



SPECTRUM

D5.3 Landscape of RIs: technologies, services, gaps

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**

e-Infrastructures, research infrastructures, compute, data storage, data transfer, software stacks, federation

This Deliverable documents and analyzes the technical characteristics 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, in particular the High-Energy Physics and Radio Astronomy fields, with a focus on the use cases described in Spectrum Deliverable D5.1 (Representative use cases: analysis and alignment), and identifies and discusses recommendations for the future evolution of the e-Infrastructure landscape in Europe.

Document Description			
D5.3 Landscape of RIs: technologies, services, gaps			
Work Package Number 5			
Document Type	Report		
Document status	Under EC Review	Version	1.0
Dissemination Level	Public		
Copyright status	 <p>This material by Parties of the SPECTRUM Consortium is licensed under a Creative Commons Attribution 4.0 International License.</p>		
Lead Partner	FZJ		
Document link	https://documents.egi.eu/document/4073		
Digital Object Identifier	https://zenodo.org/records/15647756		
Author(s)	<ul style="list-style-type: none"> • Hans-Christian Hoppe (FZJ) • Luis Cifuentes (FZJ) • Xavier Salazar (EGI) 		
Contributor(s)	<ul style="list-style-type: none"> • Jeff Wagg (CNRS) • Tommaso Boccali (INFN) 		
Reviewers	<ul style="list-style-type: none"> • Fabio Affinito (CINECA) • Thierry Bidot (NEOVIA) 		
Moderated by	<ul style="list-style-type: none"> • Patricia Ruiz (EGI) 		
Approved by	<ul style="list-style-type: none"> • Sergio Andreozzi (EGI) – on behalf of AMB 		

Revision History			
Version	Date	Description	Contributors
VO.1	08/01/2025	First draft	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ)
VO.2	21/02/2025	Second Draft	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ) Xavier Salazar (EGI)
VO.3	06/05/2025	Rearranging content, adding appendices,	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ) Xavier Salazar (EGI) Tommaso Boccali (INFN) Jeff Wagg (CNRS)
VO.4	31/05/2026	Adding further e-Infrastructure data to Annex 1; Additional contents, abstract, and executive summary, and formatting adjustments. Version ready for internal review	Luis Cifuentes (FZJ) Hans-Christian Hoppe (FZJ) Xavier Salazar (EGI)
VO.5	06/06/2025	Internal Review	Raymond Oonk (SURF) Thierry Bidot (NEOVIA)
VO.6	11/06/2025	Version incorporating the suggestions of the internal Deliverable review.	Luis Cifuentes (FZJ) Hans-Christian (FZJ)
VO.7	12/06/2025	AMB Approval	Sergio Andreatozzi (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
AES-256	Advanced Encryption Standard 256-bit
AI	Artificial Intelligence
AMB	Activity Management Board
AMD	Advanced Micro Devices
API	Application Programming Interface
ARC	Advanced Resource Connector
ARM	Acorn RISC Machine
ATLAS	A Toroidal LHC ApparatuS
BIDS	Brain Imaging Data Structure
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
CM	Communications Manager
CMS	Compact Muon Solenoid
CoP	Community of Practice
CPUs	Central Processing Units
CSCS	Centro Svizzero di Calcolo Scientifico
CSV	Comma-Separated Values
CUDA	Compute Unified Device Architecture
CWL	Common Workflow Language
DB	Database

DRAM	Dynamic Random-Access Memory
DoA	Description of Action
DocDB	EGI Document Database
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
EM	Exploitation Manager
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
FFT	Fast Fourier Transform
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
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
PD	Project Director
PDF	Portable Document Format
PFlop/s	1015 Floating point operations per second
PM	Project Manager
PMO	Project Management Office
PO	Project Objective
POSIX	Portable Operating System Interface
PUNCH4NFDI	Particles, Universe, Nuclei and Hadrons for the NFDI
QC	Quantum Computing
QoS	Quality of Service
QEMU	Quick Emulator
QRM	Quality and Risk Manager
RA	Radio Astronomy
RAID	Redundant Array of Independent Disks
RES	Red Española de Supercomputación
REST	Representational State Transfer
RHEL	Red Hat Enterprise Linux
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
SUSE	Software und System-Entwicklung
SW	Software
TByte	1012 Bytes
TFlop/s	1012 Floating point operations per second
TLS	Transport Layer Security
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	12
List of Tables	12
List of Requirements	12
Executive summary	13
1. Introduction	14
2. e-Infrastructure Analysis Template	15
2.1. Storage/Data Capabilities and Services	15
2.2. Data Transfer and Federation	15
2.3. Compute Capabilities and Services	16
2.4. Software Services, Interfaces, and Stacks	16
2.5. Compute Federation Services	16
3. Selection of Studied e-Infrastructures	17
3.1. HPC-oriented e-Infrastructures	17
3.2. HTC-oriented e-Infrastructures	20
3.3. Data-oriented e-Infrastructures	21
3.4. Cloud-oriented e-Infrastructures	21
4. Analysis Summary and Compendium of Use Case Requirements	23
4.1. Technical Analysis Summary	23
4.1.1. Storage/Data Capabilities and Services	23
4.1.2. Data Transfer and Federation	24
4.1.3. Compute Capabilities and Services	24
4.1.4. Software Services, Interfaces, and Stacks	25
4.1.5. Compute Federation Services	25
4.2. Compendium of Use Case Requirements	25
4.2.1. Storage/Data Capabilities and Services	25
4.2.2. Data Transfer and Federation	26
4.2.3. Compute Capabilities and Services	27
4.2.4. Software Services, Interfaces, and Stacks	27
4.2.5. Compute Federation Services	27
5. Requirements for Future e-Infrastructures	28
5.1. Storage/Data Capabilities and Services	28
5.2. Compute Capabilities and Services	30
5.3. Software Services, Interfaces, and Stacks	31
6. Annex 1 – E-Infrastructure Technical Analysis	33
6.1. EuroHPC JU e-Infrastructure	33
6.1.1. EuroHPC JU Petascale Systems	34
6.1.2. EuroHPC JU Pre-Exascale and Exascale Systems	38
6.1.3. Mid-Range Systems	42
6.1.4. AI Factories	42
6.1.5. Quantum Computer Systems	43
6.1.6. EuroHPC JU Federation Platform	44

6.2. GCS e-Infrastructure	46
6.3. NHR Alliance e-Infrastructure	48
6.4. RES e-Infrastructure	52
6.5. GENCI e-Infrastructure	55
6.6. CSCS e-Infrastructure	57
6.7. EPCC e-Infrastructure	59
6.8. ICSC e-Infrastructure	61
6.9. WLCG e-Infrastructure	63
6.10. NIKHEF e-Infrastructure	64
6.11. EGI HTC-oriented e-Infrastructure	66
6.12. SRCNet SKA HTC/Data-oriented e-Infrastructure	68
6.13. LOFAR – Central Processing (CEP) e-Infrastructure	70
6.14. WLCG Data-Oriented e-Infrastructure	72
6.15. EBRAINS e-Infrastructure	72
6.16. LOFAR Long-Term Archive e-Infrastructure	74
6.17. ErUM-Data-Hub e-Infrastructure	75
6.18. PUNCH4NFDI e-Infrastructure – Data/AI-Access	77
6.19. Copernicus e-Infrastructure	79
6.20. EGI Federation e-Infrastructure	81
6.21. SURF Grid/Spider e-Infrastructure	83
6.22. EOSC Federation e-Infrastructure	84
6.23. Simpl Data Federation Platform	86
7. Annex 2 – Links to Important Documents	90

List of Figures

- [Figure 1: EuroHPC JU HPC systems](#)
- [Figure 2: Simpl-Open Top-Level System Architecture](#)
- [Figure 3: Simpl-Open Implementation Roadmap](#)

List of Tables

- [Table 1: Technical Characteristics of EuroHPC JU Petascale Systems](#)
- [Table 2: Technical Characteristics of EuroHPC JU Pre-Exascale and Exascale Systems](#)
- [Table 3: Technical Characteristics of the EuroHPC Federation Platform](#)
- [Table 4: Technical Characteristics of GCS e-Infrastructure](#)
- [Table 5: Technical Characteristics of NHR Alliance e-Infrastructure](#)
- [Table 6: Technical Characteristics of RES e-Infrastructure](#)
- [Table 7: Technical Characteristics of GENCI e-Infrastructure](#)
- [Table 8: Technical Characteristics of CSCS e-Infrastructure \(Alps system\)](#)
- [Table 9: Technical Characteristics of EPCC e-Infrastructure](#)
- [Table 10: Technical Characteristics of ICSC e-Infrastructure](#)
- [Table 11: Technical Characteristics of WLCG e-Infrastructure](#)
- [Table 12: Technical Characteristics of NIKHEF e-Infrastructure](#)
- [Table 13: Technical Characteristics of EGI HTC-oriented e-Infrastructure](#)
- [Table 14: Technical Characteristics of SKA HTC/Data-oriented e-Infrastructure](#)
- [Table 15: Technical Characteristics of LOFAR – Central Processing \(CEP\) e-Infrastructure](#)
- [Table 16: Technical Characteristics of EBRAINS e-Infrastructure](#)
- [Table 17: Technical Characteristics of LOFAR Long-Term Archive e-Infrastructure](#)
- [Table 18: Technical Characteristics of ErUM-Data-Hub e-Infrastructure](#)
- [Table 19: Technical Characteristics of PUNCH4NFDI e-Infrastructure – Data/AI-Access](#)
- [Table 20: Technical Characteristics of Copernicus e-Infrastructure](#)
- [Table 21: Technical Characteristics of EGI Federation e-Infrastructure](#)
- [Table 22: Technical Characteristics of SURF Grid/Spider e-Infrastructure](#)
- [Table 23: Technical Characteristics of EOSC Federation e-Infrastructure](#)
- [Table 24: Technical Characteristics of Simpl Data Federation e-Infrastructure](#)

List of Requirements

- [Requirement #1: FAIR storage for observation/experiment data across domains](#)
- [Requirement #2: Take care in including GDPR and special category data](#)
- [Requirement #3: Automated, efficient data movement and data staging](#)
- [Requirement #4: Plan data transfer capacity according to research community requirements](#)
- [Requirement #5: Enable portability of compute tasks across accelerated compute platforms](#)
- [Requirement #6: Establish common workflow systems supported by all e-Infrastructures](#)
- [Requirement #7: Establish common SW stacks for compute applications](#)

Executive summary

This document presents the results of a comprehensive study on the technical characteristics of 20 existing or planned European e-Infrastructures that serve European research communities and data-intensive use cases with significant computing requirements. It defines an analysis template covering the different technical aspects of such infrastructures, 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 complete 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 technical characteristics are compared to research community requirements, and recommendations formulated for the future evolution and improvement of European e-Infrastructures to meet the scientific user needs.

1. Introduction

This document constitutes Deliverable D5.3 (Landscape of RIs: technologies, services, gaps) of the SPECTRUM project. It contains an analysis of the technical characteristics of the advanced compute and data resources provided currently (or within a year) by selected European e-Infrastructures, as well as a compendium of the respective current and future 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).

From a comparison of the status quo and the needs of scientific end-user communities, this document then derives and documents a set of requirements for the design, deployment, and operation of future e-Infrastructures, including a mid-term (up to three years) and long-term (up to seven years) view. These are aligned with the recommendations provided in the companion SPECTRUM Deliverable D5.2 (Interoperable Access Policies: Analysis and Recommendations).

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, 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 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 use discussed in the SPECTRUM Deliverable D5.1 has taken precedence,

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, with the analysis of each e-Infrastructure according to the template detailed in Annex 1. Section 4 provides a concise summary of the analysis, references the list of scientific use cases covered in D5.1, and compiles a compendium of their current, mid- and long-term requirements. Finally, Section 5 presents the future-looking requirements for the design, deployment, and operation of e-Infrastructures able to meet the needs of the research communities.

2. e-Infrastructure Analysis Template

This section lays out the template for analysis of the e-Infrastructures used in Annex 1 (the actual analysis) and Section 5 (the requirements for future e-Infrastructures). It focuses on the technical aspects of the capabilities and services offered, in contrast to the template in Spectrum Deliverable D5.2 (Interoperable access policies: analysis and recommendations), which looks at the access policies.

2.1. Storage/Data Capabilities and Services

- Storage systems provided and supported data abstraction
 - Principal data abstraction(s) supported by the storage system(s): files/file systems, databases, object stores, key/value stores, ...
- Locality, capacity and performance
 - Levels of storage: node/rack-local, per-system, per-site
 - Rough data capacity for metadata and volume data
 - Rough data access performance for metadata and volume read and write/modify data access: latency, bandwidth
- Access methods and SW interfaces
 - Conventional Read/Write/Modify, streaming
 - SW APIs and service interfaces: POSIX, S3, ...
- Data formats supported
 - Important scientific metadata and volume data formats (general or domain-specific) supported or optimised
- Data lifecycles and data safety
 - Data retention periods at the various levels
 - Data duplication and data backup
- Data security
 - Access control granularity: coarse-grained per user/group, fine-grained ACLs, ...
 - Data encryption

2.2. Data Transfer and Federation

- Data transfer methods and services
 - Internal to centre/infrastructure: methods/interfaces, security requirements, rough performance
 - External to end-users: methods/interfaces, security requirements, rough performance
 - For data federation services between centres/infrastructures: methods/interfaces, security requirements, rough performance

2.3. Compute Capabilities and Services

- Architecture and platforms
 - Homogeneous/heterogeneous systems
 - Kind and make of CPUs, accelerators, memory, local storage, interconnects
- Compute capabilities and performance
 - At core/node, system, infrastructure level: # of cores/CPU/GPU, operations/s, memory, and local storage latencies and throughput, interconnect latency and throughput
 - If available, specific capacity and capability measures: AI, domain-specific benchmarks
- Non-conventional compute support: QC, neuromorphic, ...
- Compute modes supported: bare metal/managed SW, batch, interactive, [Cloud-like] services, workflow orchestration, ...
- Operational aspects: runtime/size limits, QoS guarantees, ...

2.4. Software Services, Interfaces, and Stacks

- SW API and stack characteristics
 - Relevant programming languages/models and frameworks supported (optimized) by the e-Infrastructure
 - Key libraries or SW components
- Compute the services provided
 - Shell vs. service interfaces, REST APIs, portals, ...
- Workflow orchestration systems and interfaces
 - Slurm job steps and chains, other workflow engines/interfaces
- Monitoring services and interfaces
 - At the application/service/workflow and system level

2.5. Compute Federation Services

- Resource discovery & specification, job routing, load distribution
 - Within a site/infrastructure node (between different systems)
 - Between sites of an infrastructure
 - Across multiple infrastructures

3. Selection of Studied e-Infrastructures

The Spectrum DoA specifies a minimum count of 15 e-Infrastructures to be studied in Deliverables D5.2 and D5.3. The total number of such infrastructures used by European scientists across the fields is significantly larger, and even if one restricts the analysis to the HEP and RA fields, a selection would have to be made since such infrastructures exist at different levels, namely trans-national/European, national, or regional.

In addition, the services provided by e-Infrastructures roughly fall into a number of categories: HPC/AI¹ computing (designed to run closely-coupled applications/workflows with a high degree of parallelism), high-throughput computing/HTC² which targets running a large number of tasks at the same time, which themselves have only a low degree of parallelism, 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 modern, service-based interfaces (pioneered by cloud service providers) and the capability to transparently utilize resources at multiple geographic locations and potentially provided by several third parties.

Three key criteria for selection For D5.2 and D5.3 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 section presents the list of e-Infrastructures investigated by the four categories mentioned above. Unless stated otherwise, both D5.2 and D5.3 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**
 - o At the time of writing, the EuroHPC JU HPC infrastructure consisted of eight deployed systems and operational hosting sites:
 - Deucalion hosted by the Minho Advanced Computing Center (MAAC) in Guimarães / Portugal.
 - Discoverer hosted at Sofia Tech Park / Bulgaria.
 - Karolina hosted by IT4Innovations in Ostrava / Czech Republic.
 - Leonardo hosted by CINECA in Bologna / Italy.
 - LUMI hosted by CSC IT Center for Science in Kajaani / Finland.
 - MareNostrum 5 hosted by BSC in Barcelona / Spain.
 - Meluxina hosted by LuxProvide in Luxembourg.

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.

- Vega hosted by the University of Maribor / Slovenia.
- o In addition, four additional systems will become operational in 2025 or 2026 at new hosting sites:
 - Alice Recoque to be hosted 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 hosted by Linköping University / Sweden; the system is a mid-range supercomputer currently in procurement and planned to become operational in 2025.
 - Daedalus to be hosted 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.
 - JUPITER to be hosted by Jülich Supercomputing Centre (JSC) / Germany; the system is Europe's first Exascale supercomputer and is currently in the installation phase. It will become operational in 2025.
- o T5.3 (and this Deliverable) studied these in relation to the technical properties of capabilities and services offered to scientific end-users, whereas T5.2 (and D5.2) investigated the access policies.
- o The EuroHPC JU co-funds the costs required for system purchase, installation, and operation, typically by 50%. 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.
- o EuroHPC JU defines a core set of access policies and rules, which are discussed in detail in the SPECTRUM Deliverable D5.2.
- o 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, D5.3 cannot provide the same level of details on specific technology and services offered as for the HPC infrastructure.
- o Also, in Q4/2024, the EuroHPC JU awarded a contract for deploying a "**European Federation Platform**" (EFP) across their HPC/AI systems; technical information has started to become available at the beginning of 2025 – the analysis of this very relevant development in terms of access policies is therefore of a preliminary nature.
- o 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. Since these systems are co-located and integrated with HPC systems, which in turn are available under them, it is currently assumed that the access policy will be derived from and similar to one of the five above-mentioned access modalities, yet their specific access rules are not known at the time of writing.
- o In Q1/2025, the HyperCon report was published with the **summary for connectivity requirements and network design for HPC systems in Europe** with the aim to develop a roadmap for a European HPC network that meets the evolving connectivity, performance, and security needs of diverse users across Europe.
- o 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.

- 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, D5.3 can only mention this effort, yet not provide any details on specific technology and services provided.
- **Gauss Center for Supercomputing/GCS in Germany**
 - GCS is an association of the three largest German supercomputer centres (Hochleistungsrechenzentrum Stuttgart/HLRS, Jülich Supercomputing Centre/JSC, and Leibniz Rechenzentrum Munich/LRZ).
 - GCS provides access to German scientists for a collection of HPC systems, including parts of EuroHPC JU systems, across scientific disciplines:
 - Hunter at Hochleistungsrechenzentrum Stuttgart (HLRS).
 - JURECA and JUWELS at JSC in Jülich.
 - SuperMUC-NG at Leibniz-Rechenzentrum in Garching.
 - GCS will provide access to its own contingent of the JUPITER Exascale-supercomputer at JSC, in addition to the EuroHPC JU; initial availability is planned for 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; the most relevant systems are:
 - CLAIX at Technical University Aachen.
 - Barnard and Capella at the Technical University of Dresden.
 - Goethe-NHR at Goethe-University Frankfurt.
 - MOGON NHR at the University of Mainz.
 - HoreKa at Karlsruhe Institute of Technology (KIT).
- **Red Española de Supercomputación/RES³ in Spain**
 - 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 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**

³ <https://www.res.es/en/about-res/nodes>

- EPCC operates the largest UK HPC system (currently Archer 2) and makes it available for scientific and industrial use across disciplines.
- ICSC (Centro Nazionale di Ricerca in HPC, Big Data e Quantum Computing) 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. In addition to these, the ICSC enhances the infrastructure with ~150 Million Euros of site consolidation, new resources, and an improved network infrastructure. It also deploys two production-level quantum computers.
 - The ICSC allocates HPC and Cloud resources to Italian researchers and SMEs, via a granting process that enables institutions, both internal and external to the ICSC, to be assigned resources.
 - The ICSC aspires to become the default go-to point for advanced computing in Italy, and as such, is expected to be a partner for the current and future experiments of HEP and RA. For example, HEP workflows are already being partially executed on its infrastructure, and the ICSC is expected to serve the needs of future RA experiments, such as SKA.

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, 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 manner. 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 across 170 compute centers worldwide, which in 2025 comprise approximately 1.4 million 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 Tier-1 WLCG centre.
- 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 (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 comprise the SRCNet will receive, preserve, and disseminate the products generated by the SKA Science Data Processors, which are 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 architecture with a central hub and nine country nodes (Belgium, Denmark, France, Germany, Greece, Italy, Netherlands, Norway, Spain, Sweden, and Switzerland).
 - 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 and Poland
 - The LOFAR long-term archive provides access to the complete set of radio-astronomy data produced by the LOFAR instrument.
- ErUM⁴ Data Hub in Germany
 - This infrastructure is a central networking and transfer node for data related to the exploration of 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 NFDI⁵ 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 originally introduced and made popular by commercial Cloud vendors and (ii) the use of resources that are geographically distributed

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).

and/or owned and controlled by different entities. Thus, the distinction between HPC/HTC and data-oriented infrastructures is sometimes a bit blurry, and other (HPC, HTC, Data) e-Infrastructures are increasingly 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 the commercial Cloud providers. While these are being used directly in certain 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 and international 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 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, with a central EOSC EU node procured by the EC and a number of EOSC pilot nodes in different countries and for different thematic research communities. The deployment of these pilot nodes started in March 2025 and is currently in the build-up phase.
 - The EOSC Federation consists of multiple EOSC Nodes that collaborate to share and manage scientific data, knowledge, and resources across scientific disciplines and geographies.
 - Following an open call for Expressions of Interest in summer 2024, the EOSC has selected a **first wave of 13 candidate nodes** in March 2025; the organisations in charge of these candidate nodes are **BBMRI ERIC**, CERN, **CNR (Blue-Cloud 2026)**, **CNRS (Data Terra)**, **CSC-IT Centre for Science**, **CVTI SR**, Life Science Research Node (**ELIXIR**, **EMBL**, **Euro-Biolmaging ERIC**, and **Instruct-ERIC**), **ESRF (PaNOSC)**, **EUDAT**, **ICSC**, **NCN**, **NFDI SURF**. These nodes will be successively established and for the EOSC Federation.
- **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. Analysis Summary and Compendium of Use Case Requirements

This section starts with a concise summary of the detailed technical analysis results in Subsection 4.1 – the details can be found in Annex 1. Following this, it presents a compendium of the use case requirements identified in Deliverable D5.1 (Representative use cases: analysis and alignment), Subsection 4.2. Priority is given to cross-cutting requirements shared by several of the studied use cases. In addition, requirements for use cases submitted via the Spectrum Community of Practice (CoP) have been analyzed, and the results are included here.

4.1. Technical Analysis Summary

4.1.1. Storage/Data Capabilities and Services

For HPC e-Infrastructures, sites provision high-performance storage resources with parallel file systems (such as GPFS or Lustre), which can support the large numbers of sophisticated, high-performance nodes without slowing down computation. In some cases, multiple storage layers are provided, with a limited-capacity solid-state drive-based layer complemented by a larger, most often spinning disk-based layer. Most of the sites do have their own (tape-based) backup store and conduct regular backups. Node-local storage (always based on solid state devices) is rare; ad-hoc file systems which can leverage the unmatched aggregated bandwidth of such local storage have been investigated⁶, and they can be of obvious use as scratch storage for node-local data. Data access by applications is through POSIX-compatible I/O interfaces and system SW layers implementing general or domain-specific data and information structures, such as [HDF5](#) or [NETCDF](#). Data lookup is handled either in such layers or in the application space.

While the high-performance storage systems used in HPC centers provide impressive I/O bandwidths, their capacity is limited compared to general-purpose storage systems. For data-intensive use cases, input data has to be staged in from permanent storage before compute tasks on it can be run, and large volume results have to be staged out after they are generated. This applies in particular to the highest performance storage layers based on solid-state technology. Streaming of data between storage layers as computation progresses is not supported at the application level (yet might happen under the hood as part of the storage device and system software operation).

Data e-Infrastructures, on the other hand, focus on provisioning large storage capacities, and they add services to find and locate data⁷ without the end user needing to be familiar with file system directory structures, etc. Generic architectures and platforms such as [EOSC](#) and [Simpl](#) provide significant flexibility and enable the addition of domain-specific extensions and services, while e-Infrastructures like [WLCG](#) are purpose-built to support domain-specific needs. While such e-Infrastructures are based on low-level abstractions like S3 or POSIX file systems, they offer higher-level data abstractions and interfaces according to user needs.

HTC-oriented e-Infrastructures focus on achieving high throughput for data analytics and (relatively compared to HPC-style simulations) simple computation tasks, which are run in ensemble mode across many compute cores/nodes. Thus, achieving high I/O bandwidths sufficient to feed the throughput compute ensembles is critical⁸, and larger data infrastructure nodes do provision high-performance storage systems comparable to HPC centers. The key difference between HPC and HTC is that the compute tasks of the latter are independent of each other; it is possible to spread HTC ensembles across loosely-connected compute nodes, and exploitation of solid-state node-local storage is more straightforward.

⁶ See for instance the [NEXTGenIO](#) and [ADMIRE](#) projects, and the work on the [GekkoFS](#) file system.

⁷ Which is critical to ensure the "F" and "I" of the [FAIR](#) principles.

⁸ One could argue that HTC usage scenarios can have even higher requirements than HPC applications, due to performing less compute operations per unit of data.

4.1.2. Data Transfer and Federation

Since they do not generally provide long-term data storage for experiment/observation data, HPC centers require researchers⁹ to transfer data from the outside into their protected domains and stage it to their high-performance parallel file systems. Conversely, results data is typically moved from their domain to external sites/systems. The centers studied support a variety of methods for these data transfers, ranging from universally supported low-level, single-stream SSL transfer (using scp or rsync) to more advanced, multi-stream methods like Globus Connect (for CSCS), yet these are not commonly supported across the HPC e-Infrastructures and their sites. If a distinct highest-performance layer (such as based on solid-state devices) is provided, the internal data staging to/from the regular storage system is handled by the end-users or their applications/job scripts, using site-internal Linux mechanisms.

Depending on the HPC or HTC site security architecture, incoming and outgoing transfers need to pass through a number of intermediate systems, which can slow transfer down. In addition, communication between compute nodes to/from the outside Internet is often prevented, and if it is allowed, it typically goes via proxy systems. The HPC centres studied do not provide data federation services, except in some cases where they share access to parallel file systems.

Data-oriented infrastructures, on the other hand, typically provide services for searching for and locating of data according to generic or domain-specific criteria, and enable access to and transfer of that data without end users needing to know the location and details of access and transfer interfaces. Good examples are the [WLCG](#) and [EGI](#) infrastructures, which support the descriptive [FTS](#) data movement service.

Generally applicable data e-Infrastructures or platforms, such as [EOSC](#) and [Simpl](#), implement generic frameworks for defining data and information structures, catalogues for metadata, and associating data providers with storage resources; they (in particular Simpl) can be adapted to a wide variety of scenarios for finding/locating, controlling access to, and making data available to end users without participants needing to know the underlying infrastructure.

4.1.3. Compute Capabilities and Services

The traditional dominance of HPC systems using general-purpose processors (CPUs) has been successfully challenged by massively parallel accelerators (mainly GPUs), which can provide significant increases in delivered performance and energy efficiency for suitable, adapted workloads. In the general HPC field, the transition from CPU-only or CPU-optimised applications to GPU-adapted/optimized codes is in full swing.

The field of data center GPUs is today dominated by a single vendor (NVIDIA), which is pushing its own, proprietary programming model (CUDA) and offers a very capable stack of compilers, libraries, and tools around this model¹⁰. Many applications have been successfully ported and optimized to use CUDA or its related SW stack elements – yet the proprietary nature of CUDA means that the GPUs of other vendors (for instance, AMD and Intel) are either not supported, or require creation of special additional code versions.

In addition, the emergence of neural-network (NN) based ML and AI techniques is changing the specific requirements on the compute capabilities; the nature of NN-based computation strongly favors GPU architectures, reduced size data formats enable further performance and energy improvements, and a fruitful co-design between the evolving NN and ML fields and GPUs has emerged. Different from “conventional HPC”, the NN/ML field is, to a very large degree, based on using high-level programming frameworks, which can be optimized for different CPU and accelerator architectures without requiring significant application code changes.

⁹ In particular if they are not employed by the HPC center.

¹⁰ See <https://developer.nvidia.com/cuda-zone> for details.

4.1.4. Software Services, Interfaces, and Stacks

The field of workflow programming, management, and orchestration systems is very fragmented, with a number of competing standards and many mutually incompatible systems in place, and no obvious “common denominator” adopted in the HPC field. For HPC centers, the basic level of support involves chaining job steps using the Slurm system; however, this is very low-level and cumbersome for end-users, and moreover, restricted to a single center.

Many projects and activities have proposed workflow systems which provide ways to define (“program”) workflows, to submit them for execution and monitor their progress, and to orchestrate their execution across multiple systems and sites, including the required data transfers. In addition, in the field of enterprise and Cloud computing, there are several established systems and standards.

The [EuroHPC JU Federation Platform](#) (EFP) will proceed to implement and deploy a generic workflow system with the above capabilities across a large set of diverse HPC systems. It remains to be seen whether that common system can indeed satisfy the needs of the majority of use cases.

The [WLCG](#) e-Infrastructure supports a workflow system designed for the LHC use cases; it supports the above-listed functionality, orchestrating the batch-style execution on the WLCG resources in a way that is transparent to the end user. The LHC experiments (and other participants in the WLCG) can add their specific SW layers and interfaces on top to match the level of abstraction of their end users.

4.1.5. Compute Federation Services

The [EuroHPC JU Federation Platform](#) is introducing a federation of HPC end-user identities across the EuroHPC JU infrastructure and will add unified services for cross-site workflows and simplified, common interfaces for the use of the HPC systems, following previous initiatives like FENIX and Grid Computing platforms. Other HPC e-Infrastructures, such as CSCS and GENCI, provide similar federation capabilities between their constituent sites/systems. It should be mentioned that none of these support end user access privileges/quotas across all participating systems, and that automatic routing of accesses/compute tasks to the best suited systems is not available¹¹.

As mentioned in the previous subsection, [WLCG](#) federates its HTC resources for workflow execution. The [EGI federation](#) provides a brokering service for scientific end-users and makes access to public and commercial providers of Cloud resources available via a common interface.

4.2. Compendium of Use Case Requirements

4.2.1. Storage/Data Capabilities and Services

The aggregated **data volumes** collected/created by the use cases vary greatly across use cases between 100s of GBytes – several EBytes, with a clear growth curve for the large use cases until 2030 (see, for instance, the CERN LHC and SKA plans). The data used or created by the use cases is either “raw” detector/instrument data, which cannot be reproduced, or derived data coming from data analysis or simulation of experiment/detector data, which can be re-generated by re-running the analysis or simulation steps.

The first kind of data is therefore irreplaceable and must be safely stored for the duration of the instrument/experiment and the research communities working with this data. For purposes of validation and calibration, a strong argument can be made to preserve this data for an indefinite amount of time¹². For HEP and RA, this data represents the lion's share of the total aggregated data. Current practice is for the organizations operating the instruments/experiments (often partnering with research communities “using”

¹¹ The EFP promises to implement hierarchical scheduling for workflows which will address this problem.

¹² One example are the photographic astronomical observations performed by optical telescopes in since the early 20th century, which are regularly use to correlate with current observations.

the collected data) to design, deploy, and operate the required data storage infrastructure, with the WLCG being a prime example.

For the second kind of data, re-generation is certainly possible while the applications and workflows required to do that are still available and working on the available infrastructure; over time, such capability might be impacted, though, unless precautions are taken to preserve such SW artifacts and the required information to use them for data regeneration.

The scientific data of both kinds must be findable, accessible, interoperable, and reusable (FAIR) for scientific end users and, where applicable, the public. “Findable” requires establishing data catalogues and metadata structures and lookup capabilities that match the way information is structured and handled by a scientific domain. “Accessible” requires provisioning methods for end users to read and interact with the data. “Interoperable” and “Reusable” require careful definition of the inherent structure for all data, and storage of machine-independent programs and instructions for reading, writing, and extracting semantics of the data.

Data security and data safety are important for the studied use cases insofar as long-term retention and guaranteed availability, plus protection against accidental or intentional tampering of the instrument/experiment data, are concerned. The use cases do not include identifiable personal data protected by GDPR or commercially sensitive information, and the only data confidentiality requirement is focused on protecting and supporting the scientific process by enforcing an exclusivity period, which gives teams who collected/generated the data time to publish first. Keeping data confidential for long periods of time is a non-requirement. In this situation, complex encryption schemes with added overhead in performance and key management do not seem to be required.

Data volumes for compute tasks: the volume of input data required for the applications/workflows used by the research communities to analyse the data and create scientific results based on it is significantly smaller (HEP collision analysis or RA observations can work on subsets of the full instrument datasets in the area of 10s of MBytes) and produce output data volumes generally in the same size bracket.

4.2.2. Data Transfer and Federation

Data staging: HPC centers are typically not set up to store and handle the full HEP or RA instrument/experiment data; they invest heavily in high-performance storage systems supporting parallel I/O operations, which can feed their large HPC systems without slowing down computation, rather than in increasing their long-term storage capacities. For analysis/simulation/prediction applications/workflows which require HPC resources, the small input datasets today are staged in (extracted and transferred from instrument/experiment data storage sites to the high-performance storage tiers of the target HPC center), the application/workflow runs at full speed using and stores results on the high-performance storage tier, and finally, results data is staged out to the end-users system or to other data storage sites.

Bundling of several/many compute tasks and the required input data is regularly performed to amortize the per-transfer overhead and saturate the available transfer bandwidths. Such an approach works well for use cases with many independent compute steps that can be run concurrently.

Data transfer: The basic requirement of the use cases studied is to have data transfer services into and out of the high-performance storage tiers used by the HPC centres, which can be used by unattended workflows and provide satisfactory end-to-end data transfer rates. The WAN connectivity situation and issues in Europe have been studied by the EuroHyPerCon Study¹³, and improvements of the WAN infrastructure will be set in motion by the European Commission. Of equal importance is the efficient support of parallel data transfer protocols and tools across all relevant compute and data e-Infrastructures and their centers.

¹³ This study has the title “Connectivity Requirements and Network Design for HPC Systems in Europe” and was contracted by the EuroHPC JU; it is published at <https://op.europa.eu/en/publication-detail/-/publication/122922c3-fd63-11ef-b7db-01aa75ed71a1/language-en>.

Data federation: The use cases studied in detail in D5.1 do not explicitly require data federation capabilities in a specific way; however, the high-energy physics and EBRAINS use cases do in effect rely on established data federation provided by their infrastructures.

4.2.3. Compute Capabilities and Services

The use cases studied project significant increases in **compute demands**, driven by the increase in data volumes and the rising complexity of the analysis, simulation, and prediction steps. In the use cases considered, GPU computing has, amongst others, made inroads for Monte Carlo simulations (based on many independent instances of parallel computation) and for image analysis, signal detection, and instrument calibration.

HEP and RA are both mapping out how to best perform the **transition** of their applications **to GPUs** to reap the performance and energy benefits, yet it is also clear that in these fields and in other domains a significant number of important codes cannot easily be adapted/optimised for GPU-type accelerators, and that advice from experienced HPC specialists would be valued in which codes to prioritise and how to modify them.

The research communities expect that applications/workflows can easily be run across the range of compute e-Infrastructures and systems available to them; this requires **portability** of the workflows (for instance meaning that a workflow script/definition can coordinate workflow steps and data transfers across all e-Infrastructures) and the workflow steps (applications and data transfer commands). Such portability has to cover functionality (workflows and their steps run with the expected results), and it can also cover performance (the expectation is that the applications/steps run efficiently across all available systems), although this particular aspect remains a research topic.

4.2.4. Software Services, Interfaces, and Stacks

Many of the considered use cases rely heavily on multi-step **automated workflows** that combine and sequence data access, transfer, and management with analysis, simulation, and prediction computations, taking into account the graph of task dependencies, rather than manually running these steps in the correct sequence on the correct systems. It is imperative that e-Infrastructures support the definition, launching, monitoring, and accounting of (potentially complex and dynamic) workflows across their resources and the systems used by the scientific end-users.

The use cases also 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. To avoid researchers spending undue effort in configuring and adapting their workflows to accommodate such, a **common set of standard interfaces, protocols, and policies** which streamlines the deployment of workflows to multiple compute and data centers.

4.2.5. Compute Federation Services

The use cases studied in detail in D5.1 do not explicitly require federation of compute resources in any specific way; however, the high-energy physics and EBRAINS use cases do in effect rely on compute federation services.

5. Requirements for Future e-Infrastructures

This section derives technical and operational requirements for future e-Infrastructures. It considers the technical status quo summarised in Section 4, and the use case needs to be documented there, following the structure laid out in Section 2.

5.1. Storage/Data Capabilities and Services

Requirement #1: FAIR storage for observation/experiment data across domains

Across data-intensive scientific domains, agree on a common architecture and core implementation for e-Infrastructures that preserve observation/experiment data (including meta- and provenance data) and make it available to scientists and the public according to the FAIR principles.

The key requirement for such an architecture and implemented core components is to preserve the vast and rapidly growing amounts of observation/experiment data together with the relevant metadata and provenance information according to the FAIR principles. The capability to search for and locate data based on provenance and contents, and the ability to federate in an efficient way with HTC and HPC infrastructures, will be other key requirements.

Rationale

The instrument/experiment-specific solutions (such as WLCG) have shown how FAIR storage of observation/experiment data can be done in an effective way in the HEP domain, and the [SRCNet SKA](#) e-Infrastructure is deploying its own, domain-specific infrastructure for the SKA part of the RA domain. Other science domains (like Earth system sciences and neuroscience) operate their own data e-Infrastructures. This raises the danger of increasing the number of “island” solutions to be designed, implemented, and operated, and the corresponding duplication of work and costs.

The planned sharp increase in data volumes and the need to preserve such data for long periods (which will exceed the lifetime of projects and, in cases of the instrument infrastructure), and the increase in the number of data-oriented research communities and their participating scientists support such a requirement. Leveraging economies of scale would reduce the costs for implementing/deploying data e-Infrastructures and for operating these.

Such an approach would also eliminate duplicate work in adapting to the evolution of compute-oriented e-Infrastructures and the ongoing federation efforts, for instance, in the EuroHPC JU infrastructure and supporting federation between the compute and data sides, or in the context of the vision for data exploitation platforms pursued by the EGI Federation¹⁴

Implementation considerations

The established large-scale data e-Infrastructures like [WLCG](#), Destination Earth, and [LOFAR](#), and the development of generic frameworks for federated data spaces like [EOSC](#) and [Simpl](#) are clear candidates for architecture and the common software frameworks. The work of the EGI Federation on supporting Research Infrastructures to create data exploitation platforms is also a relevant activity to consider. A good balance between fully generic frameworks (that provide ultimate flexibility and maximise the circle of supported domains, yet might create performance bottlenecks) and highly customized and optimised designs must be found. It seems clear that addressing domain-specific needs will require specific configurations and very likely extensions of the core system.

The requirement does not state that there should be only one huge site operating such a data e-Infrastructure, or that there can be only one such e-Infrastructure – the expectation is that the same key components are used across sites or e-Infrastructures. A single e-Infrastructure or a federation of distinct e-Infrastructures could all be viable, as long as the implementation provides a single entry point for storing and accessing scientific data across the supported domains and research communities.

¹⁴ See presentation on the future of EGI: <https://indico.egi.eu/event/6638/sessions/5410/#20250605>

Internationality

Many large research communities are of international composition (including EU and non-EU members), and likewise, some key e-Infrastructures include EU and non-EU users, EU and non-EU origin data, and EU and non-EU resources. While the proposed data infrastructure is mainly oriented towards EU needs, it should be considered opening up to international participation. This will require to address at least three policy and one technical issue:

- Policy 1: under which conditions to admit non-EU end users (accessing data in read-only mode)
- Policy 2: under which conditions to store data from non-EU origin or ownership
- Policy 3: under which conditions to include (storage) resources outside of the EU
- Policy 4: what is the appropriate funding model
- Technical 1: any additional technical measures required to assure data safety and security

Requirement #2: Take care in including GDPR and special category data

Scientific domains in life sciences, medicine would require storing and processing of personally identifiable data in an e-Infrastructure; it will be required to minimize the impact of the required data protection measures in terms of complexity and performance, in particular with regard to use cases that do not work on such data.

Rationale

Research domains and use cases that involve personal data (in particular, special protected data according to GDPR) or use/create commercially valuable/sensitive information will require stringent data protection measures and policies to keep such data safe from unauthorized access or disclosure. Strong encryption measures for data at rest and in transit would be required here, combined with highly secure management of keys, fine-grained access control, and potentially fine-grained access logging. In the past, separate, secure data spaces were used for storing medical data.

Implementation Considerations

Integrating such measures with a common European data continuum would need to be done in a way that does not impact usability and performance for Open Science. Whether a better way than separate, secluded data spaces can be found is not clear at the moment, and the performance and usability impact of mandatory encryption of data at rest, in computations, and transfer depends on the available technology. Additional important aspects to consider are establishing the validity and security of applications acting on protected data and the impact of operating a data e-Infrastructure with full compliance with the data protection rules (which can be substantial).

A potential design decision could be to develop and evolve two architectures for e-Infrastructures: one for use cases without personally identifiable data, and a more complex, second architecture that would ensure the full extent of data protection for use cases relying on personally identifiable data. In addition, the viability of anonymization or pseudonymization techniques¹⁵ needs to be established, potentially per use case.

Requirement #3: Automated, efficient data movement and data staging

Data movement and data staging between sites in an e-Infrastructure, or between federated e-Infrastructures (such as between data and HTC/HPC infrastructures) should be handled automatically.

Rationale

Data staging is customarily performed by researchers in their workflows or job scripts using the interfaces and tools made available by the data and compute e-Infrastructures. This complicates the setup of these; HPC centers do support a multitude of data transfer mechanisms, which further complicates the researcher's task. While larger research communities have been able to automate staging in custom

¹⁵ Like foreseen by the Simpl platform.

e-Infrastructures (such as WLCG), generally applicable support for data staging at a high level is required to cover the multitude of use cases.

Implementation Considerations

Standard interfaces to express the data staging needs in a workflow must be defined and supported on all data and compute e-Infrastructures, and reliable, tuned implementations of these would need to be provided, which must be integrated with the site-specific security architectures.

A relevant proxy on how such automated support could be performed is the FTS system used by EI and WLCG, and the way it is integrated into the WLCG workflow support.

An alternative approach is being deployed by the Destination Earth¹⁶ initiative – analysis/simulation/prediction applications/workflows running on the participating HPC sites use an object-based “data bridge” interface to request the required meteorological or climatological input data, which is then transparently transferred to the HPC compute site executing the applications/workflows. Staged data is cached to avoid unnecessary re-transmissions for subsequent usage. The concept could be extended to use data streaming.

Automated, efficient data transfer into and out of the center-local storage systems raises requirements for the network and **system security architecture**: HPC centers usually restrict access to a prescribed set of entry and/or login/data transfer nodes, which must be configured to allow automated transfers into and out of the center, and in case of local storage systems at the compute nodes, efficient access to these.

Requirement #4: Plan data transfer capacity according to research community requirements

In particular, for future HPC e-Infrastructures, the evolving needs of research communities should drive the planning of data transfer and ephemeral storage capacities.

Rationale

HPC and HTC resources are expensive assets that will be operated for a significant number of years; to ensure continued effective usability, the data transfer capacities provided should be planned according to the evolving needs of the research communities using such compute resources.

Implementation Considerations

In particular, it is imperative to **adjust** the **capacity** and **performance** of inbound and outbound data transfers to the growing requirements, for instance, by sizing the pool of data transfer nodes and increasing their bandwidth to the high-performance storage tiers. For the latter, the amount of data copies required to enter the protected compute center infrastructure must be minimized. If storage local to the compute nodes is used, methods to avoid staging via an intermediate store system should be identified. In cases where end-to-end encryption is required, the impact on data transfer performance must be minimized.

5.2. Compute Capabilities and Services

Requirement #5: Enable portability of compute tasks across accelerated compute platforms

Ensure that compute tasks used by research communities can be used across different families/vendors of accelerated compute resources.

Rationale

Since the HPC field, as demonstrated for instance by the recent EuroHPC JU procurements, has already switched to heterogeneous, scalable systems that combine CPU and GPU computing elements, there is no general gap in deploying GPU compute capacity. An emerging problem is the lack of even functional portability of software codes using the CUDA programming model to accelerators of other vendors than

¹⁶ See <https://destination-earth.eu/> for details.

NVIDIA. This restricts the field of platforms on which compute tasks using CUDA or CUDA-based components can be run. Fulfilling these requirements would enable faster uptake of accelerators (mainly GPUs) across different vendors and avoid duplication of porting/adaptation work.

Implementation Considerations

Achieving full functional and performance¹⁷ portability for general HPC applications has been an elusive goal for a long time. However, functional portability has been established via a number of de-facto standards (such as MPI and [OpenMP](#)), and portability layers for accelerators are in active development (such as [SYCL](#), [OpenMP offload directives](#), or [Kokkos](#)). Even without performance portability, use of such layers or interfaces would at least enable execution of compute tasks across platforms, with later optimization according to a specific platform's characteristics.

In addition, the actual value of carefully designed domain-specific languages and environments has been validated by the emergence of AI/Deep Learning frameworks such as [Tensorflow](#) or [pytorch](#), which enable AI applications to be written at a high level of abstraction, and backend optimizations to be performed in a way mostly transparent to the application.

Establishing common programming models and tools that support accelerators independent of the vendor is of key importance. To be successful, such models must enable efficient use of different accelerator implementations[1], be applicable to common programming languages, and be usable by skilled HPC SW developers. In addition, key libraries (such as for linear algebra, FFT, data format conversions, I/O formats, etc.) must be made available.

5.3. Software Services, Interfaces, and Stacks

Requirement #6: Establish common workflow systems supported by all e-Infrastructures

Define a number of workflow systems covering the needs of key research communities and ensure support by e-Infrastructures.

Rationale

HEP experiments have developed workflow systems adapted to their requirements (see, for instance, [WLCG](#)) and have defined the interfaces required from HPC centers to run these workflows. Other science domains, likewise, have designed and implemented a number of different workflow systems for their needs. The HEP workflow system can likely be used for RA activities like LOFAR and SKA, yet it is not clear whether it will be able to cover other domains in a sufficient way.

A critical factor is, of course, establishing support for such workflow systems; WLCG uses a bespoke implementation based on [HTCondor](#), while the ongoing development of the EFP relies on a different workflow formulation (based on [LEXIS](#)). Therefore, co-design between the user communities and the e-Infrastructures is clearly necessary to establish the “right” set of workflow systems and ways to effectively support them.

Implementation Considerations

Longer-term and in collaboration with the relevant European players in compute and data federation, a next-generation workflow system including the full range of compute and data access/management/transfer requirements across a wide range of research domains should be designed (based on existing, proven systems) and implemented.

Requirement #7: Establish common SW stacks for compute applications

Establish common SW stacks for research communities deployed on the relevant e-Infrastructures.

¹⁷ Meaning achieving similar levels of performance relative to a platform's capabilities.

Rationale

HPC/HTC centers show significant differences regarding their hardware (in particular due to the emergence of accelerators) and system configuration; while the use of a Linux (or Linux-compatible) OS is standard, the provisioned SW stacks are also diverse, with different compilers, libraries, and frameworks made available. While user support is usually available, adapting code/usage practices can create significant additional effort on the side of the code developers or the end users, if execution of compute tasks is required across multiple sites of an e-Infrastructure.

Implementation Considerations

Due to its performance overhead, hypervisor-based virtualization is rarely offered by HPC centers. Instead, container technologies (Singularity and/or Apptainer) are widely supported to enable end-users to package and run their applications “anywhere”. There is currently no agreement on which of the two approaches to support, and support for Kubernetes or other container models is spotty; in addition, managing a set of containers for a research domain and making sure that optimized support for all combinations of CPU/GPU/interconnect encountered across HPC e-Infrastructures is available can be significant and lead to an explosion of distinct container files matching such combinations.

While key SW 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 HDF5 or NetCDF, I/O interfaces like POSIX and Ceph), the deployed SW environments and stacks still differ considerably, partly caused by the choice between different implementations or different versions, partly caused by differences in HW support and system configuration. Besides the danger of applications or workflows not running correctly, there is the spectre of codes running correctly yet achieving sub-optimal performance.

A number of R&D projects have proposed standardized SW stacks (see, for instance, DEEP-SEA¹⁸ and EESSI¹⁹), yet such approaches have not been taken up to the degree required.

¹⁸ See <https://deep-projects.eu/> for details.

¹⁹ See <https://www.eessi.io/> for details.

6. Annex 1 – E-Infrastructure Technical Analysis

This section presents the analysis results for the set of selected e-Infrastructures described in Section 3, as outlined in the template presented in Section 2, and provides a summary view of the analysis results.

6.1. EuroHPC JU e-Infrastructure

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. EuroHPC JU also defines a set of access policies and ancillary processes, which are discussed in SPECTRUM Deliverable D5.2.

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 about 100 PFlop/s), and pre-Exascale/Exascale systems (above 100 PFlop/s).

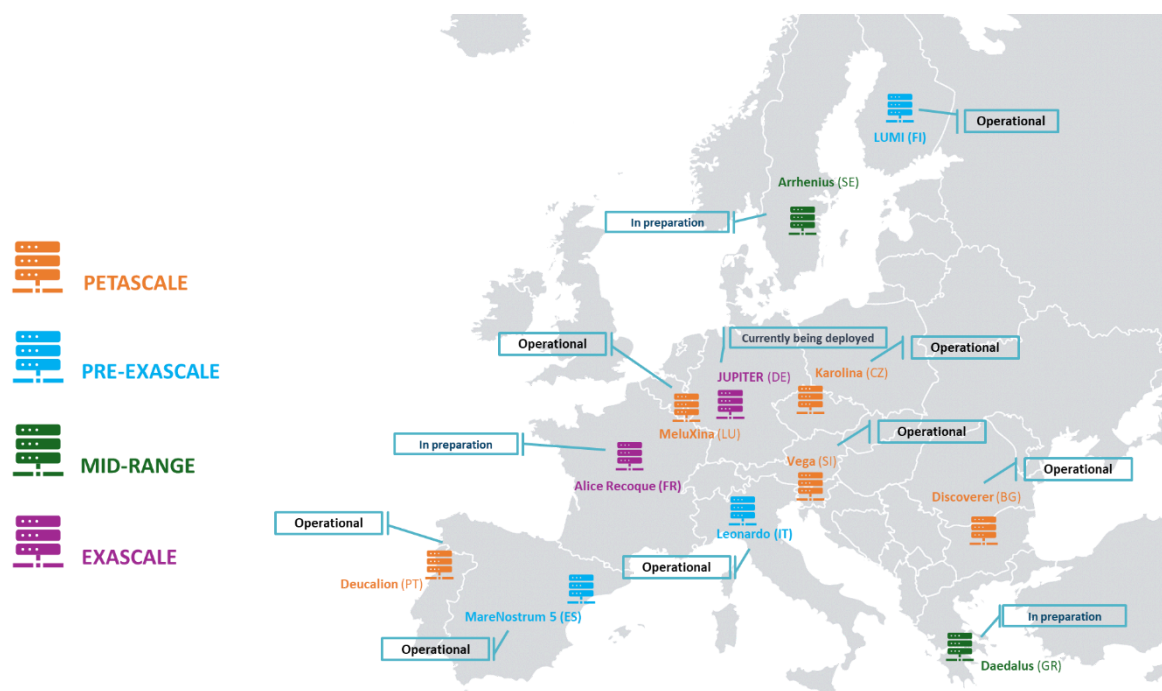


Figure 1: EuroHPC JU HPC systems

Sections 6.2.1 through 6.2.3 discuss the technical capabilities of these classes of systems.

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-optimized 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. Section 6.2.4 discusses these AI Factories – since they are in the definition/procurement phase, technical details at the time of writing this document are sparse.

At the time of writing, the EuroHPC JU had approved a total of nine Quantum Computers based on different technologies, which are in various stages of procurement and deployment. Two of these Quantum Computers were coupled to existing HPC infrastructure in the HPCQS project and will enter operation shortly. Section 6.2.5 discusses these systems.

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, as well as the FENIX federation infrastructure used by the EBRAINS infrastructure, will occur. Section 6.2.6 discusses the technical approach and capabilities of the EFP platform.

6.1.1. EuroHPC JU Petascale Systems

Table 1: Technical Characteristics of EuroHPC JU Petascale Systems

Technical Analysis – EuroHPC JU Petascale Systems					
System/site	Deucalion	Discoverer	Karolina	Meluxina	Vega
Data/Storage Characteristics and Services					
Storage abstraction	Lustre PFS, NFS and local file systems for CPU nodes	Lustre PFS and local file systems for CPU nodes	Lustre PFS and NFS file systems	Lustre PFS	Parallel file system with POSIX-compliant access (Lustre FS)
Storage levels & locality	High-speed partition (Lustre) / NAS (NFS)	Lustre PFS / WEKA (direct GPU access, optimized for AI)	Scratch PFS (Lustre) / Home NFS / Project	Local (not on CPU nodes) / Scratch / Project/ Backup	Multi-tier: high-performance parallel file system and local SSD on nodes
Storage capacity	High-Speed: 10.6 PByte raw NAS: 50 PByte raw	Lustre PFS: 7.2 PByte raw WEKA: 273 TByte	Scratch: 1361 TByte raw Home: 31 TByte raw Project: 15 PByte raw	Local: 1.9 TByte Scratch: 500 TByte raw Project: 12 PByte raw Backup: 7.5 PByte raw	High-Performance Storage Tier (HPST): 10 DDN Exascaler ES400NVX (NVMe disks), usable capacity 1 PB, Large-Capacity Storage Tier (LCST): 61 storage nodes, 23 PB Raw, 18 PB Usable
Performance	High-Speed: 340/260 GByte/s R/W		Scratch: 731/1200 GByte/s R/W Home: 1.9/3.1 GByte/s R/W Project: 29 GByte/s	Scratch: 400 GByte/s Project: 190 GByte/s Backup: 30 GByte/s	~600 GB/s bandwidth, low-latency via Infiniband HDR

Access methods & interfaces	POSIX with Lustre and NFS extensions	POSIX with Lustre extensions (CPU and GPU nodes) POSIX plus WEKA extensions (GPU nodes only)	POSIX with Lustre and NFS extensions	POSIX with Lustre extensions	POSIX, MPI-IO, HDF5, NetCDF
Data formats	Universal support for data formats which are mapped to POSIX files	Universal support for data formats which are mapped to POSIX files	Universal support for data formats which are mapped to POSIX files	Universal support for data formats which are mapped to POSIX files	Common HPC formats supported (binary, HDF5, NetCDF, etc.)
Data retention & backup	High-Speed is (potentially) long-term, backup works via NAS storage	Lustre is long-term WEKA is ephemeral	Scratch is temporary Home is backed up, yet no guarantee is given wrt. restoring old file versions	Scratch is temporary Project is long-term and backed up automatically Backup is intermediate backup target, backed by tape library	No long-term backup on Lustre; users responsible for archiving
Access control	Usual Linux ACL mechanisms	Usual Linux ACL mechanisms	Usual Linux ACL mechanisms	Usual Linux ACL mechanisms	LDAP-based user management, project-based allocations
Encryption	Specific public details are not available.				
Data Transfer & Federation					
Internal methods	Specific public details are not available.				
External methods, security	scp	rsync, scp & sftp possible but deprecated	scp, sftp, rsync Compute nodes cann access Internet resources tunneling thorough login nodes	scp, sftp, rsync	Secure transfer via GridFTP, SCP, rsync over SSH; VPN-based access
External data federation	Specific public details are not available				
Compute Capabilities and Services					
Processors/	Fujitsu ARM	AMD x86 EPYC	AMD x86 EPYC	AMD x86 EPYC	AMD x86 EPYC

Accelerators	<p>A64FX CPU</p> <p>AMD x86 EPYC "Rome" CPU</p> <p>NVIDIA A100 GPU (40/80 GByte HBM2e)</p>	<p>"Rome" CPU</p> <p>Intel Xeon "Sapphire Rapids" CPU</p> <p>NVIDIA H200 TGPU (141 HBM3)</p>	<p>"Rome" CPU</p> <p>Intel Xeon "Cascade Lake" CPU</p> <p>NVIDIA A100 GPU (40/80 GByte HBM2e)</p>	<p>"Rome" CPU</p> <p>NVIDIA A100 GPU (40 GByte HBM2e)</p> <p>NVIDIA H200 GPU (141 GByte HBM3)</p>	<p>"Rome" CPU</p> <p>NVIDIA A100 GPU (40 GByte HBM2e)</p>
Nodes	<p>ARM partition: 1632x CPU nodes with Fujitsu A64FX (48 cores) and 32 GByte HBM2 plus 512 GByte NVMe storage</p> <p>x86 CPU partition: 500x CPU nodes with 2x AMD EPYC 7742 (64 cores) and 256 GByte DRAM plus 480 GByte NVMe storage</p> <p>Accelerated partition: 33x CPU/GPU nodes with 2x AMD EPYC 7742 (64 cores) and 4xNVIDIA A100 GPU and 512 GByte DRAM plus 480 GByte SATA local SSD storage</p>	<p>CPU partition: 1128x CPU nodes with 2x AMD EPYC 7H12 (64 cores) and 256 GByte [1110 nodes] or 1 TByte [18 nodes] DRAM</p> <p>4x NVIDIA DGX2 with 2xDual Intel Xeon 8480C [56 cores] CPU and 1 TByte DRAM and 8x NVIDIA H200 T GPUs each</p>	<p>CPU partition: 720x CPU nodes with 2x AMD EPYC 7H12 (64 cores) and 256 GByte DRAM</p> <p>Accelerated partition: 72x CPU/GPU nodes with 2x AMD EPYC 7763 (64 cores) and 1 TByte DRAM and 8x NVIDIA A100 GPUs each</p> <p>Cloud partition: 36x CPU nodes with 2x AMD EPYC 7H12 (64 cores) and 256 GByte DRAM each</p> <p>Data Analytics system (SMP): 36x CPU nodes with 2x Intel Xeon 8268 (24 cores) and 24 TByte DRAM in total</p>	<p>CPU module: 573x CPU nodes with 2x AMD EPYC 7H12 (64 cores) and 512 GByte DRAM each</p> <p>Accelerated module: 200x CPU/GPU nodes with 2x AMD EPYC 7752 (32 cores) and 512 GByte DRAM and 8x NVIDIA A100 each plus 1.9 TByte SSD</p> <p>Large memory module: 20 CPU nodes with 2x AMD EPYC 7H12 (64 cores) and 4096 GByte DRAM each plus 1.9 TByte SSD</p> <p>Cloud module via 20 CPU nodes (2x EPYC 7H12 and 512 GByte DRAM), 4 CPU/GPU nodes (2xAMD EPYC 7452 (32 cores) and 768 GByte DRAM plus NVIDIA H200) and 1 CPU/GPU node (2xAMD EPYC 7452 (32 cores) and 2.3 TByte DRAM plus 8xNVIDIA H200)</p>	<p>CPU partition: 960x CPU nodes with 2x AMD EPYC 7H12 (64 cores) and 256 [768 nodes or 1 TByte [192 nodes] DRAM</p> <p>CPU/GPU partition: 60x CPU/GPU nodes with 2x AMD EPYC 7H12 (64 cores) and 512 GByte DRAM and 4x NVIDIA A100 GPUs each</p>

Interconnect	NVIDIA InfiniBand HDR (200 Gbit/s) using Fat tree topology One HDR connection to x86 CPU nodes, one HDR100 connection to ARM nodes 2 connections to CPU/GPU nodes	NVIDIA InfiniBand HDR (200 Gbit/s) for CPU nodes (one connection each), NDR (400 Gbit/s) for GPU/DGX2 nodes (3 connections each)	NVIDIA InfiniBand HDR (200 Gbit/s) using Fat tree topology One HDR100 connection to each CPU/Cloud node, 4 HDR connections to each CPU/GPU node, 2 HDR connections to Data Analytics system	NVIDIA InfiniBand HDR (200 Gbit/s) using Dragonfly+ topology One HDR connection to each CPU/Large memory node, 2 HDR connections to each CPU/GPU node	NVIDIA InfiniBand HDR (200 Gbit/s) using Dragonfly+ topology One HDR100 connection to each CPU node, 2 HDR100 connections to each CPU/GPU node
Performance	7.5 PFlop/s peak Linpack	5.9 PFlop/s peak Linpack	12.9 PFlop/s peak Linpack	18.9 PFlop/s peak Linpack	10.1 PFlop/s peak Linpack
Non-conventional compute	Qulacs quantum simulation SW	Specific public details are not available			GPU acceleration (NVIDIA A100), AI/ML optimized nodes
Compute modes	Linux shell on login and compute nodes Slurm jobs Singularity containers	Linux shell on login and compute nodes Slurm jobs	Linux shell on login and compute nodes Slurm jobs Apptainer containers Virtual machines with QEMU	Linux shell on login and compute nodes Slurm jobs Singularity containers	
Limits and guarantees	Specific public details are not available				
Software Services, Interfaces & Stacks					
OS	Rocky Linux.	Linux (specific distribution not specified).	Rocky Linux 8.9.	Linux (specific distribution not specified).	Linux (specific distribution not specified).
Languages, libs	EESSI stack, C++/Fortran/Python on Mathematical libraries (CPU and NVIDIA CUDA for GPU).	C++/Fortran/Python on Mathematical libraries (CPU and NVIDIA CUDA for GPU)	C++/Fortran/Python on Mathematical libraries (CPU and NVIDIA CUDA for GPU)	C, C++, Fortran, Python; MPI, OpenMP, CUDA, ROCm.	C/C++, Fortran, Python, MPI, OpenMP, CUDA, ROCm
Interfaces	SSH, Slurm.	SSH, Slurm.	SSH, Slurm. Easybuild and	SSH, Slurm.	SSH, Slurm (sbatch/srun),

			Spack for SW package management		JupyterHub (for some workflows)
Workflow support	Slurm job steps and chaining.	Batch scripting with Slurm.	Slurm job steps and chaining	Batch scripting with Slurm; support for containerized workflows.	
Monitoring	Slurm energy measurement functionality. Slurm job monitoring; module system for environment management.				
Compute Federation					
Resource discovery	Slurm `sinfo` provides access to partition and node availability information.				
Job routing & load distribution	Slurm scheduler routes jobs to appropriate partitions based on user-defined scripts.				

6.1.2. EuroHPC JU Pre-Exascale and Exascale Systems

Since the procurement of the Alice Recoque system started in September 2024 (under the identifier EUROHPC/2024/CD/0006) and the procurement decision is not known at the time of writing, no firm technical details about the system are available yet. The supercomputer will replace the Joliot-Curie system currently operated by GENCI (see section 6.6).

“Alice Recoque” will deliver a sustained computing power exceeding 1 EFlop/s measured by the Linpack benchmark. Like the JUPITER system, it will be based on a modular supercomputing architecture and provide both accelerated compute resources (using GPUs or other accelerators) for converged HPC or AI (training, fine-tuning) workloads and general-purpose computing resources for traditional HPC or AI (ML or inference at scale) workloads. In addition, a “high-performance data analysis” (HPDA) partition will support interactive or pre/post processing tasks that demand high-performance 3D rendering and visualization capabilities, large memory footprints, and/or access to very fast local storage. A Cloud partition will deliver compute resources for a multi-tenant IaaS/PaaS platform based on OpenStack and Kubernetes. This will enable users to host scientific tools, data, and workflows that can be exposed as services to scientific research communities or the public. CERN and SKA have submitted letters of support for Alice Recoque.

The modular architecture of the selected system should allow for the addition and integration of new hardware services or partitions based on European technologies at a later stage. Examples include hybrid quantum computing (such as 1000+ qubits NISQ solutions or early LSQ models under the French HQI national strategy), cryo- or neuromorphic computing, photonic networking, and DNA storage as a potential replacement of tapes using additional funding from future separate calls.

Table 2: Technical Characteristics of EuroHPC JU Pre-Exascale and Exascale Systems

Technical Analysis – EuroHPC JU Pre-Exascale and Exascale Systems				
System/site	JUPITER (JSC)	Leonardo (CINECA)	LUMI (CSC)	MareNostrum 5 (BSC)
Data/Storage Capabilities and Services				
Storage abstraction	IBM GPFS PFS (Spectrum Scale).	Lustre PFS and local file systems for CPU nodes.	Lustre PFS and Ceph.	IBM GPFS PFS (Spectrum Scale) and local file systems.
Storage	ExaFlash fast /	Fast tier / Metadata	Fast tier / Capacity	Fast tier / Capacity tier.

levels & locality	ExaStore main tier.	Storage / Capacity tier/ local storage (CPU nodes).	tier / Data management service/DMS (Ceph-based).	
Storage capacity	Exaflash: 29 PByte raw. Exastore: 300 PByte raw.	Fast tier + Metadata: 5.7 PByte raw. Capacity tier: 106 PetaByte raw. Local storage: 4 TByte NVMe storage per CPU node.	Fast tier: 7 PByte. Capacity tier: 80 PetaByte. Data management service: 30 PByte.	Fast tier: 4.8 PByte raw. Capacity tier: 376 PByte raw.
Performance (latency & BW)	Exaflash 3000/2000 GByte/s R/W.	Fast tier (Flash): 1400 GByte/s. Capacity tier: 744/620 GByte/s R/W.	Fast tier (Flash): 1740 GByte/s. Capacity tier: 240 GByte/s.	Fast tier (Flash): 1600/1200 GByte/s R/W. Capacity tier: 600 GByte/s.
Access methods & interfaces	POSIX plus GPFS extensions.	POSIX plus Lustre extensions.	Fast/capacity tier: POSIX plus Lustre extensions. DMS: Ceph object storage.	Fast/capacity tier: POSIX plus GPFS extensions. POSIX for local storage, purged after each job.
Data formats	Universal support for data formats which are mapped to POSIX files.			
Data retention & backup	Exaflash for scratch storage only. Exastore for long-term storage. 700 PByte tape storage.	53 PByte archive partition plus 42 GByte temporary partition Regular professional backup.	Fast tier: scratch storage only. Capacity tier: scratch and long-term (project duration) storage. DSM provides long-term (project duration) storage.	Scratch and long-term filesystems on GPFS. Local file systems on the nodes ephemeral. 400 PByte tape storage. Incremental backup .
Access control	Usual Linux ACL mechanisms	Usual Linux ACL mechanisms plus CryptoFS isolation	Usual Linux ACL mechanisms.	Usual Linux ACL mechanisms.
Encryption	Specific public details are not available.	Encryption on Fast and Capacity tier based on user keys.	Specific public details are not available.	Specific public details are not available.
Data Transfer & Federation				
Internal methods & performance	Specific public details are not available.	POSIX enabled protocols.	Fast & capacity tiers: usual Linux data transfer tools.	Specific public details are not available.

			DMS: rclone, s3cmd and restic.	
External methods, security & performance	Specific public details are not available.	Data Movers.	Fast & capacity tiers: SCP and Rsync.	SCP. No external Internet connections to/from compute nodes.
External data federation	Specific public details are not available.	Specific public details are not available.	Specific public details are not available.	Specific public details are not available.
Compute Capabilities and Services				
Processors/ Accelerators	XPU "Booster" module with NVIDIA Grace-Hopper XPU (ARM CPU/ H100 GPU). CPU "Cluster" module with SiPearl "Rhea" ARM CPU.	CPU/GPU module with Intel Xeon "Ice Lake" CPU and NVIDIA A100 GPU (64 GByte HBM2e). CPU module with Intel Xeon "Sapphire Rapids" CPU.	CPU and CPU/GPU partitions with AMD x86 EPYC "Trento" CPU. AMD MI250x GPUs (128 GB HBM2e memory).	Accelerated module with Intel Xeon "Sapphire Rapids" CPU and NVIDIA H100 GPU (64 GByte HBM2e). GPP module with Intel Xeon "Sapphire Rapids" CPU. XPU module with NVIDIA Grace-Hopper XPU (ARM CPU/ H100 GPU).
Nodes	Booster: ~ 6000 nodes with 4xNVIDIA GH200 XPU (72 cores and 120 GByte DRAM and 96 GByte HBM3). Cluster: ≥ 1300 nodes with 2xSiPearl Rhea-1 (80 cores, 64 GByte HBM2e) and 512 GByte (some to have 1 TByte) DRAM.	3456 CPU/GPU nodes with 1x Intel Xeon 8358 (32 cores) and 4xNVIDIA A100 GPU and 512 GByte DRAM. 1536 CPU nodes with 2x Intel Xeon 8480+ and 512 GByte DRAM.	2978 CPU/GPU nodes with 1x AMD EPYC 7A53 (64 cores) and 4xAMD MI250x and 512 GB DRAM. 2048 CPU nodes with 2xAMD 7763 (64 cores) and 256 GByte [1888 nodes], 512 GByte [128 nodes] and 1024 GByte DRAM.	1120 CPU/GPU nodes with 2xIntel Xeon 8460Y+ (40 cores) and 4xNVIDIA H100 GPU and 512 GByte DRAM; in addition 480 GByte NVMe storage. 6408 GPP nodes with 2xIntel Xeon 8480+ (56 cores) and 256 / 1024 [216 nodes] / 2028 [10 nodes] GByte DRAM; in addition 960 GByte NVMe storage. 72 GPP/HBM nodes with 2xIntel Xeon 8480+ (56 cores) and 128 GByte HBM; in addition 960 GByte NVMe storage.
Interconnect	NVIDIA InfiniBand NDR (400 Gbit/s) with 4 connections to each	NVIDIA InfiniBand HDR (200 Gbit/s) with 2 connections to	HPE Slingshot-11 200 Gbit/s connected to each GPU (4	NVIDIA InfiniBand NDR200 (200 Gbit/s) with 4 connections to

	Booster node and 1 connection to each Cluster node using Dragonfly+ topology.	CPU/GPU and 1 connection to CPU nodes using Dragonfly+ topology. Connections to Fast & Metadata tiers with 800 GBit/s, to Capacity tier with 400 Gbit/s.	connections by CPU/GPU node) and with 1 connection per CPU node using Dragonfly topology.	each CPU/GPU node and 1 connection to each CPU node using Fat tree topology.
Performance (Compute, memory)	>= 1 EFlop/s peak Linpac . GH200 peak BW of 4 TByte/s for HBM3, 500 GByte/s for the DRAM).	315 PFlop/s peak Linpack.	589 PFlop/s peak Linpack.	314 PFlop/s peak Linpack.
Non-conventional compute	Potential connection with the HPCQS Quantum Computer at JSC and other QC systems.	Potential connection with the EuroQCS-Italy Quantum Computer.	Specific public details are not available.	Potential connection with the EuroQCS-Spain Quantum Computer.
Compute modes	Linux shell on login and compute nodes. Slurm jobs. Apptainer, containers.	Linux shell on login and compute nodes. Slurm jobs. Singularity containers.	Linux shell on login and compute nodes. Slurm jobs. Apptainer/Singularity containers.	Linux shell on login and compute nodes. Slurm jobs. Apptainer/Singularity containers.
Limits and guarantees	Specific public details are not available.	Specific public details are not available.	Maximum job execution time: 48 hours.	Specific public details are not available.
Software Services, Interfaces & Stacks				
OS	Rocky Linux.	RedHat Linux.	HPE Cray OS.	Linux.
Languages, libs	C++, Fortran, Python. Mathematical libraries (CPU and NVIDIA CUDA for GPU).	C++, Fortran, Python. Mathematical libraries (CPU and NVIDIA CUDA for GPU).	C++, Fortran, Python. Mathematical libraries (CPU and AMD ROCm for GPU).	C++, Fortran, Python. Mathematical libraries (CPU and NVIDIA CUDA for GPU).
Interfaces	Easybuild for SW package management.	Spack for SW package management.	Easybuild and Spack for SW package management.	Easybuild and Spack for SW package management.
Workflow support	Slurm job steps and chaining.			Slurm job steps and chaining. PyCOMPSs.
Monitoring	Specific public details are not available.			

Compute Federation	
Resource discovery	Slurm mechanisms.
Job routing & load distribution	Queue selection and specification of job properties by the user.

6.1.3. Mid-Range Systems

The Arrhenius supercomputer at Linköping University will be operated by the National Academic Infrastructure for Supercomputing in Sweden and support scientific simulations and AI/ML on open and sensitive data. A request for proposals was issued in November 2024, and the system is expected to be installed in 2025. Currently, the procurement process is ongoing, and no details about the decision and the resulting system architecture and configuration are publicly known. Early presentations²⁰ mention a target size of approx. 1000 dual-processor CPU and 250 quad-GPU nodes, and a performance of more than 40 PFlop/s is expected²¹.

The Daedalus supercomputer will be operated by GRNET (National Infrastructure for Research and Technology) at the Lavrion Technological Cultural Park in Athens/Greece. In April 2025, a procurement decision was reached, and the system is planned to become operational early in 2026 for use. It will utilize the NVIDIA GH200, a combined CPU/GPU processor, and offer 1 PByte of high-performance NVMe storage and 10 PByte of capacity storage. The peak HPL performance of Daedalus is expected to be around 115 PFlop/s.

6.1.4. AI Factories

The 13 approved AI factories are listed below, with relevant publicly available information; since procurements have to be started or are in progress at the time of writing, the amount of hard technical data is highly limited. All AI factories will comply with the applicable European AI and data protection regulations.

- **LUMI AI Factory** (LUMI AIF, hosted by CSC IT Center for Science in Finland)
 - The LUMI AIF will focus on support for manufacturing industries, health and life sciences, communication technologies and networks, climate and weather, materials technology, and languages.
 - It will provide a large, accelerated partition utilising next-generation GPUs offering flexibility in supporting a wide range of compute workloads, from AI training to inference and classical simulations.
 - The LUMI AIF will also leverage a one-of-a-kind experimental quantum computing platform, LUMI-IQ, designed for QC+AI workloads.
- **HammerHAI** (hosted by High-Performance Computing Center Stuttgart (HLRS) in Germany).
 - HammerHAI will focus on with a particular focus on manufacturing industries, engineering, automotive, and prioritise supporting start-ups, and small and medium-sized enterprises (SMEs).
 - HammerHAI will be designed for large-scale AI model training and inference, offer “Cloud-like” usability, and comply with the General Data Protection Regulation (GDPR) and all other applicable European AI and data protection regulations.
- **Pharos** (hosted by the National Infrastructures for Research and Technology (GRNET) in Greece alongside Daedalus).
 - Pharos will be focused on health, culture & language, and sustainability.
 - It will deploy the collaborative “Pharos hub” to support startups and SMEs in AI model development, deployment, and scaling.

²⁰ See <https://www.naiss.se/wp-content/uploads/2024/02/Gert-Svensson-Arrhenius-procurement-Dec2023.pdf> for details.

²¹ See https://eurohpc-ju.europa.eu/way-open-build-eurohpc-world-class-supercomputer-sweden-2024-06-18_en for details.

- **IT4LIA** (hosted by CINECA Consorzio Interuniversitario in Italy, alongside Leonardo).
 - IT4LIA will focus on applications in agritech, cybersecurity, earth sciences (weather/climate/environment), and manufacturing.
 - It will provide support for AI research and innovation, including HPC-enabled workflows, AI model development and optimization, security testing, and regulatory compliance.
- **Luxembourg AI Factory** (L-AI, hosted by LuxProvide in Luxembourg alongside MeluXina).
 - L-AI will rely on a new AI supercomputer MeluXina-AI and focus on cybersecurity, finance, green economy and space.
 - It will support Cloud interfaces, multi-tenant and multi-site processing of dynamic workloads. In addition, it will offer highly secure processing environments for private AI and support an end-to-end computing continuum through compute and data bridges.
- **BSC AI Factory** (hosted by Barcelona Supercomputing Center (BSC-CNS) in Spain, alongside MareNostrum 5).
 - The BSC AI Factory will rely on an AI-optimised partition of its Marenostrum 5 supercomputer and target the sectors of healthcare, public administration, climate, biotech, energy, communication, finance, and legal services.
- **MIMER** (hosted by the National Academic Infrastructure of Supercomputing (NAISS) in Sweden).
 - MIMER will focus on the areas of gaming, autonomous systems, life science and materials sciences.
- **AI:AT** (hosted by the Technical University of Vienna in Austria).
 - AI:AT will focus on the sectors biotechnology, sustainability, energy, and manufacturing.
- **Brain++** (hosted at Sofia Tech Park in Bulgaria, alongside Discoverer).
 - Brain will focus on LLMs for the Bulgarian language, robotics, space observation, and product manufacturing (with an emphasis on consumer goods).
 - It will rely on a new, AI-optimised supercomputer Discoverer++.
- **AI2F** (hosted by GENCI in France, using Alice Recoque).
 - AI2F will focus on the defense, energy, aerospace, edtech, agriculture, finance, humanities, robotics, health, earth science, materials science, and mobility sectors.
 - It will rely on Alice Recoque, the second EuroHPC Exascale supercomputer.
- **JAIF** (hosted by Jülich Supercomputing Centre in Germany, alongside JUPITER).
 - For training, JAIF will rely on the JUPITER supercomputer, and add a module designed for AI inference to that system.
 - It will focus on sectors like healthcare, energy, climate change, education, media, the public sector and finance.
- **PIAST** (hosted by the Poznan Supercomputing and Networking Center (PSNC) in Poland).
 - PIADT will focus on the areas of health care, cybersecurity, robotics, and sustainable development.
- **SLAIF** (hosted by IZUM in Slovenia).
 - SLAIF will rely on a new AI-optimised supercomputer replacing the current Vega system, with a performance of 10 EFlop/s in mixed precision mode and 100 PFlop/s for double precision floating point computation.
 - It will focus on energy systems, environmental monitoring, smart agriculture, personalized medicine, drug discovery, language technologies in digital services, media and creative industries, as well as materials science.

6.1.5. Quantum Computer Systems

- 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 access policies for these two centers.

6.1.6. EuroHPC JU Federation Platform

After a public tender²², the implementation and deployment of the EuroHPC Federation Platform (EFP) did start beginning of 2025, with the objective of completing a first version (“minimum viable platform” in 2025 and starting deployment and operation to the EuroHPC JU HPC e-Infrastructure, the Quantum Systems and in perspective also the AI factories in Q1/2026, with a target of nine HPC/Quantum and AI hosting sites. The EFP is a purely software platform that utilizes the existing and evolving EuroHPC e-Infrastructure.

The first version available at the beginning of 2026 will cover the authentication and authorisation infrastructure (AAI), common services for resource allocation and management by the end-user, and a workflow manager that can handle complex, multi-site workflows and provides initial hierarchical scheduling functionality.

The EFP planning foresees another four releases up to 2029:

- Version 2 planned for end of 2026, with initial of Fenix, Simpl, and EOSC, and the new EuroHPC JU peer review platform for access submissions, to be rolled out across the full EuroHPC JU e-Infrastructure.
- Version 3 planned for end of 2027 with deeper integration of Fenix, Simpl, and EOSC, and initial integration of other data lakes/repositories (list of such to be determined).
- Version 4 planned for end of 2028, completing the integration of Fenix, Simpl, and EOSC, and the “other data lakes/repositories”.
- Version 5 planned for Q3/2029, with functionality/improvements to be determined.

The EFP is developed by a consortium of **CSC- IT Centre for Science** (Finland, consortium lead), **IT4Innovations National Supercomputing Centre** at the Technical University of Ostrava (Czechia), **GÉANT** (Netherlands), **Ghent University** (Belgium), **University of Tartu** (Finland), and **NORDUnet** under a commercial contract. The EuroHPC JU has expressed its intention to ensure support and further adaptation/extension of the EFP after the end of the contracted development/deployment.

Table 3: Technical Characteristics of the EuroHPC Federation Platform

Technical Analysis – EuroHPC Federation Platform	
System/site	The EuroHPC Federation Platform (EFP) is a software platform that will be deployed on the HPC, Quantum, and AI systems of the EuroHPC JU e-Infrastructure; the software itself will be open source and available for use outside of this e-Infrastructure.
Data/Storage Capabilities and Services	
Storage abstraction	The EFP will provide access to the storage systems and services at the EuroHPC JU sites, as described in Table 1 and Table 2 above. It does in and of itself not add any capacity, data formats, or backup capabilities, and should not have an impact on performance aspects.

²² More information about the tender (EUROHPC/2023/CD/0003) can be found at https://eurohpc-ju.europa.eu/eurohpc-federation-platform_en and <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/tender-details/15701>.

	EFP mentions that they will provide common file management interfaces as part of their interactive use component.
Access control	<p>In the first phase, data access control will continue to be handled in terms of site-local user identities and Linux or local storage system mechanisms.</p> <p>The planned integration with Simpl, EOSC, and potentially other data clouds/lakes will most likely change that, yet no details are publicly known at this time.</p>
Data Transfer & Federation	
Internal methods & performance	EFP might introduce common interfaces as an extension of the common file management interfaces in its interactive use component.
External methods, security & performance	The EFP workflow component will, over time, include support for (automated) data staging between external storage locations and the resources executing the workflow steps on that data. This will presumably also apply to staging out of results data.
External data federation	The EuroHPC JU has stated publicly ²³ that data federation is not seen as being in their purview. However, the EFP will evolve to integrate data federation as provided by Simpl and EOSC, as well as other, not yet identified data repositories/lakes. The first deployment of an initial solution is planned for the end of 2026.
Compute Capabilities and Services	
Processors/ Accelerators	The EFP can provide access to the compute systems and services at the EuroHPC JU sites, as described in Table 1 and Table 2 above. It does not, in and of itself, add any capacity or capability, with the exception of compute federation (see below in this table).
Software Services, Interfaces & Stacks	
OS	The EFP will run on the OS versions deployed at the EuroHPC JU hosting sites; given its nature, most of its components will run on head/login and management nodes and/or Cloud-like partitions at the centers.
Languages, libs	<p>The EFP will provide a “federated software catalogue” based on the results of the European Environment for Scientific Software Installations (EESSI) project²⁴. This will include precompiled binary distributions of common scientific software optimised for the local system architecture. The ultimate objective is to hide the peculiarities of heterogeneous systems from the end user.</p> <p>It should be noted that this approach will certainly support end-users relying on the pre-installed set of key applications very well; however, it does not address end-users working with their own applications or developing these.</p>
Interfaces	<p>The EFP will make a number of common interfaces available across all federated systems:</p> <ul style="list-style-type: none"> • AAI: a single-sign-on system based on MyAccessID as an OPENID Connect or SAML Identity provider; EFP does not implement an identity provider. End-users can go through MyAccessID authentication directly to the desired target resource, or create short-lived SSH certificates for access (similar to the method CSCS uses). • Resource allocation/management: based on the Waldur system, a Web portal will provide project dashboards and enable end-users to see and manage their allocations and request new allocations (a new, automated EuroHPC JU peer review system will be integrated). • HEAppE: REST API for “HPC as a Service” access with Simple application interfaces and functionality targeted to non-expert users.

²³ For instance in the federation session of the EuroHPC Summit in Krakow on March 20, 2025.

²⁴ The EESSI approach is documented in a paper at <https://onlinelibrary.wiley.com/doi/10.1002/spe.3075>.

	<ul style="list-style-type: none"> Interactive access: Web interface for new users, including remote desktop, shell access, and interactive frameworks like Jupyter notebooks.
Workflow support	<p>The EFP will provide capabilities to visually manage distributed workflows, to manage workflow data across the federated resources, and stage it to/from external data resources. APIs to important programming languages (Python and R are specifically mentioned) will be provided. The workflow management backends will interface with traditional batch systems (Slurm) and Cloud-like systems (OpenStack/Kubernetes). It will be based on LEXIS.</p> <p>Over time, smart scheduling hierarchical scheduling services will be deployed, which optimize workflow execution/orchestration according to monitoring data from the federated infrastructure.</p>
Monitoring	The EFP will create a layer for accessing monitoring data across the federated systems in a unified way; details are not publicly known at this time.
Compute Federation	
Resource discovery	The level of support provided by EFP is not clear at this time.
Job routing & load distribution	<p>Through the hierarchical scheduler, EFP will provide ways for automating the routing of tasks to resources.</p> <p>It should be noted that the EFP will operate with site-local access privileges and quotes; it is not planned to federate allocation quotas across the systems.</p>

6.2. GCS e-Infrastructure

Table 4: Technical Characteristics of GCS e-Infrastructure

Technical Analysis – GSC Systems			
System/site	Hunter (HLRS)	JUWELS Cluster (JSC)	SuperMUC-NG phase 1 (LRZ)
Data/Storage Capabilities and Services			
Storage abstraction	Multi-tiered storage architecture, high-performance parallel file systems, and object storage, optimized for GPU-centric workloads.	Hierarchical storage system with GPFS for various data types.	Tiered storage infrastructure with GPFS optimized for HPC workloads.
Storage levels & locality	Multiple storage levels, including high-speed scratch spaces and persistent storage.	Several storage tiers (\$HOME, \$PROJECT, \$SCRATCH, \$FASTDATA, \$DATA, and \$ARCHIVE: long-term storage for infrequently accessed data).	Structured storage hierarchy (\$HOME, \$WORK, \$SCRATCH tiers).
Storage capacity	The specific capabilities are not detailed; nevertheless, the system is designed to handle large-scale data associated with GPU-centric computations.	Multiple data volumes across its various file systems (70+ PB total across all tiers).	Significant storage resources (Over 50 PB combined in WORK and SCRATCH tiers) designed to accommodate the demands of HPC applications.

Performance (latency & BW)	High-speed interconnects and storage systems for data-intensive GPU workloads.	GPFS parallel file system, high I/O bandwidth for large-scale simulations and data analysis; scalable to handle ML/data analytics.	High-performance storage subsystem designed to support the intensive I/O requirements; optimized GPFS for parallel access.
Access methods & interfaces	Access via SSH.	Access via SSH, UNICORE, or Jupyter.	Access via SSH.
Data formats	Supports a wide range of data formats compatible with GPU-accelerated applications and standard HPC workflows.	Accommodates various data formats, including those used in scientific computing, data analysis, and machine learning applications.	Designed to handle diverse data formats prevalent in computational science and engineering domains.
Data retention & backup	Data retention policies and backup strategies aligned with project requirements.	Data retention policies across its storage tiers, with archival storage options for long-term data preservation.	LRZ policy-based retention guidelines, with backup mechanisms in place to safeguard critical data.
Access control	User authentication and authorization mechanisms.	Access control through user accounts, project allocations, and secure authentication protocols, with LDAP/DFN-AAI authentication.	LDAP-based; user authentication and project-based permissions.
Encryption	Encryption protocols to secure data in-transit and at-rest.	Encryption mechanisms for secure data transfer and storage via HTTPS/SSH.	Encryption strategies to protect sensitive data; in-transit SSL protocols.
Data Transfer & Federation			
Internal methods & performance	High-speed internal network infrastructure, efficient data movement between compute nodes and storage systems; NVLink/NVMe with InfiniBand.	High-bandwidth internal network, rapid data transfer between various components of the system, including compute nodes and storage tiers; Mellanox InfiniBand interconnect across tiers.	High-performance internal network designed; Intel Omni-Path; high I/O throughput among storage and compute nodes.
External methods, security & performance	Secure VPN/SSH tunneling.	Globus, SFTP, HTTP(S); firewall-secured endpoints.	B2SHARE, Globus, SCP; with firewall/NAT-secured connections.
External data federation	Supports federated identity via DFN-AAI; EUDAT integration.	Federated access through EUDAT & PRACE.	Integration with EUDAT and B2SAFE for research sharing.
Compute Capabilities and Services			
Processors/Accelerators	752 AMD x86/GPU accelerated processing units.	2511 dual-processor Intel x86 Skylake nodes) and Booster (936 dual-processor AMD x86 Epyc Rome nodes with four NVIDIA A100 GPUs each).	6480 dual-processor Intel x86 Skylake nodes.

Nodes	~700 nodes with 4 GPUs each.	Cluster: 2500 nodes; Booster: 900 GPU nodes.	~6,400 compute nodes.
Interconnect	InfiniBand HDR.	InfiniBand EDR/HDR.	Intel Omni-Path Architecture.
Performance (Compute, memory)	Petascale performance; large HBM GPU memory per node.	Cluster: 12.27 PetaFLOP/s; Booster: 73.02 PetaFLOP/s; scalable shared memory.	~26.9 PF peak; 96 GB DDR4/node.
Non-conventional compute	AI-accelerated training/inference; large shared memory nodes.	Booster GPU nodes for AI/ML.	Fat nodes with 768 GB RAM for memory-bound applications.
Compute modes	Batch and interactive via Slurm.	Slurm-based; interactive Jupyter sessions.	Batch and interactive via Slurm.
Limits and guarantees	Usage quota per project; GPU quotas.	Fair-share scheduling; project-based limits.	CPU-hour limits per project; time quotas.
Software Services, Interfaces & Stacks			
OS	Rocky Linux / RHEL-based.	CentOS 7, transitioning to Rocky Linux.	SUSE Linux Enterprise Server (SLES).
Languages, libs	Python, CUDA, C++, TensorFlow, PyTorch.	C/C++, Fortran, Python, CUDA, Intel MKL.	C/C++, Python, Fortran, Intel MKL.
Interfaces	SSH, REST, Jupyter.	SSH, UNICORE, JupyterLab.	SSH, module system.
Workflow support	Slurm + Jupyter; GitHub CI/CD integration.	Slurm, CWL, Nextflow, Snakemake.	Slurm, PBS-like scripting.
Monitoring	Grafana dashboards; job-level metrics.	JupyterLab monitoring, Nagios-based system; LLview.	Monitoring via DCDB, Grafana.
Compute Federation			
Resource discovery	Centralized directory, API-driven.	UNICORE registry, SLAM.	PRACE registry, SAML-based discovery.
Job routing & load distribution	Project-scheduled, AI-aware schedulers.	UNICORE job delegation, Slurm federation.	Manual delegation; PRACE-wide batch job routing.

6.3. NHR Alliance e-Infrastructure

Table 5: Technical Characteristics of NHR Alliance e-Infrastructure

Technical Analysis – NHR Alliance Systems					
System/site	Barnard/Capella (TU Dresden)	CLAIX (RWTH Aachen)	Goethe-NHR (GU Frankfurt)	HoreKa (KIT)	MOGON NHR (JGU Mainz)
Data/Storage Capabilities and Services					
Storage abstraction	Shared parallel file system	Multiple file systems	General-purpose storage systems;	Multiple storage systems	Parallel file system architecture, with a

	across Barnard and Capella clusters.	including project and work group spaces.	public details are not available.	including Lustre and BeeOND.	unified storage environment for all compute nodes.
Storage levels & locality	Tiered storage with fast staging storage (1 PB WekaIO) for Capella.	Tiered storage with BeeOND for temporary high-performance storage.	Tiered storage solutions, including high-performance temporary storage options for scratch data.	Tiered storage with high-performance temporary storage (BeeOND).	Tiered storage solutions, including high-performance temporary storage options like BeeOND for scratch data.
Storage capacity	A shared 40 PB parallel file system is available across all HPC systems.	GPFS-based storage system from DDN offers a capacity of approx. 4 PB; Lustre-based storage system from DDN based on Exascaler5 technology provides 26 PB of capacity.	Specific public details are not available.	Lustre-based parallel file system with a capacity of 15 PB.	Parallel file system has a capacity of 8 PB, with an aggregated bandwidth of 98 GB/s.
Performance (latency & BW)	Capella offers over 1 TB/s bandwidth, optimized for data-intensive applications.	GPFS-based storage system with a bandwidth of 80 GB/s; Lustre-based storage system with a bandwidth of 500 GB/s.	Specific public details are not available.	Specific public details are not available.	HDR Infiniband interconnects.
Access methods & interfaces	Access via shared file system; interactive use through Jupyter notebooks.	Access via SSH, JupyterHub, with a wide range of software modules.	Access via SSH; environment modules for software management.	Access via SSH, JupyterHub, with a wide range of software modules.	Access via SSH; environment modules for software management.
Data formats	Supports standard HPC data formats.				
Data retention & backup	Specific policies on data retention and backup are managed by the Center for Information Services and HPC (ZIH).	Specific policies on data retention and backup are managed by the university's IT Center.	Specific policies on data retention and backup are managed by the university's computing center.	Specific policies on data retention and backup are managed by KIT's computing center.	Specific policies on data retention and backup are managed by the JGU Center of Data Processing (ZDV).
Access control	User accounts with password authentication; SSH key management required.				
Encryption	SSH-based encryption for data in-transit.				

Data Transfer & Federation					
Internal methods & performance	High-speed internal interconnects for efficient data transfer and data-intensive applications.				
External methods, security & performance	Data transfer to and from external systems is typically conducted via secure protocols like SCP or SFTP over SSH; performance depends on external network conditions.				
External data federation	Specific public details are not available.				
Compute Capabilities and Services					
Processors/Accelerators	<ul style="list-style-type: none">• Barnard: Intel "Sapphire Rapids" CPUs.• Capella: 576 NVIDIA H100 GPUs.	Intel Xeon 8468 Sapphire CPUs; NVIDIA H100 GPUs.	<ul style="list-style-type: none">• CPUs:<ul style="list-style-type: none">– Intel Xeon Skylake Gold 6148– Intel Xeon Broadwell E5-2640 v4.• GPUs:<ul style="list-style-type: none">– AMD EPYC 7452.	HoreKa employs AMD EPYC 7763 CPUs and NVIDIA A100 GPUs.	<ul style="list-style-type: none">• CPU Nodes: Each node is equipped with 2× AMD EPYC 7713 processors, each having 64 cores, totaling 128 cores per node.• GPU Nodes:<ul style="list-style-type: none">– 2 nodes with 4× AMD MI250 GPUs.– 7 nodes with 8× NVIDIA A40 GPUs.– 11 nodes with 4× NVIDIA A100-SXM4 40GB GPUs.– 4 nodes with 8× NVIDIA A100-SXM4 80GB GPUs.
Nodes	<ul style="list-style-type: none">• Barnard: 720 compute nodes.• Capella: 144 nodes, each with 2 AMD CPUs and 4 GPUs.	632 CPUs nodes and 52 GPU nodes.	623 CPUs nodes and 112 GPU nodes.	~799 out of 813 CPU nodes are dedicated to computations and 60 GPU nodes (HoreKa Green + HoreKa Blue).	The cluster comprises 590 compute nodes, totaling ~75,000 CPU cores.
Interconnect	Specific public details are not available.	Specific public details are not available.	Specific public details are not available.	High-speed InfiniBand HDR interconnect.	Interconnected via HDR Infiniband in a non-blocking fat-tree topology.
Performance (Compute, memory)	Barnard: 104 cores and 512 GB RAM per node; 90 nodes with 1 TB RAM; Capella: 768 GB DDR5 RAM per node.	Specific public details are not available.	Specific public details are not available.	Specific public details are not available.	The system delivers a peak performance of 2.8 PFLOPS, with node memory configurations ranging from 256 GB to 2 TB.
Non-conventional	Capella	GPU nodes	GPU nodes	GPU nodes	GPU nodes support

onal compute	optimized for machine learning and deep learning applications.	support specialized workloads, including machine learning and deep learning applications.	support specialized workloads, including machine learning and deep learning applications.	support specialized workloads, including machine learning and deep learning applications.	specialized workloads, including machine learning and deep learning applications, leveraging high-performance accelerators.
Compute modes	Supports both batch and interactive computing with Jupyter notebooks.	Supports batch processing via SLURM and interactive computing through JupyterHub.	Batch processing via SLURM; interactive sessions discouraged on login nodes.	Batch processing via SLURM; interactive sessions discouraged on login nodes.	Batch processing via SLURM.
Limits and guarantees	Job restrictions and partitions managed via SLURM.				Job runtime limits vary by partition, ranging from 6 to 12 days. Specific partitions cater to different resource requirements, such as high-memory or GPU-intensive jobs.
Software Services, Interfaces & Stacks					
OS	All nodes operate on a standardized Linux distribution; not specified.	Rocky Linux 9.	AlmaLinux 9.	All nodes operate on a standardized Linux distribution; not specified.	All nodes operate on a standardized Linux distribution; not specified.
Languages, libs	Pre-installed software environments.	Wide range of software available through the module system.	Software managed via environment modules; users can install custom software.	Wide range of software available through the module system.	Wide range of software available through the module system.
Interfaces	Terminal, remote desktop, with job submission and management handled through SLURM; shared file systems with Jupyter notebooks.	Terminal, remote desktop, with job submission via SLURM and HPC JupyterHub.	Terminal, remote desktop, with job submission and management handled through SLURM.		
Workflow support	Workflow management tools compatible with SLURM.				
Monitoring	Job monitoring available through SLURM.				

Compute Federation	
Resource discovery	Available resources and partition information using SLURM commands.
Job routing & load distribution	Managed via SLURM partitions.

6.4. RES e-Infrastructure

Table 6: Technical Characteristics of RES e-Infrastructure

Technical Analysis – RES e-Infrastructure				
System/site	General (RES)	Mare Nostrum 5 (BSC)	Finisterrae (CESGA)	La Palma (IAC)
Data/Storage Capabilities and Services				
Storage abstraction	Multiple file systems depending on the site including GPFS, Lustre, Tape, Unified Storage.	Shared parallel file system across GPP and ACC (IBM Storage Scale/GPFS), NVMe, HDD, Tape).	Parallel filesystem (Lustre), NVMe, HDD, Tape.	Parallel filesystem (Lustre/GPFS), HDD, Tape.
Storage levels & locality	Depending on the site: Node-local, shared parallel FS, tape/archive.	Node-local NVMe (scratch), shared parallel FS (project, home), archive tape.	Node-local SSD, shared parallel FS (scratch, home, project), tape.	Node-local, shared parallel FS, archive.
Storage capacity	Site specific.	650 PB total: 2.48 PB NVMe (performance), 248 PB HDD (data), 400 PB tape (archive).	~20 PB (shared), multi-PB tape.	~PB scale (shared), multi-PB tape.
Performance (latency & BW)	Site specific – NVMe –low BW, low latency, Lustre high BW, infiniband.	NVMe: very low latency, high BW; HDD: high BW; InfiniBand NDR200 (100–800 Gb/s per node).	NVMe: low latency, high BW; Lustre: high BW; Infiniband.	High BW (InfiniBand), low latency (NVMe).
Access methods & interfaces	Depending on the site: SSH, SCP, SFTP, BBP, POSIX, MPI-IO.	POSIX, SSH, SCP, SFTP, GPFS tools, IBM Storage Scale, REST (some), NFS.	POSIX, SSH, SCP, SFTP, Lustre tools, NFS.	POSIX, SSH, SCP, SFTP, NFS.
Data formats	Any (POSIX), scientific formats.	Any POSIX-compliant, scientific (HDF5, NetCDF, etc.), user-defined.	Any POSIX-compliant, scientific (HDF5, NetCDF, etc.).	Any POSIX-compliant, scientific.
Data retention & backup	Depending on the site: Tape backup, daily/periodic,	Daily/periodic backup, tape archive, user-initiated,	Daily/periodic backup, tape archive, user/project-based	Tape archive, periodic backup.

	user-initiated, site policy.	project-based retention.	retention.	
Access control	User accounts with SSH key authentication.			
Encryption	At rest (optional, especially on tape/NVMe), data-in-transit (SSH/SCP).			
Data Transfer & Federation				
Internal methods & performance	Site specific: High-speed internal networks (Infiniband/OPA); parallel FS; multi-GB/s to TB/s bandwidth; node-local and shared storage.	Infiniband NDR200 fabric; internal data transfer between storages; parallel file system (GPFS); up to 1.2 TB/s aggregate write bandwidth; TB/s-class internal.	Infiniband interconnect; Lustre parallel FS; high internal bandwidth (multi-GB/s); node-local SSD for scratch.	Infiniband or high-speed Ethernet; parallel FS (Lustre/GPFS); high internal bandwidth.
External methods, security & performance	Site specific: SSH/SCP, SFTP, rsync, Globus, GridFTP; SSH-based security; external bandwidth typically 1–100 Gbps; performance depends on site and peering.	Data Transfer Services (DTS) using parallel SCP, SSHFS, SFTP, Globus-URL-Copy; tested at 500–800 MB/s on 10 Gbps links; SSH/SCP for security; external transfers controlled and logged; supports asynchronous and parallel transfers; external links up to 10–100 Gbps.	SCP, SFTP, Globus, rsync; SSH-based security; performance depends on external link (typically 1–10 Gbps); external transfers coordinated by user/project.	SCP, SFTP, rsync, SSHFS; SSH-based security; performance depends on external connection (1–10 Gbps typical).
External data federation	Most RES sites support data federation via GridFTP, Globus, CVMFS, and integration with European and global data infrastructures.	Supports integration with federated data services (e.g., CMS workflows, PIC storage); DTS manages output to external RSEs; can use CVMFS, GridFTP, and other data federation tools; supports workflow integration with external scientific data grids.	Supports data federation via GridFTP, Globus, and integration with scientific data grids; can participate in multi-site workflows.	Supports data federation via GridFTP, Globus, CVMFS; participates in distributed scientific data workflows.
Compute Capabilities and Services				
Processors/Accelerators	Site specific – including General Purpose CPUs, and Accelerated partitions with GPUs	• GPP (General purpose partition): 6 408 nodes: 2x Intel Shappire Rapids 8480+ 56 cores 2 GHz 72 nodes: 2x Intel Shappire Rapids 03H-LC 56 cores 1.7	• CPUs: Intel Xeon Ice Lake 8352Y 32 cores 2.2 GHz • GPUs: –64 nodes: 2x NVIDIA A100	Intel Xeon E5–2670/1600 20M 2.6 GHz

		GHz. • ACC(Accelerated partition): 2x Intel Shappire Rapids 8460Y+ 40 cores 2.3 GHz. 4x NVIDIA Hopper 64 GB HBM	1 node: 5x NVIDIA A100 1 node: 8x NVIDIA A100	
Nodes	Site specific – including General Purpose CPUs, and Accelerated partitions with GPUs	GPP: 6 480 nodes (725 760 cores) 1 120 nodes (89 600 cores + GPU) – ACC	273 CPU nodes (17 472 cores) and 66 GPU nodes (4 224 cores)	252 nodes (4 032 cores)
Interconnect	Site specific: High-speed internal networks (Infiniband/OPA); parallel FS; multi-GB/s to TB/s bandwidth; node-local and shared storage.	Infiniband NDR200.	Infiniband	Mellanox Infiniband FDR10
Performance (Compute, memory)	Site specific	GPP: RAM: 6 192x256 GB, 216x1024 GB, 72x160 GB Storage: 960GB NVMe per node. ACC: RAM: 512 GB/node (CPU) + 256 GB/node (GPU) Storage: 480 GB NVMe per node.	CPU – 16 FAT nodes with 2048 GB of RAM 1 OPTANE node with 8192 GB of RAM	2x Intel Xeon E5-2670/1600 20M 2.6 GHz and 32 GB of RAM
Non-conventional compute	N/A	N/A	N/A	N/A
Compute modes	Site specific – including batch & interactive computing	Batch processing via slurm. Limited interactive access for debugging, development or visualization tasks	Batch processing via slurm. Limited interactive access for debugging, development or visualization tasks	Batch processing via slurm.
Limits and guarantees	Job restrictions and partitions managed via SLURM.			
Software Services, Interfaces & Stacks				
OS	Site specific – Linux based	Red Hat Enterprise Server	Rocky Linux 8.4	Linux SuSE 12
Languages, libs	Wide range of software available.			

Interfaces	Site specific –including terminal, remote desktop, with job submission via SLURM	SSH login nodes, remote desktop with job submission and management handled through SLURM; parallel file systems & NVMe local storage	SSH login nodes, remote desktop with job submission and management handled through SLURM; parallel file systems (Lustre) & NVMe SSD local storage	Terminal, remote desktop, with job submission and management handled through SLURM.
Workflow support	Workflow execution supported via SLURM.			
Monitoring	Job monitoring available through SLURM.			
Compute Federation				
Resource discovery	Available resources and partition information using SLURM commands.			
Job routing & load distribution	Managed via SLURM partitions.			

6.5. GENCI e-Infrastructure

Table 7: Technical Characteristics of GENCI e-Infrastructure

Technical Analysis – GENCI e-Infrastructure			
System/site	Adastra (CINES)	Jean Zay (IDRIS)	Joliot Curie (CEA)
Data/Storage Capabilities and Services			
Storage abstraction	Features a Cray ClusterStor E1000 system.	Hierarchical storage system with multiple levels.	Features a Lustre file system.
Storage levels & locality	Several storage tiers (\$HOME, \$WORK, \$SCRATCH, \$STORE).	Several storage tiers (\$HOME, \$WORK, \$SCRATCH, \$STORE).	Several storage tiers (\$HOME, \$WORK, \$SCRATCH, \$STORE, \$BurstBuffer).
Storage capacity	~55 PB total.	~55 PB total.	~140 PB total.
Performance (latency & BW)	High BW via InfiniBand HDR.	Up to 1 TB/s (Lustre), low latency.	High BW, low latency (InfiniBand).
Access methods & interfaces	POSIX-compliant access, Slurm, SSH, API (AI/ML).	POSIX-compliant access, SSH, REST for AI.	POSIX-compliant access, batch system, SSH.
Data formats	User-defined, supports standard HPC data formats (e.g., HDF5, NetCDF).	User-defined, supports standard HPC data formats (e.g., HDF5, NetCDF).	User-defined, supports standard HPC data formats (e.g., HDF5, NetCDF).

Data retention & backup	Regular backups, quotas.	Periodic, user-specific.	Regular backups, retention aligned with project duration.
Access control	User authentication and project-based access.		
Encryption	SSH-based encryption for data in-transit.		
Data Transfer & Federation			
Internal methods & performance	Lustre shared FS, NVMe cache.	High-speed NVMe to Lustre.	Fast shared FS, hierarchical.
External methods, security & performance	GridFTP, SCP, Globus, VPN.	SCP, GridFTP, Globus (soon).	SCP, GridFTP, secured API.
External data federation	ELIXIR, EuroHPC integration.	EUDAT, EuroHPC pilot.	EOSC, PRACE, EuroHPC-ready.
Compute Capabilities and Services			
Processors/Accelerators	<ul style="list-style-type: none">CPU: AMD EPYC Genoa (96-core).GPU: AMD EPYC Trento – MI250X.	<ul style="list-style-type: none">CPU: Intel Cascade Lake (various generations).GPU: NVIDIA V100, A100, H100.	<ul style="list-style-type: none">CPU: Intel Skylake, AMD EPYC Rome.GPU: Intel Cascade Lake.
Nodes	<ul style="list-style-type: none">CPU Nodes: 544.GPU Nodes: 356.	<ul style="list-style-type: none">Scalar Nodes: 720.Accelerated Nodes: 823 (V100), 52 (A100), 364 (H100).	<ul style="list-style-type: none">CPU Nodes: ~3900.Accelerated Nodes: 32.
Interconnect	Slingshot-11 200.	OPA 100 / Infiniband NDR400.	InfiniBand HDR100.
Performance (Compute, memory)	~90 PFLOPS.	~125 PFLOPS.	~20 PFLOPS.
Non-conventional compute	Optimized for AI, mixed CPU-GPU.	Deep learning, hybrid computing.	AI training, hybrid simulations.
Compute modes	Supports batch processing via SLURM.	Supports batch processing via SLURM and interactive computing through JupyterHub.	Supports batch processing via SLURM, and containers.
Limits and guarantees	Time/project-based quotas.	Fair-share scheduler.	Tiered QoS, job limits.
Software Services, Interfaces & Stacks			
OS	Linux distribution; Cray Linux / SUSE.	Linux distribution; CentOS / RHEL-based.	Linux distribution; CentOS / RHEL.
Languages, libs	Wide range of software available, C, C++, Fortran, Python, ROCm libs.	Wide range of software available, C, C++, Python, TensorFlow, PyTorch.	Wide range of software available, C, C++, Python, CUDA, AI libs.

Interfaces	SSH, SLURM, modules, API (AI).	SSH, SLURM, JupyterLab.	SSH, SLURM, REST APIs.
Workflow support	Workflow management tools compatible with SLURM.	Workflow management tools compatible with SLURM + Jupyter.	Workflow management tools compatible with SLURM.
Monitoring	Job monitoring available through SLURM.		
Compute Federation			
Resource discovery	Available resources and partition information using SLURM commands.		
Job routing & load distribution	Managed via SLURM partitions.		

6.6. CSCS e-Infrastructure

Table 8: Technical Characteristics of CSCS e-Infrastructure (Alps system)

Technical Analysis – CSCS e-Infrastructure (Alps system)		
System/site	Alps.Clariden	Alps.Eiger
Data/Storage Capabilities and Services		
Storage abstraction	Utilizes a tiered storage architecture managed by the HPE Data Management Framework (DMF), enabling efficient data movement between different storage tiers.	
Storage levels & locality	Comprises multiple storage tiers, including high-performance SSDs and large-capacity HDDs, with data locality optimized for compute tasks.	
Storage capacity	HDD: 100+10 PiB, SSD: 5+1 PiB, Tape: 2x130.	Shared with Alps.Clariden.
Performance (latency & BW)	The storage system offers high throughput, with the scratch disk extension achieving 1 TB/s throughput.	
Access methods & interfaces	Supports standard POSIX interfaces and integrates with FirecREST, a RESTful API for managing HPC resources, allowing for operations like job submissions and data transfers.	
Data formats	Supports a wide range of scientific data formats compatible with various research applications.	
Data retention & backup	Implements robust data retention policies with backups stored in extensive tape libraries, ensuring data durability and disaster recovery capabilities	
Encryption	Data encryption is applied both at rest and during transit, adhering to industry-standard security practices.	

Data Transfer & Federation		
Internal methods & performance	High-speed data transfer within the infrastructure is facilitated by the HPE Slingshot-11 interconnect, providing 200 Gbps per module.	
External methods, security & performance	FirecREST enables secure and efficient data transfers to external systems, supporting various authentication methods and ensuring data integrity.	
External data federation	The infrastructure supports federated data processing, allowing integration with external data sources and collaboration across different research institutions.	
Compute Capabilities and Services		
Processors/ Accelerators	NVIDIA Grace-Hopper XPU (ARM CPU/ H100 GPU).	AMD x86 EPYC "Rome" CPU.
Nodes	2688 nodes, 4xGH200 XPU (72 cores), 128 GByte LPDDR5 + 96 GByte HBM3 memory.	1024 nodes, 2xEPYC 7742 CPU (64 cores), 256/512 GByte DDR.
Interconnect	HPC Cray Slingshot-11 interconnect, 200 Gbit/s per XPU or per CPU node.	
Performance (Compute, memory)	434.9 PFlop/s HPL.	7,200 TFlop/s HPL.
Non-conventional compute	Supports AI/ML workloads through dedicated Grace-Hopper nodes and containerized environments like Sarus.	
Compute modes	Batch, interactive via JupyterLab, and container workflows.	
Limits and guarantees	Resource allocation and usage policies are governed by CSCS's user regulations, ensuring fair access and system stability.	
Software Services, Interfaces & Stacks		
OS	Runs on a Linux-based OS tailored for high-performance computing environments.	
Languages, libs	Supports programming languages like C, C++, Fortran, and Python, along with scientific libraries such as Cray LibSci and MPI implementations.	
Interfaces	Provides command-line access, RESTful APIs via FirecREST, and web-based interfaces for job submission and monitoring.	
Workflow support	AiiDA, FirecREST, CI/CD.	
Monitoring	System monitoring and performance analysis are facilitated through tools like ReFrame and Cray's performance suite.	
Compute Federation		

Resource discovery	Users can discover available resources and services through CSCS's user portal and documentation.
Job routing & load distribution	Workload management is handled by the Slurm scheduler, ensuring efficient job distribution across the system.

6.7. EPCC e-Infrastructure

Table 9: Technical Characteristics of EPCC e-Infrastructure

Technical Analysis – EPCC e-Infrastructure	
System/site	Archer 2 (UK)
Data/Storage Capabilities and Services	
Storage abstraction	ARCHER2 employs a tiered storage architecture comprising home, work, and RDFaaS (Research Data Facility as a Service) file systems.
Storage levels & locality	<ul style="list-style-type: none"> • Home File Systems: Accessible on login and data analysis nodes; intended for critical data and source code. • Work File Systems: Available on login, data analysis, and compute nodes; optimized for high-performance parallel I/O. • RDFaaS: Provides long-term data storage solutions.
Storage capacity	<ul style="list-style-type: none"> • Home: 1 PB usable space. • Work: 10.8 PB usable space across three Lustre file systems. • Burst Buffer: 1 PB NVMe-based storage for high-speed I/O operations.
Performance (latency & BW)	<ul style="list-style-type: none"> • Memory Bandwidth: 380 GB/s per node. • Interconnect Bandwidth: 25 GB/s per node bi-directional.
Access methods & interfaces	Users access storage via standard POSIX-compliant interfaces.
Data formats	Supports common scientific data formats including HDF5, NetCDF, and MPI-IO.
Data retention & backup	<ul style="list-style-type: none"> • Home File Systems: Regularly backed up with snapshots taken hourly, daily, and weekly. • Work File Systems: Not backed up; users are responsible for data preservation.
Access control	User authentication and authorization managed via the SAFE system, ensuring secure access to resources.
Encryption	Data in transit is secured using SSH protocols.
Data Transfer & Federation	
Internal methods & performance	High-speed data movement facilitated by the Slingshot interconnect and Lustre file systems.
External methods, security &	Data transfer to/from ARCHER2 is performed via secure protocols like SCP and SFTP.

performance	
External data federation	Specific public details are not available.
Compute Capabilities and Services	
Processors/Accelerators	<ul style="list-style-type: none"> Each compute node features dual AMD EPYC 7742 64-core processors (128 cores per node). A testbed with AMD MI210 GPUs is available for development purposes.
Nodes	Total of 5,860 compute nodes: <ul style="list-style-type: none"> 5,276 standard memory nodes (256 GB RAM). 584 high-memory nodes (512 GB RAM).
Interconnect	HPE Slingshot network with a dragonfly topology, providing low-latency, high-bandwidth communication.
Performance (Compute, memory)	<ul style="list-style-type: none"> Peak performance of 28 PFLOP/s. Total system memory of 1.57 PB.
Non-conventional compute	GPU testbed comprising 4 nodes, each with 4 AMD MI210 GPUs, for exploring heterogeneous computing workloads.
Compute modes	<ul style="list-style-type: none"> Exclusive node access per job, ensuring dedicated resources. Support for various parallel computing paradigms, including MPI and OpenMP.
Limits and guarantees	Resource allocations are managed through project-specific quotas and scheduling policies.
Software Services, Interfaces & Stacks	
OS	HPE Cray Linux Environment (CLE) based on SUSE Linux Enterprise Server 15.
Languages, libs	<ul style="list-style-type: none"> Support for multiple programming languages: C, C++, Fortran, Python. Availability of scientific libraries and tools, including MPI, BLAS, LAPACK, and HDF5.
Interfaces	<ul style="list-style-type: none"> Command-line access via SSH. Job submission and management through Slurm.
Workflow support	Users can implement custom workflows using scripting languages and workflow management tools compatible with HPC environments.
Monitoring	System performance and job monitoring facilitated through EPCC's SAFE portal and integrated monitoring tools.
Compute Federation	
Resource discovery	Accessible via the SAFE portal.
Job routing & load distribution	Jobs are scheduled and managed using Slurm.

6.8. ICSC e-Infrastructure

The largest supercomputer in the ICSC e-Infrastructure is CINECA's Leonardo system, the details of which are reported in [Table 2](#) above. [Table 10](#) below reports on the INFN systems.

Table 10: Technical Characteristics of ICSC e-Infrastructure

Technical Analysis – ICSC e-Infrastructure	
System/site	A federation of HPC, Cloud, and Quantum Computing systems, connected via a Tbps network and with common patterns for access, etc. To mention a few of the systems: Leonardo (CINECA, see Table 2) and INFN.
	INFN
Data/Storage Capabilities and Services	
Storage abstraction	INFN employs a federated storage model, integrating various storage systems across its centers, providing a unified data access layer.
Storage levels & locality	The infrastructure supports multi-tiered storage, including high-speed NVMe for caching and SATA disks for bulk storage, distributed across multiple sites.
Storage capacity	Collectively, INFN's distributed infrastructure offers approximately 120 PB of enterprise-level disk space and 100 PB of tape storage.
Performance (latency & BW)	High-performance interconnects like InfiniBand NDR 400G ensure low-latency and high-bandwidth data access.
Access methods & interfaces	Data access is facilitated through standard protocols and interfaces, including POSIX-compliant file systems and web-based portals.
Data formats	User-dependent. Supports a wide range of data formats common in scientific research, ensuring compatibility with various applications.
Data retention & backup	Implements robust data retention policies with regular backups and replication across sites to ensure data integrity and availability. Permanent and scratch; RAID for data safety; backup on tape for certain use cases.
Access control	Utilizes federated identity management systems, including OpenID Connect and OAuth2, to manage user authentication and authorization.
Encryption	N/A on standard systems, possible on EPIC platform (for sensitive data). Data encryption is employed both at rest and in transit to ensure data security across the infrastructure.
Data Transfer & Federation	
Internal methods & performance	High-speed internal networks, including InfiniBand and 100 GbE connections, facilitate efficient data transfer within the infrastructure.
External methods, security & performance	Integration with national research networks like GARR enables secure and high-performance data exchange with external collaborators.

External data federation	The infrastructure is designed to interoperate with other national and international research infrastructures, supporting data federation and collaborative research. WLCG Data federation, federations from other initiatives (e.g., VIRGO, CDF).
Compute Capabilities and Services	
Processors/Accelerators	The compute resources include a mix of CPUs, GPUs (e.g., NVIDIA H100), and FPGAs, catering to diverse computational workloads. Various generations of Intel/AMD x86_64 systems. A few (but increasing) number of aarch64 servers.
Nodes	Compute nodes are equipped with high-core-count CPUs, substantial RAM (up to 1.5 TB per node), and high-speed interconnects. >100,000 cores overall.
Interconnect	Typically, nodes connected at 10 Gbps; storage systems up to 100 Gbps.
Performance (Compute, memory)	The infrastructure supports high-performance computing tasks, with nodes optimized for both compute-intensive and memory-intensive applications. ~1,3 MHS23.
Non-conventional compute	Nodes with FPGA.
Compute modes	Supports various compute modes, including batch processing, interactive sessions, and cloud-based deployments.
Limits and guarantees	Resource allocation policies ensure fair usage among users, with service-level agreements defining performance and availability metrics.
Software Services, Interfaces & Stacks	
OS	Various Linux flavours.
Languages, libs	Supports a wide array of programming languages (e.g., C, C++, Python) and scientific libraries (e.g., MPI, OpenMP).
Interfaces	Provides command-line interfaces, web portals, and APIs for job submission and resource management.
Workflow support	Integrates workflow management tools to orchestrate complex computational tasks across the infrastructure: SLURM, LSF, HTCondor.
Monitoring	Comprehensive monitoring systems track resource usage, performance metrics, and system health (internal).
Compute Federation	
Resource discovery	Federated systems enable users to discover and access resources across multiple sites seamlessly.
Job routing & load distribution	Advanced scheduling batch systems distribute workloads efficiently.

6.9. WLCG e-Infrastructure

Table 11: Technical Characteristics of WLCG e-Infrastructure

Technical Analysis – WLCG e-Infrastructure	
System/site	Worldwide LHC Computing GRID (~160 sites)
Data/Storage Capabilities and Services	
Storage abstraction	Storage seen via multiple protocols.
Storage levels & locality	Disk, tape (so warm and cold storage). In some cases, fast disks are deployed as caches, but this is nearly invisible to the users.
Storage capacity	<ul style="list-style-type: none"> • Disk: 1.2 EB. • Tape: 2.4 EB.
Performance (latency & BW)	Typical figures are with respect to the capability to feed CPUs. A requirement of 5 MB/s/Core and 5 MS/s/TB is typical even if not enforced.
Access methods & interfaces	POSIX, XrootD, SRM, WebDAV.
Data formats	ROOT files (https://root.cern/).
Data retention & backup	Managed by the experiments (centers do NOT operate transfers and deletions).
Access control	X509 and tokens.
Encryption	None.
Data Transfer & Federation	
Internal methods & performance	Mostly POSIX and Xrootd and WebDAV, but NFS not atypical.
External methods, security & performance	FTS mediated, with protocols like Xrootd, SRM, WebDAV. Security via X509 proxies, transitioning to tokens.
External data federation	Rucio-based.
Compute Capabilities and Services	
Processors/A ccelerators	Typically, x86_64 cores. More recently, Power9 and ARM cores available at some sites. Scarce use of accelerators for the moment.
Nodes	~ 1 MCores, nodes depending on density are ~ 64 times less.

Interconnect	10 Gbps enough for computing nodes.
Performance (Compute, memory)	12,8 MHS23; request is 2–4 GB/core, depending on the experiment and whether HT is on/off.
Non-conventional compute	A few FPGA nodes, but not in pledges.
Compute modes	Mostly via GRID, which lands on batch systems. LSF and HTCondor the most typical, SLURM when using HPC centers.
Limits and guarantees	Specific public details are not available.
Software Services, Interfaces & Stacks	
OS	Any recent Linux, since most of the payloads execute vis Singularity virtualization.
Languages, libs	Code is mostly C++, with some Python for configuration/analysis and Fortran for old legacy codes.
Interfaces	Specific public details are not available.
Workflow support	SLURM, LSF, HTCondor.
Monitoring	Complex monitoring system put in place by WLCG: https://monit.web.cern.ch/ .
Compute Federation	
Resource discovery	Mostly using a Pilot (late binding) system, with pilot factories (see https://iopscience.iop.org/article/10.1088/1742-6596/513/3/032086/pdf).
Job routing & load distribution	HTCondor / Panda / Dirac / Aline global queues.

6.10. NIKHEF e-Infrastructure

Table 12: Technical Characteristics of NIKHEF e-Infrastructure

Technical Analysis – NIKHEF e-Infrastructure	
System/site	Nikhef Data Processing Facility (NDPF)
Data/Storage Capabilities and Services	
Storage abstraction	Layered abstraction using dCache for disk-based storage and Enstore for tape archives. This allows transparent access to different storage media.
Storage levels & locality	Multi-tiered: hot (disk) and cold (tape) storage. Local storage at Nikhef and integration with remote Tier-0 (CERN) and other Tier-1/2 sites.
Storage capacity	The disk pool is typically in the range of 15–20 PB, and the tape archive is comparable or larger.

Performance (latency & BW)	High-bandwidth connections (100 Gbps via SURFnet). Disk-based access offers low latency; tape archival introduces higher latency but with high throughput.
Access methods & interfaces	GridFTP, XRootD, SRM (Storage Resource Manager), and WebDAV supported via dCache.
Data formats	Supports HEP-standard data formats including ROOT, HDF5, and binary/raw outputs from LHC experiments.
Data retention & backup	Critical data replicated on tape systems. Tape backup provides long-term archival; redundancy built into the system via distributed copies.
Access control	X.509 certificate-based authentication and VOMS (Virtual Organization Membership Service) for role-based access control.
Encryption	Data in transit encrypted using TLS; support for encrypted storage also present at the hardware and software levels.
Data Transfer & Federation	
Internal methods & performance	Local transfers over high-speed internal networks. Optimized I/O via dCache pools.
External methods, security & performance	Interconnected with LHCONE and SURFnet, allowing secure high-speed data movement between Tier-0, Tier-1, and Tier-2 centers using protocols like GridFTP and FTS (File Transfer Service).
External data federation	Part of WLCG data federation through FTS3 and XRootD; supports remote data access and distributed workflows.
Compute Capabilities and Services	
Processors/Accelerators	Multi-core x86 CPUs, mostly Intel and AMD-based. No significant GPU or FPGA deployments at present.
Nodes	Thousands of cores across several hundred physical servers. Supports batch and interactive workloads.
Interconnect	High-speed Ethernet (10/25/100 Gbps) used within racks and to switches; inter-site links via SURFnet.
Performance (Compute, memory)	HTC-optimized, typical memory configurations range from 4–8 GB per core.
Compute modes	Batch processing primarily via HTCondor; supports interactive and pilot job execution as well.
Limits and guarantees	Resource quotas and fair-share scheduling via HTCondor; SLAs defined through WLCG MoUs.
Software Services, Interfaces & Stacks	
OS	CentOS Stream/AlmaLinux and other RHEL derivatives; migration toward EL9 platforms ongoing.

Languages, libs	Supports C++, Python, ROOT, Geant4, and experiment-specific frameworks (e.g., CMSSW for CMS, Athena for ATLAS).
Interfaces	SSH, Grid middleware interfaces (gLite, EMI, ARC), Web UIs, and API-based access.
Workflow support	Workflow engines like PanDA, DIRAC, and CRAB supported for respective experiments.
Monitoring	Monitoring via Grafana, Prometheus, and WLCG-standard tools (e.g., Dashboard, GGUS integration).
Compute Federation	
Resource discovery	BDII (Berkeley Database Information Index), CRIC, and WLCG top-level info systems.
Job routing & load distribution	HTCondor central manager for job dispatch; integrated with WLCG workload management systems.

6.11. EGI HTC-oriented e-Infrastructure

Table 13: Technical Characteristics of EGI HTC-oriented e-Infrastructure

Technical Analysis – EGI HTC-oriented e-Infrastructure	
System/site	EGI Federated HTC Platform
Data/Storage Capabilities and Services	
Storage abstraction	EGI offers multiple storage types, including block storage (attached to VMs), grid storage (for HTC/HPC scenarios), and object storage (for cloud-native applications and archiving).
Storage levels & locality	Storage is distributed across a federation of data centers, allowing data to be stored close to computational resources, enhancing performance and reducing latency.
Storage capacity	The EGI Federation provides over 580 PB of online storage capacity across HTC and Cloud storage providers.
Performance (latency & BW)	Grid storage is optimized for high-throughput scenarios, ensuring low latency and high bandwidth access to large datasets.
Access methods & interfaces	Data can be accessed through various standard protocols and interfaces, including POSIX, WebDAV, and S3-compatible APIs, facilitating integration with diverse applications.
Data formats	Supports a wide range of data formats, accommodating the diverse needs of different scientific communities.
Data retention & backup	Data retention policies and backup strategies are managed by individual resource providers within the federation, ensuring data durability and availability.
Access control	Utilizes EGI Check-in, a federated authentication and authorization infrastructure (AAI), enabling secure and seamless access to services using institutional credentials.
Encryption	Data encryption is supported both at rest and in transit, depending on the configurations of individual resource providers and user requirements.

Data Transfer & Federation	
Internal methods & performance	EGI's Data Transfer service allows asynchronous transfer of large data sets with features like automatic retries and real-time statistics, ensuring efficient internal data movement.
External methods, security & performance	Supports secure data transfers across organizational boundaries, leveraging federated identity and access management to maintain security and performance.
External data federation	Through services like DataHub, EGI enables access to key scientific datasets, promoting data sharing and collaboration across different research communities.
Compute Capabilities and Services	
Processors/Accelerators	Provides access to a vast array of CPU cores and supports GPU acceleration for compute-intensive tasks, enhancing performance for suitable workloads.
Nodes	Comprises a distributed network of computing centers, offering a massive amount of computing power via a standard interface and virtual organization membership.
Interconnect	High-speed networking infrastructure connects the distributed resources, facilitating efficient data movement and job execution across the federation.
Performance (Compute, memory)	With over 1 million cores of installed capacity, EGI supports the execution of over 1.6 million computing jobs per day, catering to high-throughput computing needs.
Non-conventional compute	Supports containerized workloads (via singularity and udocker) and provides resources for applications requiring specialized computing environments.
Compute modes	Offers batch computing mode and orchestration with the EGI Workload Manager service, interactive sessions support is available in selected data centres of the federation, accommodating diverse computational workflows.
Limits and guarantees	Provides service-level guarantees and dedicated support for long-term access, ensuring reliability and performance for critical research tasks.
Software Services, Interfaces & Stacks	
OS	The compute environments predominantly run on RH-based Linux distributions, tailored to support the specific software stacks required for data processing. Support for containers allow users to run other OS distributions as needed.
Languages, libs	Offers a broad range of programming languages and scientific libraries, facilitating the development and execution of diverse scientific applications. EGI Software Distribution service providers access to a wide collection of software through CVMFS.
Interfaces	Provides standard interfaces and APIs for job submission, data access, and resource management, ensuring interoperability and ease of use.
Workflow support	Supports workflow management tools and platforms, enabling users to design, execute, and monitor complex computational workflows.
Monitoring	Integrated monitoring and accounting tools provide insights into resource availability and consumption, aiding in efficient resource utilization.

Compute Federation	
Resource discovery	Utilizes a federated information system based on the OGF GLUE standard, allowing users to discover and access available resources across the federation.
Job routing & load distribution	Employs workload and data management tools to efficiently manage computational tasks, distributing jobs across the infrastructure to optimize performance and resource utilization.

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

Table 14: Technical Characteristics of SKA HTC/Data-oriented e-Infrastructure

Technical Analysis – SRCNet SKA HTC/Data-oriented e-Infrastructure	
System/site	SKA Regional Centre Network (SRCNet)
Data/Storage Capabilities and Services	
Storage abstraction	SRCNet employs a federated data lake architecture, integrating diverse storage systems across its nodes. This setup utilizes tools like Rucio for data management, ensuring seamless access and interoperability among various storage backends.
Storage levels & locality	Data is distributed across multiple tiers located in regional SRC nodes and includes archives and long-term backups.
Storage capacity	The SKA is projected to generate between 500 and 700 petabytes (PB) of data annually at full operations. SRCNet nodes are scaling their storage infrastructures to accommodate this influx, with capacities expected to reach exabyte scales in the coming years (~2029).
Performance (latency & BW)	High-throughput networks, including at least two 100 Gbps links, connect SRCs via national research and education networks (NRENs) and GÉANT. Data Transfer Nodes (DTNs) are optimized for large data movements, employing technologies like IPv6 and jumbo frames to ensure efficient and secure transfers.
Access methods & interfaces	Users interact with the data through standardized protocols such as WebDAV and GridFTP. These interfaces are designed to be transparent and location-agnostic, facilitating seamless data access across the federated network.
Data formats	Data is stored in formats adhering to the FAIR (Findable, Accessible, Interoperable, Reusable) principles and Virtual Observatory (VO) standards, ensuring compatibility and ease of use across various platforms.
Data retention & backup	SRCNet maintains multiple data replicas across its nodes (at least two copies) to ensure redundancy and resilience. Backup strategies are implemented to safeguard data integrity and support long-term preservation.
Access control	A federated Authentication and Authorization Infrastructure (AAI) provides single sign-on capabilities across SRCs, enabling secure and streamlined user access.
Encryption	Data transfers are secured using encryption protocols like TLS. Storage systems support encryption at rest, ensuring data confidentiality and compliance with security standards.

Data Transfer & Federation	
Internal methods & performance	Within SRCNet, data is transferred using optimized DTNs and high-speed networks. Performance monitoring tools, such as perfSONAR, are deployed to ensure efficient data movement and to identify potential bottlenecks.
External methods, security & performance	Data exchanges with external entities are facilitated through secure protocols and agreements, ensuring data integrity and compliance with international standards. The network infrastructure is designed to handle high-volume data transfers with minimal latency.
External data federation	SRCNet collaborates with other major research infrastructures, such as CERN and the Vera C. Rubin Observatory, to enable cross-disciplinary data sharing and joint scientific endeavors.
Compute Capabilities and Services	
Processors/Accelerators	SRCs utilize a mix of high-performance CPUs and GPUs to handle diverse computational workloads, including data processing, simulations, and machine learning tasks.
Nodes	The compute infrastructure comprises a heterogeneous mix of on-premises HPC clusters, cloud-based resources, and virtualized environments, offering flexibility and scalability.
Interconnect	High-speed interconnects, such as InfiniBand and 100 Gbps Ethernet, facilitate rapid data movement within and between compute nodes, ensuring efficient parallel processing.
Performance (Compute, memory)	SRCs are scaling their compute capabilities to meet the demands of SKA data processing, with performance targets reaching petaflop scales and memory configurations optimized for large-scale data analytics.
Non-conventional compute	Exploratory initiatives are underway to integrate serverless computing models, such as Function-as-a-Service (FaaS), enabling dynamic and event-driven processing of SKA data.
Compute modes	SRCs support various compute modes, including batch processing, interactive sessions, and workflow-based executions, catering to the diverse needs of the scientific community.
Limits and guarantees	Resource allocations are managed based on user requirements and project priorities, with service-level agreements (SLAs) in place to ensure fair and predictable access to compute resources.
Software Services, Interfaces & Stacks	
OS	The compute environments predominantly run on Linux distributions, tailored to support the specific software stacks required for SKA data processing.
Languages, libs	A broad spectrum of programming languages and libraries is supported, including Python, C/C++, and specialized astronomical data analysis tools, facilitating diverse scientific workflows.
Interfaces	User interfaces encompass command-line tools, web portals, and APIs, providing flexible access points for data analysis, job submission, and resource management.
Workflow support	Workflow management systems, such as those based on the Common Workflow Language (CWL), are employed to orchestrate complex data processing pipelines across the distributed infrastructure.
Monitoring	Comprehensive monitoring solutions track system performance, resource utilization, and job statuses, ensuring operational efficiency and facilitating proactive issue resolution.
Compute Federation	

Resource discovery	A centralized registry and metadata catalog facilitate the discovery of available resources, datasets, and services across the SRCNet, enhancing collaboration and resource utilization.
Job routing & load distribution	Advanced scheduling and load-balancing mechanisms distribute computational tasks across the network, optimizing resource usage and minimizing processing times.

6.13. LOFAR – Central Processing (CEP) e-Infrastructure

Table 15: Technical Characteristics of LOFAR – Central Processing (CEP) e-Infrastructure

Technical Analysis – LOFAR – Central Processing (CEP) e-Infrastructure	
System/site	Central signal processing for the LOFAR radio-interferometry instrument. Distributed and federated data lake (LOFAR Long Term Archive: LTA) with sites in Poland (PSNC), Germany (FZJ), and the Netherlands (SURF).
	Central Processing (CEP) Facility
Data/Storage Capabilities and Services	
Storage abstraction	<ul style="list-style-type: none"> • Tier 0: CEP for initial data processing. • Tier 1: Core Long Term Archive (LTA) storage sites. • Tier 2: Access points for the scientific community.
Storage levels & locality	<ul style="list-style-type: none"> • Tier 0: Temporary storage at CEP. • Tier 1: Long-term storage at LTA sites in the Netherlands (SURF), Germany (FZJ), and Poland (PSNC). • Tier 2: Access points for researchers and scientists.
Storage capacity	As of October 2023, the Jülich LTA site in Germany hosts 1.5 PB on disk and 21.6 PB on tape, with an annual growth rate of 2 PB.
Performance (latency & BW)	High-throughput data processing is achieved through the efficient use of GPU resources, with data transfer rates reaching up to 150 Gb/s.
Access methods & interfaces	Data access is facilitated through web interfaces and APIs, enabling users to retrieve and interact with data stored in the LTA.
Data formats	LOFAR utilizes the Hierarchical Data Format version 5 (HDF5) for storing its data products, including beam-formed time-series data and transient buffer board data.
Data retention & backup	Data is retained in the LTA with backup mechanisms in place, including tape storage solutions, to ensure data integrity and availability over the long term.
Access control	Access to data is managed through authentication and authorization protocols to ensure secure and appropriate data usage.
Encryption	Data encryption is employed during transfer and storage to protect sensitive information and maintain data confidentiality.
Data Transfer & Federation	
Internal methods & performance	High-speed internal networks facilitate efficient data transfer between CEP and LTA sites, ensuring timely processing and storage.

External methods, security & performance	Data is transferred to external sites and users via secure protocols, with performance optimized through parallel data streams and bandwidth management.
External data federation	The LTA's federated architecture allows for seamless integration and access to data across multiple international sites, supporting collaborative research efforts.
Compute Capabilities and Services	
Processors/Accelerators	The CEP employs a GPU-based correlator and beamformer known as COBALT for processing data from LOFAR stations.
Nodes	The computing infrastructure includes multiple nodes equipped with GPUs to handle the intensive processing requirements of LOFAR data.
Interconnect	High-speed interconnects, such as InfiniBand, are used to connect compute nodes, facilitating rapid data exchange and processing.
Performance (Compute, memory)	The system is designed to handle several tens of teraflops of processing power, with sufficient memory to support large-scale data processing tasks.
Non-conventional compute	Utilization of GPUs for parallel processing tasks in the COBALT system represents a non-conventional computing approach within the CEP.
Compute modes	Supports both batch and real-time processing modes to accommodate various observational and data analysis requirements.
Limits and guarantees	The system is engineered to meet the demanding processing needs of LOFAR, with scalability to handle increasing data volumes and complexity.
Software Services, Interfaces & Stacks	
OS	The compute nodes operate on Linux-based systems, providing a stable and flexible environment for data processing applications.
Languages, libs	A range of programming languages and libraries is supported, including Python, C++, and specialized astronomical data processing libraries.
Interfaces	User interfaces include command-line tools, web portals, and APIs, facilitating diverse methods of interaction with the system.
Workflow support	Workflow management systems are in place to orchestrate data processing pipelines, ensuring efficient and reproducible analyses.
Monitoring	Comprehensive monitoring tools track system performance, resource utilization, and data processing metrics to maintain optimal operation.
Compute Federation	
Resource discovery	The federated architecture enables the discovery of compute resources across the LTA network, allowing for dynamic allocation based on processing needs.
Job routing & load distribution	Workloads are distributed across available compute resources to balance load and maximize processing efficiency throughout the federated system.

6.14. WLCG Data-Oriented e-Infrastructure

The information for the data-oriented side of WLCG is integrated in [Table 11](#) (Subsection 6.9).

6.15. EBRAINS e-Infrastructure

Table 16: Technical Characteristics of EBRAINS e-Infrastructure

Technical Analysis – EBRAINS e-Infrastructure			
System/site	Cluster OpenStack / CEA (France)	JSC Cloud / JSC (Germany)	CINECA's Galileo100 / CINECA (Italy)
Data/Storage Capabilities and Services			
Storage abstraction	Block and object storage via OpenStack Cinder and Swift.	S3-compatible object storage; Ceph-based backend.	Lustre parallel file system, GPFS for certain projects.
Storage levels & locality	Local SSDs, NAS, hierarchical storage with tape backend.	Local NVMe/SSD, Ceph clusters, long-term archival tiers.	Lustre tiers (scratch, work, archive); locality-aware allocation.
Storage capacity	Tens of petabytes; scalable via OpenStack.	Multiple petabytes across Ceph clusters.	Over 100 PB total including scratch and archive.
Performance (latency & BW)	Low-latency SSD; high bandwidth via InfiniBand/Ethernet.	High throughput Ceph with optimized performance.	High IOPS via Lustre; InfiniBand interconnects.
Access methods & interfaces	REST APIs, POSIX, Swift/Cinder APIs.	S3 API, POSIX via FUSE/CephFS, web portals.	POSIX, SLURM-managed file staging, GridFTP.
Data formats	HDF5, NIfTI, BIDS, JSON, custom EBRAINS formats.	HDF5, NIfTI, BIDS, JSON.	HDF5, NIfTI, NetCDF, EBRAINS-compatible.
Data retention & backup	Tiered storage with backups and archival.	Regular backups, replication, archival storage.	Policy-based retention; backup to tape archive.
Access control	RBAC via OpenStack Keystone.	IAM with federated AAI, per-user/group controls.	LDAP/SSO integration, project-based access rights.
Encryption	TLS in-transit; optional at-rest.	TLS in-transit, Ceph encryption at-rest.	Encryption in-transit; at-rest via filesystem or hardware.
Data Transfer & Federation			
Internal methods & performance	High-speed interconnects, OpenStack networking.	Optimized internal network, scientific data movement.	High-speed InfiniBand + Ethernet for intra-cluster and storage.
External methods, security & performance	Globus, SFTP, HTTPS; VPN/federated login.	Globus, HTTPS, federated AAI; secure links.	GridFTP, Globus, SCP; AAI and VPN-secured.
External data	EBRAINS KG & FENIX	EBRAINS, EUDAT/B2 federation	FENIX data federation, PRACE data

federation	integration.	support.	sharing support.
Compute Capabilities and Services			
Processors/Accelerators	x86 CPUs, optional NVIDIA GPU.	x86 CPUs, NVIDIA GPUs, optional FPGA.	Intel Xeon CPUs, NVIDIA A100 GPUs.
Nodes	Dozens to hundreds of virtual nodes.	Hundreds of VMs or containers.	Over 3,000 nodes (CPU), 160+ GPU-enabled nodes.
Interconnect	High-speed Ethernet or InfiniBand.	Internal fabric; high-speed Ethernet.	InfiniBand HDR100.
Performance (Compute, memory)	Up to 128 GB RAM, high-core VMs.	Up to 1 TB RAM VMs, scalable GPUs.	Up to 512 GB RAM per node, high-core CPUs.
Non-conventional compute	Limited FPGA testbeds.	GPU, neuromorphic support.	Neuromorphic support via SpiNNaker interface, future RISC-V nodes.
Compute modes	Batch, interactive, Jupyter, containers.	Cloud-native, JupyterHub, containers, batch.	Batch (SLURM), interactive, containerized.
Limits and guarantees	Quotas and SLA agreements.	Quotas, fair-share, dynamic scaling.	Quota-based job submission, SLA through PRACE or national calls.
Software Services, Interfaces & Stacks			
OS	Ubuntu, CentOS, Debian.	Ubuntu, Debian, CentOS.	CentOS, Rocky Linux.
Languages, libs	Python, R, C/C++, MATLAB, TF, PyTorch.	Python, R, C/C++, ML/DL libs, EBRAINS APIs.	Python, C/C++, Fortran, ML libraries, MPI, EBRAINS APIs.
Interfaces	Web portal, Horizon, CLI, API.	EBRAINS portal, CLI tools, REST APIs.	SLURM scripts, SSH access, EBRAINS GUI linkages.
Workflow support	CWL, Snakemake, Nextflow, Galaxy.	Galaxy, Nextflow, EBRAINS Workflows.	Nextflow, Snakemake, EBRAINS Workflows, SLURM pipelines.
Monitoring	Grafana, Prometheus, Ceilometer.	Grafana, internal monitors.	Grafana dashboards, SLURM-based resource tracking.
Compute Federation			
Resource discovery	EBRAINS KG & Resource Catalog.	EBRAINS catalog, FAIR metadata.	Integrated via EBRAINS Resource Registry.
Job routing & load distribution	Manual/semi-auto, orchestrators (e.g. K8s, Slurm).	Federation tools, K8s, EBRAINS schedulers.	Federated scheduler APIs, Slurm integration.

6.16. LOFAR Long-Term Archive e-Infrastructure

Table 17: Technical Characteristics of LOFAR Long-Term Archive e-Infrastructure

Technical Analysis – LOFAR Long-Term Archive e-Infrastructure	
System/site	Prototypical LTA site
Data/Storage Capabilities and Services	
Storage abstraction	dCache managed storage environment, providing file management through the dCache API and file access using WebDAV.
Storage levels & locality	Online storage that provides immediate access to data. Nearline (tape) storage from which data can be retrieved on demand. Automatic migration from online to nearline storage is required.
Storage capacity	The total network is expected to scale to 120 PB of storage (the vast majority nearline) over the next several years. Individual sites will obviously have lower capacity; a site providing less than 1 PB would likely not be viable.
Performance (latency & BW)	10 Gbit/second connections between storage and compute systems and to the outside world. Online storage should provide “instant” access. Staging from nearline storage may take hours or days.
Access methods & interfaces	Web interface and programmatic API available for data discovery and staging. Data accessed using WebDAV.
Data formats	Bulk data is stored in standard astronomical data formats (typically FITS for images, MeasurementSet for visibility data, and HDF5 for beamformed data). Accompanying ancillary data (diagnostic plots, logs, etc) are stored in a variety of standard formats (PNG, PDF, JSON, etc).
Data retention & backup	Intermediate data products ingested from the telescope are stored for 18 months while they are processed to advanced forms. Advanced data products are stored indefinitely.
Access control	See discussion above.
Encryption	None.
Data Transfer & Federation	
External data federation	Data is not normally federated across LTA sites, but specific datasets may be transferred where necessary.
Compute Capabilities and Services	
Processors/ Accelerators	Current pipelines are based on x86-64. The use of GPUs in LOFAR data processing is currently not common, but is expected to become so over the next several years; work with the LTA sites to provide appropriate resources is also expected.
Nodes	Varies by site.
Interconnect	Varies by site.
Performance	Varies by site.

Non-conventional compute	None.
Compute modes	Batch processing via SLURM. Support for Jupyter notebooks driving Dask clusters is anticipated in the (indefinite) future.
Limits and guarantees	Varies by site
Software Services, Interfaces & Stacks	
OS	Linux; details vary by site.
Languages, libs	Python, C++. Likely to include CUDA in future. A wide variety of standard libraries for numerical and astronomical computing.
Interfaces	SSH
Workflow support	SLURM
Monitoring	Varies by site. Work is underway to develop a unified LTA dashboard aggregating data from across multiple sites.
Compute Federation	
Resource discovery	SLURM, plus additional LOFAR-specific software (ATDB).
Job routing & load distribution	SLURM, plus additional LOFAR-specific software (ATDB).

6.17. ErUM-Data-Hub e-Infrastructure

Table 18: Technical Characteristics of ErUM-Data-Hub e-Infrastructure

Technical Analysis – ErUM-Data-Hub e-Infrastructure		
System/site	The ErUM-Data-Hub collaborates with several HPC centers to provide computational resources for research. To mention some of the systems: JUWELS Booster (JSC – see Table 3), SuperMUC-NG (LRZ – see Table 3), HAWK (HLRS), TOPAS (KIT).	
	HAWK (HLRS)	TOPAS (KIT)
Data/Storage Capabilities and Services		
Storage abstraction	POSIX-compliant parallel file systems (e.g., Lustre).	dCache, EOS, and GPFS for abstracted large-scale data access.
Storage levels & locality	Hierarchical: SSD (scratch), HDD, and tape for archive.	Multi-tier: hot (SSD), warm (HDD), cold (tape/archive).
Storage capacity	~26 PB usable.	Multiple PBs across systems (dCache/EOS); scalability planned.

Performance (latency & BW)	High BW (~1 TB/s aggregate); optimized for low latency.	High throughput (GB/s scale); designed for long-term ingestion.
Access methods & interfaces	SSH, SLURM, REST APIs for I/O; POSIX, GPFS.	GridFTP, WebDAV, XRootD, POSIX, REST.
Data formats	HDF5, ROOT, NetCDF, CSV, binary blobs.	ROOT, HDF5, Parquet, FITS (community-dependent).
Data retention & backup	Tiered with policy-based retention; tape backup.	Long-term via tape, short-term replication across dCache nodes.
Access control	LDAP, Kerberos, project-based ACLs.	X.509 certificates, VO/group-based permissions.
Encryption	At-rest encryption on tape/archive layers; in-transit via SSH/TLS.	TLS for transport, file-level encryption on request.
Data Transfer & Federation		
Internal methods & performance	InfiniBand EDR; MPI-IO, shared file systems.	100 Gbit/s interconnect; fast local transfers using SSD buffer.
External methods, security & performance	B2DROP, EUDAT tools, GridFTP; secure with SSH/TLS.	GridFTP, FTS3, REST, Globus; supports authentication protocols.
External data federation	Support via EUDAT B2FIND and B2SHARE; HLRS manages cross-institution workflows.	Federated access using dCache Federation and VO infrastructure.
Compute Capabilities and Services		
Processors/ Accelerators	AMD EPYC 7742 (Rome), 128 cores/node.	AMD EPYC, some GPU-accelerated nodes (A100).
Nodes	5,632 nodes total.	~400 compute nodes in federated environment.
Interconnect	InfiniBand HDR100, low-latency.	Infiniband & 10/100 Gbit Ethernet.
Performance	~26 PFLOPS (theoretical peak).	Scalable to multiple PFLOPS depending on federation load.
Non-conventional compute	GPU-enabled nodes, experimental ML workloads supported.	FPGA-based acceleration (experimental), ML/AI inference jobs.
Compute modes	Batch (SLURM), interactive, container-based (Singularity).	nteractive (JupyterHub), batch (HTCondor), containers.
Limits and guarantees	Fair-share, QoS scheduling, project limits.	Time-based quotas, data priorities per community.

Software Services, Interfaces & Stacks		
OS	CentOS/ROCKY Linux, tuned kernels.	Ubuntu/CentOS-based nodes.
Languages, libs	Python, C/C++, Fortran, OpenMPI, CUDA.	Python, R, Julia, ROOT, SciPy, TensorFlow.
Interfaces	SSH, Web UI, REST API for job/data mgmt.	Web UIs, Jupyter Notebooks, command-line, REST.
Workflow support	CWL, Snakemake, Nextflow, Slurm job scripts.	Galaxy, Airflow, Jupyter-based pipeline.
Monitoring	Grafana, Nagios, custom HLRS dashboards.	Prometheus, Grafana, ELK stack.
Compute Federation		
Resource discovery	SLURM job info, web portal for scheduling.	VO-based metadata catalogs and job dashboards.
Job routing & load distribution	SLURM scheduling, HLRS resource broker.	HTCondor with job migration and local priority policies.

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

Table 19: Technical Characteristics of PUNCH4NFDI e-Infrastructure – Data/AI-Access

Technical Analysis – PUNCH4NFDI e-Infrastructure		
System/site	The PUNCH4NFDI collaborates with several HPC centers for compute and data services for research. To mention some of the systems: SuperMUC-NG (LRZ – see Table 3), JURECA-DC (JSC), bwForCluster NEMO (University of Freiburg).	
	JURECA-DC (JSC)	bwForCluster NEMO (University of Freiburg)
Data/Storage Capabilities and Services		
Storage abstraction	Lustre parallel file system with hierarchical storage (SSD, HDD, tape).	BeeGFS parallel file system.
Storage levels & locality	Multi-tier (NVMe, HDD, archive); local SSDs per node.	BeeGFS local and shared parallel file systems.
Storage capacity	~1 PB fast storage, 20+ PB archival.	Several PB total; exact depends on configuration.
Performance (latency & BW)	High I/O bandwidth (up to 150 GB/s aggregate).	High bandwidth, optimized for I/O-heavy workloads.
Access methods & interfaces	POSIX, GridFTP, SCP, HPSS interfaces.	POSIX, SSH, SCP.
Data formats	HDF5, NetCDF, ROOT, others.	HDF5, CSV, FITS, domain-specific formats.

Data retention & backup	Long-term archival on tape; backup services.	Periodic backup, short-term archival.
Access control	LDAP-based authentication, role-based access.	bwIDM identity federation, project-based access.
Encryption	Encrypted transfers (SSH/SCP), tape encryption.	Encrypted connections (SSH/SCP); filesystem-level encryption on demand.
Data Transfer & Federation		
Internal methods & performance	Lustre striping, Infiniband support.	BeeGFS + Infiniband.
External methods, security & performance	GridFTP, SCP, Globus.	SCP, Globus Online.
External data federation	Supports EUDAT/B2SAFE integration; NFDI-compliant data publication.	Federated access via bwDataArchiv and NFDI platforms.
Compute Capabilities and Services		
Processors/ Accelerators	MD EPYC CPUs, NVIDIA A100 GPUs.	Intel Xeon CPUs.
Nodes	~768 compute nodes + 192 GPU nodes.	>700 nodes.
Interconnect	Mellanox HDR200 InfiniBand.	Intel Omni-Path.
Performance	~23 PFlop/s peak (modular architecture).	~2.8 PFlop/s peak.
Non-conventional compute	FPGA support in future modules.	None.
Compute modes	Batch, interactive, GPU-accelerated.	Batch, limited interactive sessions.
Limits and guarantees	Fair-share, project-based quotas.	Project-based compute hours, queue system.
Software Services, Interfaces & Stacks		
OS	CentOS 8 Stream, custom kernel.	CentOS, transitioning to Rocky Linux.
Languages, libs	C/C++, Fortran, Python, MPI, CUDA, OpenMP.	C/C++, Python, MPI.
Interfaces	SSH, SLURM, JupyterHub, UNICORE.	SSH, SLURM, JupyterHub.
Workflow support	CWL, Snakemake, Nextflow.	Snakemake, custom scripts.
Monitoring	Grafana, Prometheus, custom dashboards.	Ganglia, SLURM monitoring tools.

Compute Federation	
Resource discovery	Through JSC & bwHPC user portals; metadata shared with NFDI registries.
Job routing & load distribution	Via SLURM partitions and project queues; cross-cluster workload balancing not automated.

6.19. Copernicus e-Infrastructure

Table 20: Technical Characteristics of Copernicus e-Infrastructure

Technical Analysis – Copernicus e-Infrastructure			
System/site	Copernicus Data Space Ecosystem (CDSE)	WEkEO Platform	CREODIAS Platform
Data/Storage Capabilities and Services			
Storage abstraction	CDSE, WEkEO, and CREODIAS offer object storage for scalable data access, often built on cloud-native storage layers compatible with S3 interfaces.		
Storage levels & locality	CDSE utilizes a hybrid cloud model with storage located in European data centers, ensuring GDPR compliance.	WEkEO and CREODIAS also host data across multiple EU-located facilities to optimize access locality.	
Storage capacity	Petabyte-scale storage.	Multi-petabyte capacity.	Petabyte-scale, highly scalable.
Performance (latency & BW)	High bandwidth with CDN support, optimized latency.	Optimized latency with CDN and edge nodes.	SSD-backed storage, low-latency, high BW.
Access methods & interfaces	REST, OData APIs, CLI, GUI.	Harmonized API, Jupyter Notebooks, GUI.	EO Finder, REST APIs, OpenSearch.
Data formats	GeoTIFF, NetCDF, JPEG2000, SAFE, HDF5.	GeoTIFF, NetCDF, SAFE, others.	GeoTIFF, SAFE, NetCDF, JPEG2000, HDF5.
Data retention & backup	Long-term storage, regular backups.	Aligned with Copernicus policies.	Policy-compliant long-term retention.
Access control	OAuth2, role-based access.	SSO, OAuth2, user roles.	OAuth2, SSO, role management.
Encryption	HTTPS/TLS in transit, AES-256 at rest.	Encrypted transit and at rest (AES-256).	AES-256 at rest, TLS in transit.
Data Transfer & Federation			
Internal methods & performance	High-speed internal interconnects.	Efficient internal transfers in cloud infra.	High-speed internal networking.

External methods, security & performance	HTTPS, SFTP, Aspera, secured via CDN.	Aspera, HTTPS, secure, and performant.	Aspera, SFTP, HTTPS with performance tuning.
External data federation	Federated data access via standardized APIs.		Standardized access to Copernicus services.
Compute Capabilities and Services			
Processors/ Accelerators	Intel Xeon CPUs, NVIDIA GPUs (A100/V100), cloud-optimized compute nodes.	Intel Xeon CPUs, NVIDIA Tesla GPUs, suitable for EO workloads.	Intel Xeon CPUs, NVIDIA Tesla & A100 GPUs, FPGA accelerators for ML/AI.
Nodes	Virtual machines provisioned on OpenStack-based infrastructure, supporting auto-scaling and large distributed workloads in the cloud environment. Designed for EO data processing at scale.	Kubernetes-managed nodes that support containerized workflows. Designed to run Jupyter Notebooks and handle EO API requests. Offers scalable cloud environments tuned for harmonized data access.	Mix of virtual and bare-metal nodes, including GPU-accelerated and HPC nodes. Supports scalable compute environments optimized for machine learning, AI tasks, and EO data pipelines.
Interconnect	High-speed Ethernet.	Cloud-native high-speed links.	Infiniband, high-speed Ethernet.
Performance	Elastic, high-throughput compute.	Elastic compute with Kubernetes.	High-performance, low-latency compute.
Non-conventional compute	N/A.	N/A.	FPGA, AI accelerators.
Compute modes	Batch, interactive (Jupyter), serverless.	Batch, serverless, Jupyter.	Batch, Jupyter, serverless.
Limits and guarantees	99.9% uptime, user quotas.	SLAs, resource quotas.	SLAs with high uptime, usage limits.
Software Services, Interfaces & Stacks			
OS	Ubuntu, Debian.	Ubuntu.	Ubuntu, CentOS.
Languages, libs	Python, R, C++, geospatial libs.	Python, R, C++, EO libraries.	Python, R, C++, extensive EO libs.
Interfaces	Web GUI, APIs, CLI.	Jupyter, API, GUI.	GUI, APIs, CLI, OpenSearch.
Workflow support	Airflow, EO pipelines.	Workflow engines, EO integration.	Apache Airflow, CWL supported.
Monitoring	Real-time dashboards, logs.	User and system monitoring.	Dashboards, usage logs.
Compute Federation			

Resource discovery	Centralized catalog, metadata search.	WEkEO viewer for data/resources.	EO Finder, metadata catalog
Job routing & load distribution	Kubernetes-based orchestration.	Platform load balancing and orchestration.	SLURM, Kubernetes job routing.

6.20. EGI Federation e-Infrastructure

Table 21: Technical Characteristics of EGI Federation e-Infrastructure

Technical Analysis – EGI Federation e-Infrastructure	
System/site	The EGI Federation offers a federated cloud infrastructure that integrates various cloud resource providers across Europe. This infrastructure delivers a unified platform for compute, storage, and data services, tailored to support the diverse needs of scientific communities.
	EGI's Federated Cloud
Data/Storage Capabilities and Services	
Storage abstraction	EGI provides multiple storage solutions, including block storage for VM-attached volumes, grid storage optimized for High Throughput Computing (HTC) and High-Performance Computing (HPC) scenarios, and object storage suitable for cloud-native applications and data archiving.
Storage levels & locality	Data is stored across a distributed network of resource centers, ensuring data redundancy and locality optimization for performance and compliance.
Storage capacity	The infrastructure boasts substantial storage capabilities, with offerings exceeding 500 PB, accommodating the vast data requirements of various research projects.
Performance (latency & BW)	Designed for high-performance needs, the storage solutions minimize latency and maximize bandwidth, crucial for data-intensive computations.
Access methods & interfaces	Data access is facilitated through standard protocols and interfaces, ensuring compatibility and ease of integration with existing workflows.
Data formats	Supports a wide range of data formats, catering to the diverse requirements of different scientific domains.
Data retention & backup	Robust data retention policies and backup mechanisms are in place to safeguard against data loss and ensure long-term availability.
Access control	Utilizes federated identity management systems, allowing secure and streamlined access control across the infrastructure.
Encryption	Data encryption is implemented both at rest (as optional) and during transit, adhering to best practices for data security.
Data Transfer & Federation	
Internal methods & performance	High-speed networks facilitate efficient data transfer between resource centers, optimizing performance for collaborative projects.

External methods, security & performance	Secure data transfer protocols are employed for external communications, ensuring data integrity and confidentiality.
External data federation	The infrastructure supports data federation, enabling seamless integration and access to external data sources, enhancing collaborative research efforts.
Compute Capabilities and Services	
Processors/Accelerators	EGI's compute resources include a mix of CPUs and accelerators like GPUs, supporting a broad spectrum of computational tasks.
Nodes	The infrastructure comprises a vast array of compute nodes, providing scalable resources to meet varying computational demands.
Interconnect	High-speed interconnects between nodes ensure low-latency communication, vital for parallel processing and large-scale simulations.
Performance (Compute, memory)	Offers high-performance compute and memory resources, suitable for intensive computational workloads.
Non-conventional compute	Supports non-traditional computing paradigms, including containerization and serverless computing, to accommodate emerging research methodologies.
Compute modes	Provides various compute modes, including batch processing, interactive sessions, and long-running services, offering flexibility for different use cases.
Limits and guarantees	Service level agreements (SLAs) define resource limits and guarantees, ensuring predictable performance and resource availability.
Software Services, Interfaces & Stacks	
OS	Supports a range of operating systems, allowing users to select environments that best fit their application requirements.
Languages, libs	Provides a comprehensive suite of programming languages and scientific libraries, facilitating diverse research applications.
Interfaces	Offers standardized interfaces and APIs, promoting interoperability and ease of integration with user applications.
Workflow support	Includes tools and services for workflow management, enabling users to design, execute, and monitor complex computational workflows.
Monitoring	Robust monitoring tools are available, providing insights into resource usage, performance metrics, and system health.
Compute Federation	
Resource discovery	Implements federated resource discovery mechanisms, allowing users to locate and utilize resources across the distributed infrastructure effectively.
Job routing & load distribution	Employs intelligent job routing and load balancing strategies to optimize resource utilization and ensure efficient task execution.

6.21. SURF Grid/Spider e-Infrastructure

Table 22: Technical Characteristics of SURF Grid/Spider e-Infrastructure

Technical Analysis – SURF Grid/Spider e-Infrastructure		
System/site	SURF Grid	SURF Spider
Data/Storage Capabilities and Services		
Storage abstraction	Object and block storage via OpenStack Swift/Ceph.	Same as SURF Grid.
Storage levels & locality	Multi-tiered: SSDs for hot data, HDDs for cold storage, all local to compute.	Same, with potential extension for custom deployments.
Storage capacity	Petabyte-scale, expandable on demand.	Similar petabyte-scale, customizable.
Performance (latency & BW)	High throughput and moderate latency for large-scale data workflows.	Optimized for throughput, with slightly better latency due to the hybrid architecture.
Access methods & interfaces	REST APIs, POSIX-like access via FUSE or GridFTP.	REST APIs, enhanced with custom endpoints for user-specific clones.
Data formats	Open scientific data formats (HDF5, NetCDF, etc.).	
Data retention & backup	Scheduled backups with redundancy across availability zones.	Backup policies adjustable per clone instance.
Access control	Federated Identity via SURFconext, role-based access.	
Encryption	At rest and in transit using TLS and disk encryption.	
Data Transfer & Federation		
Internal methods & performance	Intra-cloud high-speed transfer (10–40 Gbps links).	Similar to customizable interconnects.
External methods, security & performance	GridFTP, HTTPS with identity federation; 1–10 Gbps for external nodes.	Same with additional endpoint control.
External data federation	WLCG Tier-1 participation, EGI Federation.	Extendable to domain-specific federations.
Compute Capabilities and Services		
Processors/ Accelerators	x86_64 CPUs, NVIDIA GPUs in some nodes.	CPUs, GPUs, FPGAs on demand.
Nodes	Thousands of heterogeneous nodes.	Modular node design; scalable from tens to thousands.

Interconnect	Infiniband, 10/40/100 GbE.	High-speed virtual networking on OpenStack, physical Infiniband optional.
Performance	Optimized for HTC: many loosely coupled jobs.	Mixed-mode: HTC + HPC elements.
Non-conventional compute	N/A.	FPGAs and containerized compute environments supported.
Compute modes	Batch jobs, pilot jobs.	Batch, pilot, container-based.
Limits and guarantees	SLA-backed; quotas and job walltime limits.	SLA-backed; flexible policies per user/group.
Software Services, Interfaces & Stacks		
OS	Linux (CentOS/AlmaLinux).	Linux (Ubuntu, CentOS), user-defined images.
Languages, libs	Python, R, C/C++, Fortran; scientific libraries.	
Interfaces	CLI, web portal, REST APIs.	Same + custom interfaces for clones.
Workflow support	CWL, Snakemake, Nextflow.	
Monitoring	Grafana, Prometheus.	Same, with additional monitoring per custom clone.
Compute Federation		
Resource discovery	EGI and WLCG tools.	Customizable; can integrate with SURF or third-party.
Job routing & load distribution	HTCondor, ARC middleware.	HTCondor, OpenStack-native orchestration.

6.22. EOSC Federation e-Infrastructure

Table 23: Technical Characteristics of EOSC Federation e-Infrastructure

Technical Analysis – EOSC Federation e-Infrastructure	
System/site	EOSC EU node
Data/Storage Capabilities and Services	
Storage abstraction	Offers multiple storage types, including block storage (attached to VMs or containers) and File Sync and Share (for seamless access to data across devices)
Storage levels & locality	Supports multiple storage tiers, including cloud-based and federated storage systems, ensuring data locality and compliance with regional data governance policies.
Storage capacity	Offers scalable storage solutions capable of handling large volumes of research data, accommodating the needs of diverse scientific disciplines.

Performance (latency & BW)	Designed for high-performance data access with optimized latency and bandwidth to support data-intensive research workflows.
Access methods & interfaces	Provides various access methods, including web portals, APIs, and command-line interfaces, facilitating seamless data interaction.
Data formats	Supports a wide range of data formats to ensure interoperability across different research domains.
Data retention & backup	Implements robust data retention policies and backup mechanisms to safeguard research data over time.
Access control	Utilizes a federated Authentication and Authorization Infrastructure (AAI) to manage user access rights and ensure secure data handling.
Encryption	Supports encryption protocols for data at rest and in transit to maintain data confidentiality and integrity.
Data Transfer & Federation	
Internal methods & performance	A Data Transfer service allows asynchronous transfer of large data sets with features like automatic retries and real-time statistics, ensuring efficient internal data movement.
External methods, security & performance	Supports secure and high-performance data transfer with external systems, adhering to international security standards.
External data federation	Enables integration with external data sources and repositories through the Open Cloud Mesh (OCM) API, promoting a federated approach to data sharing and collaboration.
Compute Capabilities and Services	
Processors/Accelerators	Provides access to a variety of computing resources, including CPUs and GPUs, to support diverse computational needs.
Nodes	Operates a distributed network of computing nodes to ensure scalability and resilience in processing workloads.
Interconnect	Features high-speed interconnects between computing resources to facilitate efficient data processing and communication.
Performance (Compute, memory)	Optimized for elastic computing tasks, offering diverse compute power and memory resources.
Non-conventional compute	Supports emerging computing paradigms, such as containerized applications and interactive notebooks, to cater to modern research methodologies.
Compute modes	Offers various compute modes, including batch processing and interactive sessions, to accommodate different research workflows.
Limits and guarantees	Implements usage quotas based on credits and service level agreements to ensure fair resource allocation and performance guarantees.

Software Services, Interfaces & Stacks	
OS	Supports multiple operating systems to provide flexibility in software deployment and execution.
Languages, libs	Offers a comprehensive suite of programming languages and scientific libraries to support diverse research applications. Offers pre-defined templates of research applications and tools in the Tools Hub for on-demand deployment by users.
Interfaces	Provides user-friendly interfaces, including web portals and APIs, to facilitate access to services and resources.
Workflow support	Enables the design and execution of complex research workflows, integrating various tools and services within the EOSC ecosystem.
Monitoring	Implements monitoring tools to track resource usage, performance metrics, and system health, ensuring transparency and reliability.
Compute Federation	
Resource discovery	Facilitates the discovery of computing resources across the federated infrastructure, promoting efficient resource utilization.
Job routing & load distribution	Employs intelligent job scheduling and load balancing mechanisms to optimize computational workloads within each data centre.

6.23. Simpl Data Federation Platform

Simpl (Smart Middleware Platform) is a secure middleware SW platform that supports federated data access and interoperability in European data initiatives which is being implemented by a consortium of [Eviden](#) (Belgium subsidiary), [Aruba](#) (Italy), [Capgemini](#) (Netherlands subsidiary), [Engineering International Belgium](#) (Belgium), [IONOS](#) (Germany), and [COSMOTÉ Global Solutions](#) (Belgium) under contract²⁵ with the European Commission.

Simpl consists of three main components:

- Simpl-Open is an open-source software stack that enables the setup and operation of federated data spaces. It is offered by the European Commission as a free SW suite to both the public and private sectors.
- Simpl-Labs is an environment for data spaces to assess their level of interoperability with Simpl. It also supports sectoral data spaces to experiment with the deployment, maintenance, and support of Simpl-Open before deploying it for their own needs.
- Simpl-Live are specific instances of Simpl-Open deployed for specific sectoral data spaces where the European Commission plays an active role in their management

The top-level architecture of Simpl is illustrated in [Figure 2](#). Participants can have four different roles: a *consumer* looks up data in the metadata stored at the central catalogue (4) and accesses the data stored by the data space at the *infrastructure provider* (7). The *data provider* uploads (6) and publishes metadata plus infrastructure and application/service information in the catalogue (3). The *governing agency* configures the central catalogue (1) and onboards the other participants (2).

²⁵ For tender details see <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/tender-details/12922>.

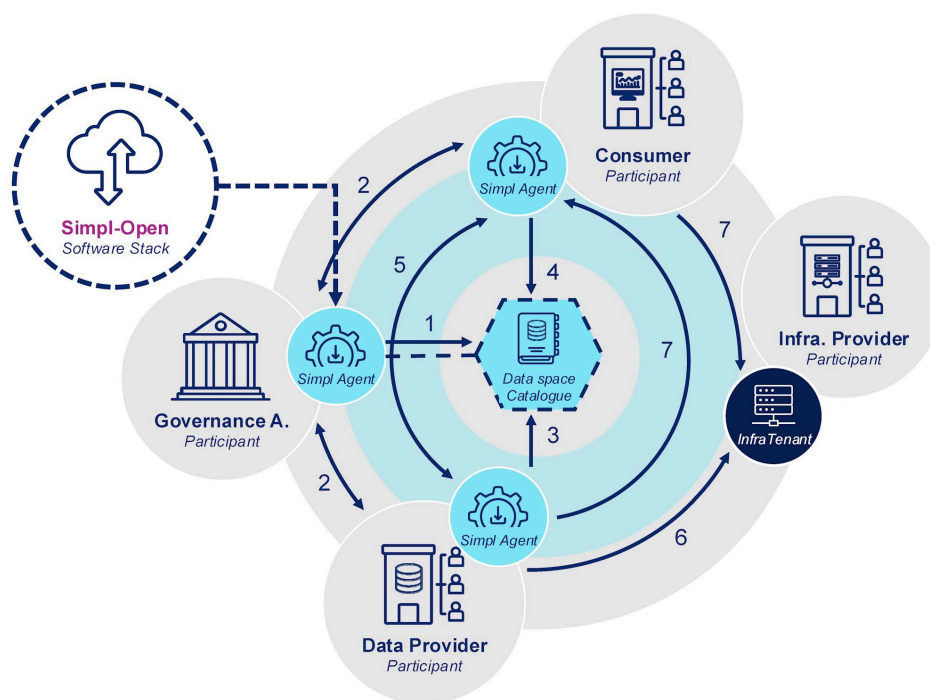


Figure 2: Simpl-Open Top-Level System Architecture

Simpl-Open is still in the early stages of development. A first “minimum viable product” (MVP) version was released at the end of January 2025 with this functionality

- Onboarding of users by the governance authority.
- Catalogues for infrastructure & data, usage policies, and quality rules.
- UI and API for adding to the catalogues and for validating entries.
- UI and API for basic and extended searches within a single data space.
- Resource selection, request to use that resource, and establish a basic contract between the data provider and the consumer.
- Infrastructure deployment, data set access, and initial types of data processing.
- Secure communication between Simpl agents, logging of basic metrics for all user actions.

The next release, planned for June 2025, will add

- UI and APIs for the governance authority to link identity attributes to an onboarding procedure template, which will be reviewed before approval and assigned to the onboarded user.
- Automated validation/rejection of onboarding requests by the governance authority according to configurable rules.
- Renewal and revocation of participants’ credentials by the governance authority.
- Publication of infrastructure deployment scripts and association with the resource description for the deployment scripts to be run once a user requests the resource.
- Structure for the data addresses according to the technical data transfer process requirements, based on the data sharing method (like, for instance, S3), enabling automated data sharing.
- Decommissioning of infrastructure resources previously provisioned for a consumer.
- Dashboard showing health and availability status of all agent services in real-time.
- Streamlining and standardisation of APIs.

As can be seen from the list of functionality and improvements, the Simpl-Open system is not yet ready for actual employment and use. The Simpl-Open Roadmap shown in [Figure 3](#) indicates that important capabilities (such as support/ticketing/helpdesk, HPC integration, and data analytics) will be worked on until the end of 2026.

The first release and deployment of Simpl-Labs is expected for 2025.

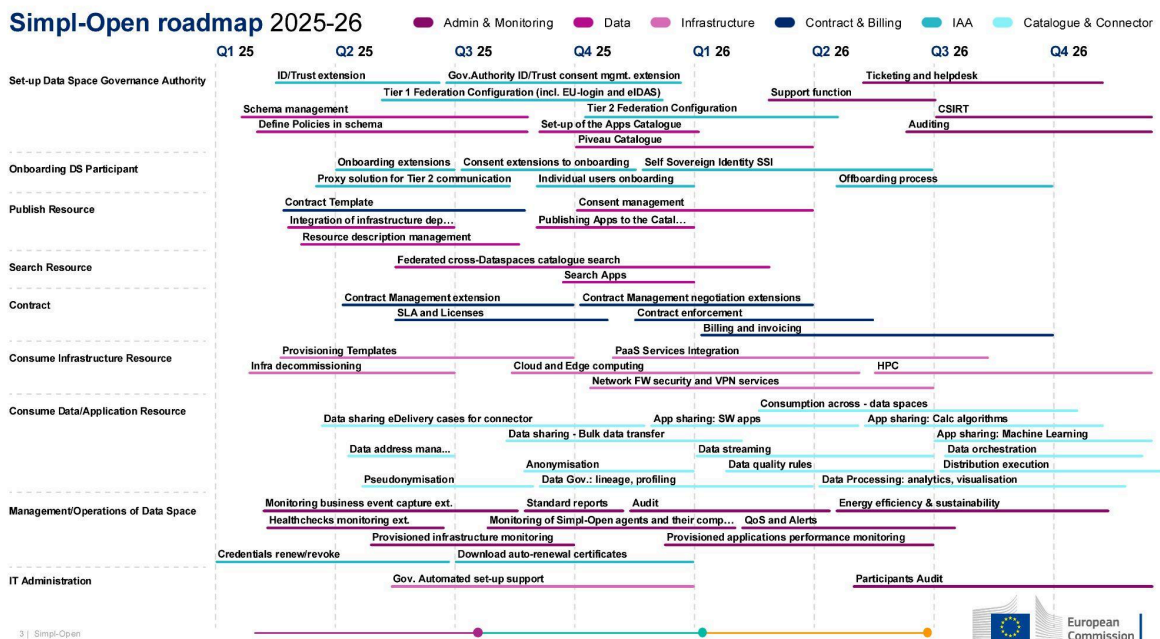


Figure 3: Simpl-Open Implementation Roadmap

Table 24: Technical Characteristics of Simpl Data Federation e-Infrastructure

Technical Analysis – Simpl Data Federation Platform	
System/site	Simpl-Open is a software middleware platform for providing federated data access and interoperability in European data initiatives; it is available for free use as open source software.
Data/Storage Capabilities and Services	
Storage abstraction	Simpl is designed to be flexible regarding the storage infrastructures supported and the data abstractions provided to consumers.
Storage levels & locality	Simpl will implement support for HPC sites, which will presumably include support for high-performance storage tiers; locality of access can be supported by configuring multiple data locations (i.e., infrastructure providers) and data transfer policies that select the nearest provider for each request.
Storage capacity	Simpl itself does not provide any capacity or capability to store the payload data. A deployed Simpl infrastructure uses the storage capabilities of the data providers and the infrastructure providers. Simpl does manage its own central catalogue infrastructure for data, infrastructure, and applications.
Performance (latency & BW)	N/A at this point in time.
Access methods & interfaces	Specific public details are not available.
Data formats	Simpl is data format agnostic.

Data retention & backup	While Simpl could drive data duplication through suitably configured policies, backups will remain the duty of the data providers and the infrastructure providers.
Access control	Simpl will implement mechanisms for anonymization/pseudonymization of sensitive data; access rules can be
Encryption	For data at rest, encryption will be the duty of the data/ infrastructure providers and the consumer; communication between the different Simpl SW agents is encrypted, and policies can be set up to mandate encrypted file transfers
Data Transfer & Federation	
External methods, security & performance	Simpl can select and drive data transfer mechanisms between the infrastructure provider and consumer or data provider, depending on the configuration of data sharing methods and the actual data addresses; this provides a high degree of flexibility and hides the details from the consumer.
External data federation	Simpl federates data through the central catalogue infrastructure, which can accommodate a multitude of data providers and infrastructure providers. Consumers access the catalogue to search for data, and the Simpl system identifies data locations (at infrastructure providers) and handles the requested access according to the configured policies.
Software Services, Interfaces & Stacks	
OS	Simpl is required to use open standards and implementations of building blocks as far as possible; in the tender requirements, it is suggested to base the work on tools and components from the Linux Foundation and, in particular, its sub-foundation, the Cloud Native Computing Foundation .
Languages, libs	Simpl relies on OpenStack/Kubernetes and uses Java to code the interfaces to participants.
Interfaces	<ul style="list-style-type: none"> • AAI for participant identities is handled via an interface compliant with eIDAS, implemented on top of openSAML. • Keycloak is used to manage and verify credentials within the Simpl system.
Workflow support	Specific public details are not available.
Monitoring	Simpl implements its own logging and monitoring system, providing dashboards for various roles within the infrastructure.

7. Annex 2 – Links to Important Documents

e-Infrastructure	Link
EuroHPC JU	https://eurohpc-ju.europa.eu/access-our-supercomputers/access-policy-and-faq_en https://eurohpc-ju.europa.eu/selection-first-seven-ai-factories-drive-europe-s-leadership-ai-2024-12-10_en
EuroHPC JU Mid-Range & Petascale Systems	https://rnca.fccn.pt/en/deucalion https://sofiatech.bg/en/petascale-supercomputer https://www.it4i.cz/en/infrastructure/karolina https://www.luxprovide.lu/meluxina https://izum.si/en/vega-en https://www.naiss.se/news/2024/11/request-for-proposals-for-new-supercomputer-arrhenius
EuroHPC JU pre-Exascale Systems	https://leonardo-supercomputer.cineca.eu https://www.lumi-supercomputer.eu/about-lumi https://www.bsc.es/marenostrum/marenostrum-5
EuroHPC JU Exascale Systems	https://www.fz-juelich.de/en/ias/jsc/jupiter https://eurohpc-ju.europa.eu/new-call-procure-european-exascale-supercomputer-alice-recoque-2024-09-09_en
EuroHPC Federation Platform	https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en https://cdn.sanity.io/files/461i44gu/production/fa594fa2ef4982bb63af608107b0cc7a2f2369e5.pptx https://onlinelibrary.wiley.com/doi/10.1002/spe.3075
GCS	https://www.hlr.de/de/loesungen/systeme/hunter https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/juwels https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/jureca https://doku.lrz.de/supermuc-ng-10745965.html
NHR	https://help.itc.rwth-aachen.de/service/rhr4fjjutttf/article/fbd107191cf14c4b8307f44f545cf68a/ https://tu-dresden.de/zih/hochleistungsrechnen/hpc

	https://csc.uni-frankfurt.de/wiki/doku.php?id=public:usage:goethe https://hpc-en.uni-mainz.de/2023/03/15/new-high-performance-computer-in-augurated-mogon-nhr-south-west https://www.nhr.kit.edu/userdocs/horeka
RES	https://www.res.es/en/access-to-res https://www.bsc.es/sites/default/files/public/u2416/20230220_marenostum5_sergi_girona_bsc.pdf https://arxiv.org/html/2503.09917v1 https://indico.ilab.org/event/459/contributions/11634/attachments/9680/14119/20230511_CMS-BSC-integration%20(1).pdf
GENCI	https://dci.dci-gitlab.cines.fr/webextranet/architecture/index.html http://www.idris.fr/eng/jean-zay/cpu/jean-zay-cpu-hw-eng.html https://www-hpc.cea.fr/en/Joliot-Curie.html
CSCS	https://www.cscs.ch/computers/alps
EPCC	https://www.archer2.ac.uk/about/hardware.html
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 Archive	https://lta.lofar.eu
ERUM	https://erumdatahub.de/en
PUNCH4NFDI	https://www.punch4nfdi.de
Copernicus	https://www.copernicus.eu/en/about-copernicus/infrastructure-overview
EGI Federation	https://www.egi.eu/egi-federation
SURF	https://servicedesk.surf.nl/wiki/display/WIKI/Grid https://servicedesk.surf.nl/wiki/display/WIKI/Spider
EOSC	https://eosc.eu/eosc-federation-handbook

Simpl	https://digital-strategy.ec.europa.eu/en/policies/Simpl https://simpl-programme.ec.europa.eu/
-------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------