SPECTRUM

A brief overview of the SPECTRUM project

Eric Wulff
Staff Computing Engineer at CERN

21 January 2026
ODISSEE All-Hands Meeting

Funded by
the European Union

SPECTRUM is funded by the European Union - Grant Agreement Number 101131550

Project overview

The Exascale Challenge: A New Era for Science

High-Energy Physics (HEP) and Radio Astronomy (RA) are entering the Exascale era, driven by next-generation instruments

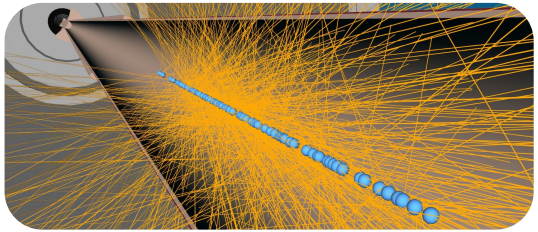

High-Luminosity Large Hadron Collider (HL-LHC)

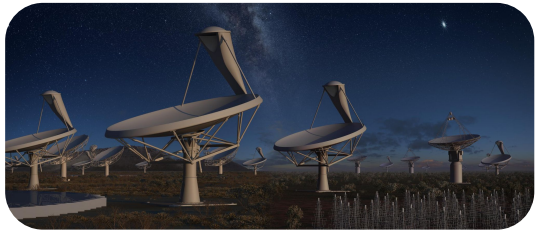

Square Kilometre Array (SKA)

LOw Frequency ARray (LOFAR) 2.0

Volume
Exabytes of data requiring storage and analysis

Velocity
Data generated at terabytes/petabytes per second, demanding real-time processing

Variety
Complex, heterogeneous data from millions of sensors or thousands of antennas

SPECTRUM

Elevating Data-Intensive Science in Europe

CHALLENGE

The **computing** for the **next generation of scientific instruments** in High-Energy Physics and Radio Astronomy is an **unsolved problem**

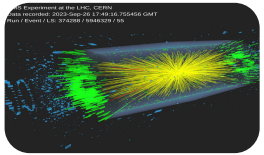

GOAL

Deliver a **Strategic Research, Innovation and Deployment Agenda (SRIDA)** and a **Technical Blueprint** for a **European compute and data continuum**.

IMPACT

- **Common vision** on future of infrastructures for data-intensive research
- **Aligning efforts** and reduce fragmentation among initiatives
- **Optimising use** of common digital infrastructures for research

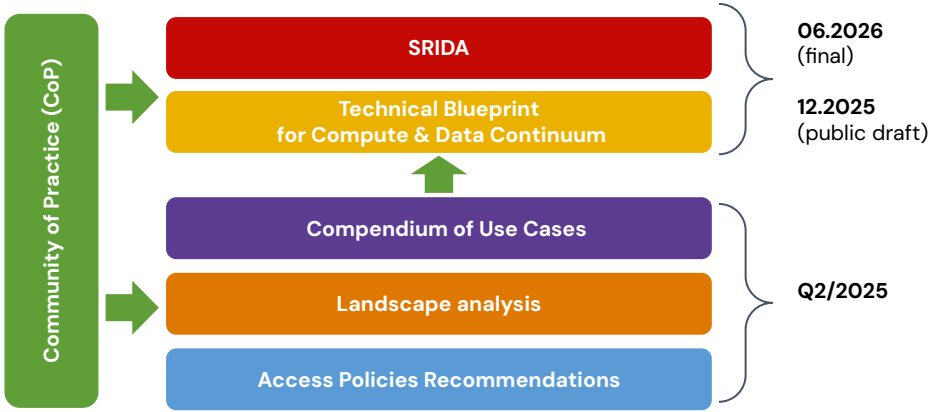

Research Communities

e-Infrastructures

The SPECTRUM Project is funded by the European Union Grant Agreement Number 101131550 - **More Info:** spectrumproject.eu

SPECTRUM Strategic Foresight Framework

Current Results

https://www.spectrumproject.eu/publications

Community of Practice (CoP)

- Cross-disciplines WGs with experts from HEP, RA and digital infras (link)
- Knowledge Base (link)
- Collaboration platform

Current Working Groups

- WG1: Data Management and Access
- WG2: Workflow management and organization
- WG3: Compute Environment
- WG4: SW tools
- WG5: Scientific Use cases
- WG6: Facilities

Compendium of Science Use Cases (14)

- From science case to technical challenges, requirements, gaps
- Both technical and policy aspects

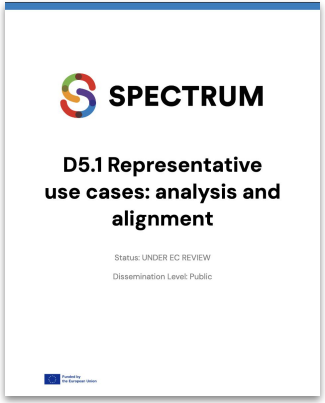

- High Energy Physics (6)
- Radio Astronomy (5)
- Chemistry (1)
- Meteorological Science (1)
- Neuroscience (1)

Current Results

https://www.spectrumproject.eu/publications

Access Policies Recommendations

- Existing access policies across the continuum
- Development of recommendation for Interoperable access

Landscape Analysis

- Existing approaches, services, technical solutions and policies for the federation of data and compute infrastructures

Formulating the Final Results

Technology Track

Capability Mapping → Technology Assessment → Capability Alignment

Technical Blueprint

Strategic Track

Priority Areas → Investment Areas → Policy alignment

SRIDA
(Strategic Research, Innovation and Deployment Agenda)

A Technical Blueprint for a European Compute and Data Continuum

Goals and context within SPECTRUM

Goals of the Technical Blueprint:

- Define **key capabilities** for an interoperable compute and data ecosystem
- Integrate **scientific, technical, and policy requirements** from the High Energy Physics (HEP) and Radio Astronomy (RA) Communities
- Provide **clear and actionable recommendations** for future e-infrastructure design and research funding programmes

Builds on previous work within SPECTRUM:

- Community of Practice (CoP) insights and survey findings
- Use cases, access policies, infrastructure landscape, and gap analysis
- Strategic Research, Innovation & Deployment Agenda (SRIDA)

Key priorities: interfaces, co-design, portability, data federation, workflows, AAI/security, AI/ML, and long-term provisioning.

Outcome: a set of actionable technical and policy recommendations for Europe.

Leadership and Timeline

Editorial team

Eric Wulff
(CERN)

Maria Girone
(CERN)

Jeff Wagg
(CNRS)

Publication timeline

Draft
Technical Blueprint for Compute & Data Continuum

12.2025
(public draft)

Final
Technical Blueprint for Compute & Data Continuum

Q2 2026
(final)

Technical Activities and Recommended Actions (1/3)

ACTIVITY	DESCRIPTION	ACTIONS
01 **Standardization of interfaces**	Fragmented APIs and job-submission interfaces hinder site integration and multi-site workflows	Develop and adopt common interfaces, metadata standards, and portability layers aligned with the EuroHPC Federation Platform, enabling seamless execution across computing sites.
02 **Infrastructure Co-design**	Early HEP/RA involvement in infrastructure design is essential	Engage EuroHPC/national programs for “strategic access.” Represent HEP/RA requirements in early system design (e.g. express need for double precision compute).
03 **Software Portability & Maintenance**	Codes tuned for CPUs are not optimized for GPUs, FPGAs, or emerging hardware	Review codebases for accelerator readiness. Fund pilot porting projects using portable frameworks.

Technical Activities and Recommended Actions (2/3)

ACTIVITY	DESCRIPTION	ACTIONS
04 **Data Management & Network Performance**	Lack of long-term storage and limited bandwidth to compute sites constrain scalable data workflows; standard, high-throughput data transfer solutions required	Run joint SRCNet-WLCG–HPC data challenges.
05 **Workflow Adaptation & Optimization**	Lack of a standardized, production-ready integration layer that allows existing workflow systems to interface seamlessly with site schedulers and policies	Connect existing workflow systems to HPC resource managers for native submission and monitoring, including end-to-end orchestration, and provenance capture.
06 **Security & AAI**	Multiple incompatible authentication and authorization systems hinders cross-continuum access	Establish a federated AAI framework allowing users to authenticate once and access multiple sites securely, with consistent authorization policies. Harmonize site authorization to accept federated identities and allow for service accounts.

Technical Activities and Recommended Actions (3/3)

ACTIVITY	DESCRIPTION	ACTIONS
07 **AI/ML Integration & Computational trends**	Rapid growth in AI workloads and hardware requires strategic software refactoring and skills development. Low-precision AI-optimized hardware conflicts with double-precision requirements for (parts of) of HEP/RA	Integrate distributed AI/ML services (training, inference, model repositories) directly into the continuum to accelerate data-driven discovery. Engage HEP and RA communities in HPC co-design processes to influence system architecture, ensuring sustained double-precision performance alongside AI-optimized capabilities. Identify and fund ML-suitable tasks. Train domain scientists in ML.
08 **Longer-term Resource Provisioning & Operational Sustainability**	Short-term resource allocations hinders long-term viability	Adopt multi-year resource-allocation frameworks. Investigate renewable allocation models aligned with experiment lifecycles.

Strategic Research, Innovation and Deployment Agenda (SRIDA)

Leadership and Timeline

Editorial team

Sergio Andreozzi
(EGI)

Xavier Salazar
(EGI)

Publication timeline

Draft
Strategic Research Innovation & Deployment Agenda (SRIDA)

12.2025
(public draft)

Final
Strategic Research Innovation & Deployment Agenda (SRIDA)

Q2 2026
(final)

Strategic Pillars of the European Compute & Data Continuum for Research

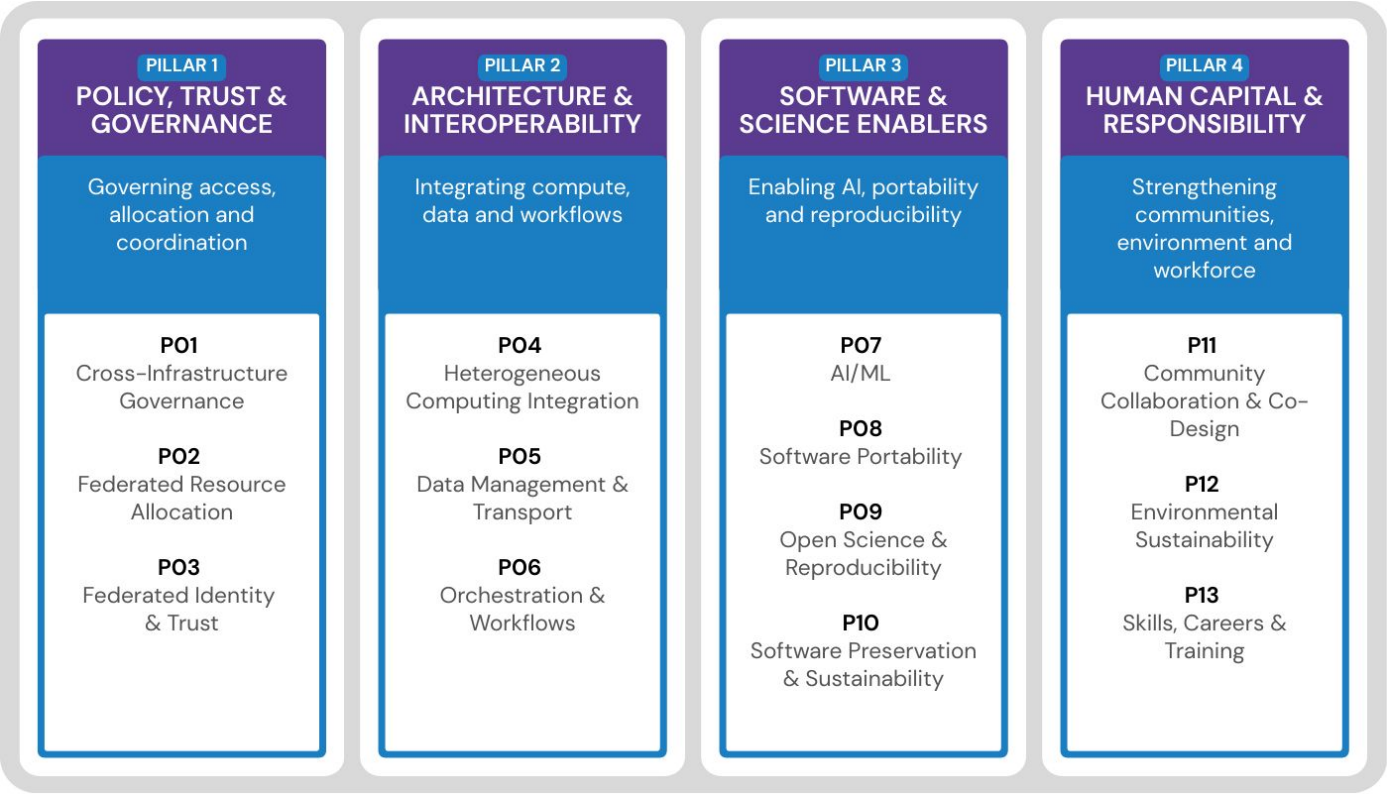

Conclusion and Next Steps

Next steps

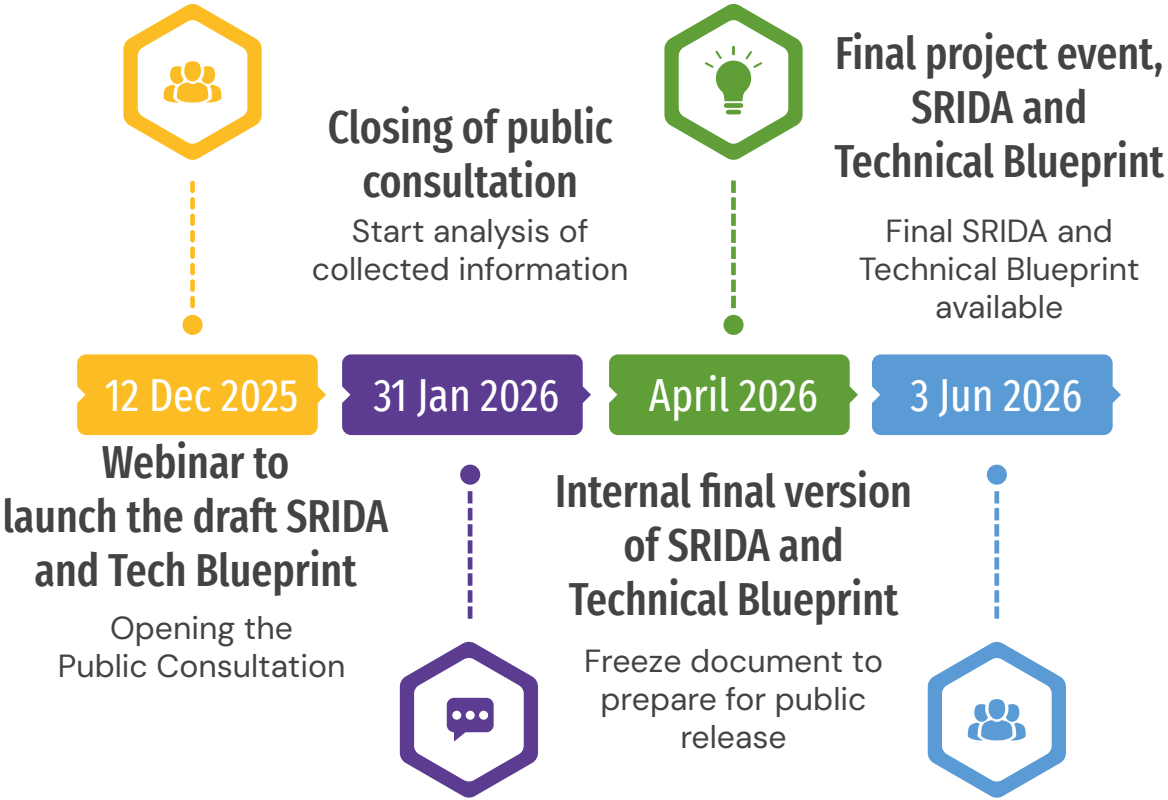
12 Dec 2025
Webinar to launch the draft SRIDA and Tech Blueprint
Opening the Public Consultation
31 Jan 2026
Closing of public consultation
Start analysis of collected information
April 2026
Internal final version of SRIDA and Technical Blueprint
Freeze document to prepare for public release
3 Jun 2026
Final project event, SRIDA and Technical Blueprint
Final SRIDA and Technical Blueprint available

Help shape the SRIDA and Technical Blueprint

Join the consultation process

By following this link, you can access a survey that enables you to share your feedback with us.

The survey includes direct access to the draft versions of both the SRIDA and the Technical Blueprint, and its length adapts to the time you have at your disposal.

Your input is essential to help shape the final documents.

The consultation will close on 31 January 2026.

spectrumproject.eu/spectrum-consultation

Thank you!

Backup

Technical Blueprint ToC

1. Introduction
2. Preliminary work
    a. The SPECTRUM CoP
    b. Use cases
    c. Landscape analysis
    d. Access policy recommendations
3. Sustainability
4. Technical Activities and Recommended Actions
5. Conclusion
6. Annex: Capability map for the compute and data continuum

Key capabilities

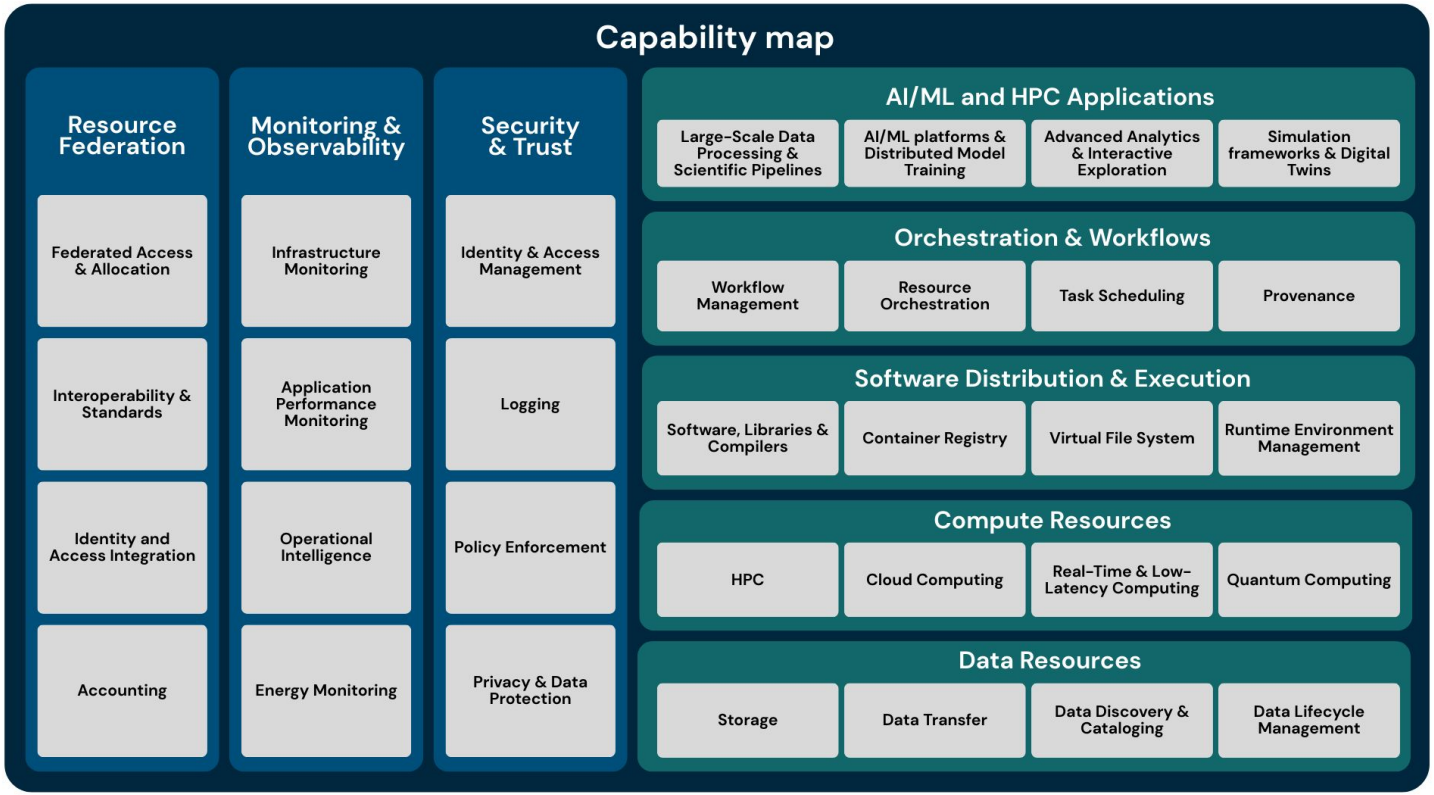

SRIDA ToC

1. Introduction
   - Purpose and Scope
   - Vision and Strategic Goals
   - Target Audience
2. Integrated Strategic Foresight Framework and Methodology
3. Strategic Context and Trends
   - European Context
   - The Evolving Landscape
4. Strategic Priorities
   - 4 pillars, 13 priorities
5. Investment Areas
6. Multi-Annual Roadmap
7. Conclusion
8. References

Annexes

- Annex A: Priorities
- Annex B: Priority Description Template
- Annex C: Priority Evidence and Validation

Pillar 1: Policy, Trust & Governance

PRIORITY	ADDRESSES	DELIVERS
PO1 **Cross-Infrastructure Governance**	Governance fragmentation: thematic RIs and pan-European infrastructures lack systematic coordination	Structured coordination mechanisms connecting scientific communities with infrastructure providers (e.g. coordination forum)
PO2 **Federated Resource Allocation**	Short-term allocation cycles incompatible with multi-decade research infrastructure programmes	Multi-year strategic allocations aligned with research infrastructure timelines
PO3 **Federated Identity & Trust**	Multiple authentication barriers blocking seamless access and automated workflows	Seamless federated authentication; Automated long-running workflows

Pillar 2: Architecture & Interoperability

PRIORITY	ADDRESSES	DELIVERS
P04 **Heterogeneous Computing Integration**	Growing heterogeneity of CPU, GPU and accelerator resources across federated sites	Unified interfaces enabling seamless workflow execution across HPC, HTC and Cloud
P05 **Data Management & Transport**	Exabyte-scale data volumes requiring coordinated federation across scientific domains	Federated FAIR data management with automated large-scale data movement across HTC, HPC, and Cloud
P06 **Orchestration & Workflows**	Multiple workflow systems with limited cross-facility interoperability and resilience	Interoperable workflow orchestration with provenance tracking across infrastructures

Pillar 3: Software & Science Enablers

PRIORITY	ADDRESSES	DELIVERS
P07 AI/ML	Transformative AI potential not yet systematically integrated into production workflows	Production deployment with training infrastructure, distributed inference, and FAIR model registries
P08 Software Portability	Proprietary software ecosystems limiting portability and hardware flexibility	Deploy portable frameworks enabling cross-platform execution
P09 Open Science & Reproducibility	Multi-decade experimental timescales with evolving software environments	Analyses preserved in re-executable form with full provenance
P10 Software Sustainability	Fragmented funding threatening long-term software preservation and availability	Open Science software repositories with software management plans

Pillar 4: Human Capital & Responsibility

PRIORITY	ADDRESSES	DELIVERS
P11 **Community Collaboration & Co-Design**	Inconsistent co-design between scientific communities and infrastructure providers	Permanent coordination mechanisms with structured community feedback loops
P12 **Environmental Sustainability**	Growing environmental impact requiring lifecycle-aware infrastructure approach	Sustainability metrics (PUE/CUE), HEPscore/watt procurement and extended hardware lifetime
P13 **Skills, Careers & Training**	Inadequate career recognition driving skilled personnel to industry	Established Research Software Engineer (RSE) career paths with competitive recognition and salaries