



Strategic Research, Innovation and Deployment Agenda (SRIDA)

ANNEXES

MAY 2026

**This document contains the Annexes of the SPECTRUM SRIDA
published at the following URL: <https://zenodo.org/records/20182879>**

Last updated June 1st, 2026

SPECTRUM is funded by the European Union Grant Agreement Number 101131550

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

Table of Contents

Annexes	4
Annex A: Priorities	4
Pillar 1: Policy, Trust & Governance	4
P01 Cross-Infrastructure Governance	4
P02 Federated Resource Allocation	8
P03 Federated Identity and Trust	12
Pillar 2: Architecture & Interoperability	16
P04 Heterogeneous Computing Integration	16
P05 Data Management and Transport	20
P06 Workflow orchestration and management	24
Pillar 3: Software & Science Enablers	27
P07 AI	27
P08 Software Portability	32
P09 Open Science and Reproducibility	35
P10 Software Preservation and Sustainability	38
Pillar 4: Human Capital & Responsibility	41
P11 Community Collaboration and Co-Design	41
P12 Environmental Sustainability	44
P13 Skills, Careers and Training	48
Annex B: Priority Description Template	51
Priority Description Template	51
Annex C: Priority Evidence and Validation	52
C.1 Desk Research and Evidence Sources	52
C1.1. SPECTRUM CoP – Early Findings	52
C1.2 Representative use cases	54
C1.3. Interoperable access policies	55
C1.4 Landscape, Requirements and Gap Analysis	57
C1.5 Feedback & Alignment with the Technical Blueprint (D6.1)	59
C.2 Feedback & Alignment from External events	61
C2.1 EGI 2025 interactive session	62
C2.2 Interactive session at International Supercomputing (ISC 2025 BoF)	63
C.3 Community of Practice and Advisory Board Validation	67
C.4 Priority to Technical Blueprint Mapping	69

Annexes

Annex A: Priorities

Pillar 1: Policy, Trust & Governance

P01 Cross-Infrastructure Governance

Strategic Rationale and Evidence Base

Strategic Importance: The compute and data continuum vision requires research workflows to span resources governed by different entities: thematic research digital infrastructures (e.g. WLCG for HL-LHC, SRCNet for SKAO), pan-European multi-disciplinary digital infrastructures (e.g. EuroHPC, EOSC, EGI), nationally funded facilities, and community-governed federations. Each governance model reflects legitimate objectives. Without coordination mechanisms bridging these structures, researchers experience governance boundaries as friction that limits infrastructure utility.

Governance coordination underpins both interoperability and efficient resource use. Following the European Interoperability Framework [49], interoperability operates at four layers: legal (compatible access policies and liability), organisational (aligned governance processes and roles), semantic (shared accounting metrics and vocabularies), and technical (standardised interfaces and protocols). Beyond interoperability, governance coordination enables complementary investment across national and European programmes, planning certainty for multi-decade scientific programmes, and alignment of infrastructure evolution with scientific requirements.

Current State: European compute and data infrastructure operates through multiple governance structures:

- **Thematic digital infrastructures** serve specific communities through community-driven governance. WLCG has sustained federated operations for over twenty years across 160 sites in more than forty countries, coordinated through 65 Memoranda of Understanding with national funding agencies. SKA Observatory is building comparable governance through SRCNet. LOFAR ERIC, established in 2024, is building its computing governance model as LOFAR 2.0 operations approach, with growing data volumes that will require coordinated multi-year commitments from member states.
- **Pan-European digital infrastructures** provide cross-domain services through different models: EuroHPC coordinates thirty-five states through direct co-funded infrastructure investment, with governance oriented toward industrial competitiveness and technological sovereignty alongside research; EOSC federates existing infrastructure under common FAIR principles; EGI Federation provides federated HTC and cloud services governed by its member institutions; and the eInfrastructure Assembly coordinates horizontal initiatives including EUDAT, OpenAIRE, PRACE and GÉANT.

National programmes provide the majority of HTC and cloud capacity, whilst the EC contributes to HPC via EuroHPC and supports transnational and virtual access to research infrastructure services through Horizon Europe INFRA-SERV calls. Alignment across these funding streams is necessary for a coherent compute continuum.

The WLCG-EGI cooperation demonstrates that coordination across governance models is achievable when mechanisms exist. Structured engagement between data-intensive science communities and EuroHPC, however, remains underdeveloped. National and European funding routes serve different purposes: national programmes provide persistent operational capacity, whilst European programmes support coordination,

innovation, and cross-border access. Bridging these complementary roles, rather than treating them as competing alternatives, is a governance challenge this priority addresses.

Community Evidence: D5.2 identifies Gap #2 (lack of coordinated governance frameworks) and Gap #16 (sustainability models not aligned across infrastructures). The JENA Computing White Paper recommends that "requirements of communities must be considered in design and procurement of future HPC systems." The WLCG Strategy notes "there is no WLCG common funding" for federating services, demonstrating that federated governance requires explicit coordination investment.

Community consultation on this priority confirmed that governance scope should focus on alignment and co-design rather than centralised control. Participants identified three topics requiring coordination beyond existing mechanisms: aligning procurement with data-intensive science requirements; coordinating innovation for full-chain workflow optimisation (as distinct from classical HPC algorithmic optimisation); and bridging the structural mismatch between operational computing needs and project-based funding.

Technical Foundation: D6.1 Section 5.2 (Co-Design of Computing, Data, and Networking Infrastructure) provides the rationale: *"HEP and RA must engage with funding agencies during the design and procurement phases of HPC systems... The most effective way to be part of the co-design process in Europe would be to obtain a recognized status within EuroHPC."* D6.1 Section 5.8 and Section 4.6 describe interoperable standards and transparent accounting across federated resources.

European Alignment: Cross-infrastructure governance connects directly to four European strategies:

- **EuroHPC:** governance coordination provides the mechanism for data-intensive science communities to engage with EuroHPC procurement and access processes, ensuring exascale investments serve scientific as well as industrial objectives.
- **EOSC:** the coordination forum aligns thematic RI data governance with EOSC's federated data infrastructure and FAIR principles, bridging domain-specific practices with the EOSC Web of FAIR Data and Services.
- **ESFRI:** research infrastructures on the ESFRI Roadmap operate on 20 to 30 year timelines; governance coordination provides the stability these multi-generational programmes require and aligns with ESFRI's 2025 recommendations on long-term engagement, funding synergies, and community development. The ESFRI thematic clusters (including ESCAPE for astronomy and particle physics) provide a coordination model that this priority extends, particularly as the Horizon Europe 2026–2027 Work Programme introduces "Data, Computing and Digital Research Infrastructures" as a sixth ESFRI domain with its own cluster. ESFRI is also positioning for a more proactive coordination role across the research and technology infrastructure landscape, advocating for consolidation through clustering and a light overarching governance model [72].
- **ERA:** the forthcoming European Research Area Act [68] will establish a binding regulatory framework for cross-border research cooperation, realising the "fifth freedom" of free movement of researchers, knowledge, and technology. The European Parliament resolution of March 2026 calls for strengthened ERA governance, enhanced coordination between national and EU R&I initiatives, and improved accessibility of research infrastructure. The ERA Policy Agenda 2025–2027 includes active actions on research infrastructures (25 Member States participating), research careers, and enabling open science. Governance coordination for the compute and data continuum directly supports these objectives by providing an institutional model for cross-border resource sharing, joint planning, and coordinated investment.

Target Outcomes and Impact Assessment

Strategic Outcomes: This priority delivers governance coordination enabling effective use of European compute and data infrastructure, benefiting thematic research communities, pan-European infrastructure providers, and funding agencies.

Outcome 1: Coordination Forum with Substantive Agenda A coordination forum brings together thematic research digital infrastructures (WLCG, SRCNet, LOFAR ERIC, and others), multidisciplinary digital infrastructures (EuroHPC, EOSC, EGI), and policy maker representatives from national funding agencies and relevant EC units. The forum addresses interoperability across four layers (legal, organisational, semantic, and technical), coordinates the innovation agenda for data-intensive science (full-chain workflow

optimisation, real-time processing), and develops policy recommendations bridging national and European funding models. Formal representation replaces ad hoc engagement dependent on personal relationships.

Outcome 2: Coherent European Infrastructure Investment Governance coordination aligns infrastructure development with scientific requirements across all participants. Thematic research infrastructures gain influence over procurement and architecture decisions affecting multi-decade programmes. Pan-European infrastructures gain structured feedback from demanding users, enabling capability development that serves documented scientific needs. Funding agencies gain a coordinated evidence base for investment decisions, reducing duplication and enabling complementary use of national and European funding streams.

Success Indicators

- Coordination forum operational with participation from thematic RI leadership, pan-European infrastructure representatives, and policy maker observers
- Active involvement of scientific communities through topical networks
- Formal engagement channels established with EuroHPC, EOSC, and EGI governance structures
- Joint working groups addressing cross-cutting technical challenges (federated identity, computing interfaces, data management standards, accounting interoperability)
- Coordinated innovation agenda published, distinguishing data-intensive science requirements from classical HPC optimisation
- Policy recommendations on the operational-versus-project funding mismatch delivered to national and European stakeholders
- Documented evidence of coordination influencing infrastructure procurement or access decisions

European Impact

Governance coordination enables Europe to realise scientific return on distributed infrastructure investments and strengthens European digital sovereignty by making European infrastructure more accessible than commercial alternatives. As the ERA Act establishes a binding framework for cross-border research cooperation, governance coordination for the compute and data continuum provides a concrete implementation model for research infrastructure interoperability and coordinated investment.

Implementation Pathway and Stakeholder Coordination

Implementation Phases

Short term (1–3 years):

- Establish a coordination forum bringing together thematic and multidisciplinary digital infrastructures with policy maker representation
- Define the forum's substantive agenda: procurement alignment, innovation coordination, and bridging the operational-versus-project funding mismatch
- Map existing coordination mechanisms and identify gaps the forum addresses
- Develop shared metrics for infrastructure utilisation and cross-infrastructure workflow performance

Mid term (3–5 years):

- Consolidate coordination to deliver cross-border access to data and compute services, piloting integrated access across thematic and digital infrastructure clusters
- Develop policy recommendations on the operational-versus-project funding mismatch, aligned with the ERA Act framework
- Assess options for institutional evolution, including formal RI representation within EuroHPC and EOSC governance and coordination through the emerging ESFRI digital cluster
- Establish joint working groups on cross-cutting technical challenges connecting the forum to Pillar 2 implementation

Long term (5+ years):

- Assess impact of governance coordination and determine whether a formalised institutional structure is warranted for multi-decade infrastructure planning
- Align governance arrangements with the multi-decade timescales of the scientific programmes they serve
- Coordinate European positions in international collaborations spanning multiple continents

Stakeholder Roles

- **Scientific Communities and Research Infrastructures**
 - Responsibilities: Publish multi-year resource plans (5+ year horizons). Designate representatives for the coordination forum. Participate in co-design with e-infrastructure providers, bringing innovation requirements (e.g. full-chain optimisation, real-time processing) alongside capacity needs. Coordinate positions across thematic infrastructures.
 - Benefits: Voice in decisions affecting multi-decade programmes; access to diverse resources through coordinated allocation; reduced cross-provider friction.
- **Computing and Data Service Providers / e-Infrastructures**
 - Responsibilities: Establish liaison mechanisms with thematic RIs including named contacts and periodic reviews. Participate in the coordination forum and working groups. Accommodate operational access models alongside project-based allocations. Provide technical coordination on interface standardisation, identity, and accounting.
 - Benefits: Demanding users driving capability development; demonstrated scientific impact justifying investment; validated federation models.
- **Policy Makers and Funding Bodies**
 - Responsibilities: Ensure data-intensive science requirements inform EuroHPC and EOSC procurement. Address the operational-versus-project funding mismatch through multi-year allocation mechanisms (PO2), strategic access designations, and operational funding streams. Align national and European RI investments, consistent with the ERA Act framework. Participate in the coordination forum.
 - Benefits: Coherent European investments; reduced duplication; clearer accountability for scientific outcomes.

Resource Needs

- Coordination capacity: liaison effort (estimated 2–3 Full-Time-Equivalent –FTE) contributed in-kind by participating organisations or through existing institutional frameworks, rather than requiring a new coordination body
- Policy analysis capacity: expertise on the operational-versus-project funding gap and EC engagement
- Pilot activities: short-term project funding (e.g. Horizon Europe INFRA-DEV calls) for forum establishment and integrated access pilots
- Funding timeline: Horizon Europe INFRA calls (2026–2027) for short-term; FP10 or institutional integration for medium-term; structural funding model for long-term sustainability

Critical Dependencies

Cross-priority dependencies with:

- P02 (Federated Resource Allocation): allocation mechanisms depend on governance agreements; the forum provides the institutional context for negotiation
- P03 (Federated Identity and Trust): trust framework governance depends on cross-institutional agreements the forum facilitates
- P04 (Heterogeneous Computing Integration): technical integration requires governance alignment on interfaces, policies, and access conditions
- P11 (Community Collaboration and Co-Design): the CoP feeds requirements into the coordination forum

P02 Federated Resource Allocation

Strategic Rationale and Evidence Base

Strategic Importance: Europe invests substantially in research computing and data infrastructure through EuroHPC, national programmes, and federated e-infrastructures. Yet allocation mechanisms fragment this investment. Scientific computing is increasingly data-intensive, with exploitation of exabyte-scale datasets relying on a diverse and heterogeneous compute continuum spanning HPC, HTC, cloud, and emerging accelerator systems. Research infrastructures increasingly operate on multi-decade timelines, yet face annual application cycles, system-specific allocations that prevent workload-appropriate resource selection, and accounting frameworks that cannot track usage across governance boundaries. The result is less utilised infrastructures, repetitive submission burden, extra overhead e.g. in proposal writing, data collection, steering, and planning uncertainty that undermines long-term scientific programmes.

Federated resource allocation addresses how compute and data resources are allocated, accounted, and funded across the continuum:

- **Compute resources:** CPU cycles, GPU hours, emerging quantum computing allocations across HPC, HTC, and cloud systems, and associated requirements in terms of I/O and memory
- **Data resources:** Storage capacity (HL-LHC requires 4–6 EB tape archival by mid-2030s),
- **Network:** bandwidth for data transfer (SKA and WLCG distributes several hundreds of petabytes per year to regional centres)
- **Data lifecycle management,** including preservation and curation costs.

Balanced and flexible allocation across Europe is essential: governance coordination and federated identity enable access, but without allocation mechanisms that match scientific planning horizons and support diverse user communities, researchers cannot reliably integrate heterogeneous resources into sustainable workflows.

Current Status: Science is outcome driven, whereas allocations are project driven; consequently, predictability in resource allocation is difficult. Whilst a substantial part of the scientific community has (WLCG) or plans for multi-annual, MoU-based access allocations, a significant part of European e-infrastructures generally provide grant-based access with short-term allocation cycles (months to one year)¹. This creates several problems:

- **Repetitive submission burden:** Researchers submit repetitive allocation requests, creating administrative overhead for both users and review panels
- **Planning uncertainty:** Multi-decade experiments cannot reliably plan workflows when allocation is uncertain beyond 12 months
- **Fragmentation of resource allocation and accounting:** The allocations are generally valid for specific hosts/sites instead of the whole e-infrastructure; this limits the overall resource utilisation
- **Utilisation inefficiency:** Short-term allocations discourage investment in workflow optimisation for specific systems

Community Evidence: The analysis of the interoperable access policies (D5.2) and the technical blueprint for compute and data continuum (D6.1) prepared by the SPECTRUM project identifies two gaps:

#1: Long-term assured access and planning of resource allocations. Long-term (3–5 years) evolution and planning are not substantially informed by the needs of research communities that have been involved in the SPECTRUM project.

#3: Resource allocation and accounting that are valid across the whole e-Infrastructure. In many cases (such as EuroHPC JU), allocations are good for specific hosting sites/systems only, while users have voiced

¹ A notable exception being the EuroHPC Strategic Access framework

the necessity to use resources across an e-Infrastructure to profit from burst capacity, maximise the efficiency of systems for parts of their workload, and increase resiliency.

Technical Foundation: D6.1 Section 5.8 provides the technical framework: *"Multi-year or renewable annual guaranteed allocations would enable experiments to fully leverage such investments and integrate HPC resources more effectively into their computing models."*

D6.1 4.6 (Resource Federation) specifies requirements for federated access and allocation: *"Usage data, such as CPU/GPU hours, storage consumption, and data transfer volumes, must be collected and reported in a consistent way across all sites. This not only supports fair allocation and equitable participation but also enables funding bodies and infrastructure providers to verify contributions and assess impact."*

Key technical requirements:

- Standardised and transparent accounting framework for heterogeneous resources (CPU, GPU, storage, network) utilisation across governance boundaries
- Cross-infrastructure resource discovery and allocation

European Alignment: Resource allocation management is critical to broader European strategies. For EuroHPC, multi-year allocations for storage and compute would enable better integration of HPC resources into long-term scientific planning, with efficient resource allocation mechanisms minimising the risk of underutilising the resources. A federated resource allocation mechanism benefits interoperability, harmonisation, integration and synergies among the ESFRI research infrastructures, as well as within the EOSC web of fair data and services.

Target Outcomes and Impact Assessment

Strategic Outcomes:

- **Adopt long-term, flexible resource allocation processes:** processes are implemented that allow scientific communities and research infrastructures to obtain long-term access privileges and quotas in a flexible and scalable manner, federated across infrastructures, with transparent accounting and monitoring. [D5.2, R2]
- **Enable e-Infrastructure-wide use of resource allocations and quotas:** Enable e-infrastructure-wide allocations and harmonized accounting, so projects and users can burst across systems and providers while ensuring fair contributions and usage transparency. [D5.2, R1]
- **Align ESFRI roadmaps with e-infrastructure roadmaps**
- **Align e-infrastructures with the European Data Strategy:** interconnect sectoral data spaces by deploying robust e-infrastructures and advanced interoperability standards

Success Indicators:

- **Metrics and checkpoints** are established for multi-year allocations
- **Standardised accounting framework** deployed across major European facilities
- **Multi-year allocation agreements** operational for data-intensive research infrastructures

European Impact: The European strategy for data aims at creating a single market for data that will ensure Europe's global competitiveness and data sovereignty. This will lead to the creation of Common European Data Spaces. The European leadership in AI requires cooperation and a seamless, integrated ecosystem combining AI Factories, Data Spaces and HPC resources. In this context, Federated resource allocation and standardised resource accounting systems enables research communities to access compute and data infrastructures across organisational boundaries, increasing utilisation of European investments and reducing reliance on non-European providers. Standardised accounting enables to measure the impact on science of EU investments on computing, allowing future better alignment between investments and needs. This strengthens European digital sovereignty by making European infrastructure easier to use than commercial alternatives.

Implementation Pathway and Stakeholder Coordination

Implementation Phases:**Short term (Years 1–3) :**

- Continue the conversations between stakeholders, in order to close the identified gaps (Long-term assured access and planning of resource allocations; Resource allocation and accounting that are valid across the whole e-Infrastructure)
- get stakeholders involved to establish an interoperability framework for resource accounting; set up and agree on relevant metrics and checkpoints; identify tasks and responsibilities for the implementation of the federated accounting system; deploy first testbeds and feasibility studies
- Identify appropriate EC calls to get funding for the design and prototyping phases of the federated accounting system (e.g. ENSURE project in INFRA-TECH-01)
- Identify roles of research communities/RIs, resource providers/e-Infrastructures, funding bodies/policy makers in establishing memoranda of understanding

Mid term (Years 3–5):

- Pilot studies of e-Infrastructure-wide use of resource allocations and quotas
- Testbeds for the accounting system are deployed and progressively integrated into resource providers, infrastructures, research communities.
- Memoranda of Understanding are put into place

Long term (Years 5+):

- A federated, flexible resource allocation process is widely adopted
- The standardised resource accounting system is routinely used in production

Stakeholder Roles:

- **Scientific Communities, Research Infrastructures**
 - Responsibilities: establish clear computing resource needs over a timescale of several years, in terms of workflows to execute, storage needed, data access patterns, local- and wide-area network requirements; consume the information aggregated by the accounting system
 - Benefits: researchers obtain long-term access privileges and quotas in a flexible and scalable manner
- **Computing & Data Service Providers / e-Infrastructures**
 - Responsibilities: Enable e-Infrastructure-wide use of resource allocations and quotas; establish MoUs with funding bodies and scientific communities/RIs; agree with Scientific Communities and Research Infrastructures on federated accounting metrics and checkpoints, implement the associated sensors and aggregate information
 - Benefits: underutilisation of resources is minimised
- **Policy Makers / Funding Bodies**
 - Responsibilities: establish clear and multi-annual funding plans with the goal of matching eInfrastructure technological evolution with scientific needs; establish MoUs with computing & data service providers and scientific communities/RIs
 - Benefits: agencies contributing funding to e-Infrastructures can gauge their return on investment by a more efficient use of the invested funding tailored to scientific community needs

Resource Needs:

- **Funding is needed for**
 - the technical aspects of the accounting system and the multi-annual, federated resource allocation mechanism
 - developing and finalising the policy aspects of federated resource allocation (a pool of people from all stakeholders)
- **Funding timeline:**
 - Short-term, project-related funds for designing/prototyping/testbeds/pilots
 - medium-long term structural funds for commissioning and operations

Critical Dependencies:

Cross-priority dependencies with

- P01-Cross-Infrastructure Governance,
- P03- Federated Identity and Trust,
- P09-Open Science and Reproducibility

P03 Federated Identity and Trust

Strategic Rationale and Evidence Base

Strategic Importance: Federated identity and trust infrastructure enables researchers to access heterogeneous computing and data resources across organisational and national boundaries. Without seamless, secure authentication and authorisation, the compute and data continuum vision remains unattainable. This priority provides the technical identity layer that makes cross-infrastructure workflows executable and data FAIR-compliant. The challenge spans both compute and data access:

- **Compute access:** Authenticating to HPC, HTC, and cloud systems; authorising job submission and resource consumption; enabling automated workflow execution across sites.
- **Data access:** Authenticating to storage systems and data repositories; authorising data read/write/transfer operations; enabling FAIR data access (data cannot be FAIR if it cannot be accessed); protecting sensitive data under GDPR whilst enabling cross-border research.

The scale is substantial: enabling 2 million researchers across 800+ universities [66] to efficiently utilise heterogeneous computing and data infrastructures whilst maintaining security, compliance, and operational efficiency. This priority describes the importance of a sound federation of HPC, HTC and cloud systems and network of trust on the example of the HEP and Radioastronomy community but is equally relevant and important for many other communities as well.

Current State: D6.1 Section 5.6 (Security, Authorization, and Authentication) documents the authentication fragmentation: *"HEP collaborations rely on the WLCG Trust model, which enables sites to grant access to users and workflows based on their membership in federated Virtual Organizations (VOs). This approach simplifies user management but presents challenges when integrating with HPC centers that impose stricter authentication and authorization policies."*

Current landscape:

- **Federated infrastructures** (EGI, WLCG) have developed technical solutions based on AARC Blueprint Architecture.
- **HPC centres** typically require independent accounts linked to individual individuals or federated identities incompatible with VO-based access control.
- **No federated service accounts:** HPC centres provide accounts linked to identifiable individuals, creating barriers for automated workflows. HPC centers do not support service accounts or robot accounts for automated workflows.

Proven technical foundations demonstrate that federated AAI works at scale. AARC Blueprint Architecture has achieved global adoption across five continents, validating its viability as a European standard. D6.1 Section 4.8 reports current tool usage is "almost equally shared between Indigo-IAM, CERN-SSO and EGI Check-in (with MyAccessID being used by the EFP)" (D6.1). EGI Check-in successfully translates between four protocols, demonstrating technical feasibility of protocol bridging [59]. eduGAIN operates globally with national federations spanning 40+ countries. EuroHPC Federation Platform is developing federated identity management based on MyAccessID service provided by GÉANT [82]. WLCG demonstrates cybersecurity coordination at federation scale through joint OSG/EGI security working groups, providing an operational model for trust framework governance [164].

Community Evidence: D5.2 identifies critical gaps: **Gap #5:** Lack of federated identity management across European e-infrastructure; **Gap #6:** Unattended execution of long-running workflows, e.g. reauthentication or MFA which creates a "human-in-the-loop" bottleneck preventing unattended execution of long-running workflows or Token-based authentication challenges for command-line and long-running workflows among other blockers.

AARC BPA 2025 (AARC-GO80) [1] provides best practice recommendations for federated AAI specifically addressing research infrastructure needs [66]. Also the AARC-Policy-Development-Kit [3] defines policies baseline expectations for all infrastructure members, so that trust can be established. The JENA Computing White Paper identifies standardised policies across the "5-domain gap" (AAI, data transfers, networking, edge services, container support) as essential for HPC integration [115].

Technical Foundation: D6.1 Section 4.8 (Security and Trust) identifies key capabilities: *"The European Compute and Data Continuum will be a federated infrastructure spanning multiple countries, institutions, and administrative domains. In such an environment, trust and security are not add-ons but foundational enablers."*

Key capabilities from D6.1 4.8:

- **Identity and Access Management (IAM):** Federated solutions supporting multiple authentication sources, account linking, transparent attribute delivery. Once authenticated, user attributes (collaboration membership, VO roles, project grants) travel with the identity so that access policies can make decisions (D6.1).
- **Trust Framework:** Governance and technical mechanisms for mutual recognition across institutional boundaries. Sites must demonstrate compliance with minimum security baselines, privacy regulations, and facilitate the capabilities necessary for enabling the FAIR principles (D6.1). OIDC Federation could serve as a concrete technical accelerator for the Trust Framework, e.g. by adopting AARC BPA as the architecture baseline, and evaluating OIDC Federation as the scalable technical mechanism to implement the Trust Framework for OIDC/OAuth ecosystems.
- **Policy Enforcement:** Dynamic, context-aware access decisions. D6.1 provides concrete examples: "data from telescope X can only be stored in GDPR-compliant facilities within the EU" or "users must re-authenticate with multi-factor authentication when accessing Tier-O resources" (D6.1).
- **Privacy & Data Protection:** GDPR compliance throughout the authentication chain, with technical safeguards including encryption at rest and in transit (D6.1).

D6.1 Section 5.6 specifies the requirement: *"The foundation of such a solution would be a federated Authentication and Authorization Infrastructure (AAI) based on industry standards. This would enable users to authenticate once and access multiple sites securely, with consistent authorization policies."*

European Alignment: Federated identity directly enables European strategic initiatives:

- **EOSC and especially EOSC Web of FAIR Data:** Authentication federation is fundamental, as in general data cannot be FAIR if it cannot be accessed [66], except public open data, which can be FAIR without AuthN.
- **EuroHPC Federation Platform:** Unified access to EuroHPC systems requires federated identity as foundation [82]
- **Digital Decade 2030 [54]:** Secure data spaces and digital public services depend on cross-border authentication
- **eIDAS 2.0 [63]:** European Digital Identity framework provides long-term alignment opportunity

Target Outcomes and Impact Assessment

Strategic Outcomes: Federated identity and trust infrastructure delivers value across three dimensions: seamless researcher access, automated scientific workflows, and FAIR data accessibility.

Outcome 1: Unified Authentication and Workflow Automation

Researchers authenticate once with home institution credentials and access any participating European compute or data infrastructure without additional authentication steps. This eliminates the current fragmentation of passwords, certificates, and tokens whilst maintaining appropriate security standards.

Scientific workflows require extended unattended execution across multiple infrastructures. Token delegation mechanisms and risk-based authentication profiles enable progressive extension of unattended workflow duration, from days to weeks to months, aligned with the continuous processing requirements of HL-LHC and SKA Observatory operations.

Outcome 2: FAIR Data Access and Trust Framework Governance

Data cannot be FAIR if it cannot be accessed [66]. Federated authorisation enables FAIR data access across repositories whilst maintaining GDPR compliance for cross-border transfers. Research collaborations manage access rights in real-time through standardised group and role mechanisms that propagate across connected infrastructures.

Trust framework governance extends the proven WLCG OSG/EGI security coordination model [164], coordinating incident response, establishing mutual recognition across institutional boundaries, and managing liability for federated operations.

Success Indicators

Within control of trust and identity management :

- AARC Blueprint Architecture deployed across EuroHPC Tier-0 systems and major national facilities.
- Specific and realistic protocol translation operational for heterogeneous compute and data systems.
- Unattended workflow execution extended from days to weeks to months.
- Authorisation changes propagating within hours across federated infrastructure
- Trust framework governance coordinating incident response across participating countries out of control of trust and identity management.
- Time-to-first-access reduced from days to hours for new researchers – (AAI aspect).
- Providing interoperability with other major infrastructures.

European Impact: Unified access management transforms European research competitiveness by removing artificial barriers to infrastructure utilisation. Researchers recover time previously lost to authentication procedures. FAIR data principles become achievable at scale, with federated authorisation enabling cross-border access whilst maintaining appropriate protections.

European digital sovereignty is strengthened as authentication remains within European governance. The AARC Blueprint Architecture, originating from European research infrastructure collaboration, has been adopted as a global standard across five continents, demonstrating European leadership in federated identity.

Implementation Pathway and Stakeholder Coordination

Implementation Phases

- **Short term** (1–3 years)
 - Trust Governance bootstrap.
 - Operational KPI agreed.
 - Review of site-local policy and risk acceptance barriers to the use of agentic access (i.e. service accounts).
- **Mid term** (3–5 years)
 - Trust framework hardening (onboarding/offboarding procedures, audit expectations across infrastructures).
 - Enable agentic access and community-coordinated authorisation, including any requisite compensatory controls or trust framework.
 - Establish a trust and security coordination framework that includes federated community participants, including continuous readiness evaluation.
- **Long term** (5+ years)
 - Make GDPR a normal operational capability rather than project-by-project integration (including auditing and policy enforcement).

Stakeholder Roles

- **Scientific communities and Research Infrastructures:** Research Communities and Research Infrastructures including end-users (e.g. community managers, workflow owners, data stewards)
 - Responsibilities: The scientific communities including research infrastructures need to identify the specific requirements for accessing a federated HPC, HTC and cloud infrastructure and agree with other communities on European standards in collaboration with the service providers. In addition, use cases showing the full potential and benefit of an integrated HPC, HTC and cloud infrastructure will be essential to ensure funding for the integrated services and equally important compute time as well.
 - Benefits: Defining European cross domain standards will help to build a sound European Compute Infrastructure. Demonstrating the impact on research and innovation in Europe

will help to secure political support (nationally and European) and required funding for services and compute resources.

- **Service providers and facilities:** Especially HPC HTC and Cloud provider, including Trust Anchor Operators/Federation Authorities, Federated AAI Service Operators (e.g. Proxy/Broker operators), Legal/Data Protection roles (policy owners, DPO), Security Operations & CSIRT.
 - Responsibilities: Service providers need to closely cooperate and exchange with the scientific communities in order to define the new federated access policies and adapt their local AAI policies to be compatible with the enhanced needs of the scientific communities (respecting their security policies).
 - Benefits: By enhancing the access possibilities to European HPC, HTC and cloud services a new European federated compute ecosystem could be created and new communities and users will be able to exploit the established HPC, HTC and cloud infrastructure. A wider and more efficient use of the compute infrastructure would boost European competitiveness and innovation.
- **Policymaker (European-level):**
 - Responsibilities: Policy makers should analyse the potential of a European federated HPC, HTC and cloud Infrastructure and recognise the benefit for research and innovation. Sufficient funding needs to be planned both for developing the federated services across disciplines and the compute resources accessible via the federated infrastructure.
 - Benefits: The most effective use of the available compute resources will help to create synergies in Europe and thus strengthen European competitiveness. The federated compute infrastructure will also help to use synergies and operate a more cost efficient compute infrastructure.

Resource Needs

- **Funding:**
 - AAI operations capacity (sustained staff for IAM/AAI not only "project-mode" work) (short term).
 - Policy capacity specialist that can align legal/privacy constraints (short term).
 - Pilots and validation across multiple facilities (testbeds, compliance validation campaigns, integration sprints) (medium-term).
 - Training programme (certification-like) (short term).
 - Lifecycle sustainability with funding line for federation core services (not project-only). (long-term).
- **Infrastructure**
 - Security-by-design setup (short term) including security engineering time to define minimum baseline controls (logging, auditing, key management, token signing, emergency revocation).
 - Authorisation at scale (medium term), including specialist to standardise entitlement semantics and reduce reliance on VO-specific bespoke mappings.
 - Adapted access to services and facilities (short, mid and long term).

Critical Dependencies

Governance agreements. PO3 cannot scale without an agreed trust governance approach: who recognises whom, under what baseline, with what incident response, etc. – e.g. resource allocation for community access.

Standardised authorisation semantics. group/role/entitlement definitions and lifecycle management must be interoperable.

Cross-priority Dependencies with:

- PO1-Cross-Infrastructure Governance
- PO4 Heterogeneous Computing Integration

Pillar 2: Architecture & Interoperability

PO4 Heterogeneous Computing Integration

Strategic Rationale and Evidence Base

Strategic Importance: The compute continuum (seamless integration of HPC, HTC, Cloud, edge, and emerging quantum resources) is essential for data-intensive science operating at exabyte scales. HL-LHC will generate 1 EB/year requiring processing across heterogeneous resources (D5.1), whilst SKA Observatory distributes 700 PB/year to regional centres requiring diverse computational capabilities (D5.1). Scientific workflows increasingly require combinations of high-throughput batch processing, high-performance simulation, interactive analysis, and accelerated computing that no single resource type can optimally provide.

European research communities face a fragmented landscape where HPC, HTC, and Cloud resources operate under different allocation models, access policies, and technical interfaces. HEP communities have historically optimised for HTC workloads (e.g. WLCG sustained 7.4B HTC CPU hours in 2024 [59]), whilst emerging requirements for ML training, simulation, and real-time processing demand HPC and GPU capabilities that require new integration approaches.

This priority describes the importance of an integration of heterogeneous computing (continuum computing) including HPC, HTC, cloud, edge and emerging QC systems on the example of the HEP and Astronomy community but is equally relevant and important for many other communities, e.g. numerical weather prediction and foundation model training as well.

Current State: D6.1 Section 5.1 (Standardization of Interfaces) identifies the core integration challenge: *"Today, each supercomputing site uses its own portals, schedulers and software stacks. Therefore, an experiment wishing to run on multiple HPC centers must custom-adapt to each... Without a standardized approach, scaling the number of participating HPC sites remains a labor-intensive challenge."*

D6.1 Section 2.3.2.5 (Compute Federation Services) documents the compute federation landscape from D5.3 (D5.3): *"It should be mentioned that none of these [federation platforms] 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 available."*

Current landscape shows partial federation within resource types, but limited cross-type integration:

- **HPC federation:** EuroHPC Federation Platform [79] currently under development, will integrate EuroHPC supercomputing facilities, AI factories and Quantum Computers) and will provide AARC-BPA-compliant AAI across EuroHPC systems, with interoperable services for cross-site workflows [82]. Integration with HTC federations is not yet addressed and how Gigafactories will be incorporated remains also unclear.
- **HTC/Cloud federation:** WLCG federates HTC resources for workflow execution via experiment-specific layers, with centres providing access points mediated by batch systems (D6.1). EGI Federation federates HTC and Cloud resources across European e-infrastructures via common AAI, accounting, and service catalogues [59]. Similarly to WLCG, SRCnet is currently being built, the federated digital infrastructure to process, store and distribute SKA Telescopes scientific data.

The JENA Computing White Paper identifies seven technical gaps between HPC and HTC: edge services, federated access and AAI, workflow management, wide-area networking, data management, software deployment, and programming models [115]. The JENA HPC Working Group specifies six technical areas requiring standardised policies across HPC centres: AAI, data transfers, networking, edge services, container support, and software deployment [116].

Significant progress has been made through European projects. The interTwin project [111] delivered working HPC federation solutions across Cloud, HTC, and HPC testbeds with 10 pilot Digital Twin applications, now serving as "foundation for EGI Federation's HPC integration strategy" [59]. EGI's strategic transition explicitly commits to moving from "compute-focused roots" to "data-and-compute-centric infrastructure" [59], with plans to formalise HPC as a fully-fledged offering. The EXTRACT project [87] has been working on solutions for workflow execution and orchestration.

EuroHPC's Federation Platform addresses unified access to 23+ heterogeneous systems including 8 pre-exascale/exascale HPC systems, 8 quantum computers, and (at the time of writing) 13 AI Factories [82]. However, integration with existing HTC federations (e.g., WLCG with 160 sites across 40+ countries) and emerging Cloud resources remains incomplete.

In the context of BDVA, DataLabs initiative [18] brings together key stakeholders in the digital ecosystems for the implementation of the European Data Union Strategy [78] following the AI Continent Action Plan [9].

Community Evidence: D6.1 Section 2.5 consolidates requirements from WP5 analysis, identifying key gaps relevant to heterogeneous computing integration:

- **Gap #3: Interoperable allocation and accounting.** Allocations and quotas are bound to specific HPC centres, limiting flexibility. Accounting practices differ across infrastructures. Recommendation: Enable e-infrastructure-wide allocations and harmonised accounting, so projects can burst across systems whilst ensuring fair contributions and usage transparency (D6.1).
- **Gap #4: High-level and portable software environments.** Software stacks differ widely between sites; portability across CPUs, GPUs, and accelerators is inconsistent. Scientists often face brittle interfaces. Recommendation: Provide standardised, efficient software stacks (containers, CVMFS) and high-level domain-specific interfaces, enabling reproducibility and accelerator portability (D6.1).
- **Gap #5: Workflow execution and orchestration.** Workflow engines are fragmented across communities (e.g., HTCondor in HEP vs LEXIS Platform in EuroHPC). Job granularity and scheduling flexibility represent fundamental differences between HTC and HPC that must be addressed. Cross-facility workflows lack resilience and provenance capture. Recommendation: Co-design and converge on common APIs for workflow systems, with support for unattended execution, provenance tracking, and cross-infrastructure orchestration (D6.1).
- **Gap #14: Interactive access to significant-scale computing.** Interactive access to small/medium partitions is possible, but problems include delays in availability and potential underuse of resources. Recommendation: Extend scheduling/orchestration to support interactive compute use cases (D6.1).

Technical Foundation: D6.1 Section 4.1 (Compute Resources) provides the architectural framework: "*The Compute Resources capability is a foundational layer of the European Compute and Data Continuum, providing seamless access to heterogeneous computational resources distributed across HPC centers, community and commercial clouds, edge infrastructures, and emerging quantum platforms.*"

D6.1 Section 5.5 (Workflow Adaptation and Optimization) recommends: "*Integrate existing application specific workflow systems with HPC resource managers to enable native submission, data staging, and monitoring across federated infrastructures to allow for orchestration of complex workflows across grids and supercomputers*" (D6.1).

Multiple approaches exist for addressing these gaps, including meta-schedulers and platform-level orchestration as well as API-based access mechanisms are essential for workflow engine integration.

European Alignment: The ESFRI Landscape Analysis 2024 identifies that "systems of linked RIs are likely to be necessary" for complex scientific challenges, with distributed RIs requiring "central coordination" and European Partnerships expected to "reduce fragmentation of the research and innovation landscape" [71]. This provides policy-level justification for heterogeneous computing integration.

EuroHPC targets federation across its investment in pre-exascale and exascale systems, with the Federation Platform (under development) designed to provide unified access whilst preserving national sovereignty [82]. EGI's 2025+ strategic direction commits to the data-compute continuum transition [59], aligning with EOSC's federated infrastructure model [66]. Both infrastructure providers and user communities have responsibilities in achieving federation goals.

JENA identifies standardised policies across HPC centres for AAI, data transfers, networking, edge services, and container support as prerequisites for effective technical integration [13]. In order to create the compute continuum, policies should focus on integrating existing initiatives rather than forcing individual HPC or HTC centers to change.

In addition it is also critical to introduce a shared language and ontologies in order to enable resource sharing. This is a first-stage requirement that enables many other technical solutions. Ontologies are politically easier to agree on as they do not change existing habits, only require precise definition of needs.

Target Outcomes and Impact Assessment

Strategic Outcomes: Heterogeneous computing integration delivers value across two dimensions: seamless workflow execution and standardised integration infrastructure.

Outcome 1: Standardised Integration Policies Across European HPC, HTC Centres and Cloud Providers

European pre-exascale and exascale centres implement standardised policies for AAI, data transfers, networking, edge services, and container support, enabling plug-and-play integration for data-intensive science workflows. Cross-platform accounting metrics extend established benchmark methodologies (e.g., HEPSCORE) to unified CPU+GPU+quantum measurement, enabling comparable resource tracking across heterogeneous infrastructure. Policies also need to connect with Data Infrastructures and Data Providers

Outcome 2: Seamless Workflow Execution Across Compute Continuum

Researchers submit workflows that automatically decompose across optimal resource types (HTC for throughput, HPC for tightly-coupled simulation, Cloud for burst capacity, GPU for ML training) without manual resource management. The unified submission interface supports automatic resource selection across federated sites, enabling scientific workflows to span organisational boundaries whilst maintaining data locality optimisation and provenance tracking.

Success Indicators:

For outcome 1:

- Integration coverage: Sites meeting standardised integration criteria (AAI, containers, data staging, network access).
- Workflow validation: Representative HEP and RA workflows demonstrated on multiple resource types.
- Orchestration adoption: Production workflows using cross-infrastructure orchestration.
- Accounting interoperability: Common metrics (e.g., HEPSCORE-based) reported across HPC/HTC/Cloud sites.

For outcome 2:

- Workflow portability: Percentage of scientific workflows executable across multiple resource types without modification.
- Resource utilisation efficiency: Utilisation rate of heterogeneous resources under unified allocation.
- Time-to-science: Reduction in time from workflow submission to results for multi-resource workflows.
- Integration coverage: Number of sites implementing standardised integration policies.
- Accounting harmonisation: Common accounting framework operational across federated HPC/HTC/Cloud resources.

European Impact: Heterogeneous computing integration maximises return on Europe's distributed infrastructure investments. Rather than requiring researchers to choose between HPC for simulation, HTC for throughput, or Cloud for flexibility, integrated infrastructure enables optimal resource matching for each workflow component. This increases utilisation efficiency, reduces time-to-science, and enables new computational approaches combining traditional physics simulation with ML-accelerated analysis.

Implementation Pathway and Stakeholder Coordination

Implementation Phases

- **Short term (1-3 years):**
 - Agree on a common (descriptive) language – ontologies – over the domain specific ontologies. Once having them, then they can be mapped to HPC systems.

- Reliance on and starting with existing initiatives, with e.g. EuroHPC Federation Platform (EFP) establishing a European federated and secure supercomputing, AI and quantum computing service.
- Start discussion about policies around existing initiatives.
- Analyse the established policies, develop suitable prototypes and start with the implementation.
- **Mid term (3–5 years):**
 - Show first implementations and API-based access across initiatives EuroHPC, EGI and Cloud providers.
 - Demonstrate different use cases including HTC, batch execution, AI training, and inference with different requirements.
 - Defining next steps based on lessons learned.
- **Long term (5+ years):**
 - Unified federation across different HPC, HTC and cloud providers based on lessons learned from earlier phases.

Stakeholder Roles

- **Scientific Communities, Research Infrastructures**
 - Responsibilities: Community workflows and codes must be adapted for the efficient use of a federated compute continuum infrastructure, especially HPC resources.
 - Benefits: Provide full ability to use the compute continuum.
- **Computing & Data Service Providers / e-Infrastructures**
 - Responsibilities: Align policies and protocols, ensure federation fits funding models.
 - Benefits: Enable the compute continuum, offer services to a wide range of users.
- **Policy Makers / Funding Bodies**
 - Responsibilities: Provide political agreement, legal framework, and funding support.
 - Benefits: Most efficient use of funding, to help the EU keep being competitive. Realise a federated compute and data infrastructure.

Resource Needs

- A federated European compute continuum involving all relevant services providers requires first support by politics and funding agencies.
- An efficient approach to establish the federated compute continuum would be dedicated funded projects both for service providers and scientific communities.
- Both compute service providers and scientific communities need to agree and engage in the joint projects. To build the legal framework and technical solutions is a basic requirement, however, compute resources (HPC, HTC, Cloud, etc) for scientific communities and research infrastructures are needed as well in order to allow the federated use of compute resources.

Critical Dependencies

First, a common ontology for the different stakeholders (scientific communities, research infrastructures and services providers) needs to be defined. Then, the legal & political framework can be elaborated and defined. Also, a political discussion and decision with a firm funding commitment is needed. The continuum computing priority has of course also cross-dependencies to almost all other priorities in this document.

Cross-priority dependencies with:

- PO1: Cross-Infrastructure Governance
- PO2: Federated Resource Allocation
- PO3: Federated Identity and Trust
- PO5: Data Management and Transport
- PO6: Workflow Orchestration and Management.
- PO8: Software Portability.
- P11: Community Collaboration and Co-Design.

P05 Data Management and Transport

Strategic Rationale and Evidence Base

Strategic Importance SPECTRUM use cases are based on the acquisition and processing of very large amounts of data. The structure of the data (their architecture) and their management are fundamental to European e-Infrastructure competitiveness. The capability to store data and efficiently transport them to HPC centres is essential for HEP and RA communities, whilst native management according to FAIR principles enables cross-domain research and Open Science.

Current State: The management of large amounts of data around HTC computing centres is reasonably addressed by existing tools within each domain. Pan-European research networks (GÉANT) and dedicated overlays (e.g. LHCONE, LHCOPN) provide high-capacity transport, though sustained investment is needed for exabyte-scale data movement. Gaps exist primarily between different types of infrastructures (e.g., data infrastructures and HPC facilities) and across different domains (e.g., standardisation of FAIR storage) (D6.1).

Community Evidence: Three main gaps have been identified in the data domain (D6.1):

- Gap #6 (Data federation and FAIR principles): No unified architectural approach or core implementation for FAIR storage across domains, including standardised metadata and provenance tracking
- Gap #7 (Data transfer and staging): Transfer of large datasets between long-term archives and transient compute resources is highly fragmented and often relies on ad-hoc tools. Remote access in some HPC is not possible as outside connections from nodes is blocked for security reasons.
- Gap #9 (Sensitive data handling): Complexity of integrating GDPR-sensitive and special category data threatens to impact usability and performance for open scientific data use cases

Technical Foundation: D6.1 4.2 (Data Resources) defines four capabilities comprising the Data Continuum (D6.1):

- Federated Storage: Distributed storage presenting unified namespace to users and applications, with data replication, integrity management, and embedded access controls
- Data Transfer: High-performance data movement services (e.g., FTS, Rucio, XRootD) orchestrating bulk transfers across research network backbones (e.g. GÉANT, LHCONE)
- Data Discovery and Cataloguing: Metadata catalogues and registries enabling data findability, with provenance tracking supporting FAIR principles
- Data Lifecycle Management: Automated policies for retention, tiered storage migration, integrity verification, and long-term preservation

These capabilities transform fragmented storage silos into a seamless, policy-compliant ecosystem for exabyte-scale scientific data.

European Alignment: ESFRI identifies data sharing across Research Infrastructures as essential for complex scientific challenges [71]. EOSC provides the policy framework through its federated infrastructure model [66], whilst EuroHPC requires efficient data access to exploit European compute capacity for data-intensive workloads [82].

Target Outcomes and Impact Assessment

Strategic Outcomes

Outcome 1: Federated FAIR Data Management

Cross-domain standards for federated FAIR data management enable Research Infrastructures to share data whilst maintaining custodian control. This includes common metadata catalogues, provenance tracking, and access interfaces built on progressive maturity frameworks. Flexible architectures decouple sensitive and non-sensitive data repositories, enabling GDPR compliance without impacting performance for open scientific use cases.

Outcome 2: Automated Data Movement at Scale

Large-scale data movement infrastructure operates with advance planning capabilities across HTC facilities, HPC centres, and cloud providers. High-performance transfer services (e.g., FTS, Rucio) integrate with HPC storage systems, enabling autonomous data staging for workflows spanning federated European infrastructure. Multi-protocol support ensures compatibility with diverse storage backends.

Success Indicators

- FAIR adoption: Research Infrastructures implementing common FAIR data management frameworks
- Transfer automation: Production workflows with fully automated data pre-staging across infrastructure types
- Integration breadth: HPC centres, cloud providers, and HTC sites integrated via common data movement protocols
- Data accessibility: Representative datasets available for analysis within target timeframes
- Interoperability validation: Cross-domain workflows demonstrated using federated data services

European Impact

Federated FAIR data management strengthens European data sovereignty by enabling Research Infrastructures to share data under common governance whilst retaining custodian control. Automated data movement maximises return on EuroHPC investment by enabling data-intensive workloads to exploit European compute capacity. Shared data infrastructure reduces duplication across Research Infrastructures, improving cost efficiency and sustainability.

Implementation Pathway and Stakeholder Coordination**Implementation Phases****Short term (1–3 years):**

- Perform a gap analysis to select pilot use cases related to FAIR storage, high-performance data movement and access, and sensitive data handling, with sufficient coverage for different topics.
- Identify critical interfaces between:
 - Research Infrastructures (RI) generating or stewarding large datasets
 - e-Infrastructures providing compute, network and storage services
 - Community and cross-domain data-access interfaces / APIs (including metadata)
- Identify and evaluate viable architectures for federated FAIR storage, metadata cataloguing, and provenance tracking.
- Pilot one or more representative use cases, such as for example:
 - FAIR-compliant cross-RI data federation
 - Automated data staging from archives into HPC environments
 - Handling of mixed sensitive/non-sensitive datasets

Mid term (3–5 years):

- Prototype and demonstrate integrated FAIR data management across multiple RIs and e-Infrastructure providers.
- Validate that architectures are fit for purpose for exabyte-scale data volumes.
- Expand scope of supported use cases:
 - High-frequency data ingestion (e.g., radio astronomy streaming or GW low-latency data)
 - Multi-protocol access to federated storage
 - GDPR-sensitive workflows using secure enclaves or logically separated data zones
- Develop standards and practices governing:
 - FAIR metadata schemas
 - Access control & identity integration
 - Operational interfaces for cross-domain data movement
 - Compliance requirements for RIs and e-Infrastructure providers
- Draft a Blueprint Architecture for the European Data Management & Transfer ecosystem.

Long term (5 years):

- Roll-out federated data management and automated transport capabilities across European e-Infrastructures.

- Onboard more communities adopting federated FAIR data services, including new scientific domains.
- Ensure sustainability and continual evolution through collective governance, coordinated investment, and community-driven standards.

Stakeholder Roles

• Scientific Communities, Research Infrastructures

Responsibilities:

- Co-design federated FAIR data architectures, including interoperable metadata standards and provenance models.
- Provide domain-specific requirements for data retention, curation, sensitivity, and access.
- Adopt harmonised interfaces and participate in pilot and validation campaigns.

Benefits:

- Seamless access to distributed data holdings with uniform FAIR-compliant interfaces.
- Increased scientific productivity due to reduced time-to-access and improved dataset reuse.
- Greater visibility, interoperability, and cross-domain collaboration.

• Computing & Data Service Providers / e-Infrastructures

Responsibilities:

- Deploy scalable federated storage systems with unified namespaces and policy-controlled access.
- Provide high-performance, automated data-transfer capabilities integrated with workflows.
- Adopt high-level data-access APIs
- Provide secure environments for sensitive and special-category data, aligned with GDPR and ethical frameworks.
- Monitoring services

Benefits:

- Maximised utilisation of storage, network, and compute resources through coordinated data movement.
- Reduced operational overhead thanks to automation and standards-based interfaces.
- Broader user base as more communities adopt shared data services.

• Policy Makers / Funding Bodies

Responsibilities:

- Ensure continued investment in high-capacity pan-European research networks (e.g., GÉANT).
- Support development of FAIR-aligned infrastructure and long-term preservation facilities.
- Coordinate regulatory and governance frameworks for cross-border data management (GDPR, FAIR metrics, EOSC policies).

Benefits:

- Strengthened European data sovereignty and reduced duplication across scientific domains.
- More efficient use of public infrastructure investments.

Resource Needs

- Compute, network, and storage resources scaling across phases from pilot to exabyte-level deployment.
- Persistent high-speed connectivity through GÉANT and domain overlays (LHCOPN, LHCONE, or equivalents).
- Monitoring services
- Long-term archival and nearline storage; data orchestration tools including intelligent tiering systems; metadata catalogues.
- Engineering resources within RIs and e-Infrastructures to develop, port, test, and integrate FAIR and data-transfer capabilities.
- Governance structures for standards, certification, and service evolution.

Critical Dependencies

- Commitment from infrastructure providers to expose harmonised and interoperable data interfaces.
- Adoption of common metadata standards, provenance models, and access protocols across domains.

Cross-priority dependencies with:

- PO1 (Cross-Infrastructure Governance)
- PO2 (Federated Resource Allocation), PO3 (Access, Security & Trust)
- PO4 (Heterogeneous Computing Integration)

PO6 Workflow orchestration and management

Strategic Rationale and Evidence Base

Strategic Importance: Scientific workflows are becoming increasingly sophisticated. Advanced workflow management systems capable of orchestrating multi-step processes across distributed resources and domains, ensuring data locality, and optimizing resource utilization, are therefore needed. Unifying workflow environments is crucial for sustaining and expanding innovative approaches such as AI-based workflows. Automated workflow tools are needed to streamline updates and deployments of e.g. AI models in real-time applications. Mixed workloads that combine embarrassingly parallel tasks with tightly coupled AI/ML processes must also be supported. Modern systems must adapt to support this diversity by providing heterogeneous architectures that integrate CPUs, GPUs, and high-speed interconnects. Streamlining workflow software and using automated workflow management for orchestrating, monitoring, accounting, is also important to reduce operational overhead and maximise resource utilisation

Current State Well-developed, automated tools enable workflow management for complex tasks including submission, monitoring, accounting of workloads, and their data management. Integrating these management systems into HPC centers has required careful tailoring to match the computational resources, configuration, and policies. Standard mechanisms for integrating HPC centers into workload management systems are required to maximize HPC adoption. This has been tackled by european projects such as EXTRACT project [88]

Community Evidence The analysis of the interoperable access policies (D5.2) and the technical blueprint for compute and data continuum (D6.1) prepared by the SPECTRUM project identifies following gaps related to workflow orchestration:

- **#5: Workflow execution and orchestration:** Workflow engines are fragmented across communities.. Cross-facility workflows lack resilience and robust, standardised provenance capture. .
- **#15: High-level end-user interfaces:** Most e-Infrastructures do not offer high-level, abstracted interfaces. Domain scientists should not need to know system details and interact with brittle application interfaces.
- **#16: Unattended execution of long-running workflows** Ensure reliable and unattended execution of long-running workflows and user-operated services.

Technical Foundation Workload Management relies on the following supporting SPECTRUM Compute and Data Continuum capabilities (see D6.1 Section 4.4): Workflow Management, Resource Orchestration, Task Scheduling, Provenance

European Alignment Workflow management is crucial to broader European strategies for sustaining and expanding innovative approaches such as AI-based workflows, and mixed workloads that combine embarrassingly parallel tasks with tightly coupled processes, such as those used for the application of AI/ML. Resource providers will benefit from automated workflow management that makes efficient usage of their resources and their exposure to user communities. The EOSC web of fair data and service benefits from a unified workflow management that is independent of the implementation details and that enables seamless access for European researchers to compute and data resources.

Target Outcomes and Impact Assessment

Strategic Outcomes

- Workflow orchestration evolves into a cross-facility capability with built-in resilience and capturing of provenance.
- Workflow scheduling and orchestration supports interactive compute use cases
- High-level general and domain-specific user interfaces are deployed to lower the threshold to a broader research community.

Success Indicators

- Number of use cases/domains supported by a standardised workflow execution environment (to be tracked per environment and framework)
- Number of resource providers enabling cross-facility capabilities
- Number of supported use cases and associated number of users

European Impact Standardised workflow management enables researchers to access compute resources and execute their workflows efficiently and independently on the implementation details, thus decreasing the time-to-completion of their research projects and boosting the competitiveness of the European Research Area.

Implementation Pathway and Stakeholder Roles

Implementation Phases

- **Short-term (Years 1-3).**
 - Continuation of the gap analysis that started with the landscape analysis of SPECTRUM for a selection of use cases.
 - Identification of critical interfaces with 1) research infrastructures, 2) e-Infrastructures, 3) (general) user interfaces/API's.
 - Identification and continuation of work for viable orchestration & workflow management frameworks for supporting cross-facility & cross-domain scientific workflows.
 - Pilot one or more use cases (HEP, SRCnet, LOFAR).
- **Medium-term (Years 3-5)**
 - Prototyping & demonstration of selected use cases with selected e-Infrastructure providers.
 - Validation of being fit-for-purpose.
 - Exploration of scope of use cases that can/should be supported (possibly resulting in multiple classes of cases, each requiring a specific solution).
 - Drafting standards and practices that (each) stakeholder must adhere to (e.g. Blueprint Architecture).
- **Long-term (Years 5+)**
 - Roll-out across European e-infrastructures, grow communities & research infrastructures that build on/benefit.

Stakeholder Roles

- **Scientific Communities, Research Infrastructures, Computing & Data Service Providers / e-Infrastructures**
 - Responsibilities: Co-design and converge on a limited set of interoperable workflow systems, with support for unattended execution, provenance tracking, and cross-infrastructure orchestration.
 - Benefits: (for scientific communities / research infrastructures) scientific output is accelerated; (for Computing & Data Service Providers / e-Infrastructures) resource utilization is optimised.
- **Scientific Communities, Research Infrastructures, Computing & Data Service Providers / e-Infrastructures**
 - Responsibilities: Extend scheduling/orchestration to support interactive compute use cases.
 - Benefits: (for scientific communities): being able to perform interactive data analysis tasks on a scalable infrastructure, with predictable, resilient behaviour and sufficiently low latency/performance; (for Computing & Data Service Providers / e-Infrastructures) efficient use of owned resources.
- **Computing & Data Service Providers / e-Infrastructures**
 - Responsibilities: support deployment of/integration with domain-specific user interfaces by e.g. providing service hosting and integration APIs.
 - Benefits: diverse use of infrastructure; (for users) the high-level part of the scientific workflows can be scaled up to ever-increasing dataset dimensions

Resource Needs

- Compute, network, (temporary) storage resources scaling up through the phases (pilot/demonstration/full deployment).
- Engineering resources for porting/building/demonstrating capabilities (embedded in research infrastructures & e-infrastructure providers).
- Long-term resources for supporting & maintaining shared capabilities.

Critical Dependencies

- Coordination between research infrastructures & e-infrastructures towards achieving fit-for-purpose unification of orchestration & workflow frameworks.
- Commitment from e-infrastructure providers to provide the required unified capabilities through harmonized interfaces.

Cross-priority dependencies with:

- PO1 (Cross-Infrastructure Governance),
- PO2 (Federated Resource Allocation),
- PO3 (Access, Security & Trust),
- PO4 (Heterogeneous Computing Integration),
- PO5 (Data Management and Transport),
- PO8 (Portability – Software Portability)

Pillar 3: Software & Science Enablers

PO7 AI

Strategic Rationale and Evidence Base

Strategic Importance: AI methods are reshaping HEP and RA workflows across simulation, reconstruction, calibration, analysis, and instrument operations: pre-trained models at the LHC trigger classify collision events in real time, AI at telescope edges detects astrophysical transients within milliseconds, and digital twins model detector radiation damage, alignment shifts, and telescope-array calibration drift. ML-based fast simulation has demonstrated two-to-three orders of magnitude computational cost reduction for high-impact HEP and RA use cases whilst maintaining scientific accuracy [22]. The European AI moment is open: the AI Continent Action Plan commits over €20 billion to AI infrastructure, AI Factories enter operations in 2025–2026, AI Gigafactories are proposed at €100 million-plus per site, and the AI Act sets the regulatory baseline. With a Commission-estimated European compute capacity gap of around 20 GW by 2036 and three hyperscalers holding roughly 70% of the European cloud market, sovereign-supply development is a load-bearing dimension of the build-out. HEP and RA must be part of this effort or risk being structurally side-lined; GPU resources are not yet accepted as "pledged" allocations for LHC experiments, exposing a structural gap between AI demand and the resource accounting frameworks that fund scientific computing.

AI agents that plan and execute scientific workflow steps under human oversight (analysis triage, anomaly investigation, calibration response) extend AI integration from training and inference into operational decision-making. This requires PO3 service-account access, PO6 orchestration, and PO9 audit and provenance.

Current State: Europe possesses substantial but fragmented AI infrastructure. EuroHPC has procured supercomputers across Europe, including exascale systems, and 19 AI Factories have been selected across three waves (December 2024, March 2025, October 2025) [7]. AI capabilities are not yet integrated as first-class infrastructure across most facilities, which privilege batch-mode access over the elastic processing that production AI workflows require. Access modalities at EuroHPC sites (different AAI stacks, batch-oriented allocation) are not yet aligned with the AI lifecycle. A strategic tension also exists at the hardware layer: industry GPU evolution prioritises low-precision (AI-optimised) computation at the expense of FP64 capability that remains critical for Lattice QCD, parts of HEP detector simulation, and high-precision astrophysical workflows (D6.1 §5.7). Without active engagement at HPC procurement and design stages, future European systems will support either AI workflows or HEP/RA precision needs, not both.

Community Evidence: D6.1 identifies a key gap directly relevant to AI integration [5]: Gap #11: AI/ML integration. AI/ML workloads are increasingly central and are being investigated at every step of data-processing pipelines, from data acquisition to simulation and analysis. Recommendation: enhance services in AI/ML platforms to support distributed training, inference deployment, and access to foundation models, integrated with HPC and domain workflows.

Community surveys reveal that 90% of respondents report software performance improvement as an issue for AI workloads, yet 58% lack the resources to address it [12], quantifying the software investment gap. D6.1 §5.7 sets out the AI-vs-precision tension and three associated actions: identify analysis and simulation tasks suitable for ML acceleration; communicate precision and performance needs to HPC hardware planners; train physicists and astronomers in ML theory, frameworks and tools to fully exploit AI Factories.

Technical Foundation D6.1 Section 4.5 (AI/ML and HPC Applications) defines four capabilities for data-intensive research [5]. PO7 builds on each as follows:

- Advanced Analytics and Interactive Exploration. Jupyter, RStudio and Dask-based platforms linked to the data federation that allow seamless scale-out from a researcher's laptop to large external

resources. Target: hypothesis-generation workflows that start interactive at PB scale and transition to offline batch without per-site re-engineering.

- **AI/ML Platforms.** Distributed, GPU-rich training environments spanning HPC, cloud, and edge; scalable inference services; FAIR-compliant model registries that work as searchable "model zoos" enabling, for example, an astrophysicist to find a transient-classifier and fine-tune it with local data. Federated learning for privacy-preserving training across distributed observatories where bulk data movement is impractical (e.g. RA observatory-level streams; selected HEP collaboration-internal datasets).
- **Simulation Frameworks and Digital Twins.** Domain simulation (GEANT4 for HEP detectors, N-body and fluid solvers for astrophysics) combined with AI-assisted prediction. Digital twin platforms such as itwinai allow domain experts to use multi-node training and optimisation without dedicated ML engineering. Concrete targets: digital twin of at least one HL-LHC sub-detector for radiation and alignment modelling; digital twin pilot for an SKA precursor station for calibration prediction.
- **Large-Scale Data Processing Pipelines.** Reconstruction, Monte Carlo production, RA data reduction. AI components integrated into existing pipelines without per-site customisation, running on heterogeneous compute through standardised access protocols (PO4).
- **Agentic AI.** AI agents that plan and execute scientific workflow steps under human oversight, extending AI integration from training and inference into operational decision-making.

European Alignment: Engagement with the European AI policy stack is the strategic backbone of this priority:

- **Regulatory.** AI Act (Regulation (EU) 2024/1689) sets transparency, documentation, dataset-quality and risk-classification obligations that affect AI used in safety-relevant scientific workflows (detector control, telescope operations), as well as the model-registry and provenance work in this priority.
- **Coordination.** The AI Office (DG CNECT, 2024) is the central EC coordination point for AI strategy and AI Act enforcement, including scientific-advisory channels.
- **Funding.** The AI Continent Action Plan (2025) commits over €10 billion for AI Factories and over €20 billion for AI infrastructure overall. The Apply AI strategy and the Cloud and AI Development Act proposal (CAIDA, target publication 27 May 2026, EU technological sovereignty package) target a sovereign EU AI and cloud stack.
- **Infrastructure.** AI Factories (EuroHPC, 19 sites across three waves [7]) carry an explicit obligation under their governing regulation to support the official European Data Spaces. AI Gigafactories proposed at scale capable of training 1 billion-parameter-class foundation models. IPCEI-CIS provides a co-funding pathway for sovereign federated cloud.
- **Science.** The RAISE programme coordinates AI capabilities for science across the EU via SCIANCE, its AI-in-science landscape-mapping project. The SCIANCE working group for Astronomy and Fundamental Physics provides the route for the RA and HEP communities; substantive engagement and use-case contribution through that working group is the actionable EC route in this priority.
- **Data.** EOSC enables FAIR-compliant model and dataset sharing across communities [6].
- **Strategic frameworks.** ESFRI identifies the need for Research Infrastructures to adopt emerging technologies including AI to maintain scientific competitiveness [27].

Target Outcomes and Impact Assessment

Strategic Outcomes:

Outcome 1: AI-Enabled Scientific Workflows. AI capabilities are integrated into scientific workflows as first-class infrastructure rather than add-ons. GPU-rich platforms support distributed training, inference deployment, and access to foundation models across HPC, HTC, and cloud resources. AI surrogate models achieve two to three orders of magnitude computational cost reduction for high-impact HEP and RA use cases, opening research questions previously computationally intractable.

Outcome 2: Sustainable AI Capability. FAIR principles apply not only to training data but also to training processes and models. Federated learning enables model training where data cannot be centralised due to sovereignty, GDPR, or volume constraints. Energy-aware scheduling addresses AI energy consumption whilst building on European Green500 leadership.

Outcome 3: HEP and RA Strategically Positioned in the EU AI Programme. HEP and RA hold formal representation in the RAISE AI for Science pillar. HEP and RA precision requirements are incorporated in AI Factory and AI Gigafactory procurement specifications, preserving FP64 capability alongside AI-optimised acceleration.

Success Indicators:

- Resource accounting reform. GPU and AI-Factory allocations to HEP and RA experiments accepted as pledged resources within the WLCG and SRCNet accounting frameworks within the mid-term horizon.
- Operational AI workloads on AI Factories. A representative subset of European AI Factories with operational HEP or RA pilot workloads.
- FAIR AI lifecycle. Training data, processes, and models documented and preserved following FAIR principles; shared model registries operational with measurable cross-community reuse.
- Precision-capable AI infrastructure. Lattice QCD (or equivalent FP64 reference workload) sustained at acceptable performance on at least one AI-Factory-class system.
- Programme participation. HEP and RA representation secured in the AI for Science pillar of the RAISE programme; recurring engagement with the AI Office on scientific AI compliance.
- Workflow integration. AI methods integrated into scientific workflows through common orchestration frameworks (PO6).
- Energy efficiency. AI workloads scheduled with energy awareness; reporting integrated with P12 sustainability accounting
- Agentic AI workflows. Agent-mediated scientific workflow execution (analysis triage, anomaly investigation, calibration response) operational across HEP and RA collaborations under a documented oversight regime; agent activity audited within WLCG and SRCNet operational reporting.

European Impact: AI integration is a strategic contribution to European digital sovereignty: indigenous AI capability for scientific research reduces dependence on external providers, and the FAIR-compliant model and dataset infrastructure provides a sovereign alternative to commercial model marketplaces. Engagement with the EU AI policy stack positions HEP and RA as evidence-base contributors rather than peripheral users. Workforce uplift through AI training builds capacity that flows back into European industry. Coupled with sustained FP64 capability in procurement specifications, the priority safeguards the long-term feasibility of precision-dependent science within the AI-led infrastructure transition.

Implementation Pathway and Stakeholder Roles

Implementation Phases:

- In the short-term:
 - Formal engagement with the RAISE programme to secure HEP and RA representation in the AI for Science pillar. Identify SPECTRUM-side interlocutors and a governance route.
 - Per-experiment assessment of which analysis and simulation tasks are suitable for ML acceleration and which require sustained FP64 (per D6.1 §5.7 Action 1).
 - Pilot validation of AI platforms, federated storage and orchestration on a small set of AI Factory sites (PO4, PO5).
 - Communicate FP64 precision and performance needs to EuroHPC procurement bodies so future procurements specify balanced hardware mixes (per D6.1 §5.7 Action 2).
 - Pilot federated-learning deployment for one cross-observatory RA use case and one HEP collaboration-internal use case.
 - Engage the AI Office on AI Act application to scientific AI (model documentation, dataset transparency, risk classification of operational AI, including the agent-class profile).
 - Define oversight framework for agentic AI in scientific workflows: autonomy scope, approval requirements, audit requirements, interfaces with PO3 service-account access and PO9 provenance.

- Pilot one HEP analysis-triage agent and one RA transient-follow-up agent under bounded oversight: PO3 service-account access, PO6 orchestration, PO9 audit and provenance.
- Mid-term:
 - Integrating cloud and HPC infrastructures, with the cloud providing a unified access model, can lower the entry barrier for scientists and accelerate the adoption of larger-scale AI workloads across HEP and research infrastructures. This approach leverages cloud interoperability while providing access to HPC-grade compute for the most demanding workloads.
 - GPU resources accepted as pledged within WLCG and SRCNet accounting frameworks.
 - Agent-mediated workflow execution operational across multiple HEP and RA workflows; oversight framework documented and adopted by participating infrastructures.
- In the long term:
 - Sustained, interoperable European continuum with production-grade federated AI/ML environments, automated provenance tracking, formalized co-design and governance models, and embedded AI/ML as strategic capabilities underpinning Europe's exascale science leadership.
 - Agentic AI as a routine capability for HEP and RA; controls, audit trails, and incident reporting integrated with WLCG and SRCNet operational tooling; AI Act compliance baseline applied to scientific AI agents.

These phases rely on standardization, co-design with scientific communities, advanced resource orchestration, sustainable policies, and a skilled workforce.

Stakeholder Roles

Scientific Communities

- **Responsibilities:** Identify workloads suitable for AI acceleration and those requiring sustained FP64; produce workload-mix specifications as input to HPC and AI Factory procurement; drive FAIR adoption for models and training processes; engage with RAISE governance via designated SPECTRUM interlocutors.
- **Benefits:** AI-enabled workflows opening intractable research questions; secured AI Factory and Gigafactory access; recognition of GPU contributions in accounting frameworks.

Research Infrastructures

- **Responsibilities:** contribute HEP and RA use cases and requirements to the SCIENCE Astronomy and Fundamental Physics working group; coordinate federated-learning pilots across member sites; provide foundational datasets and instrument-edge integration points for AI deployment.
- **Benefits:** improved utilisation; AI-assisted detector calibration, radiation and alignment modelling, and transient detection at lower marginal cost.

Computing & Data Service Providers / e-Infrastructures

- **Responsibilities:** integrate AI platform capability across the continuum (EuroHPC supercomputers, AI Factories, national HPC, EGI-federated grid and cloud sites, WLCG, SRCNet); reconcile AAI across community AAls and AI Factory access; provide elastic allocation modes for AI workloads alongside batch; implement model-registry hosting and federation; coordinate energy-aware scheduling.
- **Benefits:** standardised interfaces, improved utilisation, reduced operating costs through shared FAIR model infrastructure, and positioning of EGI/WLCG/SRCNet alongside AI Factories as the operational substrate for scientific AI.

Policy Makers / Funding Bodies

- **Responsibilities:** steer coherent EU and national investment across the AI continuum, covering AI Factories, EuroHPC and national facilities, horizontal HTC/Cloud federations like EGI or community federations like WLCG and SRCNet; include HEP and RA precision needs in procurement and provide multi-year compute allocations for science (PO2); sustain HEP and RA contribution to SCIENCE; ensure cross-DG coordination (DG RTD, DG CNECT, DG GROW).
- **Benefits:** coherent return on coordinated EU and national AI investment; visible scientific impact across European AI policy and the SCIENCE programme; sovereign capability for high-value science; alignment of AI Act objectives with operational scientific practice.

Resource Needs

- Compute. Annual GPU-hours across the continuum (AI Factories for training, national HPC and EGI-federated grid sites for inference and analysis, WLCG and SRCNet for community workflows), with sustained FP64 capability allocation alongside AI-optimised allocation.

- Storage. Training-data staging capacity and model-registry storage.
- Workforce. ML engineers and ML-physicist hybrid roles within RIs and infrastructure providers, aligned with the P13 framework.
- Software development. Integration of AI capabilities into existing reconstruction, simulation, and analysis pipelines, aligned with PO8 and PO9.
- Engagement. Sustained coordination capacity for SCIANCE working-group, AI Office and AI Factory engagement.

Critical Dependencies

- SRIDA priorities: PO2 (multi-year AI compute allocations), PO3 (cross-AAI and service-account access for agentic workflows), PO4 (HPC/HTC/cloud integration), PO5 (training-data movement and federated-learning access), PO6 (orchestration of AI and agentic execution), PO8–P10 (FAIR model lifecycle, reproducibility, audit and provenance), P12 (energy-aware scheduling), P13 (ML and hybrid-role capacity).
- Community and infrastructure: WLCG and SRCNet accounting frameworks accepting GPU and AI-Factory allocations as pledged resources; EGI Federation onboarding AI workload patterns (elastic allocation, model-registry hosting); per-experiment HEP and RA workload assessment for ML suitability versus FP64.
- Policy and timeline: AI Act enforcement (AI Office and Member State authorities, including agent-class profile); AI Factory rollout (EuroHPC, 19 sites scheduled); AI Gigafactory selection (proposed, outcome not yet decided); SCIANCE working-group structure and RAISE governance evolution.
- Risk: if the post-2027 EuroHPC procurement cycle drops FP64 in favour of AI-optimised throughput, the precision-dependent fraction of HEP and RA computing loses its European home.

P08 Software Portability

Strategic Rationale and Evidence Base

Strategic Importance SPECTRUM envisions a European Compute and Data Continuum, which will enable research infrastructures and researchers to access and use the most appropriate data centre for any given task in hand. For example, the choice of centre may depend on data locality, connectivity, availability of specialist hardware, or simply on the distribution of load across the continuum. Users will benefit from seamless access to the most effective and efficient e-infrastructures to address their needs.

Taking advantage of these capabilities requires that the software tools and workflows required by the user are available on the appropriate e-infrastructure. This requires that they be *portable* across centres: tools developed in one centre should run without modification in other centres, and, where possible, should automatically take advantage of the special capabilities each centre offers, such as accelerator hardware.

In this way, the SPECTRUM Compute and Data Continuum will *maximise the efficiency* of the use of compute resources while simultaneously *eliminating the overhead* imposed on researchers and research infrastructures by a need to modify software depending on locality and *improving reproducibility and sustainability* of the software by decoupling it from specific execution environments.

Current State We identify two essential aspects of the software portability ecosystem: whether the code itself is engineered in such a way as to make it easily portable, and whether there exist software delivery mechanisms that make it possible to deploy, manage and execute the software on multiple platforms. We discuss these aspects independently.

With respect to the first aspect—software engineering—we note that most contemporary HEP and RA software packages are based on large and highly-specialized codebases that have been developed over many years ly deployed.primarily to target the dominant x86_64 CPU architecture. Software is typically optimized at compile-time to the specific characteristics of the system on which it will be executed, and is dependent on specific features of the execution environment, such as system libraries or configuration files, and portability has not been prioritised. In many cases, work is now underway to adapt existing code to accelerator hardware, notably GPUs, in order to take advantage of their superior energy efficiency and computational throughput. While this work is a step towards portability, it often targets specific hardware environments that are not generically available.

Similarly, although individual compute and data centres generally provide ways—such as modules or containers— for users to configure their execution environment and choose from available software, there is little consistency across the ecosystem. The variety of different approaches and technologies easily leads to brittle environments, poor software portability, and concomitant confusion.

However, we note that this situation is beginning to change. The scientific community is maturing in its approach to software development. In particular, we note the success of projects such as EVERSE [65] that are promulgating best practices in research software engineering across the European scientific landscape, while the FAIR Principles for Research Software (**FAIR4RS** [89]) provide a roadmap for continued improvement.

At the same time, a range of tools that can help are becoming available. These range from low-level abstractions over HPC hardware (e.g. Kokkos [1214]) through containerized systems that encapsulate software together with its dependencies (e.g. Apptainer [18]) to systems like **CernVM-FS** [36] and **EESSI** [90] that can distribute and manage complex software environments.

We therefore identify the coming years as a critical period: the principles and tools that can address the current challenges are available, but coordinated, focused effort across the scientific and technical community is necessary to fully take advantage of them.

Community Evidence The SPECTRUM Technical Blueprint specifically identifies portability as a key requirement (**D6.1 Section 2.4** recommendation 4). This requirement is derived from **D5.3** Requirements 5 (“Enable portability of compute tasks across accelerated compute platforms”) and 7 (“Establish common SW stacks for compute applications”).

Technical Foundation This priority is captured in the following SPECTRUM Compute and Data Continuum capabilities (D6.1 Section 4):

- Software & Execution;
- Compute Continuum.

European Alignment Software portability is critical to the effective use of HPC/HTC resources at European scale. For example, EESSI is supported through the EuroHPC Centre-of-Excellence MultiXscale [137], and is used by the European Federation Platform [79] as the basis for its [Federated Software Catalogue](#) [90]. Meanwhile, EOSC specifically supports the deployment of containerized applications through its [Cloud Container Platform](#) [38].

The issues surrounding software engineering quality and its impact on portability are also widely acknowledged through major European projects, including, for example, [EVERSE](#) [65] and [EXTRACT](#) [87], while the ESCAPE Project's **Open-source Scientific software and Service Repository** [70] demonstrated the potential of centralized software distribution systems.

Target Outcomes and Impact Assessment

Strategic Outcomes

- Flexible transfer of software and workflows between data and compute centres with minimal adaptation.
- Minimal implementation effort for software to take advantage of a wide range of compute accelerators.
- Increases overall efficiency of the Compute and Data Continuum and reduces implementation overhead for software authors.

Success Indicators

- Number of standardized software repositories.
- Existence of commonly-agreed and applied standards enabling portability of software across multiple data centres.
- Number of codes published making use of hardware abstraction layers like Kokkos.
- Number of software packages made available through tools like CernVM-FS and EESSI.
- Number of data centres on which commonly-used HEP and RA processing software can be executed without modification.
- Number of software packages following FAIR4RS guidelines.
- Number of ESFRIs and other research infrastructures including software portability in their coding standards or developer guides.

European Impact Increased homogeneity and portability over different compute systems and architectures increases the overall efficiency of the scientific process, at every level from code implementation to workflow execution. The development of common standards here will spread to other domains, leading to increased reuse and efficiency across the European scientific ecosystem

Implementation Pathway and Stakeholder Roles

Implementation Phases

Short term

- Gap analysis of existing tools and systems supporting portability.
- Establish and document common software stacks for major research infrastructures.
- Increase adoption of well established technologies such as containerization for software distribution.

Medium term

- Based on the gap analysis, ensure that adequate tooling is available to support appropriate

- abstractions over hardware at all levels.
- Document commonly-accepted standards for portability over the SPECTRUM Compute and Data Continuum.
- Publish software stacks through common software repositories.
- Provide training to researchers and software engineers on adopting the new technologies.
- Establish a governance model in support of sustainable long-term support for portable software practices by research infrastructures and computing & data service providers.

Long term

- Require and expect that all appropriate software is published through the established standards to the relevant repositories.
- Ensure a process is in place for modernizing and updating standards as well as the supporting services and capabilities with time.
- Provide long-term, sustained support for development and maintenance of key enabling technologies.

Stakeholder Roles

Scientific Communities

- Responsibilities: Define scientific requirements and use cases. Implement software tools following best practices for portability and reuse. Distribute software through standardized repositories.
- Benefits: Reduced overhead in software implementation across multiple e-infrastructures. Seamless access to compute and accelerator capacity through the Compute and Data Continuum.

Research Infrastructures

- Responsibilities: Incentivize and support the adoption and development and adoption of portable software. Provide training. Invest in appropriate support infrastructures such as software repositories.
- Benefits: Able to make use of the federated compute infrastructure for efficient processing of data from the RI. Enhanced scientific impact.

Computing & Data Service Providers / e-Infrastructures

- Responsibilities: Facilitate the development of common standards and tools for software portability. Ensure their service offering complies with relevant standards and supports standard environments.
- Benefits: A common standard for both interfacing with community software and exchanging software with peer facilities reduces support overhead and increases efficiency, ultimately leading to reduced costs. Make the e-infrastructure more accessible for end users.

Policy Makers / Funding Bodies

- Responsibilities: Define and incentivize strategic goals that recognize the importance of software portability. Ensure the availability of structural funding for key enabling technologies.
- Benefits: Increased impact of resources spent on computing due to improved efficiency. More transparent scientific process due to benefits for reusability and reproducibility.

Resource Needs: Funding for the development of key standards and conventions as well as the necessary supporting and enabling technologies. Infrastructure for establishing software repositories and building and storing software and other artefacts. Development resources where the gap analysis shows existing tools to be inadequate.

Critical Dependencies

- [P04 Heterogeneous Computing Integration](#)
- [P09 Open Science and Reproducibility](#)
- [P10 Software Preservation and Sustainability](#)
- [P11 Community Collaboration and Co-Design](#)
- [P13 Skills, Careers and Training](#)

P09 Open Science and Reproducibility

Strategic Rationale and Evidence Base

Strategic Importance Open Science and reproducibility are fundamental pillars of modern scientific practice, ensuring that results can be independently verified, reused, and built upon across domains. As European scientific communities move toward exascale data volumes and increasingly complex AI-enhanced workflows, the ability to preserve, share, and reproduce results becomes both more challenging and more essential. In an environment where scientific discovery is data-driven, collaboration-intensive, and reliant on heterogeneous infrastructures, reproducibility emerges as a prerequisite for trust, transparency. Evidence from major SPECTRUM research communities shows that only infrastructures with FAIR-by-design principles, robust provenance tracking, and transparent workflow environments can support long-term scientific value and compliance with open science mandates. For SPECTRUM, Open Science and reproducibility are core enabling conditions for a truly interoperable, cross-domain compute and data continuum. They ensure that scientific knowledge generated in Europe remains accessible, verifiable, and re-usable across communities and generations of infrastructure and crucial way to ensure European sovereignty for research outputs including data, research software and workflows

Current State While many European e-Infrastructures have adopted elements of FAIR data management and workflow packaging, the landscape remains fragmented. Metadata standards vary significantly by community, workflow environments are often tied to specific compute centres, and reproducibility across heterogeneous HPC/HTC/Cloud platforms is not consistently guaranteed. Storage and software preservation practices differ strongly between national and thematic infrastructures and cross-domain interoperability remains limited.

Community Evidence Findings from the SPECTRUM Technical Blueprint, particularly Section 2.5 (software stacks, workflow systems, portability, and reproducibility gaps), reinforce that reproducibility challenges are widespread. Key gaps highlighted by communities include: inconsistent metadata and provenance capture across infrastructures, lack of standardised workflow systems supported by multiple HPC/HTC centres; divergent software stacks and insufficiently portable execution environments; limited long-term preservation of workflows, containers, and software artefacts; absence of domain-agnostic interfaces enabling reproducible execution on heterogeneous architectures.

Technical Foundation This priority is captured in the following SPECTRUM Compute and Data Continuum Capability map (D6.1 Section 4.3 Software Distribution and Execution; Section 4.2 Data Resources)

European Alignment Promotion of Open Science best practices stands at the core of EOSC mission. OpenAIRE provides the research graph connecting publications, datasets, and software with persistent identifiers, enabling discovery and citation tracking across European research outputs. EuroHPC JU is currently developing a federation platform aiming to bring together the diverse computational and data analysis capabilities, advanced software environments, and high-performance storage solutions available on the hosting sites. ESFRI roadmaps highlight long-term data preservation and reproducibility as essential infrastructure requirements for large-scale research facilities. EU Digital Decade prioritises trustworthy digital infrastructures, transparency, and secure data spaces, all of which depend on reproducible, open, verifiable scientific workflows. There are currently several European initiatives such as EVERSE [86], OSCARS [141] directly working towards consolidating Open Science best practices across European Research Infrastructures and Scientific Communities. A number of open data alliances have been introducing open community standards (e.g., the IVOA (international virtual observatory alliance [112]), the IHDEA (international heliophysics data environment alliance [110]), or the IPDA (international planetary data alliance) [113].

Target Outcomes and Impact Assessment

Strategic Outcomes Ensure scientific workflows and results remain reproducible over 5–10 year timescales through preservation of complete computational context. This includes standardised metadata, provenance

tracking, and FAIR-compliant practices for software and workflows alongside data. On a policy side pathways towards career-recognition based in following Open Science good practices

Success Indicators Qualitative measures showing progress toward outcomes

- Increased adoption of FAIR-aligned data services across communities.
- Increased transparency and verifiability of scientific publications.
- Growth in cross-border, cross-domain scientific collaborations enabled by reproducible environments.
- Reduced operational overhead for infrastructure providers.
- Stronger alignment with European Open Science policy frameworks.

European Impact A Europe-wide adoption of reproducibility practices enhances scientific credibility, supports evidence-based policy making, and accelerates joint scientific innovation. It positions Europe as a global leader in trustworthy, transparent, and FAIR science and contributes to strategic objectives under the EOSC, EuroHPC JU, ESFRI, and Digital Decade initiatives as explained in the European alignment.

Implementation Pathway and Stakeholder Roles

Implementation Phases

Short-term (1-3 years)

- Consolidation of guidelines and best practices for implementation of FAIR and reproducible workflows under Open Science
- Implementation specific tools to increase cross-domain collaboration such as the establishment and provision of execution environments streamlining the whole data cycle from findability, analysis, reproducibility and provenance by harmonising meta-data standards, workflow- systems and software stacks across European infrastructures.
- Training & upskilling all stakeholders on Open Science and FAIR paradigms
- Establishment of Open Science-based career recognition best practices and policies

Medium-term (3-5 years)

- Consolidation of aligned interoperability standards between major EU initiatives (EOSC, EuroHPC JU, ESFRI)
- Consolidation of FAIR-by-design services available across borders and across scientific disciplines
- Develop plans for both data & software & results preservation & reproducibility as an integral part of the preparation of scientific research activities & programmes and promoting the reuse of already existing community standards
- Promotion of Open Science part of educational curricula in university programmes

long-term actions (5+ years)

- Enforcement of FAIR standard policies for publicly funded research
- Provision of long-term funding for data and software preservation across scientific disciplines

Stakeholder Roles

Scientific Communities and Research Infrastructures

- **Responsibilities:** Implement FAIR and reproducible workflow guidelines following Open Science best practices, This includes contributing with domain-specific metadata, standards, and workflow definitions and sharing the best practices across scientific communities, the provision of domain-specific FAIR repositories with provenance tracking that ensure long-term curation of data, software artefacts and metadata.
- **Benefits:** Increased trust in results, easier collaboration, reduced technical debt, enhanced global visibility. Stronger integration between European infrastructures and enablers for long-term sustainability.

Computing & Data Service Providers / e-Infrastructures

- **Responsibilities:** Provide execution environments for software and containers that enable reproducibility and provenance capture. This includes the provision of common workflow engines and the integration of reproducibility checks into compute environments.
- **Benefits:** Reduced operational overhead, simplified user support, increased cross-site interoperability.

Policy Makers / Funding Bodies

- Responsibilities: Provide long-term funding for data and software preservation. Require and enforce the use of FAIR standards in publicly funded research. Support governance frameworks for Open Science and reproducibility standards. and align national and European priorities and ensure regulatory compliance.
- Benefits: Greater return on investment and increased impact of scientific outputs from publicly funded scientific projects is not hindered due to lack of FAIRness. Increase transparency in the scientific process due to the increased reusability and reproducibility.

Resource Needs

- Updated compute and data infrastructure able to store FAIR scientific data and provide the execution environments
- Manpower including infrastructure operators, software developers to develop and maintain the SW workflows, data stewards / specialists to support the data curation across the whole data-cycle,

Critical Dependencies

Cross-priority dependencies:

- P04- Heterogeneous computing integration,
- P05-Data Management and Transport,
- P08-Software Portability,
- P10-Software Preservation and Sustainability,
- P13 Skills, Careers and Training

P10 Software Preservation and Sustainability

Strategic Rationale and Evidence Base

Strategic Importance: Software is as essential as data in enabling scientific discovery of data-intensive and compute-intensive domains such as HEP and RA. Software environments are becoming increasingly complex, tightly coupled to specific hardware, and subject to rapid technological change. Without long-term preservation strategies, scientific results cannot be reproduced, validated, or reused. SPECTRUM's vision of a federated, cross-domain compute and data continuum depends on stable, portable, and durable software artefacts. Software preservation ensures continuity across successive generations of HPC/HTC/Cloud architectures and enables reproducible science at scale. It reduces technical debt, enhances operational efficiency, and supports scientific autonomy by maintaining Europe's ability to re-run, verify, and extend past analyses

Current State: While some communities (e.g., HEP) have mature practices for preserving code, configurations, and containerised environments, the European landscape remains fragmented. This maturity difference is also reflected when comparing WLCG infrastructure for HEP, which has been already operating for years and SRCnet which is being set up and not yet in operation at the time of writing this document. Many research infrastructures lack systematic approaches for archiving workflows, capturing software dependencies, or ensuring long-term accessibility of container images and modules. Preservation responsibilities are often decentralised, unfunded, or reliant on short-term projects. There is limited standardisation across domains, and current practices do not always support the transition to heterogeneous, accelerated computing environments or new data spaces.

Community Evidence: The SPECTRUM Technical Blueprint—particularly Section 2.5—highlights wide disparities in software stack management across e-Infrastructures; urgent community requirements for common, portable, persistent software environments; gaps in workflow preservation, container longevity, and cross-architecture compatibility; inconsistent support for portable build systems and container runtimes; lack of long-term archival infrastructure for software artefacts; and risks associated with replacing or decommissioning HPC systems without preservation strategies. This has been highlighted by ESFRI related Research Infrastructures and there is an expectation to get those needs supported by compute and data infrastructure providers. On the user side the current legacy know-how spanning from the HEP community in collaborating from many years is providing pathways and learning for the set up of the systems for RA.

Technical Foundation: This priority is captured in the following SPECTRUM Compute and Data Continuum building blocks (D6.1 Section 4): Software & Execution; Compute Continuum.

European Alignment: Software presentation is part of the policy objectives for promoting FAIRness of scientific outputs being pursued under EOSC. EuroHPC needs to ensure portability and continuity across sites and successive upgrades of the supercomputing facilities. ESFRI priorities also focus on long-term sustainability, reproducibility, and cross-RI interoperability. This priority is strongly aligned and supported by the work done under the EVERSE project which can bring solutions and recommendations towards faster achievement of priority outcomes. Advocacy groups such as DORA [56] or ADORE from Research Software Alliance [8] or ESCAPE ESFRI Cluster have also been actively promoting solutions on SW sustainability

Target Outcomes and Impact Assessment

Strategic Outcomes: A sustainable, Europe-wide ecosystem in which scientific software, workflows, execution environments, key enabling SW technologies and containerised stacks are archived, portable, and reusable across generations of HPC/HTC/Cloud systems.

Success Indicators: Qualitative measures showing progress toward outcomes

- Major HPC/HTC/Cloud centres adopt harmonised software preservation policies, including documentation as integral part for of preservation policies / SW management plan
- Community-standard container images and software stacks remain executable across hardware generations.
- Increased reuse of preserved software and workflows across projects and infrastructures.
- Reduced effort required to port legacy codes to new architectures.

European Impact: Software preservation enhances Europe's scientific autonomy, reduces risk during technology refresh cycles, and lowers long-term operational costs

Implementation Pathway and Stakeholder Roles

Implementation Phases :

Short-term (1-3 years)

- Development and consolidation of guidelines, best practices and policies for research software development that include software preservation to be included as integral part of SW management plans and related documentation for organizations and initiatives
- Promotion of open source/open standards policies to prevent sovereignty issues and avoid proprietary vendor lock-in
- Establishment and provision of execution environments offering harmonising software stacks across European infrastructures.
- Deploying knowledge hub - in which know-how of different clusters is being placed including harmonized training programs on research software best practices across communities

Medium-term (3-5 years)

- Consolidation of standardized community-based container images and software stacks that can remain executable across next generations of heterogeneous hardware systems.
- Develop plans for software sustainability as an integral part of the preparation for scientific research programmes (including simulations, data analysis) and of the preparation of major experiments in particular
- Training programs on research software best practices integrated into European university curricula

Long-term actions (5+ years)

- Ensure the provision of long-term funding is kept for software preservation across scientific disciplines
- Career paths for research software engineers clearly defined, recognised under standard tools like ORCID.

Stakeholder Roles : Stakeholder group responsibilities and benefits

Scientific Communities

Responsibilities: Define domain-specific preservation needs, share best practices, adopt common containers and workflow formats.

Benefits: Long-term stability, reduced porting effort, reproducible science, stronger global collaboration.

Research Infrastructures (ESFRI RIs)

Responsibilities: Provide domain repositories, guarantee metadata quality, integrate software and workflows into preservation systems.

Benefits: Improved sustainability, FAIR compliance, interoperability across Europe.

Computing & Data Service Providers / e-Infrastructures

Responsibilities: Maintain harmonised software stacks, build and preserve container environments, offer long-term archival services.

Benefits: Reduced user support burden, stable environments, easier migration to new HW.

Policy Makers & Funding Bodies

Responsibilities: Provide long-term funding, define sustainability mandates, ensure cross-programme alignment (EOSC, EuroHPC, National).

Benefits: Greater return on investment by fostering reusability of scientific software. Increasing European competitiveness,

Resource Needs: Investment dimensions, intensity levels, and funding alignment

- Heterogeneous compute infrastructure with fully updated software stacks able to provide proper execution environments to keep research software operational over generations.

- Manpower including infrastructure operators, software developers to operate and maintain software stack and to develop and maintain the research software

Critical Dependencies: Key prerequisites and coordination requirements

- Cross-priority dependencies:
 - P04- Heterogeneous computing integration,
 - P08-Software Portability,
 - P09-Open Science and Reproducibility,
 - P13 Skills, Careers and Training

Pillar 4: Human Capital & Responsibility

P11 Community Collaboration and Co-Design

Strategic Rationale and Evidence Base

Strategic Importance

Community management is strategically important because it enables HEP, RA and other data-intensive communities to align around shared priorities, move from fragmented initiatives to a coherent roadmap, and support innovation, interoperability and long-term sustainability in the European compute and data continuum. Building on the SPECTRUM CoPe and existing networking structures, it must go beyond coordination alone by also addressing human-capacity needs, including stable roles, recognised responsibilities, organisational incentives, and retention of key engineering and software expertise. It should also ensure continuity beyond individual projects and maintain structured links with international partners, as major infrastructures operate in a global ecosystem.

Current State

The current landscape offers strong potential for closer convergence across data-intensive domains, but this has not yet been fully translated into coordinated action. Beyond technical integration, the main challenges lie in aligning priorities, collaboration models and governance across diverse communities, while also making European opportunities easier to understand and engage. A more structured coordination mechanism is therefore needed to connect communities, map EU opportunities, and support coherent engagement with international partners while preserving domain-specific needs.

Community Evidence

Key gaps relate not only to technology, but also to coordination and governance. These include: limited long-term resource planning aligned with community needs; insufficiently structured feedback loops between user communities, research infrastructures, service providers and policy stakeholders (such as EOSC and EuroHPC); insufficient involvement of all relevant stakeholders in strategic discussions and decision-making; and the need to better connect shared priorities with domain-specific requirements. Consultation feedback also underlines the importance of maintaining explicit links with European and international initiatives so that collaboration leads to operational convergence rather than parallel or fragmented efforts.

Technical Foundation

The increasing data and computing demands of HEP and RA create a strong basis for deeper collaboration across scientific communities, research infrastructures and e-infrastructures. Although the two domains differ in maturity and implementation pathways, they increasingly require a common set of capabilities, including sustainability, security, federation, orchestration, distributed software, and support across the compute and data continuum. In this context, the SPECTRUM Technical Blueprint provides not only a technical reference point, but also a shared framework for dialogue, coordination and joint planning across communities working with heterogeneous infrastructures.

European Alignment

The scientific use cases of the SPECTRUM project (WP5) are rooted in major ESFRI projects in HEP and RA, notably SKAO and HL-LHC, as well as ERICs such as LOFAR and JIVE. In this context, community management has an important role in connecting these communities with the broader European landscape, including key initiatives such as EOSC and EuroHPC, as well as related programmes and networks (e.g. APPEC, ASTRONET, RadioNet, Opticon, Europlanet, Solarnet, JENA, ESCAPE). It is also important to address the still not fully articulated alignment between EOSC and EuroHPC, which remains a source of complexity for communities seeking coherent European support across data and compute services. Strengthening these links through structured dialogue, coordination and mutual visibility is essential to help communities better understand, access and shape the European opportunities relevant to their needs (cf. Strategic

Outcomes section). This also supports closer alignment between research infrastructures, e-infrastructures and scientific communities, while reinforcing Europe's capacity to act coherently in the international data-intensive science landscape.

Target Outcomes and Impact Assessment

Strategic Outcomes Beyond identifying requirements for a sustainable European compute and data continuum to support the scientific exploitation of major European RIs, the effort in community management seeks to build bridges with similar existing and forthcoming multi-disciplinary initiatives, fostering long-term alignment and collaboration. Effective community management depends not only on technical coordination, but on stable human capacity, clear representation, explicit scope, and durable cross-initiative interfaces, since organisational, political, and workforce factors are often as decisive as the underlying technology.

Success Indicators Successful community management can be identified through clear indicators such as: sustained participation of diverse stakeholders across the HEP and RA domains; the development of a collaboration path with other scientific communities; the adoption of a shared community charter ensuring long-term sustainability; effective feedback loops that influence technical and policy outcomes; ease of access; and the creation of structured links with related European and international initiatives. Additional indicators include the establishment of stable community-facing roles and retention pathways for key technical expertise, and the production of clear stakeholder mapping describing how communities are represented.

European Impact Addressing the numerical challenges of HEP and RA science use cases is a critical step towards maximising the impact of major ESFRI and ERIC projects in these domains. The SPECTRUM Community of Practice is a central instrument in this effort, providing a visible and structured framework for coordination across scientific communities, research infrastructures and e-infrastructures. Through this mechanism, the disruptive data and computing demands of infrastructures such as SKA and the HL-LHC can be articulated collectively, connected more effectively with the EOSC and EuroHPC ecosystems, and used to reinforce the bridges between them. In doing so, the SPECTRUM CoP helps create alignment, foster synergies and consolidate Europe's leadership in data-intensive science.

Implementation Pathway and Stakeholder Roles

Implementation Phases

- **Short-term (Years 1-3):** Strengthen and support the established CoP by developing further the charter developed during the SPECTRUM project, ensuring its continuity and long-term sustainability, and, where relevant, aligning or merging activities with existing networks and consortia. In the computing domain, the CoP should also develop shared metrics, tools and methodologies to assess usage, performance, efficiency and best practices, thereby enabling the definition of common indicators, the collection and analysis of operational data, and the exchange of knowledge across data-driven research communities. This shared knowledge base should be considered a key output in its own right.
- **Medium-term (Years 3-5):** Foster interaction, coordination, and potential convergence with similar initiatives operating in related research domains. Medium-term actions should also include formalisation of cross-initiative interfaces, including a maintained inventory of relevant European and international programmes, infrastructures and policy processes, together with named liaison functions and periodic review points.
- **Long-term (Years 5+):** Maintain continuous coordination and strategic alignment across the broader network of scientific communities.

Stakeholder Roles

- **Scientific Communities** - **Responsibilities:** dedicate adequate resources to identifying technical needs, collaborate with other scientific communities as a well structured CoP, work with e-infrastructures on technical developments, and designating clearly identified individuals to coordinate these activities - **Benefits:** establish a clear pathway to solutions for computing and

data requirements, while providing a well-defined interlocutor for technical, political, and financial discussions

- **Research Infrastructures** – Responsibilities: empower community-led initiatives by allocating dedicated expertise and fostering operational collaboration across RIs – Benefits: unlock the full potential of existing infrastructures, ensuring efficiency, complementarity, and long-term sustainability through coordinated bottom-up actions
- **Computing & Data Service Providers / e-Infrastructures** – Responsibilities: actively engage in evaluating operational and funding schemes in collaboration with RIs and the scientific community – Benefits: establish a clear pathway to fulfilling their mission of advancing open science, enabling high-performance computing and data-intensive research, and supporting scientific progress
- **Policy Makers / Funding Bodies** – Responsibilities: facilitate coordination between scientific communities, RIs, and Computing & Data Service Providers, supporting joint development activities through dedicated and long term funding schemes – Benefits: ensure the optimal use of investments, maximising scientific and technological return while minimising inefficiencies and financial waste.

Resource Needs: To fulfill the objectives listed above and make community management efficient, scientific communities require a combination of resources: dedicated human expertise to identify needs and liaise with partners; financial support through flexible and sustained fundings; access to infrastructures for co-design and validation. Resource needs also include recognised career paths and retention measures for research software engineers, infrastructure engineers, and community coordinators, as consultation feedback repeatedly identifies workforce instability and competition with industry as a strategic risk to sustainability. Together, these resources enable scientific communities to address computing and data requirements effectively, coordinate across domains, and act as credible interlocutors in technical, political, and financial discussions

Critical Dependencies Effective community management has critical dependencies on long-term human and technical resource commitments and interoperable tools. A continuous dialogue between RIs, e-infrastructures, and policy makers must be ensured, supported by formal agreements, cross-community working groups, and structured feedback loops. Together, these elements enable alignment with European initiatives like EOSC and EuroHPC, fostering synergies, avoiding duplication, and ensuring sustainable impact. A representative community structure would help establish a clearer and more transparent channel of communication with European institutions, addressing the lack of visibility and coherence that is often perceived today. Critical dependencies also include clarity of mandate and scope, so that participating communities understand whether the collaboration framework is intended for HEP and RA specifically or for a broader cross-domain landscape, and how this common framework accommodates legitimate domain-specific differences.

Cross-priority dependencies: PO1 Cross-Infrastructure Governance

P12 Environmental Sustainability

Strategic Rationale and Evidence Base

Strategic Importance: Energy use has recently become a limiting factor for HPC due to the difficulty of procuring MW, cost and cooling issues. It is, however, just the tip of the iceberg as, in the current context of global warming, collapse of biodiversity and gradually exceeding planetary limits, environmental impacts of IT are significant and need to be reduced at the European scale to fulfil the Paris agreement. Conversely, large science projects, in the HEP and RA domains for instance, call for increased use and growth of HPC infrastructures and will face the challenge of doing more with fewer impacts to be environmentally, politically and socially acceptable. This is not simply a matter of improving the efficiency of the processes, although this will help, but of gaining control over the different impacts of IT in order to allow it to fit in schemes to globally reduce mankind's impacts. A holistic approach is required as science moves into the era of big data and because of the exponential adoption of AI to ensure scientific computing can continue responsibly in spite of the growing environmental pressure.

Current State

The energy challenge and GHG emissions of datacentres have been gaining traction and moves towards the use of renewable energy and more efficient cooling have been taken. While this allows for reducing the environmental impacts of the use phase, it does not allow for addressing them globally. Life-cycle assessment methodologies (ISO 14040, ISO 14044, product environmental footprint) and databases, from the cradle to the grave, have been emerging to quantify the impacts and improve the design, fabrication and overall management of the IT resources, as well as avoid the transfer of impacts between phases, impact types or geographical regions. To date, however, tooling and policies to forecast and manage impacts are still in their infancy and coordination is lacking to gain control on the whole and first stall their growth, then reduce them overall. Impacts can be local (e.g. water use) or global (e.g. climate change) and are generated by diverse communities (end users, datacenter managers, hardware suppliers, waste managers etc.) at the national or international scales, which makes them complex to address consistently. In a similar way, preliminary mitigation of negative social impacts is under development with approaches based on the more recent ISO 14075 norm.

Community Evidence

The optimisation paradigm adopted by RA and HEP is centred on weak scaling: the ability to handle larger problems using a larger amount of resources. In this frame, section 2.5 gap #10 in the Technical Blueprint identifies the need to monitor workflows, performance and energy usage across sites to efficiently exploit a multi-infrastructure compute and storage continuum. This focus on efficiency, as it allows for pushing the limits of what is technically feasible, has and will lead to further increases of storage, network and computing resources in contradiction with the imperative to reduce the environmental impacts. This is something known as the rebound effect or Jevon's paradox.

The D6.1 Technical Blueprint identifies a number of gaps relevant to environmental sustainability. Gap #10 in section 2.5 resulting from WP5 analysis relates to monitoring and observability with the lack of consistent, cross-site monitoring of workflows, performance, and energy usage (D6.1). This forms part of the inefficient resource usage gap in section 3. The same section points out that execution is carried out irrespective of environmental impacts which implies that the latter are only determined a posteriori, ie. too late to be avoided or reduced. A third gap in this section corresponds to significant digital waste due to poor data and software reuse and of failing software. The fragility of operations is finally pointed out, which in the light of the increasing pressure on energy and cooling is susceptible of leading to increased waste or the inability to operate due to shortages of resources.

Technical Foundation

Sections 3.1 and 3.2 of the Technical Blueprint propose a path for reconciling weak-scaling optimisation efforts with environmental sustainability through the development of tools allowing to view them through the prism of environmental impacts. These provide a technical basis for sustainability policies to be enforced at the continuum and datacentre levels to address their impacts with a holistic approach. Central

to this approach is using environmental impacts as credits allocated to users at all levels. This will make them visible and promote responsible behaviours (as opposed to abstract quantities like core hours). It will also allow for framing their expenses within a global environmental impact envelope defined by higher-level policy-makers so that the environment is indeed protected. It is expected that datacenters will arbitrate between procurement, maintenance and usage impact expenses on this basis. In doing so, the latter will flow down to end users who will then need to optimise and prioritise their uses accordingly. While the tooling is end-user oriented, these policies are intended to be set by datacenters, governing bodies and higher-level policy-makers in order to define a large-scale comprehensive framework for the management and use of European IT infrastructures within determined environmental impact envelopes. Guidelines for defining these policies are provided below but precise definitions based on metrics and setting thresholds is beyond the scope of the project as they are the responsibility of datacenters, governing bodies and high-level policy makers, each at his own level.

In the light of the gaps summarised above, the D6.1 Technical Blueprint recommends the development of software-hardware optimisation to optimally use resources and distribute data and compute, tailoring execution to sustainability in terms of forecasting impacts and providing end users with means to map and prioritise their jobs so they can be scheduled in an impact-aware manner, FAIR data and software publication, in direct relation to Priority PO5 (Data Management Continuum), fine software quality assessment prior to the use of production-level resources to reduce the probability and consequences of failures and improvements in the IT's adaptation capabilities, notably through the resilience of software execution to varying or failing hardware resources.

This approach is complemented with policies addressing the other phases of the IT lifecycle.

European Alignment

The European Green Deal aims to make Europe the first climate-neutral continent. This is a significant ambition, a top-level objective addressing climate change globally across the union. Recommendations for technical developments in the Technical Blueprint and policies in this document aim to empower policy-makers in this sense and allow for this objective to flow down in quantitative terms to datacenters and end-users.

EuroHPC JU Strategic Commitment: Europe has demonstrated leadership in sustainable HPC through [82]:

- JUPITER JEDI #1 Green500 ranking
- MareNostrum 5 PUE < 1.08 industry-leading efficiency
- SEANERGY €16.8M EU investment in energy-efficient workload management

EOSC and Broader European Strategies: Environmental sustainability aligns with broader European strategies. The Digital Decade Policy Programme 2030 emphasises sustainable digitalisation [54], and EOSC principles include environmental responsibility for research infrastructures [66].

JENA Federated Computing Vision: JENA Computing White Paper positions environmental sustainability as one of the central challenges for federated computing, requiring transparency, lifecycle optimisation, and cultural shift towards sustainable computing practices [115].

Target Outcomes and Impact Assessment

Strategic Outcomes

1. **Forecast and Reporting of Environmental Impacts:** Operational capability to simulate, forecast and report environmental impacts of computing workloads, enabling informed resource allocation and scheduling decisions.
2. **Improved Software Quality:** Software development practices that prioritise energy efficiency and resource optimisation alongside performance and scientific accuracy [115], contributions to libraries to increase the lifetime of software, help it mature and standardise practices to facilitate co-design.
3. **Improved Efficiency of Execution:** Multi-objective optimisation of computing workloads balancing performance with environmental impact (carbon emissions, energy consumption, resource utilisation) (D6.1).

4. **Lifecycle Extension:** Procurement and operational policies that extend hardware lifetimes, accounting for fabrication impacts (upstream) and end-of-life disposal (downstream). In the European energy mix, fabrication phase outweighs usage phase, making replacement for efficiency gains rarely beneficial (D6.1).

Success Indicators

1. **Environmental Transparency:** Environmental monitoring (impact budgets, PUE/CUE etc.) operational at major European datacentres with public reporting
2. **User Engagement:** Sustainability metrics visible through dashboards to users during job submission
3. **Procurement Integration:** Sustainability metrics (embedded carbon, water usage, repairability, HEPscore/watt etc.) adopted as standard procurement criteria alongside performance metrics
4. **Lifecycle Extension:** Hardware operational lifetimes extended across SPECTRUM infrastructures
5. **Cultural Shift:** Training programmes on sustainable computing practices operational [115], awareness to environmental impacts and ecologically friendly ways of designing and using IT infrastructures through communities of practice
6. **Policy Alignment:** Integrated sustainability policies, accounting and management from the projects to the infrastructures with coordinated and overarching supervision by European HEP/RA institutions and EuroHPC.

European Impact

Europe already leads in green HPC, with JUPITER JEDI ranked #1 on the Green500 and MareNostrum 5 achieving PUE below 1.08 [82]. Reducing GHG emissions from research infrastructures contributes to European climate neutrality by 2050 and ensures continued public support for large-scale investments such as HL-LHC and SKA. Energy efficiency improvements also reduce operational costs, while lifecycle extension provides greater economic and environmental benefit than frequent hardware replacement (D6.1).

Implementation Pathway and Stakeholder Coordination

The Technical Blueprint recommends a number of tools to support the approach described here. Although part of these tools have been the focus of research activities, achieving the reliability and efficiency required for these tools to be used in production, that is for the effective management of IT resources, requires this effort to be intensified on the short term before the tools can be made mature through consolidation based on their use through an experimental stage described as part of the phases below.

Implementation Phases

Given the European scientific and computational ambitions, the Paris agreement and the corresponding European Union policies call for immediate action with the available means and for their extension. The paragraphs below describe the programmatic aspects while additional detail useful to make this actionable is provided in Annex C.3.

- **Short-term (Years 1-3):** This first phase has two facets. One consists in putting in practice available solutions to achieve a reduction of IT's environmental impacts and provide some relief to the environmental crises in the short term with limited disruption. The other is preparatory to further this approach. Its purpose is to plan, develop and experiment with the solutions required for the change of paradigm fully integrating the consideration of environmental impacts as a constraining aspect for the development, management and use of the data and compute continuum.
- **Medium-term (Years 3-5):** In the medium term, as a result of the experimentations carried out in the first phase, the different measures should be consolidated and put into wide-scale practice. It will also be time for experimenting with the more complex measures for which the first phase has laid a basis.
- **Long-term (Years 5+):** Given the urgency of the environmental crises and the expected growth of IT, all of the different proposed measures (or adaptations thereof) should be consolidated and put into practice in this time frame. Given the likely continued rapid evolution of IT, this will also require continuously monitoring their effectiveness and adapting them to the hardware and software systems, as well as to the environmental situation.

Stakeholder Roles

- **Scientific Communities** – Responsibilities: train users on software development for low environmental impacts and mindful use of resources (priorities, experimentation plan) – Benefits: reduction of digital waste, resilient and adaptive execution, improved software quality
- **Research Infrastructures** – Responsibilities: contribute to sustainability governance, contribute to data and software reuse, incorporate IT environmental impacts in the wider assessment of those of the research infrastructure – Benefits: improved coordination between user communities, detailed accounting of environmental impacts resulting from computing and storage
- **Computing & Data Service Providers / e-Infrastructures** – Responsibilities: contribute to sustainability governance, manage the overall environmental impacts of infrastructures, allocate end-user shares such impacts as credits, support users and policy makers in tracking and managing them – Benefits: sustainable operations, detailed accounting of environmental impacts
- **Policy Makers / Funding Bodies** – Responsibilities: define environmental impact budget at the scale of the compute and storage continuum, evaluate sustainability governance – Benefits: meet obligations resulting from the Paris agreement, sustainable operations, detailed accounting of environmental impacts

Resource Needs

Given the importance of embedded impacts resulting from the fabrication of IT equipment, the transition to a more sustainable management of IT infrastructures throughout their life-cycle relies primarily on the involvement of personnel, both in a dedicated way as part of the sustainability governance staff and the development of the underlying tooling and in terms of specialisation in eco-design of infrastructures and software at the datacenter and end-user levels.

The development of the tooling recommended as part of the Technical Blueprint, although it will build on existing solutions, will require a significant amount of effort to be tailored to environmental impacts. This tooling consists mostly in software.

Critical Dependencies

The management of environmental impacts involves actors at all levels of the IT hierarchy and is by nature transverse as it involves all phases of the life cycle, hardware and software. A holistic, systems engineering view is required to avoid excessive disruption as well as the transfer of impacts. For this reason, coordination is seen as a critical dependency for this effort to be successful.

P13 Skills, Careers and Training

Strategic Rationale and Evidence Base

Strategic Importance: HEP and RA research has become computing intensive, at least since three decades. Large infrastructures and large investments are needed in order to collect, store and analyze the datasets produced. This requires dedicated teams of experts with specific skills and capabilities to design and operate the required computing systems and software stacks; failing that, the expected impact is a reduced physics exploitation capability for the planned billion Euro investments. In many cases critical pieces of software depend on individuals or small groups of people often in non-permanent positions. Together with workforce ageing, migration to industry and lack of recognition for research software engineering work, this poses a risk to research.

Training, supporting and retaining computing experts is of critical importance for current and next generation initiatives, even more so in a scenario where the same capabilities are in great demand from industry. In several cases researchers are lacking training in research software engineering and will receive the necessary training only when joining a research collaboration. Community established training in this respect is even more vital.

Current State: Computing experts (either by formation or trained inside the initiatives) suffer from inadequate career paths. They are seldom offered permanent positions, and those that do rarely have a strong career path. Academic positions in physics are not open to them, and in the best case they are retained as “technical staff”, typically on “soft money” (renewed year on year without long term job security). Computing experts also face an attractiveness deficit relative to industry (which offers higher salaries and more stable positions) and structural under-recognition in scientific outputs (rarely first author on the resulting publications, often not credited as authors at all). The work of research software engineers or researchers who code also often lacks recognition of their contributions within the communities which hampers career progression in addition.

Community Evidence See D6.1 Section 2.5 (Consolidated Gaps and Requirements), Gap #13 (Workforce sustainability and development): the problem of “Workforce, funding, and sustainability” is listed as one of the most critical gaps; the highlighted solution consists of the “... development of a coordinated European funding/governance models, in alignment with national contributions, and in the investment in training and career pathways for research software engineers and data stewards”. The Physics Briefing Book “Input for the 2026 update of the European Strategy for Particle Physics” [142], Section 12.6, explains the problem from the angle of High Energy Physics. In the conclusions, the problem of career paths is detailed as “... in the long term, this is more worrying than technological trends, and finding solutions should be a field priority”.

Technical Foundation: The problem is not strictly technical, but rooted in the historical way in which computing related tasks have been executed in the target physics domains. Up to 3 decades ago, computing was a post-facto solution, designed after data taking and long after the initial design. This was possible due to its (much) reduced complexity with respect to other technical issues (e.g. building and operating a detector). Today, the size and complexity of scientific computing has increased to a level of criticality comparable with other technical aspects, with the overwhelming majority of experts acknowledging this. Nevertheless, research and academic institutions have failed to keep pace and continue to undervalue scientific computing experts, treating them as replaceable. The result is a fast turn-around of less experienced computing collaborators, who leave after contributing critical software to their communities. Worse still is when this happens at a later career stage, where computing contributions are much more significant. Entire systems need to be replaced, without the knowledge and expertise of the original expert. This issue directly affects the sustainability of large scientific initiatives, and their capability to produce scientific results. It also hampers European strategic efforts to find synergies across scientific computing domains, as the few experts that do exist are too busy mitigating these issues. Finally, while it's desirable to have some early career individuals leave academia for industry, for those who leave later in their careers the transition is typically much less efficient for industry and more stressful on the individual.

European Alignment The career problem for science is now recognized in European initiatives, including the EuroHPC JU and the EOSC, and is leading to promoting positions as technical specialists and data stewards. Still, at the academic level, the career paths offered to computing experts (either by formation or by specialization) are less attractive than, for example, in the construction of equipment and in data analysis.

Sovereignty of computing initiatives and skills is an important aspect which needs to be preserved, in a global scenario where major non-European industries drive the technology innovation. The EOSC EVERSE project, as one example, strives to establish best practices in research software engineering underpinned with services to provide a coherent and connected infrastructure for training materials and events as well as providing recognition for research software engineers and trainers.

Target Outcomes and Impact Assessment

Strategic Outcomes Inform and reiterate among research and academic institutions the need and importance of research personnel fluent in large scale computing; prepare and fund appropriate career paths for such experts and provide means for the recognition of research software engineering work. Feedback to universities and engineering schools on training needs in scientific collaborations

Success Indicators Lower total costs for the computing aspects of European scientific initiatives, driven largely by the ability to create longer-term, more sustainable computing plans, both during the design and operations phases. Improved knowledge sharing in computing across scientific domains thanks to the existence of a small cohort of career computing experts providing a stable and strategic vision for European scientific computing. Of course this cannot be realized by SPECTRUM alone, but must come from a domain-level realization of the issue and its possible solutions. With well-defined career paths, smoother transitions between academia and industry mean individuals destined for industry will start their careers earlier.

European Impact A solid cohort of domain scientists with extended computing skills, able to confront the data deluge expected from next generation scientific initiatives. In perspective, since not all of them will be able to remain in the public research domain, a cohort of scientists is moving to the private / productive sector, and participating in its innovation.

Implementation Pathway and Stakeholder Roles

Implementation Phases

- Short-term actions (deliverable under current initiatives: RI Train, HPC Train, EuroHPC Centres of Excellence) raise awareness of academic and research partners to the problem of providing good career paths to computing experts. Medium-term actions (deliverable through new project proposals or integration into university curricula) target structural change in education and recognition.
- Long term actions require a reshaping of the funding and personnel policies in academic and research institutions.

Stakeholder Roles

Scientific Communities

- **Responsibilities:** clarify their present / future needs in computing, and the needs in expert personpower; clarify the risks attached to their vacancy. Communicate the challenges to be solved, and the possibility to directly influence scientific results.
- **Benefits:** secure long term computing positions, for the benefit of decades-long initiatives. Increase the potentially interested community.

Research Infrastructures

- **Responsibilities:** ensure sufficient staffing and smooth operations (for research and beyond).
- **Benefits:** long-term planning of resources allows to participate in decades-long initiatives, and to have personnel with a deep knowledge in the systems and thus able to react / plan in a more effective way.

Computing & Data Service Providers / e-Infrastructures

- **Responsibilities:** Provide comprehensive training on the efficient use of heterogeneous resources, covering hardware acceleration (e.g. GPUs), software best practices (e.g. containers, portability), and sustainable operations; offer secondment opportunities for Research Software Engineers to embed within operational centres. Facilitate contacts between the scientific user communities and partners from industry, in order to show further career possibilities.
- **Benefits:** Improved infrastructure efficiency through optimised user workflows; reduced support

burden; and a user community capable of exploiting advanced platform features.

Policy Makers / Funding Bodies

- Responsibilities: fund computing research positions at academia and research institutes. Ensure career paths for researchers designing / operating computing systems.
- Benefits: a more predictable landscape in computing allows for a long-range planning of initiatives. A solid scientific computing infrastructure in specific scientific domains allows for easier support for the long tail of research. The realization of a cohort of experts which can partially fuel innovation in the productive system.

Resource Needs

- More funds dedicated to computing experts (by formation, training or inclination) in academic and research institutes. Long-term plans guaranteeing sustained support.

Critical Dependencies

All the other priorities, but especially highlighted for

- P07 AI
- P08 Software Portability
- P09 Open Science and Reproducibility
- P10 Software Preservation and Sustainability
- P11 Community Collaboration and Co-Design

Annex B: Priority Description Template

Each of the thirteen SRIDA priorities follows a standardised template that ensures consistent treatment and enables systematic comparison across priorities. The template organises information into three sections addressing rationale, outcomes, and implementation.

Strategic Rationale and Evidence Base establishes why the priority matters for European research infrastructure. This section identifies the strategic importance, describes the current state and gaps, references community evidence from the Technical Blueprint and other sources, identifies the technical foundation within the Compute and Data Continuum, and demonstrates alignment with broader European strategies.

Target Outcomes and Impact Assessment defines what success looks like through capability-focused outcomes rather than activity metrics. This section presents strategic outcomes as capability statements, identifies qualitative success indicators, and articulates the contribution to European research competitiveness.

Implementation Pathway and Stakeholder Roles shows how the priority will be achieved and who is responsible. This section defines implementation phases across short-term (1-2 years), medium-term (3-5 years), and long-term (5+ years) horizons, identifies stakeholder responsibilities, outlines resource needs, and notes critical dependencies.

Priority Description Template

Strategic Rationale and Evidence Base Establish WHY this priority matters for European research infrastructure

Strategic Importance One-paragraph explanation of why this priority is essential for SPECTRUM's vision and European research competitiveness

Current State Brief description of existing capabilities and identified gaps in European landscape

Community Evidence For technical priorities, reference to the SPECTRUM Technical Blueprint (in particular mention the gaps/requirements in Section 2.5); reference to other relevant documents also possible

Technical Foundation Reference to supporting SPECTRUM Compute and Data Continuum capabilities

European Alignment Connection to broader European strategies (EOSC, EuroHPC, ESFRI, Digital Decade)

Target Outcomes and Impact Assessment Define WHAT success looks like through capability-focused outcomes and realistic progress indicators

Strategic Outcomes Capability statements describing what will be achieved

Success Indicators Qualitative measures showing progress toward outcomes

European Impact Contribution to European research competitiveness and strategic objectives

Implementation Pathway and Stakeholder Roles Show HOW the priority will be achieved and WHO is responsible for implementation

Implementation Phases Short-term, medium-term, and long-term actions

Stakeholder Roles Stakeholder group responsibilities and benefits

Resource Needs Investment dimensions, intensity levels, and funding alignment

Critical Dependencies Key prerequisites and coordination requirements

Annex C: Priority Evidence and Validation

The thirteen strategic priorities emerged from systematic analysis of SPECTRUM deliverables and validation by the Community of Practice and External Advisory Board. This annex documents the evidence base, validation process, and traceability to the Technical Blueprint.

C.1 Desk Research and Evidence Sources

The strategic priorities of the SRIDA emerged from systematic analysis across SPECTRUM Work Packages 3, 5, and 6:

- **D3.1 (SPECTRUM CoP Community of Practice – Early Findings)** describes the main gaps and findings from the work performed by the SPECTRUM CoP in the first period of the SPECTRUM project.
- **D5.1 (Representative Use Cases)** examined fourteen science use cases (mainly from HEP and RA), identifying common requirements for cross-facility workflow orchestration, federated data access, heterogeneous computing support, and automated resource provisioning.
- **D5.2 (Access Policy Analysis)** investigated access mechanisms and authentication frameworks across European e-infrastructures, identifying sixteen critical gaps including lack of long-term resource planning, fragmented identity management, and barriers to unattended workflow execution.
- **D5.3 (Landscape Analysis)** assessed twenty European e-infrastructures, characterising current capabilities, service offerings, and operational practices, whilst identifying gaps that limit effective federation and cross-infrastructure collaboration.
- **D6.1 Section 2.5 (Consolidated Gaps and Requirements)** synthesised WP5 findings into sixteen consolidated gaps ensuring traceability from identified challenges to strategic interventions.

On top of the project deliverables, feedback was also collected in 2 interactive sessions at EGI 2025 conference and at ISC 2025 and by the interviews of community experts organised by WP4. The following section provides a summary of the main findings and the relation with SRIDA priorities. For more detailed analysis, readers should consult these deliverables directly. The following sections document how this evidence was validated and mapped to priorities.

C1.1. SPECTRUM CoP – Early Findings

This section describes the main gaps and findings from the work performed by the SPECTRUM CoP as reflected in D3.1 it includes (1) the analysis of results from a survey, (2) a review of the WP3 Knowledge Hub, populated with official documents from the communities, as provided by the CoP collaborators, and (3) trends and directions as extrapolated from all the collected material. Following are the main findings:

Table C.1.1.1: Gaps and findings from the Community of Practice.

Key Areas & Gaps	Finding Description	Evidence for SRIDA
Careers and skills. Lack of stable career paths for scientific software and data engineering roles	European research infrastructures depend heavily on specialised RSEs, data engineers, and domain-specific computing experts. However, most positions remain short-term or project-based, leading to high attrition to industry. This threatens Europe's ability to design, maintain, and scale the complex systems needed in the coming decade.	Strategic Driver: Workforce Priorities: P13: Skills, Careers and Training

Key Areas & Gaps	Finding Description	Evidence for SRIDA
Software evolution Insufficient evolution of scientific software for heterogeneous and emerging architectures	Many scientific codes are not optimised for multicore CPUs, vectorisation, GPUs, FPGAs, or heterogeneous workflows. Portability and performance engineering remain uneven across domains. The lack of modern tooling, parallelisation strategies, and portable abstractions limits the ability to exploit next-generation systems.	Strategic Driver: Preservation Priorities: P04 Heterogeneous Computing Integration, P08. Software Portability
AI Integration Limited adoption and integration of AI/ML across scientific workflows and tools	Although AI/ML is recognised as transformative, its uptake remains inconsistent across communities. There is a need to embed AI both as a user-facing capability (e.g., inference, analysis, surrogate modelling) and as an optimisation tool within infrastructures. Without structured integration, scientific domains risk falling behind global trends.	Strategic Driver: AI Priorities: P07 AI/ML
FAIR, open and cross-domain collaboration Fragmented data/software practices and insufficient adherence to FAIR principles	Data and software are often siloed, inconsistently documented, or incompatible across communities. Limited interoperability and collaboration hinder reuse, co-development, and joint innovation. Shared standards, repositories, and cross-domain toolchains are needed to reduce fragmentation and increase impact.	Strategic Drive: Preservation Priorities: P01 Cross-Infrastructure Governance, P08 Software Portability, P09 Open Science and Reproducibility
Cloud, HPC and virtualisation Uneven support for cloud-native workflows and portable deployment	Many e-infrastructures operate cloud and virtualisation capabilities, but integration across providers varies. Software distribution systems (CVMFS, EESSI) ensure consistent execution environments across sites. Portability across heterogeneous architectures and cloud-HPC integration remain inconsistent.	Strategic Drive: Scale Priorities: P04 Heterogeneous Computing Integration, P06 Orchestration & Workflows, P08 Software Portability
Network and node connectivity Insufficient external networking capabilities for data-intensive workloads	A key bottleneck is the availability of high-bandwidth, external network connectivity from compute nodes. Centres express willingness to improve connectivity, but real-world capability must be validated against realistic scientific use cases, particularly for distributed or streaming workloads.	Strategic Drive: Scale Priorities: P05 Data Management and Transport
Long-Term Resource Allocation Short-duration compute/storage allocations that hinder sustained workflows	Many scientific programmes require multi-year continuous access to compute and data resources, but current allocations are often short-term. Although some centres can offer longer periods, dedicated policy and funding engagement is required to institutionalise extended allocation models.	Strategic Drive: Scale Priorities: P02 Federated Resource Allocation

C.1.2 Representative use cases

D5.1 (Representative Use Cases) examined fourteen use cases across HEP, Radio Astronomy, and adjacent domains, identifying common requirements for cross-facility workflow orchestration, federated data access, heterogeneous computing support, and automated resource provisioning. Use cases included major HEP experiments (ATLAS, CMS, LHCb, ALICE), AI applications in accelerator science, SKA Observatory and other radio astronomy initiatives (LOFAR, fast radio bursts), and data-intensive applications from meteorology, molecular dynamics, and neuroscience.

Table C.1.2.1: Challenges and Recommendations from Use Case Analysis.

Challenges	Recommendation	Evidence for SRIDA
Unprecedented Data Growth	Develop exabyte-scale data management architectures, expand long-term storage capacity, and invest in scalable data lifecycle management strategies to prevent bottlenecks across scientific domains.	Strategic Drive: Scale Priorities: P05 Data Management and Transport
Heterogeneous Computing Evolution	Design flexible, modular compute architectures that integrate CPUs, GPUs, and emerging AI accelerators; support software portability and optimization to fully exploit heterogeneous systems.	Strategic Drive: Scale Priorities: P04 Heterogeneous Computing Integration
Complex Workflow Requirements	Deploy advanced workflow orchestration frameworks capable of managing multi-step pipelines across HPC, cloud, and distributed sites; promote workflow co-design with scientific communities.	Strategic Drive: Scale Priorities: P06 Orchestration & Workflows
Data Federation	Establish robust, standards-based data federation frameworks enabling cross-site discovery, access, replication, and system-level metadata interoperability.	Strategic Drive: Scale Priorities: P05 Data Management and Transport, P09 Open Science and Reproducibility
Resource Allocation Planning	Adopt multi-year compute and storage allocation models that align with long-term scientific experiments and enable predictable planning for large-scale research programmes.	Strategic Driver: Scale Priorities: P02 Federated Resource Allocation
Standardisation Needs	Implement and enforce common facility, interface, and operational standards to reduce integration complexity, enhance interoperability, and lower maintenance costs.	Strategic Driver: Scale Priorities: P01 Cross-Infrastructure Governance, P04 Heterogeneous Computing Integration, P11 Community Collaboration and Co-Design
Authentication Evolution	Evolve authentication and authorization infrastructures toward federated, automation-friendly models (e.g., token-based systems) that balance security with usability for	Strategic Driver: Trust Priorities: P03 Federated Identity and Trust

Challenges	Recommendation	Evidence for SRIDA
	large, distributed workflows.	
Need for Architectural Flexibility	Build infrastructure supporting diverse workloads simultaneously (simulation, AI/ML, streaming, and data-intensive pipelines) within a unified but flexible exascale-ready architecture.	Strategic Driver: Scale, AI Priorities: P04 Heterogeneous Computing Integration P07 AI/ML
Scaling Data Management to Exabyte Levels	Invest in next-generation data management solutions with high-throughput ingest, automated distribution, and transparent multi-site access, ensuring scientific teams can operate across distributed environments.	Strategic Drive: Scale Priorities: P05 Data Management and Transport
Importance of Cross-Community Collaboration	Promote shared tools, AI models, data standards, and co-design initiatives between HEP, RA, and other scientific communities to increase efficiency and accelerate innovation.	Strategic Drive: Workforce Priorities P11 Community Collaboration and Co-Design

C.1.3. Interoperable access policies

D5.2 (Access Policy Analysis) investigates access mechanisms and authentication frameworks across European e-infrastructures, identifying sixteen critical gaps including lack of long-term resource planning, fragmented identity management, and barriers to unattended workflow execution. This section summarises the main findings and recommendations emerging from this deliverable.

Table C.1.3.1: Findings and recommendations from interoperable access policy analysis.

Key Areas	Finding Description and Recommendation	Evidence for SRIDA
Obtaining Access		
Lack of long-term, assured resource access	Grant-based allocations last 1–2 years and may not be renewed; this prevents long-term planning for experiments requiring stable 3–5 year access to compute and storage. Recommendation: Adopt longer-term and more flexible allocation processes aligned with multi-year scientific planning.	Strategic Drive: Scale Priorities: P02 Federated Resource Allocation
Allocations limited to single sites within an e-Infrastructure	Many systems (e.g., EuroHPC JU) only honour allocations on specific hosting sites, preventing users from shifting workloads for burst capacity, resiliency, or efficiency. Recommendation: Enable e-Infrastructure-wide allocation models where quotas apply across all participating sites.	Strategic Drive: Scale Priorities: P01 Cross-Infrastructure Governance
Access Tracks & Modalities		

Key Areas	Finding Description and Recommendation	Evidence for SRIDA
Limited interactive access at significant scale	<p>Interactive workloads face delays and inefficiencies due to batch scheduling constraints; scalable interactive sessions are difficult to obtain.</p> <p>Recommendation: Extend orchestration/scheduling systems to support on-demand, interactive compute use with access to substantial node resources.</p>	Strategic Drive: Scale Priorities: P06 Orchestration and Workflows
Lack of high-level, user-friendly science platforms	<p>Most centres lack portals or workflow interfaces (like LOFAR or PanDA) that hide low-level scheduler, module, and environment complexity from domain scientists.</p> <p>Recommendation: Introduce high-level, general and domain-specific user interfaces enabling scientists to run workflows without deep system knowledge.</p>	Strategic Drive: Scale Priorities: P06 Orchestration and Workflows
Access Management & Security		
Fragmented identity management across infrastructures	<p>End-users working across multiple federated infrastructures must manage separate identities. Lack of common AAI systems limits interoperability and efficiency.</p> <p>Recommendation: Deploy interoperable, European-wide AAI services supporting federated identities and service accounts.</p>	Strategic Drive: Scale Priorities: P03 Federated Identity and Trust
Difficulty running unattended long-running workflows under Multi-Factor Authentication	<p>MFA-based authentication methods expire frequently (often daily), breaking multi-day workflows that cannot reauthenticate interactively.</p> <p>Recommendation: Ensure reliable, unattended execution by enabling token lifetimes, service credentials, or automation-safe authentication mechanisms.</p>	Strategic Drive: Scale Priorities: P03 Federated Identity and Trust
Rules, Assurances & Monitoring		
Inconsistent low-level software environments across sites	<p>Software stacks differ significantly between centres and architectures; portable builds require standardised interfaces and widely supported distribution methods (EasyBuild, Spack, containers).</p> <p>Recommendation: Provide efficient, standardised low-level software interfaces across sites, including shared build/distribution frameworks.</p>	Strategic Drive: Workforce Priorities: P08 Software Portability, P10 Software Preservation and Sustainability
Lack of structured feedback and co-design loops	<p>While some communities (e.g., WLCG, SRCNet) use collaborative planning models, most infrastructures lack systematic mechanisms for shared roadmapping and evaluation.</p> <p>Recommendation: Establish formal feedback, improvement, and planning loops between research communities and e-Infrastructures.</p>	Strategic Drive: Workforce Priorities: P11 Community Collaboration and Co-Design

C1.4 Landscape, Requirements and Gap Analysis

D5.3 (Landscape Analysis) assessed twenty European e-infrastructures, revealing persistent challenges in WAN performance, software portability across heterogeneous architectures, and environmental sustainability monitoring. This section summarizes the key findings and recommendations of the deliverable. The infrastructures studied span High Performance Computing (HPC), High Throughput Computing (HTC), dedicated Data-oriented infrastructures, and Cloud e-Infrastructures.

- **HPC-oriented infrastructures**

The EuroHPC Joint Undertaking (JU) [80] forms the backbone of this landscape, comprising operational petascale systems [85] such as Deucalion [51], Discoverer [55], Karolina [123], Leonardo [126], LUMI [132], MareNostrum 5 [134], Meluxina [136] and Vega [162], with exascale systems JUPITER (Germany) [121], Arrhenius (Sweden) [14], Daedalus (Greece) [50] and Alice Recoque (France) [12] planned for 2025-26. National centres such as **ICSC** (Italy) [109], **GCS** [97] / **JSC** [114] / **HLRS** [106] / **LRZ** [131] / **NHR Alliance** (Germany) [138], **RES** (Spain) [148], **CSCS** (Switzerland) [43], **EPCC** (UK) [67] and **GENCI** (France) [98] complement these resources and will support future HEP and RA experiments.

- **HTC-oriented infrastructures**

The WLCG [164] serves the global HEP community with about 1.4 million CPU cores and 1.5 Exabytes of storage (2025). WLCG Tier-1 centers, the EGI HTC-oriented infrastructure [62] provide access for diverse scientific communities, together with upcoming HTC platforms that include regional centres such as the SKA SRCNet [152].

- **Data-oriented infrastructures**

They include the WLCG's data lake (across the world and with CERN in Switzerland as the central hub), the SKA SRCNet [152], the EBRAINS neuroscience platform [58], the LOFAR Long-Term Archive (LTA) [129], German's ErUM4 Data Hub [69], the ICSC National Center in Italy [109], the PUNCH4NFDI [146] initiative and the Copernicus Earth-observation data space [41].

- **Cloud-oriented infrastructures**

In contrast to the three categories of infrastructures above, Cloud-oriented e-Infrastructures often provide a combination of compute (HTC and increasingly HPC) and data services. This category include for example: the EGI FedCloud [61], which offers a federated cloud interface that unifies resources across Europe; SURF's Grid [161] and Spider [162] platforms mix cloud and HTC/HPC services the European Open Science Cloud (EOSC) Federation [158], which federates national nodes and thematic pilots (the first 13 pilot nodes were announced in March 2025²) and the Simpl Data Federation Platform [156], built under a commercial contract for the European Commission, which aims at unifying access to diverse European data spaces.

Table C.1.4.1: Requirements from landscape analysis.

Key Areas	Requirements Description	Evidence for SRIDA
Storage / Data Capabilities & Services		
FAIR storage for observation/experiment data across domains	Create shared interfaces and core implementations that preserve rapidly growing scientific datasets (including metadata and provenance) according to FAIR principles. Must enable discovery, search by provenance/content, public access, and seamless interoperability with HPC/HTC systems, while supporting international users.	Strategic Drivers: Preservation Priorities: PO5 Data Management and Transport, PO9 Open Science & Reproducibility
Support for GDPR-protected and restricted data	Provide secure, standardised frameworks for storing and processing personal or commercially sensitive data. Must include encryption, secure key management, auditing, application vetting, and possibly sandboxing. Must avoid imposing restrictive	Strategic Drivers: Trust Priorities: PO3 Federated Identity and Trust

² <https://www.eosc-beyond.eu/pilots>

Key Areas	Requirements Description	Evidence for SRIDA
	security overheads on open scientific data and ensure separation of sensitive vs. open workflows.	
Data Transfer & Federation		
Automated and efficient data movement and staging	Automate intra-site and inter-site data transfers using high-performance, multi-stream mechanisms. Standard interfaces (e.g., FTS-like) should abstract protocol details. Third-party transfer capabilities and alignment with HPC centre security policies are required for large-scale workflows.	Strategic Drivers: Preservation. PO5 Data Management and Transport
Plan transfer capacity based on community needs	Provision high-throughput data movement and ephemeral storage according to evolving domain requirements. Data movement is increasingly unpredictable; centres must be prepared to adapt transfer policies and bandwidth to dynamic workflow demands over system lifetimes.	Strategic Drivers: Preservation. PO5 Data Management and Transport
Compute Capabilities & Services		
Portability of compute tasks across accelerated platforms	Ensure research codes can run across diverse accelerators (GPUs, FPGAs, vendor families). Avoid vendor lock-in (e.g. CUDA). Promote portable programming models (SYCL, Kokkos, Alpaka, oneAPI) and transparently optimised backends (PyTorch, TensorFlow). Recognise intrinsic limits in portability for architectures like FPGAs.	Strategic Drivers: Preservation. Priorities: PO4 Heterogeneous Computing Integration, PO8 Software Portability
Software Services, Interfaces & Stacks		
Common workflow systems supported by all e-Infrastructures	Identify and support a reduced set of workflow management systems usable across domains. Avoid duplication of domain-specific frameworks. Provide interfaces to HPC fabric schedulers (e.g., external API access to Slurm). Aim to converge toward a small, interoperable workflow ecosystem co-designed with communities.	Strategic Drivers: Scale. Priorities: PO6 Orchestration and Workflows
Common software stacks for compute applications	Provide unified, container-ready software stacks optimised for combinations of CPUs, interconnects, and accelerators. Standardise low-level interfaces across heterogeneous systems to simplify portability, virtualisation, and performance optimisation.	Strategic Drivers: Preservation. Priorities : PO8 Software Portability
Compute Federation Services		
Federated compute services with interoperable identity, catalogues, and access	Support federated identity (AAI), cross-site workflow execution, and simplified user interfaces across HPC sites. Current models (EuroHPC federation platform, CSCS/GENCI federations, WLCG, EGI) demonstrate partial capability but lack unified quotas, automatic workload routing, and full compute federation. Systems like Simpl provide initial catalogues, governance onboarding, secure communication, and federated data processing workflows.	Strategic Drivers: Trust Priorities: PO1 Cross-Infrastructure Governance, PO3 Federated Identity and Trust

C1.5 Feedback & Alignment with the Technical Blueprint (D6.1)

D6.1 Section 2.5 (Consolidated Gaps and Requirements) synthesises WP5 findings into sixteen consolidated gaps with explicit mapping to SRIDA priorities, ensuring traceability from identified challenges to strategic interventions. We report here as reference and link to SRIDA priorities.

Table C.1.5.1: Consolidated Gaps/Requirements and Recommendations.

Gap/Requirement	Recommendation	Evidence for SRIDA
Long-term resource provisioning Current grant-based allocations are short (1–2 years) and fragmented by site; HEP and RA need assured access on decadal timescales.	Establish long-term (3–5 year) flexible allocation frameworks, federated across infrastructures, with transparent accounting and monitoring. Policy activity will require dedicated funding.	Strategic Drivers: Scale Priorities P02 Federated Resource Allocation
Federated access and identity management End-users face fragmented identity systems across e-infrastructures, complicating multi-site workflows. Authentication often interrupts unattended, long-running jobs.	Deploy interoperable AAI frameworks across European e-infrastructures, supporting single sign-on, attribute-based authorization, and persistent tokens for long workflows.	Strategic Drivers: Trust Priorities P02 Federated Identity and Trust
Interoperable allocation and accounting Allocations and quotas are bound to specific HPC centers, limiting flexibility. Accounting practices differ across infrastructures.	Enable e-infrastructure-wide allocations and harmonized accounting, so projects can burst across systems and providers while ensuring fair contributions and usage transparency.	Strategic Drivers: Scale Priorities P01 Cross-Infrastructure Governance, P02 Federated Resource Allocation
High-level and portable software environments Software stacks differ widely between sites; portability across CPUs, GPUs, and accelerators is inconsistent. Scientists often face brittle interfaces.	Provide standardized, efficient software stacks (containers, CVMFS) and high-level domain-specific interfaces, enabling reproducibility and accelerator portability.	Strategic Drivers: Preservation Priorities: P04 Heterogeneous Computing Integration, P08 Software Portability
Workflow execution and orchestration Workflow engines are fragmented across communities (e.g. HTCondor in HEP vs. LEXIS in EuroHPC). Cross-facility workflows lack resilience and provenance capture.	Co-design and converge on a set of common APIs for workflow systems, with support for unattended execution, provenance tracking, and cross-infrastructure orchestration.	Strategic Drivers: Scale Priorities: P06 Orchestration and Workflows
Data federation and FAIR principles No unified approach to FAIR storage, metadata, and provenance across domains; irreplaceable datasets risk fragmentation.	Establish interoperable federated data management with FAIR-compliant storage, metadata catalogs, and efficient interfaces to HPC/HTC, ensuring long-term access and reproducibility.	Strategic Drivers: Scale, Preservation Priorities: P05 Data Management, P09 Open Science

Gap/Requirement	Recommendation	Evidence for SRIDA
<p>Data transfer and staging</p> <p>Data movement across sites often relies on ad-hoc scripts and site-specific tools; transfer capacity planning is uneven.</p>	<p>Automate large-scale, third-party enabled, efficient data movement with standardized protocols, and plan data transfer/storage capacity jointly with community needs.</p>	<p>Strategic Drivers: Scale, Preservation</p> <p>Priorities P05 Data Management and Transport</p>
<p>Real-time and low-latency processing</p> <p>A small number of RA use cases (e.g. transient detection) and some HEP triggers require near-real-time data analysis via pipelines.</p>	<p>Extend continuum capabilities with real-time, low-latency processing.</p>	<p>Strategic Drivers: Scale</p> <p>Priorities P04 Heterogeneous Computing Integration</p>
<p>Security, privacy, and sensitive data handling</p> <p>Personal or commercially sensitive data requires stronger safeguards; GDPR compliance adds complexity.</p>	<p>Provide flexible architectures: robust encryption, key management, auditing, anonymization/pseudo-anonymization, and possibly separate high-protection infrastructure paths.</p>	<p>Strategic Drivers: Trust</p> <p>Priorities P03 Federated Identity and Trust</p>
<p>Monitoring and observability</p> <p>Current infrastructures lack consistent, cross-site monitoring of workflows, performance, and energy usage.</p>	<p>Deploy monitoring and observability frameworks that capture standardized metrics describing infrastructure health, application performance, provenance, and sustainability.</p>	<p>Strategic Drivers: Scale</p> <p>Priorities P01 Cross-Infrastructure Governance</p>
<p>AI/ML integration</p> <p>AI/ML workloads are increasingly central and are being investigated in every step of the data-processing pipelines, from data-acquisition to simulation and analysis.</p>	<p>Enhance services in AI/ML platforms to support distributed training, inference deployment, and access to foundation models, integrated with HPC and domain workflows.</p>	<p>Strategic Drivers: AI/ML</p> <p>Priorities P07 AI/ML</p>
<p>Co-design and feedback loops</p> <p>Systematic co-design between science communities and infrastructure providers is patchy; planning often mismatches user needs.</p>	<p>Establish structured feedback loops (like WLCG/SRCNet) across domains, embedding co-design in infrastructure roadmaps, procurement, and evaluation.</p>	<p>Strategic Drivers: Workforce</p> <p>Priorities: P11 Community Collaboration and Co-Design</p>
<p>Workforce for scientific software development, funding, and sustainability</p> <p>Long-term sustainability depends on stable funding, governance, and skilled personnel. Current models are fragmented.</p>	<p>Develop coordinated European funding/governance models, align with national contributions, and invest in training and career pathways for research software engineers and data stewards.</p>	<p>Strategic Drivers: Workforce</p> <p>Priorities: P13 Skills, Careers and Training</p>
<p>Interactive access to significant-scale computing</p>	<p>Extend scheduling/orchestration to support interactive compute use</p>	<p>Strategic Drivers: Scale</p>

Gap/Requirement	Recommendation	Evidence for SRIDA
Interactive access to small/medium partitions is possible. Problems include delays in availability, and potential underuse of resources.	cases (incl. those that have significant resource needs) Compromise to be reached between quick availability, performance guarantees and efficient use of resources.	Priorities: P02 Federated Resource Allocation
High-level end-user interfaces Most e-Infrastructures do not offer high-level, abstracted interfaces. Domain scientists should not need to know system details and interact with brittle application interfaces.	Introduce high-level general and domain-specific user interfaces. (like workflow portals, Jupyter, or CARTA)	Strategic Drivers: Workforce. Priorities: P11 Community Collaboration and Co-Design
Unattended execution of long-running workflows and services Two-factor authentication can impact execution of long-running workflows and services. Access established after interactive or MFA authentication has a limited period of validity. Workflow steps need to interactively re-authenticate.	Ensure reliable and unattended execution of long-running workflows and user-operated services.	Strategic Drivers: Scale Priorities: P02 Federated Resource Allocation, P06 Orchestration & Workflows

C.2 Feedback & Alignment from External events

During June 2025 public feedback was sought in external events, during the EGI 2025 conference and a BoF session at ISC 2025. The main objective was to start assessing the landscape and main challenges and gaps to be addressed in the Technical Blueprint and SRIDA. For those a brief presentation of the project was delivered backed by an interactive questionnaire organized during the event using the mentimeter platform.

The questions asked:

- What are the main challenges and gaps in compute/data intensive science?
- How current technologies and policies should evolve to meet the identified challenges/gaps?
- What are the main priorities to tackle over the next 5–7 years?

C.2.1 EGI 2025 interactive session

The interactive session at the EGI conference³ was reasonably attended – ca. 25 participants (ca. 80% room capacity) from which over 19 responses were gathered on the mentimeter platform that were followed by a live Q&As and discussion. Next tables collect a summary of the main findings and associate them with the different SRIDA priorities.

What are the main technical and non-technical challenges in compute/data intensive science?

Table C.2.1.1: Technical and non-technical challenges from EGI2025 interactive session.

Technical Challenges	Non-Technical Challenges
Workflow Portability (P08) Heterogeneous Infrastructures (P04) Dataflow Management (P05) Software Development Practices (P10) Interface Standardization Distributed Computing Knowledge (P13) System Configuration Variability (P06) Efficient Resource Utilisation (P02)	Community Alignment (P11) Real Needs Assessment (P11) Global Coordination across communities (P11) Role separation (data generation & data usage) (P01) Collaboration dynamics (P11) Training and lack of mutual understanding across user communities and infrastructures (P01, P11) Awareness of available resources (P13) Access to funding

How current technologies and policies should evolve to meet the identified challenges/gaps?

Table C.2.1.2: Technology and policy evolution from EGI2025 interactive session.

Technology-related	Policy Related
Enforce open standards & interoperability (P09) Data-centric & FAIR infrastructures (P09) Lower technical barriers to entry (P11) Adopt proven industry solutions (P01) Update resource efficiency metrics (P12) Stay technology-agnostic & adaptable (P01)	Prioritize training & career paths (P13) Encourage collaboration & common solutions (P11) Foster EU/global coordination (e.g., through EOSC) (P01) Support inclusivity & accessibility (P02) Address cultural/organizational gaps (P11, P13) Increase transparency of metrics (P01)

What are the main priorities to tackle over the next 5–7 years?

Table C.2.1.3: Main priorities from EGI2025 interactive session.

Priority Area	Policy Related
Adopt a Data-Centric Approach	Adopt, scale, and support data-driven models
Expand and Diversify Communities of Practice	Engage diverse, real scientist input including underrepresented groups
Strengthen Collaboration and Interoperability	Build bridges, foster dual-skilled experts
Workflow Mapping & Reuse	Map workflows and promote reuse solutions to avoid redundancy
Respond to External Disruptions	Monitor and adapt to external political and technological factors (e.g. geopolitics and GenAI)
Breaking Silos	Connect science, IT, and policy communities

³ EGI 2025 event page: <https://indico.egi.eu/event/6638/contributions/20528/>

C2.2 Interactive session at International Supercomputing (ISC 2025 BoF)

The Birds-of-a-feather interactive session⁴ at the International Supercomputing Center was very well attended (over 70 people), from which 33 responses were collected in the Mentimeter platform. The following tables provide a summary of the content gathered.

What are the main challenges and gaps in compute/data intensive science?

Table C.2.2.1: Technical and non-technical challenges from ISC2025 interactive session.

Technical Challenges	Non-Technical Challenges
Infrastructure reliability (P03) Data sharing & collaboration (P05) Data transport (P05) Interoperability (systems, models) (P04) Data security (P03) Data management & lifecycle Documentation (P05) Access patterns for compute/data (P03) Keeping up with technology (P04)	Training & upskilling (P13) Ethics & compliance (P03) Reproducibility & sustainability (P10) Trusted research environments (P03) Fast, pragmatic action over high bureaucracy (P01) Community awareness (P11)

How current technologies and policies should evolve to meet the identified challenges/gaps?

Table C.2.2.2: Technology and policy evolution from ISC2025 interactive session.

Technology-related	Policy Related
Federated access for data/compute (P03) Cross-platform software deployment (P08, P10) Scalable Data Management (P05) Open Federation Standards (P09) Cloud Evolution (P04)	Common European policies & harmonization (P01) Integrated, cross-border access frameworks (P02) Agility over perfection: prioritize timely policy action (P01) Foster cross-domain agreements on shared challenges (e.g. metadata, provenance) (P01) Coordinated, sustainable funding strategies

What are the main priorities to tackle over the next 5–7 years?

Table C.2.2.3: Main priorities from ISC2025 interactive session.

Priority Area	Policy Related
Advance Scientific Data & Metadata Management	Adopt robust, FAIR-aligned, long-term practices
Sustainable Data Storage	Plan and fund scalable, long-term storage solutions
AI Competitiveness	Invest in infrastructure and skills to keep pace with AI
Plan Strategic Funding & Partnerships	Strategic planning with experiment and infrastructure upgrades
Usability for Non-Experts	Simplify access and use of compute/data infrastructure
Reuse & Integration	Build on existing solutions; break down silos

⁴ ISC BoF page:

<https://isc.app.swapcard.com/widget/event/isc-high-performance-2025/planning/UGxhbm5pbmdfMjU4NiExMg==>

What are the main challenges and gaps in compute/data intensive science?

Table C.2.3.1: Main challenges and gaps from CoP Interviews.

Challenge Area	Related Gaps
Compute Availability & Suitability (PO4, PO7)	<ul style="list-style-type: none"> – Shortage of suitable compute resources for scientific workloads; EU investments are skewed toward AI-optimised architectures not broadly usable by most sciences – Need for general-purpose HPC, not only GPU- or AI-centric systems; risk of misalignment with scientific needs driven by political agendas (e.g., AI Act) – Operational stability: scientific digital-twin infrastructures require near-continuous operation, but supercomputers have downtime and unpredictable maintenance cycles
Data Scale, Storage, and Mobility (PO5)	<ul style="list-style-type: none"> – Data volumes exploding (multi-exabyte scale in SKA; HL-LHC era; RA/astronomy; DestinE) causing challenges for storage, transfer, and long-term preservation – Long-term storage remains unfunded and under-prioritised despite being essential and costly – Data mobility bottlenecks, especially cross-site and cross-border transfers; lack of a mature federation framework
Fragmentation of Infrastructure & Lack of Interoperability (PO1, PO4, PO9)	<ul style="list-style-type: none"> – Extreme fragmentation: many domain-specific clusters, bespoke tools, disconnected policies; a single continuum is not realistic but greater coherence is needed – Metadata gaps: many communities not ready to standardise metadata; data without structure is useless
Skills, Workforce, and Retention (P13)	<ul style="list-style-type: none"> – Severe shortage of RSEs, HPC experts, and cross-skilled profiles (all interviews reaffirm this) – High attrition to industry due to salary gaps and unclear career paths – Training bottlenecks: porting to accelerators, heterogeneous architectures, data science methods; many scientists lack sustainable coding practices
Misaligned Funding & Sustainability (PO1, P10)	<ul style="list-style-type: none"> – Project-based allocations do not support environments requiring continuous access – Infrastructure sustainability—power cost, energy efficiency, operational cost escalation
Slow Adaptation to Emerging Technologies (PO4, PO8, P10)	<ul style="list-style-type: none"> – AI growing fast but unevenly adopted; many workflows not prepared for GPUs or ML – Quantum computing considered impactful only long-term – Technology lock-in (especially NVIDIA CUDA) creates major porting challenges

How current technologies and policies should evolve to meet the identified challenges/gaps?

Table C.2.3.2: Challenges and Recommendations from CoP Interviews.

Challenge Area	Recommendations
Infrastructure & Architecture Evolution (P01, P04)	<ul style="list-style-type: none"> - Build a federated but not monolithic European data-compute continuum, enabling 80% common services and 20% domain-specific features - Invest in data-centric design: data locality, distributed storage replication, high-speed cross-border networks - Diversify architecture choices beyond quantum and GPUs, e.g., neuromorphic, FPGAs - Increase interoperability with container-oriented, cloud-native, and Kubernetes-like environments
Policy & Governance Changes (P01)	<ul style="list-style-type: none"> - Shift from project-based to sustainable long-term operational funding models for research infrastructures - Align European and national strategies to avoid duplication and fragmentation - Ensure data sovereignty, avoid commercial lock-in, and keep scientific data in open-access, public infrastructures - Integrate energy-efficiency mandates into software and hardware funding criteria
Skills, Workforce & Training Evolution (P13)	<ul style="list-style-type: none"> - Professionalise career paths, create stable long-term contracts, and recognise computing roles - Scale training in heterogeneous computing, GPU programming, modern software engineering, metadata practices - Adopt cross-community hands-on formats (hackathons, scaling workshops, co-design labs) that have proven effective
Technology Adoption and Co-Design (P07, P11)	<ul style="list-style-type: none"> - Co-develop tools and workflows with communities, not impose top-down systems - Promote shared software stacks (e.g., Rucio, PanDA) across scientific domains - Boost applied AI/ML integration in simulation, modelling, anomaly detection, and digital twins

What are the main priorities to tackle over the next 5-7 years?

Table C.2.3.3: Main priorities from CoP Interviews.

Priority Area	Policy Related
Skills, Careers & Human Capital (P13)	<ul style="list-style-type: none"> • Establish permanent RSE/HPC/Data Science career tracks at national & EU levels • Provide sustainable funding for long-term staff, not only project contracts • Create European training programmes on accelerators, heterogeneous architectures, software sustainability • Encourage mobility networks, shared curricula, and incentives for technical career recognition
Data Infrastructure: Storage, Federation & Mobility (P05)	<ul style="list-style-type: none"> • Create a federated European data framework, not a monolithic system • Implement policies for long-term, funded scientific data archiving • Support cross-border high-speed networks and guaranteed data-mobility capabilities • Enforce metadata and interoperability standards across domains
Software Portability, Modernisation &	<ul style="list-style-type: none"> • Direct EU investment toward software & workflow optimisation, not only hardware procurement

Priority Area	Policy Related
Sustainability (P08, P10)	<ul style="list-style-type: none"> • Incentivise multi-architecture portability (avoid CUDA lock-in) and support for compiler/translation frameworks • Fund shared software stacks (e.g. Rucio, PanDA) as European public goods.
Heterogeneous Compute Architectures (P04)	<ul style="list-style-type: none"> • Balance EuroHPC policy to include general-purpose systems, not only AI/GPU-heavy ones • Create procurement guidelines to support diverse architectures (e.g., FPGA, neuromorphic) • Require that hardware investment includes associated data, software, and operational budgets • Ensure continuous-access policies for infrastructures needing predictable uptime
Large-Scale AI Capacity & Integration (P07)	<ul style="list-style-type: none"> • Fund EU-wide AI training facilities for scientific datasets • Accelerate adoption of AI-driven workflows in scientific codes • Establish ethical and open-access frameworks for scientific AI datasets and models • Promote cross-community AI co-design labs to produce reusable tools
Funding, Governance & Sustainability (P01)	<ul style="list-style-type: none"> • Shift from project-based HPC access to stable operational funding framework • Mandate energy-efficiency requirements at both software and facility levels. • Create mechanisms for EU-national co-funding to maintain shared infrastructures • Limit fragmentation by restricting funding for isolated, domain-specific clusters
Cross-Community Collaboration and Co-Design (P11)	<ul style="list-style-type: none"> • Promote RI-HPC co-design programmes (workshops, pilots) over top-down solutions • Support European initiatives that integrate multiple communities (e.g. ESCAPE-like models) • Build policies for standardised APIs, shared workflow tooling, and federated access models • Encourage joint applications and coordination between RIs to avoid parallel infrastructures

C.3 Community of Practice and Advisory Board Validation

SPECTRUM CoP, comprising approximately 65 participants across six Working Groups, provided continuous validation throughout 2024–2025. Each priority was assigned to one or more Working Groups who developed the priority specifications, ensuring domain expertise and operational knowledge informed priority content. Validation occurred through regular plenary meetings, Working Group sessions, and a community survey capturing quantitative data on resource utilisation patterns, technical requirements, and policy preferences.

The External Advisory Board provided strategic guidance refining the priority framework. Key recommendations from the second EAB meeting (post-M18) included: focusing on 2–4 main innovation avenues capable of delivering concrete services within 2–3 years (identifying P02, P05, P08, P09 as strategic focus priorities); creating priorities addressing reproducibility over 5–10 year timescales (P09) and code survivability across 20–30 year instrument lifetimes (P10); adding heterogeneous computing integration as a foundational priority (P04); and refining priority scope to distinguish ESFRI pan-European infrastructure needs from generic simulation projects. These recommendations shaped the final 13-priority framework.

Table C.3.1: Assignment of SRIDA Priorities to CoP Working Groups.

Working Group	SRIDA Priorities Developed
WG1: Data Management and Access	P05 Data Management and Transport, P07 AI/ML & Accelerated Computing (data aspects)
WG2: Workflow Management and Organisation	P02 Resource Allocation, P06 Orchestration & Workflows
WG3: Compute Environment	P04 Heterogeneous Computing Integration, P07 AI/ML & Accelerated Computing (compute aspects)
WG4: Software Tools	P08 Software Portability, P10 Software Preservation and Sustainability
WG5: Scientific Use Cases	P11 Community Collaboration and Co-Design, P13 Skills, Careers and Training
WG6: Facilities	P01 Governance & Funding Coordination, P03 Access, Security & Trust, P12 Environmental Sustainability

P09 (Open Science and Reproducibility) emerged from cross-cutting discussions spanning WG1, WG4, and WG5, receiving particular emphasis following External Advisory Board feedback.

P12 Environmental Sustainability had an extended contribution on expected policy implementation. Following is the summary:

We describe two sets of measures, ones which can to some extent already be applied and one which need to be developed, tested and then gradually applied. The latter is presented in order of complexity, the first being susceptible to be applied on the short term while the last is for a longer term.

The environmental impacts below are understood to result from multi-criteria attributional (when evaluating the impact of a system) and consequential (when evaluating changes) life-cycle analyses and apply at the local and global scales. As such, they are not limited to the current phase in the life of the resource but consider all of them.

Applicable measures :

- Train users to impact-oriented software development with an introduction to life-cycle analysis and environmental impacts, together with hardware and software levers to minimise them.

- Favour the emergence of eco-design communities for hardware and software solutions :
 - create a software eco-design CoP including users and datacenter support teams as a forum to discuss good practices and share experience, as well as foster collaborations on code to reduce their environmental impacts,
 - create a similar community to address hardware aspects at the datacenter level, including how to design, operate, maintain and procure infrastructures with minimal environmental impacts,
 - have datacentres provide resources to facilitate hardware-software co-design (software environments automating fine monitoring of execution in a transparent manner to point to portions of software requiring improvement, share hardware characteristics with users as well as platform-specific good practices),
 - propose access to experimental partitions composed of older nodes (whose environmental impacts have been amortized), of diverse types and with elevated privileges, to foster exploration of bottlenecks and software fragilities which are more obvious in such contexts for development and testing.
- Require users to provide an experimentation plan as part of their resource request with a level of detail matching its approximate expected environmental impacts.
- Power manage infrastructures :
 - concentrate use of resources on fewer nodes and switch off nodes made idle,
 - finely manage the powering of components in nodes (eg. unused CPU, accelerators, USB subsystem),
 - cool to higher temperatures,
 - power cap datacentres globally and compare strategies for power distribution to nodes (eg. selectively switch off nodes versus reducing consumption via dynamic voltage and frequency scaling).
- Require reporting IT resource usage in scientific publications (eg. as an extension to the required acknowledgements to resource providers).
- Institute sustainability governance at the datacentre and governing bodies levels through dedicated committees and people in charge of developing and evaluating sustainability in an independent way across all 3 tiers (datacenters, network, user terminals) to avoid impact transfer. This governance will also be in charge of continuously evaluating sustainability practices versus the evolving technological landscape and environmental crises.

Measures to be developed :

- Develop, test, evaluate and collect feedback from users for the tools proposed in the Technical Blueprint to consolidate them and de-risk transition to impact-driven resource management.
- Promote the evaluation of the environmental impacts of computing, transfers and storage in scientific publications.
- Develop a process for accessing resources to frame use and software quality in accordance to environmental impacts (tools are addressed in the Technical Blueprint):
 - regulate compute usage by enforcing use of selected partitions for selected categories of work (eg. development/prototyping, interactive, optimisation, test runs, production, long runs) and define criteria enabling categories of work related to functional, performance, scaling, sufficiency and reliability assessments,
 - manage hardware allocations to partitions depending on distribution of categories of work,
 - regulate storage usage based on performance of storage media versus software I/O performance and define criteria related to I/O usage and behaviour of software (file formats, access patterns, data volumes) versus storage technology,
 - develop active monitoring of execution through live tracking of progress and quality and adaptively control it based on its usefulness to adjust or interrupt it,
 - require crash testing of applications : handling of exceptions and errors, exceeding available resources (eg. memory usage, time or impact allocations) and hardware failures (eg. through checkpointing) to trigger adaptive software behaviour or graceful termination of execution,
- Enable impact-oriented datacentre-level management of resources :
 - given the importance of environmental impacts during fabrication, use hardware for extended durations mainly by using failure-prone hardware for functions which are not mission or impact critical and increasing the resilience of execution,

- track datacenters' environmental impacts,
- develop comparative life-cycle analyses for existing resources to estimate usage, maintenance, disposal, replacement and extension impacts in a prospective manner,
- include environmental impacts as a central component of decision-making,
- include sustainability requirements in tenders as environmental impacts should condition feasibility as cost and performance do.
- Extend the impact-oriented management of the previous two items from the datacenter level to the compute and storage continuum.
- Current evaluation of the use of IT resources focuses on resource usage both at the datacenter (pressure on resources, occupancy of partitions) and end-user (core hours) levels. This encourages growth while the relevance of such use is not precisely evaluated. Develop and evaluate metrics rewarding sufficiency and favouring minimal use for a given purpose.
- Transform resource requests to allow allocating environmental impacts :
 - require a description of the work's relative contribution to the wider scope of the scientific project,
 - require an experimentation plan with a level of detail matching the requested environmental impacts,
 - guide users in identifying resources meeting their technical needs at the datacenter and continuum levels (compute, storage) and in forecasting the environmental impacts related to different usage scenarios (in relation to the experimentation plan),
 - allow users to request access to technical resources and obtain an overall environmental impact credit for the proposed work (in replacement for core hours on predefined partitions),
 - support staged accesses (to different partitions depending on the categories of work) and rewarding sufficiency,
 - ensure user-allocated environmental impacts fit in the datacenters or continuum's overall impact credits,
- monitor environmental impacts at the user, datacenters and continuum levels when operating under this paradigm to support them in managing their environmental credits and build a detailed accounting of the corresponding expense.

C.4 Priority to Technical Blueprint Mapping

The table below provides traceability between SRIDA priorities and D6.1 Technical Blueprint components, showing which consolidated gaps each priority addresses, which capabilities each priority implements, and which technical activities provide implementation guidance.

Table C.4.1: Mapping of SRIDA Priorities to D6.1 Technical Blueprint

Priority	D6.1 Gaps Addressed	D6.1 Capabilities	D6.1 Activities
P01 Governance & Funding Coordination	#2 Coordinated governance, #16 Sustainability models	4.6 Resource Federation, A.8 Security & Trust	5.11 Governance & Sustainability
P02 Resource Allocation	#1 Long-term provisioning, #3 Allocation & accounting	4.6 Resource Federation, 4.7 Monitoring & Accounting	5.8 Resource Provisioning & Accounting
P03 Access, Security & Trust	#5 Identity management, #6 Unattended workflows, #7 Cross-border data protection	4.8 Security & Trust, 4.6 Resource Federation	5.10 Security, Trust & Data Protection

Priority	D6.1 Gaps Addressed	D6.1 Capabilities	D6.1 Activities
P04 Heterogeneous Computing Integration	#13 Heterogeneous computing	4.1 Compute Resources, 4.4 Orchestration & Workflows	5.1 Heterogeneous Computing Integration
P05 Data Management and Transport	#4 WAN performance, #8 Intelligent data placement, #10 Data mover capabilities	4.2 Data Resources, 4.7 Monitoring & Accounting	5.3 Data Federation & Transport, 5.4 Intelligent Data Management
P06 Orchestration & Workflows	#9 Cross-facility orchestration, #11 Resource-aware scheduling	4.4 Orchestration & Workflows, A.1 Compute Resources	5.5 Workflow Orchestration & Automation
P07 AI/ML & Accelerated Computing	#12 AI/ML with HPC, #13 Heterogeneous computing	5.5 AI/ML, A.1 Compute Resources	5.6 AI/ML Infrastructure Integration
P08 Software Portability	#14 Software portability, #15 Domain repositories	4.3 Software Distribution & Execution	5.2 Software Portability & Distribution
P09 Open Science and Reproducibility	Cross-cutting: #8, #14, #15	4.3 Software Distribution & Execution, A.2 Data Resources	3.1 Data & Software Preservation, 5.9 Open Science Practices
P10 Software Preservation and Sustainability	#15 Domain repositories	4.3 Software Distribution & Execution	3.1 Data & Software Preservation, 5.2 Software Portability
P11 Community Collaboration and Co-Design	#2 Coordinated governance	4.6 Resource Federation	5.2 Co-Design Processes
P12 Environmental Sustainability	Cross-cutting: sustainability	4.7 Monitoring & Accounting	3.2 Environmental Sustainability, 5.7 Energy-Aware Operations
P13 Skills, Careers and Training	#13 Hybrid computing expertise	4.6 Resource Federation	5.2 Co-Design Processes, 5.13 Workforce Development

This mapping serves multiple audiences: technical implementers read D6.1 for architecture then consult SRIDA priorities for stakeholder coordination; policymakers read SRIDA for strategic rationale then reference D6.1 for technical validation; research communities use both documents together to understand what infrastructure can provide and how to influence its development.