

The use cases considered rely on a large set of complex software applications, libraries, and tools that are executed across the target e-Infrastructure. Different data and, in particular, compute systems often provide different hardware and software resources, which can introduce incompatibilities and limit the set of systems a given end-user application or workflow can run on. Adopting a common set of standard interfaces, protocols, and policies which streamline the deployment of workflows to multiple compute and data centers would help researchers to avoid spending undue effort in configuring and adapting their workflows. Independent of this, users will continue to need investing in optimising their workflows for modern systems.

4.2.5. Monitoring, Evaluation, and Evolution

Research communities working with long-term deployed instruments or experiments can provide valuable input to e-Infrastructures regarding the evolution of data and compute capabilities required by them in the mid- and long-term. They could conceivably also provide feedback to tweak the operation of currently deployed resources. Vice versa e-infrastructure can inform these communities about emerging technologies and resource utilization by current workflows. What would be required is establishing mid-/long-term collaborations between the research communities and the e-Infrastructures.

4.3. Gap Analysis

4.3.1. Obtaining Access

Gap #1: Long-term assured access and planning of resource allocations

A significant part of the studied infrastructures (in particular in the HPC space) does not regularly provide such assured long-term access (via grant-based access), and it is not clear whether their long-term evolution and planning are being substantially informed by the research community's needs.

The approach of awarding resource allocations as a result of reviewing specific access submissions and awarding short-term (one or two years maximum) allocations conflicts with the much longer time horizon of many research communities (in particular research infrastructures) and their instruments/experiments and the according need for mid/long-term (3-5 years timeframe) resource allocations, and can, depending on the number of researchers involved, easily lead to a flood of repetitive access submissions.

As stated in a recent ESFRI landscape study[10], there is a need for new and flexible allocation methods to support a wider set of users than those catered for by a stringent peer-review of access proposals as commonly used in the HPC area. At the same time, the need for resource providers (and their allocation committees) to review and supervise actual use of resources should be accommodated.

Gap #2: Resource allocations that are valid across the whole e-Infrastructure

Providing the architecture and mechanism for resource federation without providing federated resource allocations solves only half of the problem - it enables users to make use of resources across the infrastructure in theory, without actually giving them the capability (in terms of usage budgets/quotas) to do so. It is acknowledged that setting up "fungible" resource allocations can be difficult (what would be a fair conversion factor?) and intrudes to a larger degree into the autonomy of the e-Infrastructure sites than installing a federation software platform; yet, the gap still remains.

This capability is implemented in some e-Infrastructures (such as WLCG or the EGI federation), yet is lacking in others (including the future federated EuroHPC HPC and AI infrastructure).

Offering such interfaces can also be seen as a way to reduce the attack surface (since end-users can only invoke a limited set of services with well-defined arguments, rather than having a full shell CLI access) and thereby improve system security – end-users can only invoke the supported tasks, and the parameters/input data can be checked for problems; this is much harder to do once a user has obtained shell access.

4.3.2. Access Tracks and Modalities

Gap #3: Interactive access to significant-scale computing resources

While the HPC e-Infrastructures studied do provide interactive access to their head/login nodes or to compute nodes (using Slurm allocations), this does not fully address the requirements as stated in Section 4.2.2. Head/login nodes are shared resources with competition from other users, and allocating compute nodes using in particular a batch-oriented scheduler[11] is subject to scheduling and queuing decisions and can involve significant, unpredictable waiting times. Additionally, many centers do not permit internet access from the computer nodes, which complicates running tools like Jupyter Notebook. While it is possible to arrange reservations of such nodes at a prescribed date and time with the resource providers, this requires pre-planning well in advance.

For computational steering, such as in neuroscience use cases, the ability to "attach" to a running job and interact with it is not universally supported.

Gap #4: High-level end-user interfaces

Higher-level, general, or domain-specific interfaces for running often-used applications are not offered by most e-Infrastructures; this misses an opportunity to simplify the lives of domain scientists and also the potential upside in system security discussed in Section 4.2.2.

However, many of the studied e-Infrastructures rely on providing low-level APIs and CLI interfaces, which require domain scientists to acquire and maintain significant, sometimes even system-specific skills to make effective use of the e-Infrastructure's resources.

4.3.3. Access Management and Security

Gap #5: Federation of end-user identities (also across e-Infrastructures)

For many of the studied e-Infrastructures, particularly in the HPC space, current practice is to require the use of local user identities. The roll-out of the EuroHPC JU federation platform is expected to address this

[10] See https://landscape2024.esfri.eu.

[11] Even Cloud resources relying on more flexible orchestrators will cause waiting times depending on te degree of subscription of the available resources.

issue for HPC within the next year[12]. At a higher level we have the Authentication and Authorisation for Research and Collaboration (AARC) initiative to address the increased need for federated access and to develop and pilot cross-disciplinary authentication and authorization frameworks building on existing AAIs. The research IT community has made great strides towards the implementation of the AARC guidelines and blue print architecture. However, adoption and integration of these is still lagging, this is felt in particular by smaller research collaborations using smaller infrastructures/institutions, and needs continued investment and expertise support in the coming years. This lack of general adoption impacts researchers who regularly use different e-Infrastructures.

Gap #6: Unattended execution of long-running workflows

Having to "unlock" a certificate or key by typing a password or going through an MFA (multi-factor authentication) step requires an actual "human in the loop"; it complicates the unattended use of a resource, which is required, for instance, for executing automated workflows. Common resolutions include keeping keys unlocked for a specified period or issuing access tokens after MFA, which are valid for a certain number of hours. Both mechanisms can be integrated with automatic workflow execution, yet the time period within which any given workflow will be executed cannot be determined with sufficient certainty in the general case.

4.3.4. Rules and Assurances

Gap #7: Efficient provision of standard, low-level SW interfaces

While key software interfaces have been standardised for a long time (examples include programming languages, programming models like OpenMP and MPI, mathematical libraries such as BLAS, I/O and data format libraries such as FITS, ROOT, HDF5 or NetCDF, I/O interfaces like POSIX and Ceph), the software environments and stacks deployed by the e-Infrastructures[13] still differ considerably, partly caused by the choice between different implementations or different versions, partly caused by differences in hardware support and system configuration. Besides the danger of applications or workflows not running correctly, there is the spectre of code running correctly yet achieving sub-optimal performance.

Several R&D projects have proposed examples of standardized software stacks to be provided by the e-Infrastructures (see, for instance, DEEP-SEA[14] and EESSI[15]); however, such approaches have not (yet) been adopted to the required degree and do not necessarily contain full software ecosystem required by end users in the data intensive sciences. The CernVM File System[16] (CVMFS) provides a scalable, reliable, and low-maintenance software distribution service. Software packaging tools (such as Easybuild and Spack) for managing software component installation are in widespread use; yet, they do not fully address the underlying challenge of agreeing on a set of common interfaces and software components that work across different system architectures and configurations.

4.3.5. Monitoring, Evaluation, and Evolution

Gap #8: Organised feedback/improvement/planning loop

Long-term collaboration between research communities and HPC/HTC/data e-Infrastructures is happening in certain areas (like, for instance, between Destination Earth and EuroHPC JU, LHC and WLCG, LOFAR LTA and SURF/FZJ/PSNC, and SKA and SRCNet), and the e-Infrastructures solicit feedback from end-users, carefully monitor the operation of their resources, and improve their operation accordingly.

Outside of the examples noted above, organised efforts to collect and consolidate feedback and future resource needs with all relevant user communities, to improve the operation of an e-Infrastructure based on this feedback, and furthermore, collaborate with the relevant user communities in planning the mid- and long-term evolution of an e-Infrastructure are not commonly undertaken.

[12] EuroHPC also indicates that the federation is to be extended to the AI factories.

[13] While there are differences in SW components and libraries across different usage domains, this refers to the basic software stacks underpinning the domain-specific components.

[14] See https://deep-projects.eu/ for details.

[15] See https://www.eessi.io/ for details.

[16] See https://cernvm.cern.ch/fs/ for details.

5. Access Policy Recommendations

This section derives recommendations for evolving and improving e-Infrastructures to better meet the requirements of research communities today and in the mid-term future. It is based on the gap analysis in the previous section.

5.1. Obtaining Access

Recommendation #1: Adopt long-term, flexible resource allocation processes

As stated in **Gap #1** above, there is a need for e-Infrastructures to implement processes that allow researchers to obtain long-term access privileges and quotas in a flexible and scalable manner. This has to accommodate the needs of the research communities, the operational and security requirements of the resource providers, and the necessity of scientific agencies contributing funding to e-Infrastructures to gauge their return on investment. “Flexible” refers to the ability to handle researchers obtaining and relinquishing access privileges and quotas dynamically over time, and “scalable” relates to the potentially large number of researchers requesting such access.

Potential implementations

To meet the needs of larger research communities, e-Infrastructures should consider using membership/role-based authorization based on virtual/collaborative organizations, as discussed in Section 4.1.1, instead of requiring the handing in of specific access proposals. This would enable access based on membership of researchers in such organisations. The actual negotiation of access quotas and terms would be handled by the virtual/collaborative organisations, potentially including pledges of funding to resource providers.

Alternatively, allocations based on specific proposals should offer flexibility in duration and the number of end-users covered – a proposal by a PI for a research community should be allowed to include a potentially large number of researchers, if possible allow adding/deleting researchers during the term, and facilitate the extension of an allocation in accordance with the lifetime of the research community.

Recommendation #2: Enable e-Infrastructure-wide use of resource allocations and quotas

As stated in **Gap #2** above, e-Infrastructures should provide resource allocations and quotas to end-users which are valid across the nodes/systems of the infrastructure, enabling the end-user to make use of any resource of his/her choice, or enabling a transparent automated resource selection/brokering system to target the best suitable resources. This recommendation primarily targets compute resources.

Potential implementations

It is acknowledged that the realities of funding schemes (which can involve different sources) and the diversity of resources in some e-Infrastructures make implementing this recommendation a complex task. Automated resource selection schemes could involve mechanisms for balancing the actual use of resources according to the funding sources.

Since compute resources in an e-Infrastructure can vary substantially according to their capability and costs (operation, depreciation), a “currency conversion” might be required – as an example 1000 GPU hours on a NVIDIA A100 (9.7 TFlop/s for 64-bit arithmetic) could be converted to 285 GPU hours on a NVIDIA H100 (34 TFlop/s).

5.2. Access Tracks and Modalities

Recommendation #3: Extend scheduling/orchestration to support interactive compute use cases

As stated in **Gap #3** above, there is a need to improve support for interactive use cases on (batch) compute resources, while better balancing this with the traditional batch (queue-based) usage of HPC and many HTC resources.

The key recommendation here is for compute e-Infrastructures and resource providers to enable researchers to request exclusive, interactive use of fractional, single, or multiple nodes with assured, short-term availability; such allocations must be connected to high-speed storage.

Potential implementations

The existing resource management/scheduling tools (in the HPC world, based on e.g., Slurm) should be made more flexible and should also be considered to support modern orchestration systems, such as Kubernetes, on a subset of the HPC system, as demonstrated by the Meluxina and Karolina systems. A key requirement is the effective coupling of such "Cloud" nodes to the high-performance storage tiers at the HPC (supercomputing) centers. That this is possible has already been shown at HTC centers.

Recommendation #4: Introduce high-level general and domain-specific user interfaces

Gap #4 above discussed the benefits of introducing high-level user interfaces that better support research communities by matching the level of abstraction to the actual scientific tasks and activities. From this, a three-fold recommendation follows:

- Identify opportunities for introducing domain-specific, high-level interfaces (which can be APIs, services, or Web portals) for research communities that do not currently use them and agree on a set of such interfaces, and then work with the e-Infrastructures currently used to implement, deploy, and support them.
- Similarly, collaborate between end-user communities and e-Infrastructures to define, implement, and deploy general-purpose APIs, services, or Web portals for using compute and data infrastructures in a secure and intuitive way (learning, for instance, from initiatives such as FireCREST[17] or FTS[18]),
- Evaluate opportunities to integrate and unify domain-specific high-level interfaces to create cross-domain solutions with larger user bases and collaborate with the e-Infrastructure landscape in Europe to develop, maintain, and deploy these.

5.3. Access Management and Security

Recommendation #5: Introduce common AAI services across European e-Infrastructures

As discussed in Gap #5 above, introducing federation of identities and the systems handling authentication and authorization is a critical requirement for all e-Infrastructures, and there is considerable progress in the HPC space (which was lagging behind HTC and data e-Infrastructures in this respect) in the shape of the EuroHPC JU federation platform; the non-EuroHPC JU infrastructures without a federated AAI system should follow suit.

The longer-term recommendation here is to continue the adoption of a common guidelines and blueprint architecture (with AARC as a prime candidate) whilst stimulating more collaboration between e-Infrastructures using federated identity managements to further improve the interoperability between their AAI platform implementations, thereby enabling end users to easily switch between e-Infrastructures. In addition, broadening support for and uptake of identity federations (e.g., eduGAIN) would reduce the effort required by end users to acquire and maintain their identities across multiple providers.

Recommendation #6: Ensure Reliable and Unattended execution of long-running workflows

Gap #6 above discussed the need to enable reliable and unattended execution of (long-running) workflows and their steps across different nodes in an e-Infrastructure. It is recommended to adopt authentication methods or elements that do not require re-authentication through human interaction before a workflow has been completed.

[17] A REST interface for interacting with HPC systems and associated storage developed by CSCS (see https://www.cscs.ch/services/products/firecrest for details).

[18] A file transfer system used by CERN and EGI - see https://fts.web.cern.ch/fts/ for details.

Potential implementations

Methods that provide a "grace period" time interval of unattended access after initial authentication require reliable knowledge of the workflow execution time and the ability to set long enough "grace periods". One example of alternative approaches would be tokens that are only valid for a limited number of uses, matched to the requirements of a specific workflow. Since workflows could have been initiated on a different site within a federated e-Infrastructure, such mechanisms need to be managed by a central identity management/authentication instance.

An alternative approach can be to define higher-level interfaces to resources that restrict the actions an end-user can take and limit execution privileges to the minimum necessary.

5.4. Rules and Assurances

Recommendation #7: Provision efficient standard, low-level SW interfaces

As discussed in Gap #7 above, users who switch between different nodes of compute e-Infrastructures depend on being able to move their specific applications without recoding, and preferably without significant loss of achieved performance[19].

The short-term recommendation is for compute e-Infrastructures to ensure portability of applications (preferably binary, yet recompilation could also be sufficient) between their different nodes and systems; in the mid-term, it is recommended to deploy mechanisms that also enable portability between e-Infrastructures.

Potential implementations

One way to implement the recommendation is to continue efforts to create standard software stacks that facilitate the transfer of HPC and ML/AI applications between different systems, thereby avoiding the need for costly code and configuration changes. This can build on the results of, amongst others, the EESSI project.

Alternatively, containers can be used to package applications with the required SW stack elements and avoid conflicts with pre-installed stack elements. The HPC and HTC centers studied all support containers (mostly using Apptainer or Singularity), and the performance impact of a well-crafted container is considered negligible.

However, the increasing variety in CPU and accelerator architectures, along with the large number of combinations, creates a container management challenge. One can either create very large and complex containers with bespoke software support for all CPU/accelerator/network combinations, or manage a large number of small containers, one for each. A potential solution here can be layering of containers, such as proposed by the Sarus[20] system of CSCS.

5.5. Monitoring, Evaluation, and Evolution

Recommendation #8: Establish feedback/improvement/planning loops for research communities and e-Infrastructures

Gap #8 above discussed that organised efforts to close the loops between research communities and e-Infrastructures would be valuable in optimising operation and planning, and providing the best possible services.

The recommendation is to establish and maintain close, long-term collaboration between the European e-Infrastructures and the circle of research infrastructures using their resources.

[19] Relative to peak performance.
[20] See https://www.cscs.ch/services/products/sarus for details.

6. Annex 1 – Detailed Access Policy Analysis

This section presents the analysis results for the set of selected e-Infrastructures listed in Section 3, according to the template outlined in Section 2. A concise summary of the analysis results is provided in Section 4.1 above.

The content is derived from information made available online via the Web presences or service portals maintained by the e-Infrastructures covered, as well as published in papers and presentations, or obtained through direct communication with e-Infrastructure providers. Section 7 (Annex II) contains a list of links for each e-Infrastructure.

6.1. EuroHPC JU e-Infrastructure – General

The EuroHPC JU (Joint Undertaking) acts on behalf of the European Commission, with the main mission of establishing and sustaining a world-class ecosystem of HPC, AI, and Quantum Computing systems for European end-users from science and (for certain systems) industry. In this, the EuroHPC JU closes "hosting site contracts" with Hosting Entities that will operate these systems, provides co-funding for purchase and operational costs of these systems, and works with the hosting site in driving the system procurement. The EuroHPC JU also defines a set of access policies and ancillary processes, which are discussed in this subsection and further explored in Subsections 6.2 through 6.6 for more specific aspects.

It is important to stress that EuroHPC JU at the time of writing does approve access only to specific systems identified by the party requesting access; although a federation platform (see subsection 6.6) is being rolled out, access permissions and quotas are not "fungible" across EuroHPC JU sites.

In addition, access to systems co-funded by the EuroHPC JU is also available through alternative, mainly national paths- an example is access to the Jupiter system provided by GCS (see Subsection 6.2).

6.1.1. HPC Systems

Figure 1 shows the eight EuroHPC JU HPC systems as currently deployed, along with four systems being installed or prepared (with a 2026 time horizon for introduction into operation). The EuroHPC JU distinguishes between mid-range systems (up to a few PFlop/s), Petascale systems (up to approximately 100 PFlop/s), and pre-Exascale/Exascale systems (above 100 PFlop/s).

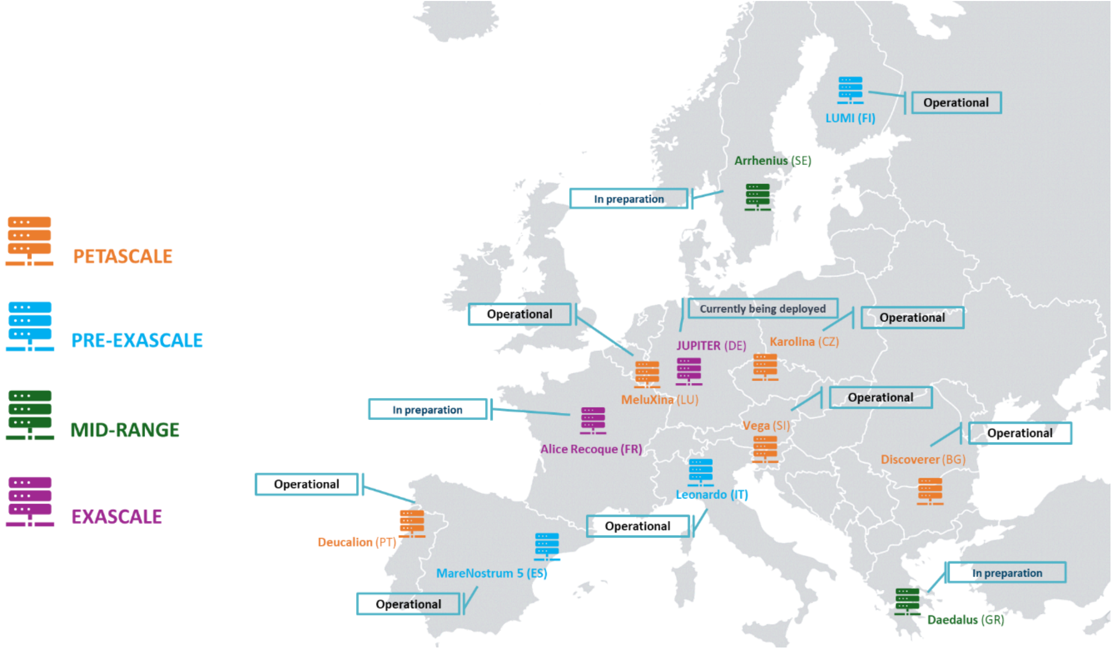

Figure 1: EuroHPC JU HPC systems

The EuroHPC JU Petascale systems are:

- Deucalion operated by the Minho Advanced Computing Center (MAAC) in Guimarães / Portugal with CPU nodes using the ARM and x86 architecture (1632 and 500 nodes), accelerated nodes combining x86 CPUs with NVIDIA A100 GPUs (33 nodes, 4 GPUs each), and an aggregated peak HPL performance of 7.5 PFlop/s.
- Discoverer located at Sofia Tech Park / Bulgaria with x86 CPU nodes (1182) and 4 NVIDIA DGX with 8 H200 GPUs each, and an aggregated peak HPL performance of 5.9 PFlop/s.
- Karolina operated by IT4Innovations in Ostrava / Czech Republic with x86 CPU nodes (720), accelerated nodes combining x86 CPUs with NVIDIA A100 GPUs (72 nodes, 8 GPUs each), and an aggregated peak HPL performance of 12.9 PFlop/s; in addition, a 36-node x86 SMP large-memory system for data analytics, and a "Cloud" module running virtual machines and OpenStack.
- Meluxina operated by LuxProvide in Luxembourg with x86 CPU nodes (573) and accelerated nodes combining x86 CPUs with NVIDIA A100 GPUs (200 nodes, 8 GPUs each) and an aggregated peak HPL performance of 18.9 PFlop/s; in addition, 20 large-memory CPU nodes and a "Cloud" module running virtual machines and OpenStack.
- Vega operated by the University of Maribor / Slovenia with x86 CPU nodes (960) and accelerated nodes combining x86 CPUs with NVIDIA A100 GPUs (60 nodes, 4 GPUs each) with an aggregated peak HPL performance of 10.1 PFlop/s.

The current pre-Exascale systems are:

- Leonardo located at CINECA in Bologna / Italy, with x86 CPU nodes (1536), accelerated nodes combining x86 CPUs with NVIDIA A100 GPUs (3456 nodes, 4 GPUs each) with an aggregated peak HPL performance of 315 PFlop/s.
- LUMI hosted by CSC IT Center for Science in Kajaani / Finland, with x86 CPU nodes (2048) and accelerated nodes combining x86 CPUs with AMD MI250x GPUs (2978 nodes, 4 GPUs each) with an aggregated peak HPL performance of 589 PFlop/s.
- MareNostrum 5 operated by BSC in Barcelona / Spain, with x86 CPU nodes (6408) and accelerated nodes combining x86 CPUs with NVIDIA H100 GPUs (1120 nodes, 4 GPUs each) with an aggregated peak HPL performance of 314 PFlop/s.

In addition, three additional systems will become operational in 2025 or 2026 at new hosting sites:

- JUPITER to be operated by Jülich Supercomputing Centre (JSC) / Germany; the system is Europe's first Exascale supercomputer, and its XPU partition is currently in the installation phase. It will become operational in 2025. JUPITER will have approx. 6000 ARM/GPU nodes (each with 4 NVIDIA GH200 XPUs) and approx. 1300 CPU nodes (plan is to use SiPearl Rhea CPUs).
- Alice Recoque to be operated by GENCI (Grand Equipement National de Calcul Intensif) at a CEA (Commissariat à l'énergie atomique et aux énergies alternatives) site close to Paris/France ; the system is the second European Exascale system and will be procured in 2025 and become operational in 2026.
- Arrhenius to be operated by Linköping University / Sweden; the system is a mid-range supercomputer currently in procurement and planned to become operational in 2025.
- Daedalus to be operated by GRNET (National Infrastructures for Research and Technology) in Athens/Greece; the system is a mid-range supercomputer, with installation and operation planned for 2025.

6.1.2. AI Factories

In addition to these systems, the EuroHPC JU will co-fund a total of (at the time of writing) thirteen "AI factories", each of which will include an AI-optimised supercomputer system, AI-focused programming/usage interfaces, support for end-users from science and industry, and training services. Deployment and entry into operation are expected to begin in 2025, with initial operational capability anticipated by the end of 2025, according to EuroHPC.

The approved AI factories are:

- LUMI AI Factory (hosted by CSC IT Center for Science in Finland).
- HammerHAI (hosted by High-Performance Computing Center Stuttgart (HLRS) in Germany).
- Pharos (hosted by the National Infrastructures for Research and Technology (GRNET) in Greece alongside Daedalus).
- IT4LIA (hosted by CINECA Consorzio Interuniversitario in Italy, alongside Leonardo).

- L-AI Factory (hosted by LuxProvide in Luxembourg alongside MeluXina).
- BSC AI Factory (hosted by Barcelona Supercomputing Center (BSC-CNS) in Spain, alongside MareNostrum 5).
- MIMER (hosted by the National Academic Infrastructure of Supercomputing (NAISS) in Sweden).
- AI: AT (hosted by the Technical University of Vienna in Austria).
- Brain++ (hosted at Sofia Tech Park in Bulgaria, alongside Discoverer).
- AI2F (hosted by GENCI in France, using Alice Recoque).
- JAIF (hosted by Jülich Supercomputing Centre in Germany, alongside JUPITER).
- PIAST (hosted by the Poznan Supercomputing and Networking Center (PSNC) in Poland).
- SLAIF (hosted by IZUM in Slovenia).

No detailed information on access policies for these AI factories is publicly available at the time of writing; here, it is assumed that the "AI and Data-Intensive Applications access policy by EuroHPC will serve as a base. Initial ideas communicated include

- Enabling initial, small allocations as "playground access" with a time-to-approval of two work days and a duration of up to three months
- Processing access requests for middle-level allocations (for instance, up to 50000 GPU hours) within four days
- Granting large-scale access (for instance, >= 50000 GPU hours) within 10 days and two cut-off dates per month

6.1.3. Quantum Computing Systems

At the time of writing, the EuroHPC JU had approved a total of nine Quantum Computers based on different Quantum Computing technologies, which are in various stages of procurement and deployment:

- HPCQS analog quantum simulator 1 (hosted by GENCI in France and to be coupled to Alice Recoque) using neutral atom technology
- HPCQS analog quantum simulator 1 (hosted by Jülich Supercomputing Centre in Germany and coupled with the JUWELS HPC system) using neutral atom technology
- LUMI-Q (hosted by IT4Innovations National Supercomputing Centre in the Czech Republic and integrated with Karolina) using supercomputing qubits in a star-shaped topology with one-to-all connectivity
- EuroQCS-France (hosted by GENCI in France and to be coupled to Alice Recoque) using a quantum-dot-based single photon source and a programmable quantum interferometer
- Euro-Q-Exa (hosted by Leibniz Supercomputing Centre in Germany) using superconducting, frequency-tunable qubits and couplers in a square-lattice topology
- EuroQCS-Italy (hosted by CINECA Consorzio Interuniversitario in Italy) using neutral atom qubits.
- EuroQCS-Poland (hosted by the Poznan Supercomputing and Networking Center (PSNC) in Poland) using trapped ions
- EuroQCS-Spain (hosted by Barcelona Supercomputing Center - BSC-CNS - in Spain), based on analog quantum computing technology
- EuroSSQ-HPC (hosted by SURF in the Netherlands) using semiconductor spin qubits

The two HPCQS systems are deployed and will enter operation shortly. They are coupled to the HPC systems of GENCI and JSC and are likely to utilize the existing EuroHPC and applicable national access policies for these two centers. We will assume here that the EuroHPC JU will utilize its existing access policy concept for the other Quantum Computers, including Quantum Computing-specific access requests and approval processes, and potentially create a bespoke quantum computing (QC) access policy. No details are publicly known at the time of writing.

6.1.4. General EuroHPC JU Access Policy Information

Table 1 presents the analysis of general aspects of access for the EuroHPC JU e-Infrastructure[21]; details on specific access tracks are given in subsections 6.1.5 through 6.1.8. In addition, the national parts of the EuroHPC infrastructures are also available via national access routes (in this document, we cover GCS for Germany, RES for Spain, and ICSC for Italy).

[21] Via the EuroHPC JU route, see https://eurohpc-ju.europa.eu/supercomputers/supercomputers-access-policy-and-faq_en.

Table 1: Access Policy Analysis – EuroHPC JU (General)

Access Policy Analysis – EuroHPC JU (General)	
Obtaining Access	
Targeted users	Scientists, public administrations, and industrial users (PIs and team members) working for organizations located in an EU member state or a third country associated with Horizon 2020 or Horizon Europe; 75% of the capacity is reserved for science, 20% for industry, and 5% for public administrations.
Process to obtain access	Submit access proposals to access calls, with subsequent peer-review and approval/rejection; evaluation criteria are excellence, innovation & impact, and quality and efficiency. Details vary per access track.
Reporting requirements	PI has to submit a final report within three months of the completion of an allocation, describing the results obtained, as well as qualitative feedback on the use of the resources to the EuroHPC peer review. Applicants must acknowledge EuroHPC JU in all publications that describe results obtained using EuroHPC resources.
Access Tracks and Modalities	
Access tracks	Currently supported tracks are benchmark access, development access, extreme scale access, regular access, and AI & data-intensive applications access; details of the last five tracks are discussed in subsections 6.1.5 through 6.1.8.
Access modalities	The most common technical access modalities are using login shells via SSH, protected by SSL certificates and increasingly multi-factor authentication (MFA), and SSL-based data transfer (scp et al.) with the same protections. Some EuroHPC sites provide additional access modalities, such as Jupyter notebooks. The EuroHPC JU Federation Platform (EFP), currently under development, will provide unified Access modalities, which include web-based access to applications, workflow execution, and Jupyter notebooks.
Summary of available resources	
Compute resources	At the time of writing, nine HPC supercomputers ranging in the single-digit to low-hundreds of PetaFlop/s performance range; the first European ExaFlop/s system is planned to become operational in 2025, as are three additional Petascale systems. In 2026, the second European ExaFlop/s supercomputer will follow. Additional AI-optimized compute capability (in thirteen AI factories) will become available in late 2025/2026, and a total of nine Quantum Computing simulators and systems will be deployed.
Data resources	High-performance parallel file systems (PFS) attached to the HPC systems, with on-site backup support - plus in some places, separate conventional storage systems and file systems. PFS holds pre-staged input data for HPC jobs, and output data created by these; in most centers, output data needs to be staged out eventually.
Data transfer resources	All centers support SSL-based data transfers like scp; some support additional, performance-oriented data transfer mechanisms such as UnicoreFTP, FTS, or Nodeum data mover. The EuroHPC JU HPC centers are connected via (mostly) NReN and GEANT-provisioned high-bandwidth WAN links.
Access management and security	
Identity management and	Currently, local identity management using vendor, home-cooked, or open source systems; all centers accept SSL certificates and keys (password-protected) for login and data access, with an ongoing trend to mandate the use of multi-factor authentication.

AAA	Some centers support community-specific identity management and AAA systems like Unicore or FENIX (based on MyAccessID) for certain projects. The EFP solution will create a federated identity management based on MyAccessID; from what is known, centres will have the option to request MFA, in which case EFP will issue short-term SSL certificates to enable automated (non-MFA) access for, amongst others, workflow execution.
Security methods and processes	Centers institute different rules for passwords and password updates; MFA is increasingly requested. From the Internet, a set of login nodes is accessible; other "worker" nodes are only visible from the login nodes and are made available through a batch scheduler (Slurm). Support for encrypted data storage is provided, yet at the time of writing, it is not enforced. Unencrypted data transfer (such as FTP or Telnet) is prevented; instead, users must rely on protected methods, like SCP. Direct access to "worker" nodes from the Internet is not permitted, and likewise, access to the Internet from worker nodes is restricted, in most cases requiring staging of data to storage followed by data transfer from a login node.
Rules and assurances	
End-user rules and policies	EuroHPC and the local hosting sites impose end-user policies mandating fair use of resources in general, outlawing criminal and malicious activities, prohibiting the processing and/or creation of personal data protected by GDPR, and disallowing the sharing of login credentials. Provision of resources follows a "best effort" paradigm – no hard guarantees are given as to resource availability, performance, and waiting times.
Fair use, security, data protection	Local hosting sites impose and track quotas for use of compute and data storage resources, aggregated to user or project identities. Batch queues with different priorities and/or available resource quotas (in space and time) support different end-user needs, and the batch scheduling system (Slurm) automates scheduling and allocation according to site-defined fairness rules. Regular end-users are not able to acquire root privileges. Data is protected by Linux and PFS mechanisms, including, in many cases, access control lists. Encryption of data at rest is not enforced; unencrypted data transfer mechanisms are blocked. The EuroHPC JU has recently started a Computer Emergency Response Team (CERT), which includes EuroHPC staff plus security experts from the Hosting Entities.
End-user support	Centers maintain groups of end-user support specialists who help with access and system issues; most centers can also provide assistance with system SW and a range of commonly used applications. End-user support is based on the promise of "best effort".
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Each hosting site monitors a significant number of indicators related to the use of its resources, details about the system behavior, and logs of excursions or errors. These are as of the time of writing, not shared with each other or with the public, and serve local analysis and according to improvements. The EuroHPC JU has access to a subset of that monitoring data and is using it in conjunction with monitoring data related to the approval process to ascertain efficient use of resources, track potential problems, and assess the quality of service: • Volume of resources offered vs. volume of resources requested • Number of applications vs. number of awarded projects • Share by country of the total number of awarded projects • Share by country of the total awarded resources • Share of requested resources per domain • Share of awarded resources per domain • Number of applications vs. number of awarded projects led by industry • Volume of resources requested vs volume of resources awarded to industry-led projects • Volume of resources awarded to SMEs

Engagement with end-users	EuroHPC JU has instituted a "User Forum" process to engage with all end-users, sharing EuroHPC updates, soliciting feedback and formulation of end-user requirements, and discuss/address infrastructure challenges. A coordination group made up of domain scientists oversees the User Forum, and Slack/Discourse channels are provided.
Evaluate, improve, and evolve infrastructure	EuroHPC pledges to take end-user feedback seriously and evolve its infrastructure accordingly.

6.1.5. EuroHPC JU e-Infrastructure – Benchmark Access

Table 2: Access Policy Analysis – EuroHPC JU (Benchmark Access)

Access Policy Analysis – EuroHPC JU (Benchmark Access)	
Obtaining Access	
Targeted users	End-users wanting to test/benchmark their applications on EuroHPC JU systems prior to applying for an Extreme Scale and/or Regular Access.
Process to obtain access	Submit an access proposal to a continuously open submission. Cut-off dates for review on the first day of every month, passing through two stages of evaluation (administrative check, technical assessment). The expected duration of review is up to 2 weeks, with access provided within three weeks of the cut-off date.
Summary of available resources	
Compute resources	Depending on the system requested, allocations can comprise 2000-3500 CPU node hours, and/or 200-600 GPU node hours. Access is valid for three months.

6.1.6. EuroHPC JU e-Infrastructure – Development Access

Table 3: Access Policy Analysis – EuroHPC JU (Development Access)

Access Policy Analysis – EuroHPC JU (Development Access)	
Obtaining Access	
Targeted users	Code development and targeted optimization activities
Process to obtain access	Submit an access proposal to a continuously open submission. Cut-off dates for review on the first day of every month, passing through two stages of evaluation (administrative check, technical assessment). The expected duration of review is up to 2 weeks, with access provided within three weeks of the cut-off date.
Summary of available resources	
Compute resources	Depending on the system requested, allocations can comprise 3500-4500 CPU hours, and/or 400-900 GPU hours. Access is valid for six to twelve months. Allocations cannot be changed or extended.

6.1.7. EuroHPC JU e-Infrastructure – Regular Access

Table 4: Access Policy Analysis – EuroHPC JU (Regular Access)

Access Policy Analysis – EuroHPC JU (Regular Access)	
Obtaining Access	
Targeted users	Research applications requesting large allocations. Each PI can only have up to one awarded Regular Access proposal at any time.
Process to obtain access	Submit an access proposal (two parts) to a continuously open submission. Cut-off dates for review occur twice a year, and have to pass through six stages of evaluation (administrative check, technical assessment, rapporteur reporting, access resource committee meeting, resource allocation panel, governance board approval). Access is to be provided within four months of the cut-off date. It is possible to submit a continuation proposal during the period of validity of the original proposal; this goes through the same review process and requires a progress report for the original allocation to be attached.
Summary of available resources	
Compute resources	Regular access calls can target each EuroHPC JU system. For the Petascale supercomputers, the minimum allocation is 60000 node hours for CPU nodes, and 25000 node hours for GPU nodes; this translates to 60000-120000 CPU and GPU hours. No upper limit is specified. For the pre-Exascale systems, minimum allocation is 60000 for CPU nodes, and between 20000 and 25000 GPU node hours; this translates to 120000 CPU hours and between 80000 adn 100000 GPU hours; upper limits are 120000-230000 CPU node hours (240000-460000 CPU hours) and 150000-220000 GPU node hours (600000-880000 GPU hours). Depending on the system requested, allocations can comprise 60000 CPU hours, and/or 20000-25000 GPU hours. Access is valid for twelve months and can, on special request, be extended by a maximum of an additional three months if the allocated resources cannot be consumed within the original allocation period. The resource allocation itself cannot be modified.

6.1.8. EuroHPC JU e-Infrastructure – Extreme-Scale Access

Table 5: Access Policy Analysis – EuroHPC JU (Extreme-Scale Access)

Access Policy Analysis – EuroHPC JU (Extreme-Scale Access)	
Obtaining Access	
Targeted users	Research applications requesting extremely large allocations in terms of compute time, data storage, and support resources. Each PI can only have up to one awarded Extreme Scale Access proposal at any time.
Process to obtain access	Submit an access proposal (two parts) to a continuously open submission. Cut-off dates for review occur twice a year, and have to pass through eight stages of evaluation (administrative check, technical assessment, scientific evaluation, response &rebuttal, rapporteur reporting, access resource committee meeting, resource allocation panel, governance board approval). Access will be provided within six months of the cut-off date. It is possible to submit a continuation proposal during the period of validity of the original proposal; this goes through the same review process and requires a progress report for the original allocation to be attached.

Summary of available resources	
Compute resources	Extreme-scale access requests can target the JUPITER, Leonardo, LUMI, and MareNostrum 5 supercomputers. Depending on the system requested, allocations can comprise a minimum of 130000-245000 node hours, with no upper limit given. This corresponds to 260000-980000 CPU and 640000-980000 GPU hours minimum. Access is valid for twelve months and can, on special request, be extended by a maximum of an additional three months if the allocated resources cannot be consumed within the original allocation period. The resource allocation itself cannot be modified. Notice must be given should a situation occur in which the use of the allocated resources is no longer possible.

6.1.9. EuroHPC JU e-Infrastructure – AI & Data-Intensive Applications

Table 6: Access Policy Analysis – EuroHPC JU (AI & Data-Intensive Applications Access)

Access Policy Analysis – EuroHPC JU (AI & Data-Intensive Applications Access)	
Obtaining Access	
Targeted users	Ethical Artificial Intelligence, Machine Learning, and in general, data-intensive applications, with a particular focus on Foundation Models and Generative AI (e.g., Large Language Models).
Process to obtain access	Submit an access proposal (two parts) to a continuously open submission. Cut-off dates for review occur six times a year, and have to pass through three stages of evaluation (administrative check, technical assessment, scientific evaluation). Access will be provided within one month of the cut-off date. It is possible to submit a continuation proposal during the period of validity of the original proposal; this goes through the same review process and requires a progress report for the original allocation to be attached.
Summary of available resources	
Compute resources	AI and Data-Intensive access requests can target the GPU nodes of the Leonardo, LUMI and MareNostrum 5 supercomputers. Depending on the system requested, allocations can comprise 20000-90000 node hours, which corresponds to 20000-180000 CPU and 80000-360000 GPU hours, depending on the system targeted. Access is valid for twelve months. The resource allocation cannot be modified. Notice must be given should a situation occur in which the use of the allocated resources is no longer possible.

6.1.10. EuroHPC JU e-Infrastructure – Federation Platform

Through the public tender EUROHPC/2023/CD/0003 published in October 2023, the EuroHPC JU did solicit offers for implementing and deploying a federation platform which would cover the EuroHPC JU computing resources; in December 2024, a consortium led by CSC-IT Centre for Science in Finland was awarded the contract for the EuroHPC Federation Platform (EFP). A first EFP release and deployment on nine or more Hosting Entities is scheduled for Q1/2026. In Q4/2026, a second release is foreseen, and in addition, all Hosting Entities and their HPC, AI factory, and Quantum Computing systems will be integrated. At that point in time, an initial coupling with Simpl and EOSC and the FENIX federation infrastructure as used by EBRAINS will happen[22]. Three more releases and support of the EFP infrastructure are scheduled until the end of 2029.

[22] See for instance the presentation at the EuroHPC HPC Summit 2025: https://www.google.com/url?q=https://cdn.sanity.io/files/461i44gu/production/fa594fa2ef4982bb63af608107b0cc7a2f2369e5.pptx&sa=D&source=docs&ust=1747382003648103&usg=AOvVaw0uuaskSBGO3dClOxAnyxC4.

The main EFP features are:

- Federated identity management and single-sign-on (SSO) capability for all supported services and systems.
- Integrated resource allocation, management, and monitoring interface covering all systems.
- Interactive use of systems via remote desktops, shell commands, Jupyter Notebooks, and bespoke AI training monitoring interfaces.
- Federated catalogue of pre-installed SW elements on all systems, and advanced discovery capabilities.
- Workflow execution relies on smart scheduling and automated data transfers between steps, complemented by graphical interfaces for creating, managing, and monitoring workflows.
- Support for direct system access on the basis of short-lived SSH certificates created through an MFA-capable login flow.

The major components of the EFP are based on readily available Open-Source SW technologies, which are in tried and trusted production already. The identity management is based on MyAccessID (as used by the FENIX[23] federation system at the heart of EBRAINS), yet EFP will not act as a MyAccessID provider. During the lifetime of the EFP project, a helpdesk and ticketing system will be provided to the end-users.

It is important to note that the EFP will provide supported federation mechanisms and services – access privileges and quotas are granted and managed by the Hosting Entities and/or system operators, and are not fungible across Hosting Entities.

In addition, while automated data transfer services for efficient execution of EFP-managed workflows will be provided, it is at this point in time unclear which functionality will be provided for integration with data infrastructures – Simpl and EOSC are listed as targets, yet no more information is given at this point in time. The EuroHPC JU has repeatedly stated that they see the establishment of data infrastructures and data federation as a problem outside of their purview, to be handled by other funding organizations.

6.2. GCS e-Infrastructure in General

The Gauss Center for Supercomputing/GCS association in Germany has the three largest German HPC centres (Höchstleistungsrechenzentrum Stuttgart/HLRS, Jülich Supercomputing Centre/JSC, and Leibniz Rechenzentrum Munich/LRZ) as its members, and it provides access to the large HPC systems of its members. There will be an overlap with EuroHPC JU systems, since access to the Jupiter Exascale system at JSC (entry into operation planned for H2/2025) and potentially to the AI factory system Hammerhai (2025/2026) at HLRS is foreseen.

Table 7: Access Policy Analysis – GCS (General)

Access Policy Analysis – GCS (General)	
Obtaining Access	
Targeted users	Researchers at German Universities and publicly funded research institutions can apply directly, and European Researchers can go through the Partnership for Advanced Computing in Europe (PRACE)[24]. Research conducted by using the GCS systems is expected to serve the "public interest", and results and findings must be made publicly available.
Process to obtain access	Applicants hand in a project application, which is peer-reviewed and then decided upon by GCS. The criteria for acceptance are scientific excellence and, for the large-scale projects, technical feasibility.
Access Tracks and Modalities	
Access tracks	**Large-scale projects** requiring large amounts of compute time (roughly a minimum of 2% or more of the system's annual compute capacity). Details are given in subsection 6.7.1. **Regular projects** with smaller compute requirements. Details are given in subsection 6.7.2.

[23] See https://fenix-ri.eu/ for details.

[24] It is not clear whether requests for GCS compute time via PRACE are still possible, given the change in role of PRACE.

Access modalities	Login shells via SSH, protected by SSL certificates and increasingly multi-factor authentication (MFA), and SSL-based data transfer (scp et al.) with the same protections. Some GCS sites provide additional access modalities, such as Jupyter notebooks. GCS does not provide any federated access services.
Summary of available resources	
Compute resources	At the time of writing, GCS offers access to • JUWELS Cluster (2511 dual-processor Intel x86 Skylake nodes) and Booster (936 dual-processor AMD x86 Epyc Rome nodes with four NVIDIA A100 GPUs each) at JSC • SuperMUC-NG phase 1 (6480 dual-processor Intel x86 Skylake nodes) at LRZ • SuperMUC-NG phase 2 (240 dual-processor Intel x86 Sapphire Rapids nodes with four Intel Ponte Vecchio GPUs each) is planned to be opened during 2025 • Hunter (752 AMD x86/GPU accelerated processing units) at HLRS
Data resources	No specific GCS data resources; successful projects will use the HLRS, JSC, or LRZ parallel file systems.
Data transfer resources	No specific GCS resources; successful projects will use the mechanisms provided by the member centres, with SSL-based data transfer (scp) as a common denominator.
Access management and security	
Identity management and AAA	The GCS centers manage their own, local user IDs and accept SSL certificates and keys (password-protected) for login and data access, with an ongoing trend to mandate use of multi-factor authentication. Some centers support community-specific identity management and AAA systems like Unicore or FENIX (based on MyAccessID). There are no GCS-level identity management/AAA services.
Security methods and processes	GCS centers institute their own distinct rules for passwords and password updates; MFA is increasingly requested. From the Internet, a set of login nodes is accessible; other “worker” nodes are only visible from the login nodes and are made available through a batch scheduler (Slurm). Support for encrypted data storage is provided, yet at the time of writing, it is not enforced. Unencrypted data transfer (such as FTP or Telnet) is prevented; instead, users must rely on protected methods, like SCP. Direct access to “worker” nodes from the Internet is not permitted, and likewise, access to the Internet from worker nodes is restricted, in most cases requiring staging of data to storage followed by data transfer from a login node.
Rules and assurances	
End-user rules and policies	The GCS centers impose end-user policies mandating fair use of resources in general, outlawing criminal and malicious activities, prohibiting processing and/or creation of personal data protected by GDPR, and disallowing sharing of login credentials. Provision of resources follows a “best effort” paradigm – no hard guarantees are given as to resource availability, performance, and waiting times. It is important to notice that GCS expects projects to make their results and findings created by the use of their HPC systems public.
Fair use, security, data protection	The GCS centers manage their own local quotas for use of compute and data storage resources, aggregated to user or project identities. Batch queues with different priorities and/or available resource quotas (in space and time) support different end-user needs, and the batch scheduling system (Slurm) automates scheduling and allocation according to site-defined fairness rules. Regular end-users are not able to acquire root privileges. Data is protected by Linux and PFS mechanisms, including in many cases, access control lists. Encryption of data at rest is not enforced; unencrypted data transfer mechanisms are blocked.

End-user support	GCS has its own HPC application support structure organized around four support levels: Level 1 is a service desk acting as the entry point for questions from end-users; it uses a ticketing system to keep track of requests. This level will provide the basic support regarding HPC system availability and use. Complex problems are escalated to levels 2 and 3, the latter of which involves cross-sectional or cross-domain teams. These levels solve complex end-user problems, and they are also involved if GCS or center monitoring data indicates consistent inefficiencies in the end-users' applications or system use. In this case, these teams will assist the user in improving scaling and per-node and I/O performance. Support level 4 aims to establish and evolve long-term collaboration between the GCS member centers, their end-users, and academia in general. The objective here is to pool the available expertise to develop new algorithms, workflows, or applications. Levels 2-4 work with and on the GCS centers, groups of end-user support specialists. End-user support for the centers is based on the promise of "best effort".
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	GCS does not disclose details on KPIs that they monitor; the member centers perform their own collection and analysis of system usage and performance data. From the end-user support scheme, it is clear that GCS staff have access to performance monitoring data to assess the efficient use of their systems.
Engagement with end-users	GCS runs a mentoring scheme, with large-scale projects being assigned a GCS mentor acting as a permanent point of contact and advisor. GCS also provides a wealth of training material and training events, in close collaboration with its member centres.
Evaluate, improve, and evolve infrastructure	No public data available.

6.2.1. GCS e-Infrastructure Large-Scale Projects

Table 8: Access Policy Analysis – GCS (Large-Scale)

Access Policy Analysis – GCS (Large-Scale)	
Obtaining Access	
Process to obtain access	Applicants hand in a project application to twice-yearly calls, which then undergoes a competitive peer-review and resource allocation process. Criteria for acceptance include clear scientific goals and milestones, and technical feasibility. In case of already active large-scale allocations for a project, the request for additional large-scale capacities for that project has to be justified. Requests for use of multiple GCS systems must also be specifically justified.
Reporting requirements	A final report must be submitted within three months of the end of the large-scale project. It has to focus on the scientific and technical outcome and detail how the granted compute time was used by the project; publications, PhD theses, and graphical material interesting for the general public should also be included. A shorter status report is required as part of a project application for extending its allocation period. A separate report for publication on the GCS Website must also be submitted within three months of the conclusion of a project, or after two years of multi-year projects. These reports should briefly describe the scientific and technical goals.

Summary of available resources	
Compute resources	**Large-scale projects** require large amounts of compute time (roughly 2% or more of the system's annual compute capacity), which translates to a **minimum** of • 45,000,000 core hours for SuperMUC-NG phase 1, equivalent to 937,500 node hours • 25,000 node hours for Hunter • 45x10^{21} floating point operations per year for JUWELS Cluster or Booster • 847x10^{21} floating point operations per year for JUPITER GCS does not specify a limit on compute resources that can be requested, and multi-year projects are also possible.

6.2.2. GCS e-Infrastructure Regular Projects

Table 9: Access Policy Analysis – GCS (Regular)

Access Policy Analysis – GCS (Regular)	
Obtaining Access	
Process to obtain access	Applicants hand in a project application (at any time for the Hawk and SuperMUC-NG systems), which undergoes a peer-review process. Criteria for acceptance include clear scientific goals and milestones, and technical feasibility. Requests for use of multiple GCS systems must also be specifically justified.
Reporting requirements	A final report must be submitted within one month of the end of the large-scale project. It has to focus on the scientific and technical outcome and detail how the granted compute time was used by the project; publications, PhD theses, and graphical material interesting for the general public should also be included. A shorter status report is required as part of a project application for extending its allocation period. A separate report for publication on the GCS Website must also be submitted within three months of the conclusion of a project, or after two years of multi-year projects. These reports should briefly describe the scientific and technical goals.
Summary of available resources	
Compute resources	Regular calls need to stay below the yearly limits of the large-scale calls detailed in subsection 6.2.1. Multi-year projects are possible.

6.3. NHR Alliance e-Infrastructure

The Nationales Hochleistungsrechnen (NHR) alliance in Germany has twelve "tier-1" HPC centres (Technical Universities of Aachen, Darmstadt, Dresden, and Kaiserslautern, Universities of Berlin, Frankfurt/Main, Göttingen, Mainz, Nuremberg/Erlangen, Paderborn, and Saarland, and the Karlsruhe Institute of Technology) as members. It provides access to mid-range HPC systems to German scientists, with a number of systems designed to support a subset of scientific disciplines.

Table 10: Access Policy Analysis – NHR (General)

Access Policy Analysis – NHR (General)	
Obtaining Access	
Targeted users	Scientists with a doctoral degree from a German accredited university can apply for HPC resources. Within the scope of the approved projects, non-doctoral scientists (of German universities) can use the resources in addition.
Process to obtain access	The PI submits an application for resource allocation at any of the NHR centers via the Web-based JARDS portal. The time between proposal submission (or cutoff date for large projects) and approval is three months. Projects can be approved for multiple years, and continuation proposals can be handed in
Reporting requirements	Each approved project must submit yearly project reports detailing the resource use and the generated scientific results and publications, and include a description of the project and its results suitable for the general public and for dissemination by NHR. Acknowledgement of the use of NHR systems must be included in all publications of scientific results generated by the use of NHR resources.
Access Tracks and Modalities	
Access tracks	Besides tracks for HPC-starters and initial tests to determine whether the selected NHR system applies to a research project, NHR offers access tracks "Normal" and "Large". Details can be found in subsections 6.3.1 and 6.3.2.
Access modalities	Login shells via SSH, protected by SSL certificates and increasingly multi-factor authentication (MFA), and SSL-based data transfer (scp et al.) with the same protections. Some NHR sites provide additional access modes, such as Jupyter notebooks or project-specific services. NHR does support transferring projects from one center to another without re-submission of a proposal; approval of such a transfer is done by the "receiving" center and contingent on evaluation of technical feasibility. Besides this, NHR does not provide any federated access services.
Summary of available resources	
Compute resources	NHR consists of no less than twelve member centers; the available computing systems are in the range of 100-1000 dual-CPU (x86) nodes for aggregated performance in the single digit PetaFlop/s range, and 100-1000 dual-CPU (x86) nodes with NVIDIA GPGPUs (A100 and H100) in the range of 10-50 PetaFlop/s. The NHR members each have special competency in distinct fields of computing/science, and offer in-depth advice and support in these fields to NHR users on their system. The scientific domains include engineering sciences, atomistic simulations, life sciences, earth system sciences, materials sciences, high-energy physics, and computational chemistry. Specific computing fields covered are data analytics and data-intensive compute, performance optimization, AI, and quantum simulations and computing.
Data resources	No specific NHR data resources; successful projects will use the local parallel and general-purpose file systems at the participating centers.
Data transfer resources	No specific NHR resources; successful projects will use the mechanisms provided by the NHR centers, with SSL-based data transfer (scp) as a common denominator.
Access management and security	
Identity management and AAA	The NHR centers manage their own, local userids and accept SSL certificates and keys (password-protected) for login and data access, with an ongoing trend to mandate use of

	multi-factor authentication. There are no NHR-level identity management/AAI services.
Security methods and processes	NHR centers institute their own distinct rules for passwords and password updates; MFA is increasingly requested. From the Internet, a set of login nodes is accessible; other “worker” nodes are only visible from the login nodes and are made available through a batch scheduler (usually Slurm). Support for encrypted data storage is provided, yet at the time of writing, it is not enforced. Unencrypted data transfer (such as FTP or Telnet) is generally prevented; instead, users must rely on protected methods, like SCP. Direct access to “worker” nodes from the Internet is not permitted, and likewise, access to the Internet from worker nodes is restricted, in most cases requiring staging of data to storage followed by data transfer from a login node.
Rules and assurances	
End-user rules and policies	The NHR centers impose end-user policies mandating fair use of resources in general, outlawing criminal and malicious activities, prohibiting processing and/or creation of personal data protected by GDPR, and disallowing the sharing of login credentials. Provision of resources follows a “best effort” paradigm – no hard guarantees are given as to resource availability, performance, and waiting times.
Fair use, security, data protection	The NHR centers manage their own local quotas for use of compute and data storage resources, aggregated to user or project identities. Batch queues with different priorities and/or available resource quotas (in space and time) support different end-user needs, and the batch scheduling system (Slurm) automates scheduling and allocation according to site-defined fairness rules. Data is protected by Linux and PFS mechanisms, including, in many cases, access control lists. Encryption of data at rest is not enforced; unencrypted data transfer mechanisms are blocked. Regular end-users are not able to acquire root privileges.
End-user support	NHR provides support during the application phase via its central organization and the participating centers; this includes advice on selecting the best-suited NHR center (including test accounts). The NHR centers support users during the project execution phase.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	The NHR centers collect a wealth of monitoring data locally; they roll up data on the use of their resources through NHR projects and pass that on to the NHR itself.
Engagement with end-users	Via the central NHR organization (during submission, approval, and reporting), and the NHR centers during project execution.
Evaluate, improve, and evolve infrastructure	No data is publicly available here.

6.3.1. NHR e-Infrastructure Normal Projects

Table 11: Access Policy Analysis – NHR (Normal)

Access Policy Analysis – NHR (Normal)	
Obtaining Access	
Process to obtain access	Applications can be submitted at any time and can start at any time (with a few exceptions for the centers at Dresden and Göttingen). They include a brief description of the scientific project and a detailed justification of the resource needs, with a focus on establishing efficient use of the resources. Resources are allocated by the central "NHR Nutzungsausschuß[25]" based on the review by the scientific committee at the respective NHR center.
Summary of available resources	
Compute resources	NHR states that "moderate" compute resource requests fall under this category, without listing specific limits. Resources available to a project can be increased by at most 25%, and the project term can be extended by up to three months without submitting a new proposal.

6.3.2. NHR e-Infrastructure Large Projects

Table 12: Access Policy Analysis – NHR (Large)

Access Policy Analysis – NHR (Large)	
Obtaining Access	
Process to obtain access	Applications can be submitted at four deadlines each year, and large projects start at fixed dates (four per year). They include a brief description of the scientific project and a justification of the resource needs, including demonstration of scalability. Resources are allocated by a scientific committee at the respective NHR center.
Summary of available resources	
Compute resources	NHR states that "large" compute resource requests fall under this category, without listing specific limits. Resources available to a project can be increased by at most 25%, and the project term can be extended by up to three months without submitting a new proposal.

[25] Meaning "usage committee".

6.4. RES e-Infrastructure

Table 13: Access Policy Analysis – RES (General)

Access Policy Analysis – RES	
Obtaining Access	
Targeted users	RES resources are open for all kinds of users within the European Union and its associated states via competitive calls for researchers performing open research, including: Academic researchers and from public administration, SME researchers for open R&D&I, and New users without prior HPC experience.
Process to obtain access	Access to the resources is through competitive calls based on the scientific excellence of the projects submitted. To submit the application, users need to access the RES intranet and properly fill out the application form. Open calls for activities are always open with cut-off dates every 4 months. Calls for tests are always open. Access has a double filter system, including an Initial peer review evaluation by the Spanish State Research Agency (Agencia Estatal de Investigación, AEI) of the general research programme, where the specific areas of the activity involving the use of RES are defined. and the Subsequent evaluation by the Access Committee, advised by a Technical Experts Panel and a Scientific Experts or Data Management Experts Panel, of the activity, which must provide detailed objectives and an approximate timeline for completion.
Reporting requirements	The RES must be acknowledged in any publications resulting from accepted projects.
Access Tracks and Modes	
Access tracks	There are three main tracks HPC - Supercomputing, AI - Artificial Intelligence, for which applications can be submitted for activities and for testing, and for DATA - Data Management, only applications for activities are available.
Summary of available resources	
Compute resources	At the time of writing the report, the total capacity of the RES supercomputers exceeds 329 PFlop/s, providing over 90 PFlop/s exclusively for RES users. These supercomputers may consist solely of processors (CPUs) or be hybrids with processors and graphics accelerators (CPUs + GPUs). This allows addressing various challenges, ranging from traditional supercomputing for complex calculations to the creation of digital twins to simulate and predict complete systems. • MareNostrum and MinoTauro at the Barcelona Supercomputing Center - Centro Nacional de Supercomputación (BSC-CNS). MareNostrum5 GPP provides a peak performance of 45.5 PFlop/s. 6 408 nodes: 2x Intel Shappire Rapids 8480+ 56 cores 2 GHz and 72 nodes: 2x Intel Shappire Rapids 03H-LC 56 cores 1.7 GHz. MareNostrum5 ACC provides 260 PFlop/s in 1 120 nodes (89 600 cores + GPU) including 2x Intel Shappire Rapids 8460Y+ 40 cores 2.3 GHz and 4x NVIDIA Hopper 64 GB HBM. • La Palma at the Canary Islands Astrophysics Institute (IAC) provides 83.85 TFlop/s peak performance in 252 nodes (4 032 cores) Intel Xeon E5-2670/1600 20M 2.6 GHz. • Altamira at the Physics Institute of Cantabria (IFCA) of the University of Cantabria (UC) offers a peak performance of 105 TFlop/s. 158 nodes (2 528 cores). IBM dx360 nodes: 2x Intel Xeon E5-2670 2.6 GHz and 64 GB of RAM. • Picasso at the University of Malaga (UMA) provides 4.34 PFlop/s peak performance. 344 nodes (38 296 cores). 26 SD530 nodes: 2x Intel Xeon Gold 6230R and 192 GB of RAM, 156 Lenovo sr645 nodes, 2x AMD EPYC 7H12 and 512GB of RAM, 24 BULL R282-Z90 nodes, 2x AMD EPYC 7H12 and 2048 TB of RAM, 34 Lenovo sr645 v3 nodes, 2x AMD EPYC 9754 and 768 GB of RAM and 4 additional nodes NVIDIA DGX nodes: 8x NVIDIA A100 GPU.

	• Tirant at the University of Valencia (UV) provided 111.8 TFlop/s peak performance. 336 nodes (5 376 cores). iDataPlex dx360m4 nodes: 2x Intel Xeon E5-2670 2.6 GHz and 32 GB of RAM. • Caesaraugusta in the Institute for Biocomputation and Physics of Complex Systems (BIFI) at the University of Zaragoza (UNIZAR) offers 537 TFlop/s peak performance. 96 nodes (6 144 cores). 2x AMD Epyc 7513A 2.6 GHz. 93 nodes thin: 256 GB RAM, 3 nodes FAT: 512 GB RAM. Storage: 480 GB SSD per node. • Caléndula at the Fundación del Centro de Supercomputación de Castilla y León (SCAYLE) offers 246 TFlop/s. 186 nodes (2 976 cores). 2x Intel Xeon E5-2670 (SandyBridge) 2.6 GHz and 32 GB of RAM. • Pirineus at the Consorci de Serveis Universitaris de Catalunya (CSUC) Pirineus III offers 427 TFlop/s peak performance. 58 nodes (11 136 cores). 50 nodes: 2x AMD Epyc 9654 2.4 GHz and 8 nodes: 2x AMD Epyc 9654 2.4 GHz AND 3TB of RAM. Pirineus III - GPU offers 1.59 PFlop/s. 29 nodes (1 856 cores + GPU). 2x AMD Epyc 9334 2.7 GHz, 2x NVIDIA H100 80 GB HBM. • Lusitania at the CénitS-COMPUTAEX offers 89.19 TFlop/s peak performance. 208 nodes (3 488 cores). 168 IBM dx360 nodes: 2x Intel Xeon E5-2670 2.6 GHz and 32 GB of RAM, 40 Fujitsu Primergy CX2550 nodes: 2x Intel Xeon E5-2660 v3 2.6 GHz, 80 GB of RAM and 256 GB SSD. • Finisterrae at the Fundación Pública Galega Centro Tecnolóxico de Supercomputación de Galicia (CESGA) FinisTerraeIII - CPU offers 1.5 PFlop/s peak performance. 273 nodes (17 472 cores). 16 FAT nodes with 2048 GB of RAM and 1 OPTANE node with 8192 GB of RAM. FinisTerraeIII - GPU - offers 2.8 PFlop/s. 66 nodes (4224 cores + GPU). 64 nodes: 2x NVIDIA A100, 1 node: 5x NVIDIA A100, 1 node: 8x NVIDIA A100. • Cibeles at the Universidad Autónoma de Madrid (UAM) offers 100 TFlop/s peak performance. 28 nodes (1 568 cores). 2x Intel Xeon Gold 6330 2.0 GHz and 512 GB of RAM. • Urederra at the Navarra de Servicios y Tecnologías, S.A.U (NASERTIC) offers 487 TFlop/s peak performance. 180 nodes (7 576 cores). 142 nodes Intel SkyLake: 2x Intel Xeon Platinum 8160 @ 2.10 GHz (6 816 cores) and 54 528 GB de RAM. 38 nodes Intel Broadwell: 2x Intel Xeon CPU E5-2640 v4 @ 2.40 GHz (760 cores) and 4 864 GB de RAM. • Xula and Turgalium at the Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT). Xula offers 135 TFlop/s peak performance. 44 nodes (1 760 cores). 2x Intel Xeon Gold 6148 2.4 GHz and 192 GB of RAM. Turgalium offers 143 TFlop/s peak performance. 40 nodes (1 440 cores). 2x Intel Xeon Gold 6254 3.1 GHz and 192 GB of RAM. There is no strict minimum or maximum limit of requested time. However, those activities requiring more than 10 million hours should provide a justification (section 5a of the application form). The amount of resources needed differs a lot depending on the research activity. Typically, RES applicants request between 50.000 and 2.000.000 hours, but smaller or larger projects are also accepted.
Data resources	The RES nodes feature over 183 PB of storage capacity, of which more than 41 PB are allocated to RES users. Data projects must have a minimum of 200TB to use the data services of the RES. As for an upper limit, there is no specified cap, but typically, data projects do not exceed 1 PB of data volume. Data resources encompass a complete storage and processing infrastructure designed to manage large volumes of information efficiently and securely: • Storage systems: Multilevel storage including disk storage for frequent access (both in files and objects) and magnetic tape storage for long-term conservation, backups, and efficient archiving of historical data. • Virtual infrastructure: Access to customizable virtual machines that allow the creation of specific environments for data processing and analysis, offering flexibility in resource configuration according to the needs of each project.

	• Complementary computing capacity: Limited computational resources for data processing-intensive tasks, enabling complex analyses, transformations, and aggregations of large datasets without the need to create a full project for HPC resources.
Data transfer resources	Each site offers its specific protocols for access, usually via SSH, SIR2, or UNIFICAT. Mare Nostrum as EUDAT service provider, BSC can facilitate the integration of B2DROP, B2SAFE, and PIDs services in data projects. Depending on the site, offer GridFTP, FTP & rsync for large volumes of data and S3.
Access management and security	
Identity management and AAA	Depending on the site, MareNostrum supports federated access protocols (EduGAIN, SIR). For other sites, mostly local accounts are needed. CESGA and BIFI can be accessed under EGI-AAI.
Security methods and processes	Depending on the site, they have a User/password system and have a Segmented Network, LDAP, and Firewalls (e.g. PaloAlto firewall + fail2ban).
Rules and assurances	
End-user rules and policies	RES includes a General Usage Terms policy restricting the use of the compute resources for exclusive research purposes and non-profit aims.
Fair use, security, data protection	A username and password to access RES facilities. They need to be strictly secret and not shared with other users, in order to avoid infiltration in the facilities. All data stored on disks is removed 1 week after the project is completed. RES is not responsible for user data. Users are encouraged to back up their own data. RES requires users to mention the use of RES facilities in any of their publications. RES may publish the subject of the research done using RES facilities, unless an NDA is signed.
End-user support	Each side provides support via a specific email. Support to DATA resources is also provided by each of the sites, ranging from 1 FTE to 5 FTE, and includes Technical Support, Data Manager, Data Engineers, or domain-specific (life sciences, quantum, etc).
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	CPU Hours with & without Priority for HPC. TB or CPU hours and VMs for Data.
Engagement with end-users	The RES organizes an annual User Conference to encourage the exchange of experiences, opinions, and ideas among users, technicians, and RES coordinators. RES supports ENS users for the organization of scientific seminars on supercomputing applications in specific scientific domains. RES nodes also organize technical formation to provide the necessary knowledge for users and technicians to use and manage the various supercomputing, AI, and data resources.
Evaluate, improve, and evolve infrastructure	Based on RES History, older versions of MareNostrum have been distributed to the different RES nodes to update their systems as newer versions of MareNostrum became available. Last time during 2024, following the inauguration of the MareNostrum5 computer.

6.4.1. RES e-Infrastructure HPC & AI standard activities

Table 14: Access Policy Analysis – RES (HPC & AI standard activities)

Access Policy Analysis – RES e-Infrastructure HPC & AI standard activities	
Obtaining Access	
Process to obtain access	The procedure for submitting an HPC resource request consists of submitting projects to open and competitive calls for researchers from Spanish and European institutions. Depending on the type of activity (ordinary, tutored, long-duration, and continuation, pre-reservation), the application form will be different, e.g., in the case of a continuation activity, the form is Simplified, or in the case of a pre-reservation. An application can be submitted at any time, and can be modified as many times as necessary before the established deadline. All the necessary information about accessing RES resources can be found in the **RES Access Protocol** document. Applications include general information on the activity type, scientific area, the description of the research project, the necessary software and numerical libraries, a description of the research team, and the needed resources - including the machine and expected processors, disk space, and requested time.
Summary of available resources	
Compute resources	If the requested activity is asking for more than 10 Million CPU hours, the applicant needs to justify the amount of resources requested for the activity.

6.4.2. RES e-Infrastructure HPC & AI test activities

Table 15: Access Policy Analysis – RES (HPC & AI test activities)

Access Policy Analysis – RES e-Infrastructure HPC & AI test activities	
Obtaining Access	
Process to obtain access	The tests are designed for users utilizing supercomputing resources for the first time in their project, allowing them to test algorithms and code for use in an HPC or AI activity. The test call is always open, and reviews will be conducted as quickly as possible to ensure fast access, but without a fixed deadline.
Summary of available resources	
Compute resources	Due to their trial nature, the tests require a reduced computing time and receive special support from the support teams.

6.4.3. RES e-Infrastructure QUANTUM activities

Table 16: Access Policy Analysis – RES (QUANTUM activities)

Access Policy Analysis – RES e-Infrastructure QUANTUM activities	
Obtaining Access	
Process to obtain access	Quantum computing resources can be accessed through the same application form as the other activities, but with specific guidelines as explained in the Quantum Resources Access Guide. The quantum machine is mainly intended for chip usage and low CPU demand. Each quantum job is associated with a 40-CPU node, which can be used to run supporting code, but the main computation must involve the quantum chip. Applications involving hybrid workflows need to request both Quantum and the appropriate classical systems (e.g., MN5-ACC, MN5-GPP). An application includes General information, including activity type, scientific area, and type of application, research project description, including the expected number of qubits required, maximum circuit depth, number of shots, gate set used, and estimated duration of the longest job. It also includes the software and libraries needed, e.g., Qibo and Qililab libraries or any specific quantum computing workflow, and the description of the research team.
Summary of available resources	
Compute resources	Users accessing Quantum Computing (QUANTUM) should expect evolving capabilities, some unexpected issues, and an opportunity to test and contribute to RES' quantum computing capabilities. The available quantum chips at the BSC cluster may change over time during the RES access period. At the time of writing this deliverable, this cluster hosts two chips: one with five qubits and another with 10 qubits. A new 20-qubit chip is expected to be installed around June/July 2025. Updated information about the cluster and available resources can be accessed here. As a reference, a typical job with a circuit depth of 1000 and 10,000 shots takes around 20 seconds to execute. Providing these details is essential for assessing the feasibility and resource impact of your proposal.

6.4.4. RES e-Infrastructure Data activities

Table 17: Access Policy Analysis – RES (Data activities)

Access Policy Analysis – RES e-Infrastructure Data activities	
Obtaining Access	
Process to obtain access	The process to access DATA resources is essentially the same as for accessing HPC resources. The main difference between the two modalities lies in the nature of the projects, and therefore, the consultation documentation differs. DATA applications must provide a Data Management Plan as explained in RES Data Services Application Guide. It needs to provide project title, information on the applicant and general description, the total desired storage (in TB), resources needs per year, host nodes, technical project description, number and description of datasets (including name and ID, volume, format, storage type main policy, back up policy and legal/ethical restrictions - such as personal data, clinical data, genomic information, etc).

Summary of available resources	
Compute resources	The minimum size required to carry out a data project is 200 TB, in order to fully take advantage of the capabilities of the RES nodes. Although there is no stipulated limit, typically, projects larger than 1 PB are the upper limit for submitted projects.

6.5. GENCI e-Infrastructure

Table 18: Access Policy Analysis – GENCI (General)

Access Policy Analysis – GENCI (General)	
Obtaining Access	
Targeted users	Principal investigators (PIs)must be scientists working in France in the academic sphere (scientific and technological establishments, higher education, research institutes, and foundations) and for public-funded research projects or parts of the industrial sphere (companies conducting research activities in France). End-users of resources can also include scientists from other countries, if they are part of a project associated with a French research institute.
Process to obtain access	PI must submit a project application to the EDARI system requesting access to the HPC resources, which will be reviewed based on the criteria of scientific and technical excellence. Decisions are made by the central GENCI evaluation and allocation committees or by the GENCI center directors. Results created by the use of GENCI resources must be published. All project end-users must separately submit requests to receive user accounts at the GENCI centres, if they do not already have these.
Reporting requirements	An activity report is required at the end of the allocation or when a project is renewed. This must list all publications produced using the GENCI resources, and DOIs must be registered with the EDARI system. Acknowledgement of the use of GENCI resources must be given in all publications.
Access Tracks and Modes	
Access tracks	GENCI supports “regular” and “dynamic” access tracks.
Summary of available resources	
Compute resources	At the time of writing, GENCI provides access to three member centers and these HPC systems: • Adastra supercomputer at CINES with an aggregated performance of 91 PFlop/s, composed of a scalar partition (2xAMD x86 Genoa, 544 nodes) with 3.3 PFlop/s, a GPU-accelerated partition (AMD x86 Trento plus 4x AMD MI250x, 356 nodes) with 74 PFlop/s, and an APU partitions (4x AMD x86/GPU MI300A, 28 nodes) with 13.7 PFlop/s. • Jean Zay supercomputer at IDRIS with an aggregated performance of 126 PFlop/s, composed of a scalar partition (2x Intel x86 Cascade Lake, 720 nodes) with 2.3 PFlop/s, an accelerated HPC AI partition (2x Intel x86 Cascade Lake plus 4x NVIDIA V100, 427 nodes 720 nodes) with 15.5 PetaFlop/s, an AMD AI-accelerated partition (2xAMD x86 Milan plus 8xNVIDIA A100, 52 nodes) with 8.2 PFlop/s, and a new accelerated HPC/AI partition (2x Intel x86 Sapphire Rapids plus 4x NVIDIA H100, 364 nodes) with 100 PFlop/s. • Joliot Curie/Irene supercomputer at CEA with an aggregated performance of 20 PFlop/s, composed of an AMD scalar partition (2x AMD x86 Rome, 64 nodes) with 2.0 PFlop/s, an Intel

	scalar HPC partition (2x Intel x86 Skylake, 192 nodes) with 6.9 PFlop/s, and an accelerated partition (2x Intel x86 Cascade Lake, 4x NVIDIA V100, 32 nodes) with 1.1 PFlop/s.
Data resources	No specific GENCI data resources; projects will use the CINES, IDRIS, and CEA parallel file systems.
Data transfer resources	No specific GENCI resources; successful projects will use the mechanisms provided by the member centres, with SSL-based data transfer (scp) as a common denominator.
Access management and security	
Identity management and AAA	The GENCI centers manage their own, local user IDs and accept SSL certificates and keys (password-protected) for login and data access; GENCI handles the process of applying for user accounts across the centers. There are no further GENCI-level identity management/AAA services.
Security methods and processes	GENCI centers institute their own distinct rules for passwords and password updates. From the Internet, a set of login nodes is accessible; other "worker" nodes are only visible from the login nodes and are made available through a batch scheduler (usually Slurm). Support for encrypted data storage is provided, yet at the time of writing, it is not enforced. Unencrypted data transfer (such as FTP or Telnet) is generally prevented; instead, users must rely on protected methods, like SCP. Direct access to "worker" nodes from the Internet is not permitted; likewise, access to the Internet from worker nodes is only possible per specific request; in most cases, staging of data to/from storage combined with data transfer to/from a login node is required.
Rules and assurances	
End-user rules and policies	GENCI prohibits appropriating the rights of use as well as access to the accounts of others by any means. It also requests that end-users limit access to resources they own (files, executables, directories, etc.) to the strict minimum possible. Deliberate propagation of viruses, the development and use of software to circumvent the security devices in place, as well as the misappropriation of resources made available for purposes other than those described in the application for award, constitute misconduct liable to professional or criminal sanctions. In addition, processing and/or creation of personal data protected by GDPR, in particular medical data, is disallowed. Provision of resources follows a "best effort" paradigm – no hard guarantees are given as to resource availability, performance, and waiting times. In exceptional circumstances, like f.i. incidents outside of the influence of the end-users, which stop long-running compute jobs, restitution of core/GPU hours lost can be requested by the project.
Fair use, security, data protection	Projects can use up to 125% of their GENCI resource allocations before being stopped from further resource access. GENCI centers expect that the resource allocation is used in a continuous way across the project period (year); priorities of jobs belonging to a project are set according to previous resource usage, with projects that have under-consumed receiving priority boosts and over-consuming projects seeing priorities reduced. Project leads are encouraged to return allocations that they cannot use, reducing their total allocation and helping the GENCI centers to support other projects from such returned allocations. The GENCI centers manage their own local quotas for the use of compute and data storage resources, aggregated to user or project identities. Batch queues with different priorities and/or available resource quotas (in space and time) support different end-user needs, and the batch scheduling system (Slurm) automates scheduling and allocation according to the GENCI fairness rules. Data is protected by Linux and PFS mechanisms, including, in many cases, access control lists. Encryption of data at rest is not enforced; unencrypted data transfer mechanisms are blocked.

	Regular end-users are not able to acquire root privileges.
End-user support	End-user support is provided by User Committees at the GENCI centers, and by regular GENCI centers support staff. The EDARI system provides information, FAQs, and access to user and project data.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	GENCI and its centers monitor the use of resources by projects and make details of the monthly consumption available via the EDARI system. The PI is also informed via email should allocations be exceeded or should significant under-consumption occur. A notional valuation in euros and in equivalent tonnes of carbon is provided to PIs for informational purposes, and the carbon footprint of each job is calculated and communicated by the GENCI centres.
Engagement with end-users	Mainly through the EDARI portal, in addition to the User Committees at the GENCI sites.

6.5.1. GENCI e-Infrastructure Regular Projects

Table 19: Access Policy Analysis – GENCI (Regular Projects)

Access Policy Analysis – GENCI (Regular Projects)	
Obtaining Access	
Process to obtain access	Proposals for regular projects can be submitted twice a year and need to include in-depth administrative, scientific, and technical information, and a detailed justification for the resources required. Decisions are made by the central GENCI evaluation and allocation committees. Time between a cutoff date and a decision can vary between two and six months, and a dynamic allocation can be provided quickly to allow a swift project start. Such projects generally run for one year. Extension requests for additional resources can be made halfway through the project and require only an updated justification for the resources required. Projects can be renewed up to two years after their initial closure. Large projects are reclassified automatically as dynamic projects if their resource requests fall below the threshold. An "on-the-fly" extension of the resource allocation can be made in exceptional circumstances.
Reporting requirements	An activity report is required at the end of the allocation or when a project is renewed. This must list all publications produced using the GENCI resources, and DOIs must be registered with the EDARI system. Acknowledgement of the use of GENCI resources must be given in all publications.
Summary of available resources	
Compute resources	Projects requesting more than 50,000 V100 GPU hours or 500,000 Intel Cascade Lake core hours are part of the "large" access track. GENCI defines conversion factors for the more modern CPUs and GPUs; for instance, AMD MI300 hours is considered equivalent to one NVIDIA H100, two NVIDIA A100 hours, and four NVIDIA V100 hours.

6.5.2. GENCI e-Infrastructure Dynamic Projects

Table 20: Access Policy Analysis – GENCI (Dynamic Projects)

Access Policy Analysis – GENCI (Dynamic Projects)	
Obtaining Access	
Process to obtain access	Proposals for dynamic projects can be submitted at any time and need to include only basic administrative, scientific, and technical information. Decisions are made by the GENCI center directors. Dynamic projects generally run for one year. An "on-the-fly" extension of the resource allocation can be made in exceptional circumstances.
Reporting requirements	An activity report is required at the end of the allocation or when a project is renewed. This must list all publications produced using the GENCI resources, and DOIs must be registered with the EDARI system. Acknowledgement of the use of GENCI resources must be given in all publications.
Summary of available resources	
Compute resources	Projects requesting less than 50,000 V100 GPU hours or 500,000 Intel Cascade Lake core hours are part of the "dynamic" access track. GENCI defines conversion factors for the more modern CPUs and GPUs; for instance, AMD MI300 hours is considered equivalent to one NVIDIA H100, two NVIDIA A100 hours, and four NVIDIA V100 hours.

6.6. CSCS e-Infrastructure

The Centro Svizzero di Calcolo Scientifico (CSCS) in Lugano/Switzerland operates the Swiss National Supercomputing resources (currently the Alps systems) and offers access to Swiss scientists; it also hosts systems and provides services for Swiss research communities and public organisations, for instance the Swiss weather service and CERN for analysis of data from the Large Hadron Collider (LHC) by special contractual arrangement.

Table 21: Access Policy Analysis – CSCS (General)

Access Policy Analysis – CSCS (General)	
Obtaining Access	
Targeted users	Researchers in Switzerland can apply for computational resources. Services and systems for Swiss research communities or public institutions can also be integrated/provided by special arrangement.
Process to obtain access	PIs submit project proposals, which are reviewed by two scientists who belong to academic establishments from around the world and two technical experts from CSCS. The final decision about resource allocations is the purview of an independent expert committee. Projects can have a duration of up to 36 months, with a progress report and separate renewal proposal being required after each year. Allocated resources can increase by a maximum of 20% from year to year. CSCS also enters into contractual arrangements to host systems and provide services to Swiss research communities and public bodies.

Access Tracks and Modes	
Access tracks	CSCS supports a range of access tracks, with production, development, and preparatory projects being relevant for the Spectrum analysis.
Access modalities	CSCS offers a range of Web-based services for end-users to access the compute and storage resources, and also supports limited-time validity SSH key pairs for SSH/SCP access. Besides the usual login shell/Slurm combination, CSCS offers restful Web interfaces and Kubernetes/OpenStack scheduling/orchestration. CSCS manages the different parts of Alps as a single system, and makes partitions available as virtual Clusters.
Summary of available resources	
Compute resources	The Alps system, run by CSCS, is distributed over different sites providing geographically redundant supercomputing services and locality to large amounts of data: • CSCS in Lugano. • EPFL in Lausanne. • Paul Scherrer Institute (PSI) in Villigen for data archiving. • ECMWF in Bologna for access to meteorological data. Alps provides an aggregate performance of more than 400 PetaFlop/s and consists of five different partitions: • 2688 nodes with 4x NVIDIA Grace-Hopper ARM/GPU combined processors. • 1024 nodes with 2x AMD x86 Rome CPUs. • 144 nodes with AMD x86 Rome CPU plus 4x NVIDIA A100 GPUs each. • 128 nodes with 4x AMD MI300A x86/GPU combined processor each. • 24 nodes with one AMD x86 Rome CPU, plus one AMD MI250X GPU each. The first two partitions are the main resources for external users; CSCS makes these available as virtual Clusters: • Clariden for the Grace-Hopper nodes • Eiger for the AMD CPU nodes
Data resources	Two hard disk-based storage systems with 10 and 100 PetaBytes, respectively, and two SSD-based storage systems with 5 PetaBytes and 1 PetaByte, respectively.
Data transfer resources	CSCS recommends the use of the Globus online endpoint or the Globus URL-copy command for file transfer to/from external sites. SCP is, of course, also supported.
Access management and security	
Identity management and AAA	CSCS users have a single account for the different services. Using Multi-factor authentication (MFA) is obligatory for most users. With their username/password and a one-time password (OTP) created during the MFA, end-users can access a single-sign-on gateway or a set of Web services. Using the SSO gateway enables users to switch between different services without re-authentication. Alternatively, end-users can log in to a SSHService web-based service (with MFA) and create SSH key pairs that are valid for 24 hours. User-generated SSH keys are not supported.
Security methods and processes	OTPs generated by Multi-factor authentication (MFA) or limited-time SSH keys provide an additional level of security; the latter should also address usability restrictions encountered by “naïve” MFA implementations. Support for encrypted data storage is provided, yet at the time of writing, it is not enforced. Unencrypted (such as ftp or telnet) data transfer is, in general, prevented; instead, users have to rely on protected methods like the suggested Globus data transfer or scp. Direct access to “worker” nodes from the Internet is not permitted, and likewise, access to the Internet from worker nodes is restricted, in most cases requiring staging of data to storage followed by data transfer from a login node.

Rules and assurances	
End-user rules and policies	CSCS imposes end-user policies mandating fair use of resources in general, outlawing criminal and malicious activities, prohibiting processing and/or creation of personal data protected by GDPR, and disallowing the sharing of login credentials. It reminds users not to flood Slurm with hundreds of jobs and commands at the same time, and not to run compute or memory-intensive applications on the login nodes. Provision of resources follows a "best effort" paradigm – no hard guarantees are given as to resource availability, performance, and waiting times.
Fair use, security, data protection	CSCS manages batch queues with different priorities and/or available resource quotas (in space and time) supporting different end-user needs, and the batch scheduling system (Slurm) automates scheduling and allocation according to the CSCS fairness rules. Alternatively, Kubernetes/OpenStack can be used within existing Slurm allocations. Data is protected by Linux and PFS mechanisms, including, in many cases, access control lists. Encryption of data at rest is not enforced; unencrypted data transfer mechanisms are blocked. Regular end-users are not able to acquire root privileges.
End-user support	The CSCS support team offers a Web-based service desk and a ticketing system.

6.6.1. CSCS e-Infrastructure Production Projects

Table 22: Access Policy Analysis – CSCS (Production Projects)

Access Policy Analysis – CSCS (Production Projects)	
Obtaining Access	
Process to obtain access	Proposals for production projects can be submitted by the PI for twice yearly cutoffs (May 19 and October 20) and need to include the scientific goals and objectives, research methods, algorithms, and code parallelization approach (including memory requirements). It should also include representative benchmarks, justification for resource requests (compute and data), and a project plan with tasks and milestones. Needs for visualization, and pre- and post-processing should be detailed, and any previous CSCS allocations requested, granted, and used in previous projects of the PI must be listed, in addition to relevant references from the literature. Each PI can only submit up to two production proposals per cutoff date, and proposals that have already been submitted to other HPC programs are inadmissible. The PI must hold a postdoc position with an academic institution. Only professors can request multi-year projects (up to three years). The review process consists of a technical feasibility analysis, a scientific review, and a final decision by a scientific committee. Applicants without prior access to Alps must first submit a preparatory project (see subsection 6.6.3). Such a project allows porting and testing of codes, as well as collecting all the information necessary for a production project proposal.
Reporting requirements	PIs of production projects accepted for longer than one year (with a max duration of 36 months) need to submit a progress report together with a resource request and justification within the deadline announced below. In the renewal, allocation requests can be increased by up to 20%.

Summary of available resources	
Compute resources	Large production projects use between 10,000 and 250,000 node hours. Resources are granted for one year, starting at the next allocation window after the cutoff date (April 1 or October 1 of each year).

6.6.2. CSCS e-Infrastructure Development Projects

Table 23: Access Policy Analysis – CSCS (Development Projects)

Access Policy Analysis – CSCS (Development Projects)	
Obtaining Access	
Process to obtain access	Proposals for large development projects can be submitted by the PI for twice yearly cutoffs (May 19 and October 20) and must include the background, significance and scientific objectives for the intended applications, a description of the algorithmic and code development planned, details about the current SW status (including benchmarks), and the approach for programming and parallelization. A project plan with milestones, a justification of the requested resources, and the requirements for the HW/SW infrastructure (including hardware and software tools) also need to be detailed. Projects can have a duration of up to two years. Each PI can only submit one production proposal per cutoff date, and proposals that have already been submitted to other HPC programs are inadmissible. The PI must hold a postdoc position with an academic institution. Only professors can request multi-year projects. The review process consists of a technical feasibility analysis, a scientific review, and a final decision by a scientific committee.
Reporting requirements	PIs of production projects accepted for longer than one year (with a max duration of 36 months) need to submit a progress report together with a resource request and justification within the deadline announced below. In the renewal, allocation requests can be increased by up to 20%.
Summary of available resources	
Compute resources	Development projects use between 10,000 and 250,000 node hours. Resources are granted for one year, starting from the next cutoff date (six months after the one the proposal was handed in for).

6.6.3. CSCS e-Infrastructure Preparatory Projects

Table 24: Access Policy Analysis – CSCS (Preparatory Projects)

Access Policy Analysis – CSCS (Preparatory Projects)	
Obtaining Access	
Process to obtain access	Preparatory projects enable users to port and test their codes before applying for a production project and/or to generate the performance analysis and resource justification to be included in a production project proposal. They are allocated for 3 months with a possible further 3-month extension on (reviewed) request. A preparatory project proposal can be submitted at any time, and it must contain a short description of the research methods, algorithms, and code, the parallelization approach, memory

	requirements, and the research goals achieved after three months to be promoted to a production project. The HPC Experience of the applicant (brief description of the HPC know-how of the PI and the group; need for support in compiling, porting, optimizing …) should also be included. Preparatory Projects will undergo a technical review by CSCS experts within one month of submission.
Summary of available resources	
Compute resources	Preparatory projects can use up to 250 node hours over their three-month duration. Up to three CSCS accounts can be used/created for their execution.

6.7. EPCC e-Infrastructure

The Edinburgh Parallel Computing Centre/EPCC in the UK EPCC operates the largest UK HPC system (currently ARCHER2) and makes it available for scientific use across disciplines. EPCC also supports industrial use of their systems under bespoke agreements.

Table 25: Access Policy Analysis – EPCC (General)

Access Policy Analysis – EPCC (General)	
Obtaining Access	
Targeted users	Compute time on the EPCC HPC systems (mainly Archer 2) can be allocated to UK researchers, using specific allocation mechanisms implemented and funded by the Engineering and Physical Sciences and Natural Environment Research Councils (EPSRC and NERC). Non-UK researchers can only apply as project co-leads from certain countries (Norway at the moment) or under the auspices of the International Institute of Advanced Systems Analysis (IIASA).
Process to obtain access	Project proposals must be submitted, with the level of detail depending on the amount of resources required. Proposals are reviewed by EPCC and/or EPSRC, and NERC.
Reporting requirements	EPCC reserves the right to publish information on and results of projects that use their HPC systems. End-users should register the DOIs for their publication in the EPCC service management system (SAFE).
Access Tracks and Modes	
Access tracks	EPCC offers two distinct access tracks, which are relevant to SPECTRUM: Pump-priming and “Access to HPC”. It is also possible for UK researchers applying to an EPSRC research grant to add requests for HPC resources (incl. ARCHER2) during the grant application; such resources can be extended on special request.
Access modalities	Interactive access via SSH, with mandatory multi-factor authentication (MFA).
Summary of available resources	
Compute resources	5860 compute nodes with 2x x86 AMD Rome CPUs.
Data resources	14.5 PBytes of work storage in 4 file systems (using the Lustre PFS).

Access management and security	
Identity management and AAA	EPCC identity management uses two levels: users need an account on the SAFE system, which then enables them to create one or multiple user accounts on ARCHER2 and manage reports and quotas. SAFE and ARCHER2 login accounts are local to EPCC and not federated with other UK centres.
Security methods and processes	SAFE accounts are protected by a password. ARCHER2 login accounts use an SSH key pair protected by a passphrase and a Time-based one-time password (TOTP) as a second factor.
Rules and assurances	
End-user rules and policies	EPCC mandates fair use of resources in general, outlawing immoral, illegal, and malicious activities and copyright infringements, disallowing sharing of SAFE or login credentials or appropriating other persons' accounts, and prohibiting access to data without the owner's permission. Users must be reachable by EPCC via the registered Email address.
Fair use, security, data protection	Provision of resources follows a “best effort” paradigm – no hard guarantees are given as to resource availability, performance, and waiting times. EPCC manages batch queues with different priorities and/or available resource quotas (in space and time) supporting different end-user needs, and the batch scheduling system (Slurm) automates scheduling and allocation. Data is protected by Linux and PFS mechanisms, including access control lists. Encryption of data at rest is not enforced; unencrypted data transfer mechanisms are blocked. Data transfer is supported using encrypted protocols; EPCC supports SSH/SCP, Globus Endpoint and GridFTP, and rclone (to/from Cloud storage providers). Regular end-users are not able to acquire root privileges.
End-user support	EPCC provides a service desk accessible via Email, telephone, and through the SAFE system, which uses an issue tracking system to manage support requests. EPCC also provides a wealth of self-service training material, user guides, and training workshops/seminars.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	EPCC monitors the use of resources (compute nodes and storage) by projects and users and makes the results available through the SAFE system, and through collated, public monthly and quarterly service reports.
Engagement with end-users	Users can interact with EPCC through the Service Desk. User satisfaction and feedback is collected and published in collated form in end-user survey reports.

6.7.1. EPCC e-Infrastructure Pump-Priming Access

Table 26: Access Policy Analysis – EPCC (Pump-Priming Access)

Access Policy Analysis – EPCC (Pump-Priming Access)	
Obtaining Access	
Process to obtain access	The ARCHER2 Pump Priming access track is open to EPSRC researchers for requesting a small amount of compute resources through a light-touch process. The main purpose of this track is to enable researchers to try, test, and scale their code on ARCHER2. Applicants must submit a technical assessment form by Email to the ARCHER2 service desk, which will check the contents and submit it to the EPSRC within 8 working days. The EPSRC will then make a decision within 2 weeks. The assessment form details the software to be used on ARCHER2, the sizes and lengths of compute jobs to be used, rough I/O and data transfer needs, and required system SW and support services.
Summary of available resources	
Compute resources	Pump-priming access provides up to 4000 CPU node hours for a maximum of six months.

6.7.2. EPCC e-Infrastructure “Access to HPC” Access

Table 27: Access Policy Analysis – EPCC (“Access to HPC” Access)

Access Policy Analysis – EPCC (“Access to HPC” Access)	
Obtaining Access	
Process to obtain access	The ARCHER2 “Access to HPC” access track is open to EPSRC researchers for requesting significant amounts of compute resources to carry out or support their research projects. Calls for this access track happen every six months (the current one closes towards the end of May 2025). Applicants must submit a proposal through a Web portal operated by UK Research and Innovation. Proposals can include collaboration with industry. Proposals can only be submitted once. The proposal must clearly describe the “vision” of the project (including relevancy, potential to advance knowledge in its field or across fields, impact on and specific benefits for research/society/economy, contributions to HPC and computational science), detail the approach to be taken (including previous work, feasibility, risks and their mitigation, details of codes and methods), and point out the (scientific and research management) experience, expertise and skills of the proposers. The compute resources must be detailed and justified, detailing why they are essential for the project and ensuring that they are appropriate to achieve the stated objectives.
Summary of available resources	
Compute resources	Proposals must request more than 4000 node hours. The upper limit is stated as “half the total available resources”. Each call allocates on the order of 3400000 node hours, with accepted proposals ranging from 12000 – 1300000 node hours.

6.8. ICSC e-Infrastructure

Table 28: Access Policy Analysis – ICSC

Access Policy Analysis - ICSC (General)	
Obtaining Access	
Targeted users	The access is open to all the participants of the ICSC project, including researchers, the public sector, and SME.
Process to obtain access	The applicants can ask for computing resources using the internal project portal. In the application, the scope of the project and the requested resources (HPC, HTC, cloud) should be clarified. The application then undergoes a technical review to check the feasibility of the project.
Access Tracks and Modes	
Access tracks	One access track. Calls are opened every six months.
Access modalities	Cloud access, batch computing, and interactive HPC services.
Summary of available resources	
Compute resources	Compute and data resources are those included in the ICSC network, including the infrastructure provided by INFN and CINECA.
Data resources	Storage facilities integrated with compute resources across the network, capable of handling Big Data workloads.
Data transfer resources	High-speed research network interconnecting HPC and data facilities for efficient transfer and access.
Access management and security	
Identity management and AAA	The access management is demanded by the resource providers (i.e., INFN and CINECA).
Security methods and processes	Standardized cybersecurity protocols, monitoring, and data protection compliance.
Rules and assurances	
End-user rules and policies	Policies on fair use, data security, and access management.
End-user support	End-user support via helpdesk, documentation, and training programs.
Monitoring, Evaluation, and Evolution	
KPIs monitored by	Usage statistics, user satisfaction, efficiency metrics.

the infrastructure	
Engagement with end-users	Regular interaction through surveys, workshops, and user forums.
Evaluate, improve, and evolve infrastructure	Continual assessment of services, infrastructure upgrades, and user feedback-driven development.

6.9. WLCG e-Infrastructure

Table 29: Access Policy Analysis – WLCG

Access Policy Analysis - WLCG e-Infrastructure	
Obtaining Access	
Targeted users	Researchers working with the LHC/CERN experiments, partner experiments, or for institutions using data from these.
Process to obtain access	Access is granted by virtue of membership of the LHC experiments, or of virtual organisations (VOs) supported by WLCG; there are VOs for each of the four principal LHC experiments (ALICE, Atlas, CMS, LHCb) plus a handful of others. End-users do not need to write or submit specific access proposals to WLCG once they have become experiment or VO members.
Access Tracks and Modes	
Access tracks	Access for developing/testing central experiment codes; access to perform end-user analyses. Typical workflows are simulation, reconstruction, and statistical analyses (also via AI methods).
Access modalities	Interactive via CERN and home institutions. For all the others, the access is Grid-based. Access to storage is certificate/token-based.
Summary of available resources	
Compute resources	The aggregated performance is 12.8 MHS23, measured using the HEPScore23 benchmark[26], which corresponds to roughly 4630 dual-CPU AMD EPYC 7H12 high-end servers used in many HPC systems.
Data resources	Disk and tape resources in the same sites, for a total of: • Disk: 1.2 EB. • Tape: 2.4 EB.
Data transfer resources	WLCG offers the File Transfer Software (FTS) system with Python, CLI, and Web interfaces for large-scale data transfers, which supports the S3, XrootD, and WebDAV protocols. A Web interface and a derived CLI interface are available to Simplify FTS use.

[26] See https://indico.cern.ch/event/1225116/contributions/5519006/attachments/2713539/4712490/GDB-13-09-2023-giordano.pdf or https://w3.hepix.org/benchmarking.html for details

Access management and security	
Identity management and AAA	Indigo-IAM.
Security methods and processes	Indigo-IAM implements the OpenID Connect standard; it works with authentication via SAML or OIDC accounts, X.509 certificates, or SSH keys and supports JSON Web Tokens using OAuth token exchange for authorizing unattended access.
Rules and assurances	
End-user rules and policies	User policies are experiment and site-dependent yet governed by the general WLCG security and acceptable use policies. These policies state that resources may only be used for the stated (research) purposes, that IP ownership and confidentiality agreements are to be respected, that end-users must protect their authentication credentials, and that appropriate acknowledgements or citations, to regain the use of WLCG resources, must be given.
Fair use, security, data protection	See above.
End-user support	Central (WLCG) + Central (Experiment level) + site level.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Six times/year review by the Large Hadron Collider Experiments Committee and the LHC Experiments Resources Review Boards. Monitored parameters are utilization, efficiency, availability, and reliability.
Engagement with end-users	Regular workshops (e.g., WLCG collaboration meetings) and user feedback channels, central via WLCG or via the experiments.
Evaluate, improve, and evolve infrastructure	Plans reach up to 2041 at least and are aligned with the LHC operational cycles. Short-term evolution plans can be found here.

6.10. NIKHEF e-Infrastructure

Table 30: Access Policy Analysis – NIKHEF

Access Policy Analysis - NIKHEF e-Infrastructure	
Obtaining Access	
Targeted users	Researchers and collaborations involved in high-energy physics and related domains, particularly those requiring HTC capabilities.
Process to obtain access	Interested users follow a defined process via NIKHEF's facility documentation portal; usually coordinated through their research collaborations or (for Grid resources) via the Dutch National Computing call[27].
Access Tracks and Modes	
Access tracks	Tracks may involve national/international research collaborations, project-based access, or institutional partnerships.
Access modalities	Access is typically mediated through virtual organization, federated identity systems, and allocation agreements.
Summary of available resources	
Compute resources	HTC clusters and grid computing resources, tightly integrated with WLCG (Worldwide LHC Computing Grid); support for large-scale batch jobs.
Data resources	Storage services include disk, tailored for large-volume data (e.g., from LHC experiments) accessible through dCache, and integrated with data management systems (e.g. Rucio and OSDF Cache)
Data transfer resources	High-throughput data transfer capabilities, typically through gridFTP, WebDAV, and other robust protocols used by the WLCG infrastructure.
Access management and security	
Identity management and AAA	Federated identity (e.g., via eduGAIN and EGI Check-in).
Security methods and processes	Adheres to EGI and WLCG security policies, including incident response teams and security coordination.
Rules and assurances	
End-user rules and policies	Policies govern acceptable use, data handling, and security compliance; users must adhere to collaboration-specific and NIKHEF-wide rules.
Fair use, security, data protection	Enforced via policy and technical controls; includes quotas, auditing, and data protection measures consistent with EU regulations.
End-user support	Support is provided via user documentation, helpdesk services, and through collaboration-specific support channels.

[27] See https://www.nwo.nl/en/calls/computing-time-on-national-computing-facilities for details.

Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Usage metrics, job success rates, data throughput, system availability, and user satisfaction are tracked.
Engagement with end-users	Involves regular feedback loops through collaboration meetings and user forums.
Evaluate, improve, and evolve infrastructure	Continuous improvement through performance metrics analysis, stakeholder feedback, and technology upgrades in alignment with scientific needs.

6.11. EGI HTC-oriented e-Infrastructure

Table 31: Access Policy Analysis – EGI HTC

Access Policy Analysis - EGI HTC	
Obtaining Access	
Targeted users	Individual Researchers, Research Infrastructures and Projects & Research Communities, and Commercial Research entities - via **regular access call**. SME's via **EGI Digital Innovation Hub**. Public authorities & Policy Makers - via sponsored access.
Process to obtain access	Via **EGI Access Call** • Proposals for scientific use cases will go through a review by independent experts from the EGI Federation. As part of this review, the main objectives, the overall level of maturity, and the feasibility of the use case will be evaluated. Applicants will be notified of the outcome of the evaluation within three weeks of the submission date. • If the outcome is positive, EGI will run a match-making process aiming to identify the most suitable service providers and allocate the required resources. These resources can be national or international, depending on the coverage of the request and its expected impact. • Experts from EGI will monitor the integration plans of the use case regularly, providing assistance and support to address technical issues and disseminate outstanding results.
Access Tracks and Modes	
Access tracks	Trial Access, Custom Access, and consultancy • Sponsored Access (European Commission funded or in-kind support on the national level) - free at the point of use, supported by EC project or national funding. • Paid Access - as defined by specific contract/**default SLA**. • Long-term Partnership - organizations with an established legal entity **to become part of EGI Federation**.
Access modalities	Fully managed batch processing in the EGI Infrastructure, enabling researchers and scientific communities to easily and efficiently run hundreds of thousands of batch computing jobs on the EGI Infrastructure - combining several EGI services (EGI HTC, Online storage, and workload manager).
Summary of available resources	
Compute	The service is provided by a distributed network of computing centres, giving you access to a

resources	massive amount of computing power via a standard interface and membership of a virtual organisation. With over 1 million cores of installed capacity, EGI can support over 1.6 million computing jobs per day. The key components of the EGI High Throughput Compute architecture are: Data Transfer service (FTS), the Online Storage services, and Computing Elements (CEs), which are compute resources made available through GRID interfaces. The most common implementations of CEs in the EGI infrastructure are HTCondor-CE and ARC-CE.
Data resources	Over 580 PB online storage capacity is offered by HTC and Cloud storage providers under EGI Online Storage service which includes Block Storage for durable data that doesn't need to be shared beside a single virtual machine (VM) that enables to attach the storage to VMs as volumes, making it easy to access your data whenever needed. HTC and Clud also embed: • Grid Storage: for storing and accessing large amounts of data quickly without having to worry about latency or downtime, ensuring that data can be accessed when needed. • Object Storage for cloud-native applications, archiving, or when data is shared between different VMs or multiple steps of processing workflows, ensuring data structures are efficiently stored and can be easily accessed.
Data transfer resources	EGI Data Transfer service allows scientists to move any type of data files asynchronously from one storage to another. The service includes dedicated interfaces to display statistics of ongoing transfers and manage storage resource parameters. The service components are: • FTS3 Server: The service is responsible for the asynchronous execution of the file transfer, checksumming, and retries in case of errors. • FTS3 REST: The RESTFul server which is contacted by clients via REST APIs, CLI and Python bindings. • FTS3 Monitoring: A Web interface to monitor transfer activity and server parameters. EGI DataHub service is a solution based on the Onedata technology that allows bringing data close to computing to exploit it efficiently. • Access data via GUI, POSIX, or REST API. • Support for CEPH, S3, GlusterFS, POSIX, and more.
Access management and security	
Identity management and AAA	Access based mainly on X509. EGI Check-in service proxy service that operates as a central hub to connect federated Identity Providers (IdPs) with EGI service providers; EGI Check-In is based on the AARC Blueprint Architecture. Check-in is compliant with eguGAIN, REFEDS RnS, and Sirtfi policies. Translates SAML 2.0, OpenID Connect, OAuth 2.0, and X.509 credentials.
Security methods and processes	User/password under EGI Check-in. Further security is possible under EGI Secrets Store service.
Rules and assurances	
End-user rules and policies	Access requires acceptance of the EGI Terms of Use also available from here Acceptable Use Policy (AUP) and Conditions of the 'EGI Applications on Demand Service'. General service conditions under the default EGI Service Level Agreement.
Fair use, security, data protection	Under the EGI Privacy notice.
End-user	Full user documentation of EGI Services.

support	User documentation and webinars specific to the EGI HTC service. Specific support and consultancy are also available as part of the EGI HTC service.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	HTC CPU Hours, Number of users, User communities. Under EGI Monitoring based on ARGO, to monitor the availability and reliability of the sites.
Engagement with end-users	EGI provides the central user support via EGI Helpdesk and coordinates support activities of EGI providers, who offer user support for the services/resources they contribute to the EGI ecosystem.

6.12. SRCNet SKA HTC/Data-oriented e-Infrastructure

SRCNet is not operational at the time of writing; hence, a finalized set of access policies has not been completely agreed upon. Table 32 below reports the anticipated access policies for the SRCNet SKA HTC/Data-oriented e-Infrastructure.

Table 32: Access Policy Analysis – SRCNet SKA HTC/Data-oriented e-Infrastructure

Access Policy Analysis - SRCNet SKA HTC/Data-oriented e-Infrastructure	
Obtaining Access	
Targeted users	Since SRCNet is still under construction, it is anticipated that researchers and students will apply for telescope observing and computing resource time at SRCNet through a proposal review process. These individuals and teams can be from anywhere in the world; however, access time will be proportional to the member state's contribution to the observatory. It is not yet clear what the reporting requirements will be. Successful applicants will be given access to the computing resources of SRCNet. It is also anticipated that those wishing to access and analyse archival data will be able to do so through the SRCNet. Any qualified registered scientific user or member of their team will be permitted to access the data through the interface. This network of heterogeneous compute nodes will be hosted by the SKA member countries and will receive data from the HPC centers located at the two telescope data processing centers located in Perth and Cape Town. Access will be provided to qualified scientists from the astronomical community who have either been awarded time to conduct research using the telescopes or those interested in accessing archival data. It is not yet decided whether there will be reporting requirements.
Process to obtain access	The SRCNet support staff, including SKAO, will likely determine who is granted access. However, the process is not yet set up.
Access Tracks and Modes	
Access tracks	SKA Regional Centers will provide access to regular scientific users from the community, as well as less restricted access granted to internal users employed as qualified staff by SKAO or the SRCNet nodes.
Access modalities	There are no specific access modalities foreseen at this moment. Resource management will be carried out by SRCNet to make the most efficient use of compute resources, and resource

	allocation will start from the project approval stage so that appropriate staging of the data will be done. Resource allocations may be made to users wishing to exploit exclusively public data (i.e., without an approved SKA project), following an approved SRCNet processing proposal.
Summary of available resources	
Compute resources	Analysis through the SRCs will be carried out over different regional, national, or supranational compute infrastructures. Computing requirements are varied and include processing of very large datasets for generation of advanced data products or running of pipelines (to create project-level data products or test appropriate workflows and parameters of observatory data products), as well as user-driven batch processing or interactive processing via, for example, a notebook. It will also be possible for users to test new workflows on small datasets and receive immediate results, thus enabling workflows to be refined prior to submitting large batch jobs.
Data resources	SRCNet data will be stored in a data lake across a heterogeneous set of nodes. This lake will be centrally managed but distributed and federated at the storage elements level. These data will be stored across nodes hosted by various member countries of the SKA project. Each of these nodes contributes storage capacity. Some countries will contribute their storage capacity from a single data center, while other countries will have a distributed data center network. The total computing requirements at full operations (2028+) are expected to be 25 PFlop/s with ~1 EByte of storage. Users will be provided with a persistent, personal file system (for example, Portable Operating System Interface (POSIX)), to which they can upload and download files at will, including from the archive, to within the per-project resource allocation. These files are available for further processing and visualisation, as well as to hold code and additional uploaded data for analysis. Processing logs will be stored along with the data for all processing operations and will record information on software and resources used and any other parameters to ensure reproducibility of new data product generation. The SRCNet data lake will be based on Rucio. Each node implements its own protocol. WebDAV is used in the Spanish SRC. XrootD is used in other nodes like the Sweden SRC, the Swiss SRC, and the UK SRC. It is currently anticipated that the SKAO will produce ~200PB/year of data products during steady state operations; however, there are significant uncertainties on these estimates..
Data transfer resources	This is still to be done, but the current assumption is that data will be transported from the two telescope sites to the SRCNet nodes using 100 Gbp/s links. Each SRCNet node site will initially require at least 1 Gbps of upload and download speed for each node, but the network requirements for the full system have not yet been defined. WebDAV and XrootD are amongst the mechanisms to be used for data ingestion. The SRCNet repository will be centrally managed but distributed and federated at the storage elements level.
Access management and security	
Identity management and AAA	SRCNet will use the SKA-IAM (Identity and Access Manager) for identity management, authentication, and authorization within SKAO[28].
Security methods and processes	All SRCNet nodes must have an authentication service entry, either a global service that connects to federated identity providers or a local service implementation connecting to the global SRCNet authentication declared services. All SRCNet nodes must integrate consistent authorisation modules to ensure a secure access system, following, e.g., the SKA data access policies.

[28] See https://stfc.github.io/IAM-Docs/ for details.

Rules and assurances	
End-user rules and policies	The SRCNet organisational structure will likely include a coordination committee who are responsible for reviewing data management policies to ensure appropriate quality of service (trading off performance and cost/capacity needs). Users are supported by staff working at the SKAO and SRCNet nodes.
Fair use, security, data protection	It is anticipated that at least two copies of each dataset delivered to the SRCNet will be stored across the nodes. The observatory data products and advanced data products produced by scientific users of the SRCNet are expected to be archived indefinitely.
End-user support	User support will be provided once the system is operational.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	There will be a committee composed of representatives from the SKAO and each member state who will be responsible for overseeing the operations of the SRCNet and tracking performance against KPIs. They will be supported by an SRC Operations Group, which will implement policies to ensure the quality of service to the end user. There will be an operations function capable of monitoring the availability and performance of services provided by each SRC node, reporting service faults, and tracking overall performance against pledged resources. Users will be able to request help by filing tickets through a Helpdesk operated by SKAO and SRCNet staff.
Engagement with end-users	A help desk will be available to monitor user feedback. It is anticipated that the SRC Operations Group (SOG) will also ensure the quality of service by monitoring the ability of each site storage element to accept data products from SKA sites; monitoring the ability of appropriate sites to accept batch processing jobs and to provide interactive sessions for users; monitoring the network link availability and performance (e.g., through continuous monitoring of links in use and liaison with network providers), and monitoring user support, tracking Helpdesk metrics and other feedback.
Evaluate, improve, and evolve infrastructure	It is not clear yet what reporting requirements will exist. There is the assumption that the user should be able to provide improvement suggestions through software tickets.

6.13. LOFAR - Central Processing (CEP)

Table 33: Access Policy Analysis – LOFAR - Central Processing (CEP)

Access Policy Analysis - LOFAR - Central Processing (CEP)	
Obtaining Access	
Targeted users	No individual users are specifically addressed. LOFAR CEP is for researchers and institutions that are involved in the LOFAR telescope.
Process to obtain access	Users are generally not granted access, except for special-purpose experiments arranged on an individual basis..
Access Tracks and Modes	
Access tracks	LOFAR CEP does not support a variety of different access tracks.

Access modalities	Access via SSH.
Summary of available resources	
Compute resources	**Central Processing (CEP) Facility**: Located at the University of Groningen's Centre for Information Technology (CIT), the CEP serves as the primary processing center for LOFAR data. It handles real-time data streams from LOFAR stations, performing initial processing tasks such as correlation and beamforming. **COBALT Correlator**: A GPU-based correlator and beamformer that replaced the earlier IBM Blue Gene/P system. COBALT processes the digitized signals from the LOFAR stations, enabling high-throughput data processing essential for LOFAR's operations.
Data resources	**CEP4 Storage Cluster**: Stores correlated output temporarily for post-processing and archiving. **COBALT (GPU Correlator)**: Real-time signal processing; outputs visibilities and beamformed data[29].
Data transfer resources	LOFAR employs a high-speed fiber network infrastructure to transmit data from its stations to the CEP facility in Groningen. This network ensures the timely delivery of large volumes of data for processing[30].
Access management and security	
Identity management and AAA	LOFAR utilizes an integrated Authentication and Authorization Infrastructure (AAI) to manage user identities and access rights across its distributed infrastructure.
Security methods and processes	Implementation of secure data transfer protocols and regular audits to ensure data integrity and confidentiality.
Rules and assurances	
End-user rules and policies	Users must adhere to LOFAR's data usage policies, ensuring proper citation and acknowledgment in publications.
Fair use, security, data protection	LOFAR enforces fair use policies to prevent misuse of resources and implements robust security measures to protect data.
End-user support	LOFAR provides comprehensive support through documentation, helpdesks, and user forums to assist researchers in data access and analysis.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	System uptime, data processing throughput, storage utilization, and user access metrics.
Engagement with end-users	Regular workshops, surveys, and feedback mechanisms to understand user needs and improve services.

[29] See https://science.astron.nl/telescopes/lofar/access-to-lofar-data/data-products/ for details.

[30] See https://arxiv.org/pdf/1305.3550 for details.

Evaluate, improve, and evolve infrastructure	Continuous assessment of technological advancements to upgrade infrastructure components, ensuring scalability and efficiency.

6.14. WLCG Data-Oriented e-Infrastructure

The information for the data-oriented side of WLCG is integrated in the table in subsection 6.9.

6.15. EBRAINS e-Infrastructure

Table 34: Access Policy Analysis – EBRAINS

Access Policy Analysis - EBRAINS	
Obtaining Access	
Targeted users	Researchers across the fields of neuroscience, brain health, and brain-related technologies.
Process to obtain access	EBRAINS provides access to high-performance computing, Cloud, and storage services, some of which are part of the Fenix research infrastructure. It also provides access to neuromorphic computing systems and the EBRAINS Collaboratory, which includes a Jupyter Lab instance where scientists can use the EBRAINS software tools with minimal overhead. The process to obtain access to HPC, Cloud and storage resources at the participating HPC centres (currently JSC, CINECA and CEA) is through the EBRAINS 2.0 project or through national or Fenix access calls, with a subsequent evaluation of the applications, which follows the peer review principles established by PRACE, based on the technical and scientific assessments. EBRAINS users have immediate access to the EBRAINS Collaboratory, including its Jupyter Lab instance. In order to access neuromorphic systems, users need to send a request for access, which is evaluated and granted directly by the hosting site (the University of Manchester for SpiNNaker and the University of Heidelberg for BrainScales).
Access Tracks and Modes	
Access tracks	The access can be used for a variety of use cases, from developing and hosting platform services on Virtual Machines (suitable for deploying platforms, for example, HBP Collaboratory, image services or neuromorphic computing front-end services), modelling, simulation, and data analysis tasks, Archival and Active Data Repositories for storing, sharing, and accessing data, as well as for training and education.
Access modalities	Scientific services can be accessed by EBRAINS users via the EBRAINS IAM. These services may use HPC, cloud, and storage resources via service accounts at different computing sites (at the moment JSC, CINECA, or CEA).
Summary of available resources	
Compute resources	To date, JSC, CINECA, and CEA are providing Cloud computing and storage resources, and JSC is also providing access to some HPC resources on the JUSUF cluster. It is currently difficult to exactly quantify the resources available to EBRAINS, as there is some flexibility in how much the centres provide to EBRAINS; nevertheless, most EBRAINS services are running on the JSC Cloud (Germany). Available through the latest Fenix call are the following resources: • CEA (France): Cluster OpenStack / 320 vCPUs (no vGPUs) with a total of 640 GB RAM. • JSC (Germany): JSC Cloud / 1500 vCPUs, 3000 GB Memory.

	Other sites are currently testing and evaluating the service and will make it available for users in the near future. Science services provide small amounts of HPC resources for specific-purpose jobs via service accounts. The EBRAINS Collaboratory via its Lab service provides a JupyterLab environment for notebooks with official releases of EBRAINS tools pre-installed. Images available are of up to 8GB of RAM and 2 VCPUs.
Data resources	Yet again, precisely quantifying the resources available to EBRAINS is currently challenging, as the level of support provided by the centres can vary; nevertheless, the data resources made available through the latest Fenix call are as follows: • CEA (France): 500TB / Not possible to access the S3 storage from the OpenStack cluster, but the storage can be used for back-up using the FENIX FTS service. • JSC (Germany): Virtual machines spawned within the JSC Cloud can have access to four different kinds of storage: Root File system (10s of GB) / NVMe (894 GB) / Dedicated Volume (100s of GB - 10s of TB) / DATA (100s of GB - 10s of TB).
Data transfer resources	FENIX FTS service.
Access management and security	
Identity management and AAA	HPC site-specific authentication and authorization services per project – FENIX. For science services, the EBRAINS IAM is used to manage identity and access rights across the ecosystem.
Security methods and processes	Passwords and password updates. The HPC sites require 2F authentication for most services.
Rules and assurances	
End-user rules and policies	Commitment to comply with the EBRAINS' current version of the Terms and Policies. Depending on the case, access also requires acceptance of the CEA, CINECA, or JSC Terms of Use and Acceptable Use Policy.
Fair use, security, data protection	Under the EBRAINS' Terms and Policies.
End-user support	EBRAINS provides a multi-tier support service, from the frontline to high-level scientific integration support. Questions related to computing and storage services are forwarded to mentors from each center who provide end-user support.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Project users, CPU Hours, User communities, dissemination material: number of publications/posters/papers and talks for granted projects.
Engagement with end-users	Engage with all end-users, sharing updates and soliciting feedback related to the use-cases and end-user requirements of the HBP/EBRAINS and the neuroscience communities. Additionally, the EBRAINS RI has an extensive offer of Education and Training activities which aim at developing interdisciplinary skills among its present and potential users.
Evaluate, improve, and evolve infrastructure	Larger amounts of computational resources will be required in the near future to fulfill the objectives of the HBP/EBRAINS and the neuroscience communities; for this reason, the end-user feedback is taken seriously, and other sites are currently testing and evaluating for further services.

6.16. LOFAR Long-Term Archive e-Infrastructure

Table 35: Access Policy Analysis – LOFAR Long-Term Archive e-Infrastructure

Access Policy Analysis - LOFAR Long-Term Archive e-Infrastructure	
Obtaining Access	
Targeted users	The astronomical community is interested in using LOFAR data. Scientists who are successful in observing proposals to LOFAR will use the Long-Term Archive (LTA) to retrieve data products resulting from their observations. Any other may apply for an account, which will enable them to download data that is not covered by a proprietary period.
Process to obtain access	The current system involves an online registration followed by contact with the Helpdesk operated at ASTRON. Within the near future, it is expected to move to a federated AAI system based on SURF Research Access Management.
Access Tracks and Modes	
Access tracks	LOFAR does not support a variety of different access tracks.
Access modalities	Users select data either interactively (through the website) or programmatically (through an API), and “stage” it for later download using HTTP.
Summary of available resources	
Compute resources	Compute resources are pledged to the LTA by the LOFAR ERIC member countries. As such, both the amount of resources available and the nature of those resources vary with time. The LOFAR2.0 operational era will run for five years from 2026. During this time, it is estimated that 1.5 billion CPU hours of computing will be required. This will be sourced from across existing (SURF, FZJ, PSNC) and new contributions from ERIC member countries.
Data resources	As with compute, data is stored across a distributed, heterogeneous network of data centres provided by the LOFAR ERIC members. During the five years from 2026, storage will gradually increase to a maximum of about 120 PB. As intermediate products are processed to their final form, this will eventually be reduced to 90 PB, which will be archived for the long term.
Data transfer resources	LTA sites are expected to support 10 Gbit/s bandwidth: • Between LOFAR central processing and the LTA site. • Between storage and compute systems within a given site. • To external users, including both end-users and other LTA sites.
Access management and security	
Identity management and AAA	Currently based on a bespoke AAI managed by ASTRON. Expect to move to a federated AAI based on SURF SRAM in the near future.
Security methods and processes	LOFAR data is not generally confidential. Basic access control is applied to proprietary data products (normally, data within one year of the date of observation).

Rules and assurances	
End-user rules and policies	No specific rules are documented. Users are expected to comply with the data policy (described below).
Fair use, security, data protection	LOFAR data is made available under the terms of the Science Data Policy of LOFAR ERIC, which specifies conditions of use.
End-user support	Users have access to a help desk managed by ASTRON.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Typical KPIs include: • Size/growth of LTA content. • Volume of content accessed in the period. • Number of data objects accessed in the period. • Access broken down by country of affiliation. (Note that these are not all publicly available).
Engagement with end-users	• Annual LOFAR Data Schools provide training and the opportunity for user feedback. • The LOFAR Users Committee represents the needs of the community to LOFAR ERIC. • The annual LOFAR Family Meeting provides a forum for the whole LOFAR community to come together and includes a session dedicated to infrastructure and telescope news, updates, and discussion.
Evaluate, improve, and evolve infrastructure	The services provided through the LOFAR LTA are under continuous evaluation and evolution, based on, amongst others: • New telescope functionality (e.g. the upgrade to LOFAR2.0) which requires new data centre functionality and/or produces data that is larger of a different type. • The availability of infrastructure provided by the LOFAR ERIC partners. • Community events and feedback. • Collaboration with and learning from peer facilities, in particular the SKA Observatory and associated Regional Centre Network.

6.17. ErUM-Data-Hub e-Infrastructure

Table 36: Access Policy Analysis – ErUM-Data-Hub

Access Policy Analysis	
Obtaining Access	
Targeted users	Researchers affiliated with the eight ErUM communities in Germany, encompassing fields such as particle physics, astroparticle physics, hadron and nuclear physics, synchrotron radiation, and neutron research.
Process to obtain access	Access is typically coordinated through institutional affiliations and collaborative projects. Researchers often gain access via their participation in specific experiments or collaborations within the ErUM framework.
Access Tracks and Modes	
Access tracks	Access is structured through a federated model, where researchers utilize resources provided by their home institutions, national computing centers, or through collaborative agreements within the ErUM communities.

Access modalities	Access modalities include direct institutional access, federated access through national infrastructures, and participation in collaborative projects that provide shared resources.
Summary of available resources	
Compute resources	The ErUM-Data-Hub collaborates with several High-Performance Computing centers to provide computational resources for research, such as those at JSC (JUWELS) or KIT (TOPAS), providing substantial computing power for data processing and analysis.
Data resources	• Data Storage: Large-scale storage solutions are provided by the collaborating HPC centers, ensuring secure and efficient data management. • Data Management: The hub supports comprehensive data management practices, aligning with FAIR (Findable, Accessible, Interoperable, Reusable) principles to enhance data usability.
Data transfer resources	• Networking: High-bandwidth networks, such as the DFN (Deutsches Forschungsnetz), facilitate efficient data transfer between institutions and computing centers. • Data Transfer Tools: Utilization of tools like GridFTP and XRootD for secure and efficient data movement across federated infrastructures. • Federated Infrastructures: The DIG-UM initiative promotes the development of federated digital infrastructures, enabling seamless data sharing across different platforms and institutions.
Access management and security	
Identity management and AAA	Implementation of federated identity management systems, such as UmbrellaID, to streamline user authentication and authorization across different platforms and institutions.
Security methods and processes	Adherence to standardized security protocols to protect data integrity and confidentiality, including regular audits and compliance checks.
Rules and assurances	
End-user rules and policies	Users are expected to comply with institutional and national guidelines regarding data usage, sharing, and publication.
Fair use, security, data protection	Compliance with data protection regulations, including GDPR, to ensure the privacy and security of sensitive information.
End-user support	Provision of user support through dedicated helpdesks, documentation, and training workshops to assist researchers in utilizing the infrastructure effectively.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Tracking the usage of computing and data resources. Assessing user feedback to improve services.
Engagement with end-users	Regular surveys and feedback mechanisms are in place to gather input from users, ensuring that the infrastructure evolves to meet the changing needs of the research community.

Evaluate, improve, and evolve infrastructure	Ongoing development and integration of new technologies and methodologies to enhance the efficiency, scalability, and user-friendliness of the e-Infrastructure.

6.18. PUNCH4NFDI e-Infrastructure – Data/AI-Access

Table 37: Access Policy Analysis – PUNCH4NDFI

Access Policy Analysis -PUNCH4NFDI (Data/EI Access)	
Obtaining Access	
Targeted users	Researchers in astroparticle, particle, hadron, and nuclear physics domains across German institutions affiliated with PUNCH4NFDI.
Process to obtain access	Specific procedures are not fully detailed, but access is likely mediated through project or institutional affiliation; more structured access mechanisms are under development.
Access Tracks and Modes	
Access tracks	Differentiated by use-case types (e.g., short-term access for small analyses, long-term project-based access).
Access modalities	Likely includes GUI, APIs, and CLI for different services; actual methods are not exhaustively specified.
Summary of available resources	
Compute resources	Federated compute resources, including high-performance computing (HPC) centers and institutional clusters.
Data resources	Distributed data storage across partner sites; large datasets from physics experiments and simulations.
Data transfer resources	Use of standard data transfer tools and protocols (e.g., GridFTP, HTTP); optimized for large scientific datasets.
Access management and security	
Identity management and AAA	Federated identity system; utilizes mechanisms like eduGAIN and Helmholtz AAI for authentication and authorization.
Security methods and processes	Based on best practices for data protection, includes secure access and auditing measures.
Rules and assurances	
End-user rules and policies	Users are expected to follow institutional and project-level policies regarding data use and access.
Fair use,	Emphasis on responsible use; policies to ensure data protection and infrastructure integrity.

security, data protection	
End-user support	Support is provided via documentation and user helpdesks; the level of support may vary depending on the service provider.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Resource usage, access logs, and user satisfaction; specifics may vary by component or partner site
Engagement with end-users	Regular feedback collection via surveys, workshops, or forums; collaboration with user communities.
Evaluate, improve, and evolve infrastructure	Continuous development based on user feedback and performance monitoring.

6.19. Copernicus Data-oriented e-Infrastructure

Table 38: Access Policy Analysis – Copernicus

Access Policy Analysis - Copernicus Data-oriented e-Infrastructure	
Obtaining Access	
Targeted users	Copernicus Data Space Ecosystem (CDSE) is designed to support a wide range of users, including: • Researchers and Scientists: Conducting studies in Earth observation and related fields. • Policy Makers and Public Authorities: Utilizing data for informed decision-making. • Commercial Entities and SMEs: Developing applications and services based on Copernicus data. • General Public: Accessing data for educational and informational purposes.
Process to obtain access	Accessing Copernicus data typically involves: • Registration: Creating an account on the relevant Copernicus portal (e.g., CDSE, WEkEO). • Data Discovery: Utilizing search and visualization tools to find relevant data products. • Data Access: Downloading data or processing it directly within the provided cloud environments.
Access Tracks and Modes	
Access tracks	Copernicus provides multiple access tracks to cater to different user requirements: • Open Access: Most Copernicus data and services are freely and openly accessible to all users worldwide. • Thematic Services: Specialized services offer data products tailored to specific themes such as atmosphere, marine, land, climate, and emergency management.
Access modalities	Users can access Copernicus data through various modalities: • Web Portals: User-friendly interfaces for data discovery, visualization, and download.

	• APIs: Programmatic access to data for integration into applications and services. • Cloud Processing Environments: Platforms like WEkEO and CREODIAS provide environments for processing data directly in the cloud.
Summary of available resources	
Compute resources	Main HPC Systems: Copernicus leverages a combination of cloud-based and high-performance computing (HPC) resources to process and analyze vast amounts of Earth observation data. Key components include: • Copernicus Data Space Ecosystem (CDSE): Provides Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS) environments, enabling users to process hosted data efficiently within a cloud computing ecosystem. • WEkEO Platform: Offers access to Copernicus and Sentinel data, allowing users to process data in a cloud infrastructure with capabilities such as Jupyter Notebooks, up to 128 virtual CPUs, 4TB RAM, and 40TB storage. • CREODIAS Platform: A commercial component of the Data Space, providing cloud computing services for satellite data processing, including compute, storage, backup, and data-related services.
Data resources	Copernicus offers a vast array of Earth observation data through its Sentinel missions and contributing missions. Key data resources include: • Copernicus Data Space Ecosystem (CDSE): Provides immediate access to large amounts of open and free Earth observation data from the Copernicus Sentinel satellites, including both new and historical Sentinel images, as well as data from contributing missions. • WEkEO Platform: Offers access to reference products for Sentinel 1, 2, 3, 5P, and 6, as well as data from Copernicus Services and other missions like EUMETSAT meteorological satellites. • Copernicus Services Portals: Specialized portals for Atmosphere Monitoring (CAMS), Marine Monitoring (CMEMS), Land Monitoring (CLMS), Climate Change (C3S), and Emergency Management (EMS) provide access to thematic data products.
Data transfer resources	Data transfer within the Copernicus infrastructure is facilitated through: • Cloud-Based Access: Users can access and process data directly within cloud environments, reducing the need for data downloads and enabling efficient data handling. • Web Portals: Various web portals allow users to discover, visualize, and download data products, supporting diverse access mechanisms tailored to user needs.
Access management and security	
Identity management and AAA	Copernicus employs a unified user management system, providing access to various services through a common identity management framework. This ensures streamlined AAA processes across platforms.
Security methods and processes	Security within the Copernicus infrastructure includes: • Secure Authentication: Ensuring that only authorized users can access data and services. • Data Protection Measures: Implementing protocols to safeguard data integrity and confidentiality. • Compliance with Regulations: Adhering to data protection and privacy regulations to protect user information.
Rules and assurances	
End-user rules and policies	Copernicus data is provided under an open and free data policy, allowing users to access, use, and share data without restrictions. Users are expected to acknowledge the source of the data in their applications and publications.

Fair use, security, data protection	While data is freely available, users are expected to adhere to fair use principles, ensuring that their activities do not negatively impact the infrastructure or other users. Security and data protection measures are in place to safeguard both the data and the users.
End-user support	Copernicus provides comprehensive support to users through: • Help Desks and Support Centers: Offering assistance with data access, processing, and application development. • Documentation and Tutorials: Providing guides and resources to help users effectively utilize Copernicus data and services. • Community Forums: Facilitating user engagement and knowledge sharing.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Copernicus monitors data availability, system uptime, user engagement, and service responsiveness.
Engagement with end-users	Regular user satisfaction surveys and feedback mechanisms are employed to gather insights and improve services.
Evaluate, improve, and evolve infrastructure	Feedback from users, including satisfaction surveys and usage data, plays a central role in identifying areas for improvement. This ongoing evaluation enables Copernicus to adapt and upgrade its systems, enhancing data accessibility, processing capabilities, and overall user experience.

6.20. EGI Federation e-Infrastructure

Table 39: Access Policy Analysis – EGI Federation

Access Policy Analysis - EGI Federation	
Obtaining Access	
Targeted users	Individual Researchers, Research Infrastructures and Projects & Research Communities, and Commercial Research entities - via regular access call. SME's via EGI Digital Innovation Hub. Public authorities & Policy Makers - via sponsored access.
Process to obtain access	Via EGI Access Call • Proposals for scientific use cases will go through a review by independent experts from the EGI Federation. As part of this review, the main objectives, the overall level of maturity, and the feasibility of the use case will be evaluated. Applicants will be notified of the outcome of the evaluation within three weeks of the submission date. • If the outcome is positive, EGI will run a match-making process aiming to identify the most suitable service providers and allocate the required resources. These resources can be national or international, depending on the coverage of the request and its expected impact. • Experts from EGI will monitor the integration plans of the use case regularly, providing assistance and support to address technical issues and disseminate outstanding results.
Access Tracks and Modes	
Access tracks	Trial Access, Custom Access, and consultancy • Sponsored Access (European Commission funded or in-kind support on the national level) - free at the point of use, supported by EC project or national funding. • Paid Access - as defined by specific contract / default SLA.

	• Long term Partnership - organizations with an established legal entity to become part of EGI Federation.
Access modalities	The EGI Cloud Compute (FedCloud) service offers a multi-cloud IaaS federation that brings together research clouds as a scalable computing platform for data and/or compute-driven applications and services for research and science. The following usage modalities are usually used; • Service hosting: The EGI Federated Cloud can be used to host any IT service, such as web servers, databases, etc. Cloud features, as elasticity, can help users provide better performance and reliable services. Example: NBIS Web Services, Peachnote analysis platform. • Compute and data-intensive applications: for those applications needing a considerable amount of resources in terms of computation and/or memory and/or intensive I/O. Ad-hoc computing environments can be created in the EGI cloud providers to satisfy extremely intensive HW resource requirements. Example: VERCE platform, The Genetics of Salmonella Infections, The Chipster Platform. • Datasets repository: the EGI Cloud can be used to store and manage large datasets, exploiting the large amount of disk storage available in the Federation. Example: OBSEA. • Disposable and testing environments: environments for training or testing new developments. Example: Training infrastructure. • Interactive computing, including Data analytics and visualisation with Jupyter Notebooks (EGI Notebooks).
Summary of available resources	
Compute resources	Trial access included up to 4 vCPU cores, 8GB of RAM, and 100GB of block storage.
Data resources	Over 580 PB online storage capacity is offered by HTC and Cloud storage providers under EGI Online Storage service which includes Block Storage for durable data that doesn't need to be shared beside a single virtual machine (VM) that enables to attach the storage to VMs as volumes, making it easy to access your data whenever needed. HTC and Cloud also embed: • Grid Storage: for storing and accessing large amounts of data quickly without having to worry about latency or downtime, ensuring that data can be accessed when needed. • Object Storage for cloud-native applications, archiving, or when data is shared between different VMs or multiple steps of processing workflows, ensuring data structures are efficiently stored and can be easily accessed.
Data transfer resources	EGI Data Transfer service allows scientists to move any type of data files asynchronously from one storage to another. The service includes dedicated interfaces to display statistics of ongoing transfers and manage storage resource parameters. The service components are: • FTS3 Server: The service is responsible for the asynchronous execution of the file transfer, checksumming, and retries in case of errors. • FTS3 REST: The RESTFul server which is contacted by clients via REST APIs, CLI and Python bindings. • FTS3 Monitoring: A Web interface to monitor transfer activity and server parameters. EGI DataHub service is a solution based on the Onedata technology that allows bringing data close to computing to exploit it efficiently. • Access data via GUI, POSIX, or REST API. • Support for CEPH, S3, GlusterFS, POSIX, and more.
Access management and security	
Identity management and AAA	Under EGI Check-in service, a proxy service that operates as a central hub to connect federated Identity Providers (IdPs) with EGI service providers; EGI Check-In is based on the AARC Blueprint Architecture. Check-in is compliant with eguGAIN, REFEDS RnS, and

	Sirtfi policies. Translates SAML 2.0, OpenID Connect, OAuth 2.0, and X.509 credentials.
Security methods and processes	User/password under EGI Check-in. Further security is possible under the EGI Secrets Store service.
Rules and assurances	
End-user rules and policies	Access requires acceptance of the EGI Terms of Use, also available from here, **Acceptable Use Policy (AUP), and Conditions of the 'EGI Applications on Demand Service'**. General service conditions under the default EGI Service Level Agreement.
Fair use, security, data protection	Under the EGI Privacy notice.
End-user support	Full user documentation of EGI Services. User documentation and webinars specific to the EGI Cloud service. Specific support and consultancy are also available as part of the EGI HTC service.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	CPU Hours, Number of users, User communities. Under EGI Monitoring based on ARGO to monitor the availability and reliability of the sites.
Engagement with end-users	EGI provides the central user support via EGI Helpdesk and coordinates support activities of EGI providers, who offer user support for the services/resources they contribute to the EGI ecosystem.

6.21. SURF Data Processing e-Infrastructure (HTC)

Table 40: Access Policy Analysis – SURF Data Processing Grid/Spider

Access Policy Analysis - SURF Grid/Spider	
Obtaining Access	
Targeted users	SURF is a research and education cooperative. Here we limit ourselves to the Grid and Spider HTC platforms and the Computing Time on National Computing Facilities call for proposals[31]. The target audience for this call is research performed by Dutch academia.
Process to obtain access	Obtaining access to Compute and Data services for Research is through fee-free access by national calls (peer-reviewed) or by purchasing access (for SURF member organizations).
Access Tracks and Modes	
Access tracks	• Research organizations can obtain access through national calls for computing by the Dutch Research Council (peer-reviewed). • SURF member organizations can also purchase access.
Access	SURF provides a variety of modalities of access to compute resources, including bare

[31] See https://www.nwo.nl/en/calls/computing-time-on-national-computing-facilities for details.

modalities	metal, managed software stacks (e.g. DIRAC), programmatic interfaces and options for both batch and interactive use. Additionally, SURF offers service interfaces and special support for workflow execution and orchestration. There is an ongoing investigation into the integration of Kubernetes (K8S) with Slurm, and a multi-tenant Kubernetes cluster. The access through commercial contracts includes service assurances such as SLAs, which vary depending on the specific service and, in some cases, the customer.
Summary of available resources	
Compute resources	SURF offers diverse compute capabilities spanning the range from supercomputing (HPC) to HTC and Cloud. Here we focus on the HTC platforms Grid and Spider that are deployed on top of (in-house) Openstack Cloud at SURF. Via the national CfP about 20.000 CPU cores and 40 GPUs are available for these HTC platforms in 2025. In a request these compute services can be combined with dedicated expertise (consultancy) and a variety of data services. Internal and external network is available and follows a fair use policy.
Data resources	A wide range of filesystems is in use, including dCache(FS), CXFS, CephFS, Ceph block storage, Ceph S3, NFS, S3FS, and mounted solutions over FUSE such as WebDAV. Storage solutions deployed include NextCloud, dCache, Ceph, Object Store, NVMe, NVMe over Fabrics (NVMEoF), and Tape. Hierarchical Storage Management (HSM) is supported. Storage capacity is supply and demand driven and constrained by funding. In the CfP the HTC platforms come with 28 PB of online disk and 14 PB of offline tape storage. In total SURF hosts about 100 PB of online disk and around 200 PB of nearline tape. Project space and home file systems feature redundancy but are not backed up or duplicated. Specific services at SURF are available for data duplication, backup, and long-term preservation. Data federation is provided through dCache (Grid) and iRODS. Federated identities are (being) integrated as part of the AAI solution. SURF operates various federative HTC services such as the global workload management system DIRAC, the file transfer service FTS, and the VO management system VOMS. For software distribution CVMFS and EESSI are supported. The NL Tier-1 Grid services for WLCG is delivered as a shared, distributed service across two locations (SURF and NIKHEF). Certain services (e.g., long-term data preservation and the SURF Data Repository) are only available via the purchase route and not part of the CfPa CoP.
Data transfer resources	SURF supports a broad range of protocols and clients for both internal and external connectivity. The specific set of supported protocols and clients may vary depending on the system architecture. The network infrastructure is designed to deliver high performance, with 1.2 Tbps capacity available at the end-of-row and a minimum of 1.2 Tbps at the data center (DC) border. Compute nodes are typically provisioned with at least dual 25 Gbps interfaces. Storage systems are architected to scale horizontally to meet growing performance and capacity demands. Aligned with the DC24 initiative, SURF successfully demonstrated 800 Gbps connectivity between Amsterdam and Geneva. Sustained box-to-box transfer rates of ca. 600 Gbps were achieved in this test, with throughput intentionally capped at this level to prevent interference with production services. Staging data services are also supported such as a local OSDF Cache instance.
Access management and security	
Identity management and AAA	Different services can have different IdM and auth2: • Compute and Storage systems work with public-private SSH key pairs, Grid certificates, and/or Tokens. • Granular privileges can be set at the Unix group, ACL, etc levels. • SURF Research Access Management (SRAM) as the national AAI system (AARC BPA compliant) is available for accessing an increasing set of services (with the aim of having

	all services accessible via SRAM). All systems report to a separate accounting system (the central budget and accounting system) that aggregates these numbers, maps these to contracts and creates reports.
Security methods and processes	Security policies and measures are harmonized as much as possible, but can differ on different systems. A variety of processes and agreements exist to ensure compliance with local and European regulation. The SURF (HTC) services are ISO 27001 certified.
Rules and assurances	
End-user rules and policies	Users must follow specific access and usage policies for compute and data services, ensuring proper authentication and permissions. The usage policies are defined in the SURF Usage Agreement for users. Access is only provided if the user signs this agreement. Usage is monitored for compliance, and audits may be conducted to ensure adherence to policies. Support services are available for technical issues, with users required to follow proper reporting channels. Non-compliance may result in the suspension or revocation of access, with SURF reserving the right to modify or withdraw access to services.
Fair use, security, data protection	Fair use of Compute, Network and Data services for research is ensured through clear access and allocation policies. Resources are allocated based on project needs and priorities, with usage guidelines in place to promote efficiency and prevent misuse. Transparency is maintained through clear terms and conditions for researchers. Security at SURF is maintained through strong authentication and role-based access control. Regular security audits help identify and address vulnerabilities, ensuring robust protection of infrastructure and user data. Data protection at SURF is aligned with national and international regulations, including GDPR. Mechanisms are in place to ensure data integrity and prevent unauthorized modifications.
End-user support	In general, support includes consultancy and expertise, helpdesk services, and online documentation. Expertise is available upon request to assist end users in improving workflow execution. All projects requested via the CfP are subject to review for technical feasibility and performance.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Resource usage (including compute, storage, and network) is monitored at the project, user, and node levels. Technical personnel analyse this data to identify and address inefficiencies. Dedicated service (and innovation) programs focus on benchmarking, performance optimisation, and energy efficiency. SURF also provides monitoring services and interfaces at system level to end users. This includes e.g., Prometheus and Grafana dashboards for monitoring cluster and node performance for computing.
Engagement with end-users	Expertise is offered to end-users to help optimize the execution and performance of their research workflows. SURF monitors the usage of compute, storage, and network resources, providing feedback to users to improve efficiency and optimize resource utilization. Ongoing support and training are provided to users, helping them maximize the effectiveness of available resources. Programs on benchmarking, performance, and energy efficiency guide users toward best practices. SURF fosters collaboration with researchers, offering resources and assistance for joint projects to integrate services with research activities smoothly.
Evaluate, improve, and evolve infrastructure	At SURF, a shared, distributed Grid service is provided in collaboration with Nikhef, the high-energy physics institute in the Netherlands. Additional federation options, such as access integration for distributed data meshes and computing, are actively being explored with partners in a variety of national and international projects (incl. EOSC).

6.22. EOSC Federation e-Infrastructure

Table 41: Access Policy Analysis – EOSC Federation e-Infrastructure

Access Policy Analysis - EOSC Federation e-Infrastructure	
Obtaining Access	
Targeted users	Researchers across Europe, including those from academia, research institutions, and potentially industry and public sectors, are the primary users of EOSC services.
Process to obtain access	Access to EOSC services is facilitated through the EOSC Portal, where users can discover and request services. Access policies are defined by individual service providers and must align with EOSC's overarching access policy framework.
Access Tracks and Modes	
Access tracks	EOSC supports multiple access tracks, including open access, restricted access, and controlled access, depending on the nature of the data and services.
Access modalities	Access modalities include interactive access via web interfaces, programmatic access through APIs, and federated access using standardized protocols.
Summary of available resources	
Compute resources	EOSC provides a broad range of compute resources, including HPC, HTC, and cloud computing services. These resources are offered through EOSC Nodes and are accessible via federated services.
Data resources	EOSC integrates FAIR (Findable, Accessible, Interoperable, and Reusable) data repositories provided by EOSC Nodes. These repositories must adhere to EOSC's data policy and are expected to be certified, e.g., through CoreTrustSeal.
Data transfer resources	EOSC facilitates data transfer through standardized protocols and services that ensure secure and efficient movement of data across different infrastructures.
Access management and security	
Identity management and AAA	EOSC employs a federated Authentication and Authorization Infrastructure (AAI) to manage user identities and access rights across services. This system ensures secure and seamless access for users.
Security methods and processes	Security within EOSC is maintained through comprehensive measures, including encryption, vulnerability management, and incident response frameworks. EOSC Nodes are required to implement these security protocols to protect against threats like data breaches and unauthorized access.
Rules and assurances	
End-user rules and policies	EOSC has established a set of Rules of Participation that all service providers and users must adhere to. These rules encompass principles of openness, FAIR data standards, ethical research practices, and interoperability.

Fair use, security, data protection	EOSC's policies ensure fair use of resources, robust security measures, and strict data protection protocols. Service providers must define clear usage terms, and users are expected to comply with these terms to maintain the integrity and trustworthiness of the EOSC ecosystem.
End-user support	EOSC offers comprehensive end-user support through helpdesks, training materials, and documentation. This support is designed to assist users in effectively utilizing EOSC services and resources.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	EOSC monitors key performance indicators to assess the effectiveness and efficiency of its services. These key performance indicators include metrics related to service usage, user satisfaction, and compliance with FAIR principles.
Engagement with end-users	EOSC actively engages with its user community through consultations, feedback mechanisms, and collaborative initiatives to ensure that services meet the evolving needs of researchers.
Evaluate, improve, and evolve infrastructure	EOSC is committed to continuous improvement and the evolution of its infrastructure. This involves regular evaluations, incorporation of user feedback, and adaptation to emerging technologies and research requirements.

6.23. Simpl Data Federation e-Infrastructure

Table 42: Access Policy Analysis – Simpl Data Federation e-Infrastructure

Access Policy Analysis - Simpl Data Federation e-Infrastructure	
Obtaining Access	
Targeted users	Researchers, public sector entities, and stakeholders involved in European data spaces and research infrastructures.
Process to obtain access	Access is facilitated through the Simpl platform, which provides a secure and interoperable environment for data sharing. Users can engage with the platform via its open-source components, such as Simpl-Open, and test environments like Simpl-Labs.
Access Tracks and Modes	
Access tracks	Simpl supports various data spaces, including the European Open Science Cloud (EOSC), eHealth, Language, and Smart Communities. Each data space may have specific access protocols and requirements.
Access modalities	The platform enables cloud-to-edge federations, allowing users to access and process data across distributed infrastructures seamlessly.
Summary of available resources	
Compute resources	Simpl facilitates access to federated computing resources, integrating cloud and edge computing capabilities to support diverse research needs.

Data resources	The platform supports the creation and management of Common European Data Spaces, providing structured and interoperable data repositories for various sectors.
Data transfer resources	Simpl ensures secure and efficient data transfer mechanisms between cloud infrastructures, systems, and applications, promoting interoperability and data sovereignty.
Access management and security	
Identity management and AAA	Simpl employs federated identity management systems, allowing users to access multiple services across different domains using a single set of credentials. This approach enhances security and Simplifies the user experience.
Security methods and processes	The platform incorporates robust security protocols, including data encryption, access controls, and compliance with EU data protection regulations, to safeguard data integrity and confidentiality.
Rules and assurances	
End-user rules and policies	Users must adhere to the terms of use defined by each data space, which outline acceptable use, data handling procedures, and compliance requirements.
Fair use, security, data protection	Simpl aligns with the EU's data protection frameworks, ensuring fair use policies are enforced and data privacy is maintained across all services.
End-user support	The platform provides comprehensive support services, including documentation, helpdesks, and community forums, to assist users in navigating and utilizing the infrastructure effectively.
Monitoring, Evaluation, and Evolution	
KPIs monitored by the infrastructure	Key performance indicators include system uptime, data transfer rates, user engagement metrics, and compliance adherence, ensuring the platform meets its operational objectives.
Engagement with end-users	Simpl fosters active collaboration with its user base through workshops, feedback mechanisms, and participatory governance models to continuously refine its services.
Evaluate, improve, and evolve infrastructure	The platform is committed to ongoing development, incorporating user feedback and technological advancements to enhance functionality, scalability, and user satisfaction.

7. Annex 2 – Links to Important Documents

Table 43: Links to important documents and information

e-Infrastructure	Link
EuroHPC JU	https://cdn.sanity.io/files/461i44gu/production/92ef531f0f8aa574a6791702e65f04bf9e0422c5.pptx https://eurohpc-ju.europa.eu/selection-first-seven-ai-factories-drive-europes-leadership-ai-2024-12-10_en https://eurohpc-ju.europa.eu/eurohpc-ju-selects-additional-ai-factories-strengthen-europes-ai-leadership-2025-03-12_en https://eurohpc-ju.europa.eu/one-step-closer-european-quantum-computing-eurohpc-ju-signs-hosting-agreements-six-quantum-computers-2023-06-27_en https://eurohpc-ju.europa.eu/new-eurohpc-quantum-computer-be-hosted-netherlands-2024-10-22_en https://eurohpc-ju.europa.eu/signature-procurement-contract-eurohpc-quantum-computer-located-germany-2024-10-15_en https://eurohpc-ju.europa.eu/access-our-supercomputers/access-policy-and-faq_en https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-regular-access-mode_en https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-extreme-scale-access-mode_en https://eurohpc-ju.europa.eu/eurohpc-ju-access-call-ai-and-data-intensive-applications_en https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-benchmark-access_en https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-development-access_en https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en https://eurohpc-ju.europa.eu/selection-first-seven-ai-factories-drive-europes-leadership-ai-2024-12-10_en https://eurohpc-ju.europa.eu/eurohpc-federation-platform_en https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en https://cdn.sanity.io/files/461i44gu/production/8cb97c89ba76cc9249c37d11b6261543ac4f9e02.pptx
GCS	https://www.gauss-centre.eu/for-users/hpc-access

	https://www.gauss-centre.eu/for-users/hpc-infrastructure https://www.gauss-centre.eu/fileadmin/user_upload/GaussCall33_Jan_2025.pdf https://www.gauss-centre.eu/for-users/user-services-and-support
NHR	https://www.nhr-verein.de/en/computing-time https://www.nhr-verein.de/informationen-zur-zentrenauswahl
RES	https://www.res.es/en/access-to-res
GENCI	http://www.idris.fr/eng/info/gestion/demandes-heures-eng.html https://www.genci.fr/sites/default/files/brique/fichier/09-2023/Modalites_dacces_1.pdf https://www.cines.fr/services/formulaires-et-textes https://www.edari.fr https://www.edari.fr/documentation/index.php/Documentation_compl%C3%A8te#VieDuProjet
CSCS	https://www.cscs.ch/services/overview
EPCC	https://www.archer2.ac.uk/support-access/access.html https://www.ukri.org/opportunity/access-to-high-performance-computing-facilities-spring-2025/ https://www.archer2.ac.uk/about/policies/ https://www.archer2.ac.uk/about/reports/ https://docs.archer2.ac.uk/quick-start/quickstart-users/
ICSC	https://www.supercomputing-icsc.it/
WLCG HTC-oriented	https://wlcg.web.cern.ch/using-wlcg/who-can-use-wlcg https://wlcg.web.cern.ch/using-wlcg/computer-security
NIKHEF	https://www.nikhef.nl/pdp/doc/facility
EGI HTC-oriented	https://www.egi.eu/services/research
SKA HTC-oriented	https://www.skao.int/en/science-users/119/ska-regional-centres https://www.uksrc.org/project-overview https://swesrc.org/about https://skach.org
EBRAINS	https://www.ebrains.eu/page/terms-and-policies
LOFAR Long-term	https://lta.lofar.eu

Archive	
LOFAR – Central Processing (CEP)	https://science.astron.nl/telescopes/lofar/access-to-lofar-data/
ERUM	https://erumdatahub.de/en
PUNCH4NFDI	https://www.punch4nfdi.de
Copernicus	https://www.copernicus.eu/en/about-copernicus/infrastructure-overview https://www.copernicus.eu/en/access-data
EGI Federation	https://www.egi.eu/egi-federation
SURF	https://www.nwo.nl/en/calls/computing-time-on-national-computing-facilities https://www.surf.nl/en/small-compute-applications-nwo https://servicedesk.surf.nl/wiki/display/WIKI/Grid https://servicedesk.surf.nl/wiki/display/WIKI/Spider https://www.surf.nl/en/access-to-compute-services
EOSC	https://eosc.eu/eosc-federation-handbook
Simpl	https://digital-strategy.ec.europa.eu/en/policies/Simpl