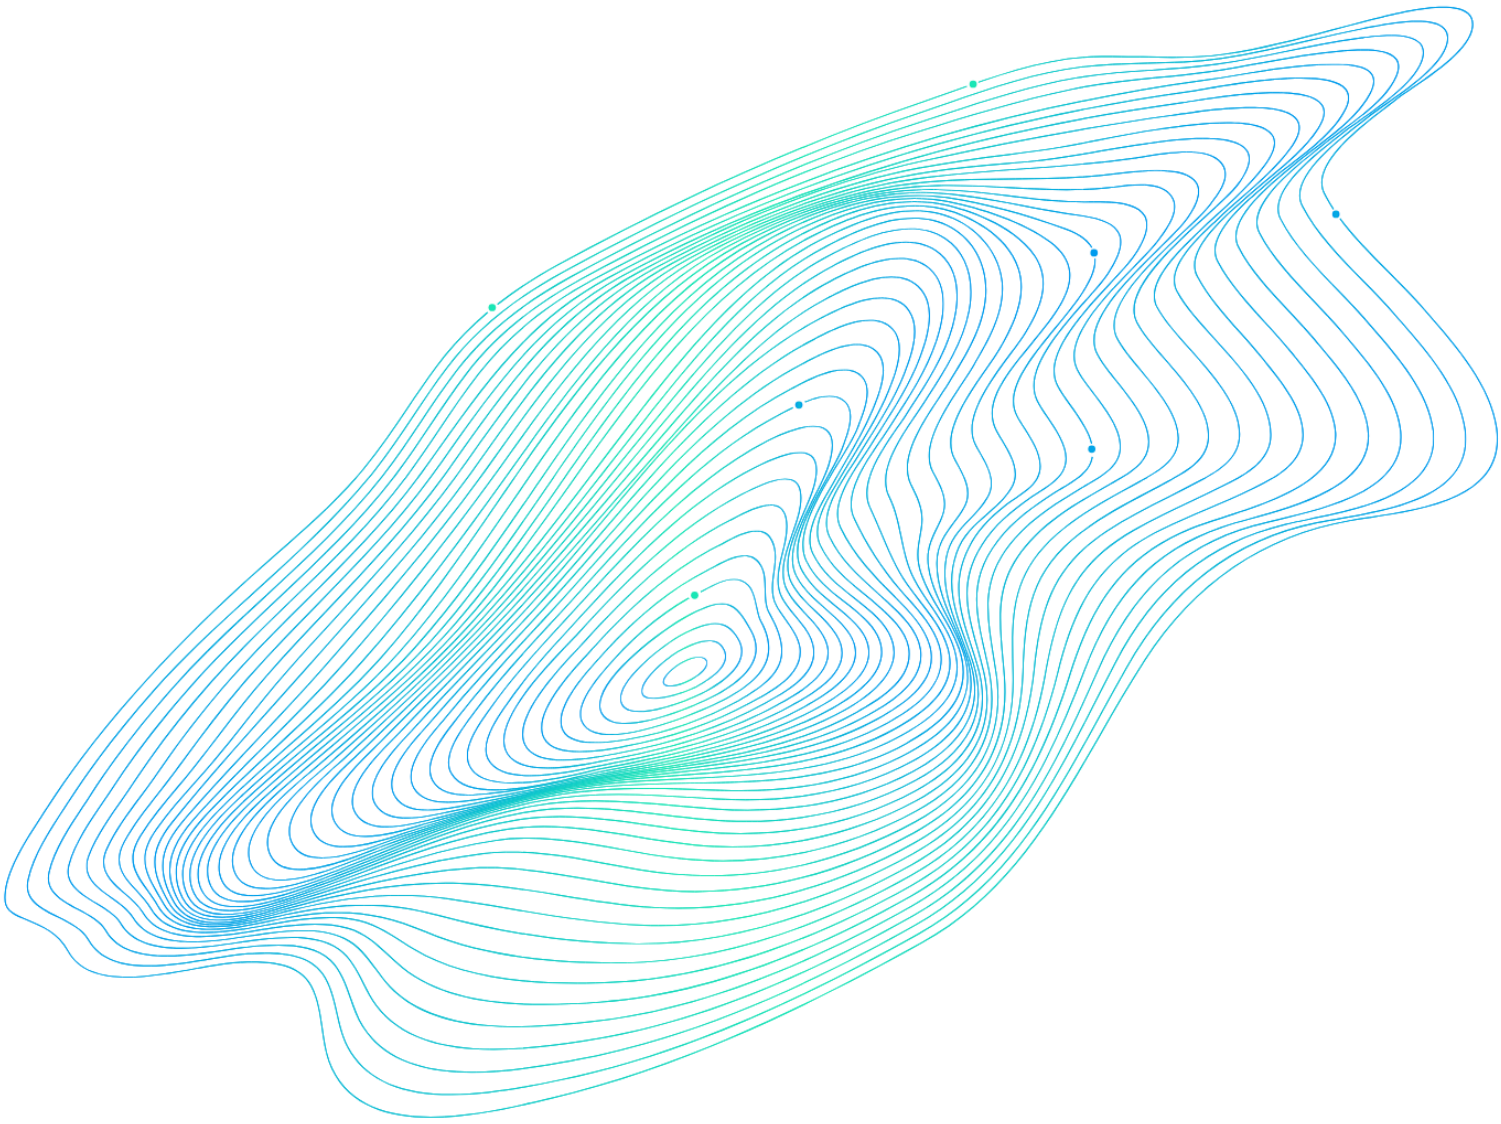
EPICURE

Unlocking European-level HPC Support

HIGH-LEVEL SPECIALISED APPLICATION SUPPORT SERVICE IN HIGH-PERFORMANCE COMPUTING (HPC)



Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No.101139786. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

EPICURE

DELIVERABLE D 4.2

Lessons learned from user requirements



Project Title	High-level specialised application support service in High-Performance Computing (HPC)
Project Acronym	EPICURE
Project Number	101139786
Type of Action	DIGITAL JU Simple Grants
Topic	DIGITAL-EUROHPC-JU-2022-APPSUPPORT-01-01
Starting Date of Project	2024-02-01
Ending Date of Project	2028-01-31
Duration of the Project	48 months
Website	epicure-hpc.eu

Work Package	WP4 Technical Coordination
Task	T4.1 Collaboration with the EuroHPC JU Peer Review Office
Lead Authors	João Barbosa (IT4I@VSB)
Contributors	Frédéric Pinel (LuxProvide) Peter Larsson (KTH)
Peer Reviewers	Samo Miklavc (IZUM) Vladislav Rusanov (SofiaTech)
Version	1.0
Due Date	31/1/2026
Submission Date	29/1/2026

Dissemination Level

X	PU: Public
	SEN: Sensitive – limited under the conditions of the Grant Agreement
	EU-RES. Classified Information: RESTREINT UE (Commission Decision 2005/444/EC)
	EU-CON. Classified Information: CONFIDENTIEL UE (Commission Decision 2005/444/EC)
	EU-SEC. Classified Information: SECRET UE (Commission Decision 2005/444/EC)

Co-funded by the European Union
EuroHPC Joint Undertaking

Version History

Revision	Date	Editors	Comments
V0.1	09/12/2025	João Barbosa (IT4I)	Table of Contents
V0.2	05/01/2026	João Barbosa (IT4I) Frederic Pinel (LuxProvide) Peter Larsson (KTH)	First complete draft of the deliverable
V0.3	12/01/2026	Samo Miklavc (IZUM) Vladislav Rusanov (SofiaTech)	Review conducted by the appointed reviewers
V0.4	19/01/2026	João Barbosa (IT4I)	Reviewers input consolidated and second draft sent to the Project Management Board (PMB)
V0.5	19/01/2026	Ville Savolainen (CSC) Virvatuli Uusimäki (CSC) David Vicente (BSC)	PMB feedback on the deliverable received
V0.6	25/01/2026	João Barbosa (IT4I)	Project Management Board feedback consolidated and final version of the deliverable made available to the PMO
V1.0	29/01/2026	Virvatuli Uusimäki (CSC)	Final quality check performed and deliverable submitted

Declaration on the use of AI

This report has been prepared with the support of AI tools which were used to assist in drafting text and checking language. All content has been thoroughly reviewed, fact-checked, and edited by the authors to ensure accuracy and alignment with the objectives of the report.

Co-funded by the European Union
EuroHPC Joint Undertaking

Glossary of Terms

Item	Description
HPC	High-Performance Computing
ESM	EPICURE Support Manager
SME	Small and Medium-sized Enterprise
GPU	Graphics Processing Unit
AI	Artificial Intelligence
L2	Level 2 support (standard centre assistance)
L3	Level 3 support (specialised/deep engagement)
PI	Principal Investigator
WP	Work Package
NCC	National Competence Centre
CoE	Centre of Excellence
KPI	Key Performance Indicator
CFD	Computational Fluid Dynamics
CS	Computer Science
IP	Intellectual Property

Co-funded by the European Union
EuroHPC Joint Undertaking

Executive Summary

Professional support services are highly sought after and exert significant influence within the EuroHPC JU ecosystem. The EPICURE project has established a successful support service, demonstrating this impact. In alignment with Milestone MS12, this deliverable offers the EuroHPC JU formal input and consolidates lessons learned from support activities conducted throughout the project duration.

Key Results:

- Supported 192 users across 8 EuroHPC JU sites with contributions from 14 partner organisations.
- Approved 90.1% of support requests with a median approval time of under 1 day, the reject projects were due to misinformation, either did not fulfil the requirements or were L1 (centre support tickets).
- Received outstanding user feedback: helpfulness (4.77/5), recommendation (4.83/5), and staff availability (4.67/5).
- Completed 23.5% of projects, with 92.3% achieving at least some of their objectives.

User Needs: Performance optimization represented the most frequent request (46.2% of projects), followed by adapting software to specific HPC systems (38.5%) and introducing GPU-based acceleration (30.8%). Users consistently valued technical expertise, prompt responses, and informative support interactions.

Key Lessons Learned:

- Modernising legacy software offers a substantial opportunity to advance European scientific computing capabilities.
- Collaboration across HPC centres facilitates knowledge sharing, although enhanced coordination mechanisms are necessary.
- Users without HPC expertise require structured onboarding processes and clear communication regarding expected timelines and results.
- Small and medium-sized enterprises (SMEs) and industrial users encounter specific challenges related to intellectual property governance and the need for streamlined access pathways.

Recommendations for EuroHPC JU:

- Make application support a permanent and broadly accessible component of the European HPC ecosystem.
- Create centralised portals for support requests while maintaining distributed delivery.

Co-funded by the European Union
EuroHPC Joint Undertaking

- Establish dedicated pathways for AI workload onboarding and SME engagement.
- Strengthen coordination and collaboration among NCCs, CoEs, AI Factories, and EPICURE application support teams, and provides value on top of the existing support channels.

The operational experience of EPICURE demonstrates that expert application support is essential for enabling European researchers to fully utilize EuroHPC JU resources. The lessons documented herein will inform the development of a more robust and sustainable European HPC Application Support service.

This report presents data collected since the EPICURE Support Manager (ESM) entered production in January 2025; prior to this date, insufficient structured data existed to support systematic analysis. Additionally, some metrics presented in this document may overlap, as individual projects or support actions can exhibit multiple characteristics simultaneously (e.g., a single project may involve both performance optimisation and GPU acceleration).

Co-funded by the European Union
EuroHPC Joint Undertaking

Table of Contents

1. Introduction 10
1.1 Purpose and Scope 10
1.2 Relation with other EPICURE outputs 11
1.3 Structure of the document 11
2. Evidence 12
2.1 Data sources 12
2.2 Data collection and processing 13
2.3 Role of the EPICURE Support Manager 14
2.4 Quantitative overview of support activity 14
2.4.1 Support Request Pipeline 15
2.4.2 Project Portfolio Overview 16
2.4.3 Supercomputer Site Distribution 19
2.4.4 Partner Effort Distribution 22
2.4.5 Temporal Trends 24
2.4.6 Process Efficiency: Approval Times 26
2.4.7 User Feedback and Satisfaction 28
2.4.8 Support Types Requested 29
2.4.9 Scientific Domains and User Affiliations 30
2.4.10 Key Performance Indicators Summary 34
3. Lessons Learned 35
3.1 User experience and satisfaction 35
3.2 Process efficiency and improvement opportunities 37
3.3 Technical landscape and challenges 38
3.4 Cross-centre collaboration 40
3.5 Industry and SME engagement 41
4. Feedback and Recommendations to EuroHPC JU 43
4.1 Sustaining application support services 43

Co-funded by the European Union
EuroHPC Joint Undertaking

4.2 Strategic infrastructure investments 43
4.3 European HPC community coordination 44
5. Conclusions 45
5.1 Summary of key lessons 45
5.2 Fulfilment of Milestone 46
5.3 Outlook 46

Co-funded by the European Union
EuroHPC Joint Undertaking

1. Introduction

1.1 Purpose and Scope

This deliverable compiles lessons learned from application support tasks completed during the EPICURE project. It provides the EuroHPC JU with structured input on service gaps, user needs, and recommendations to strengthen the European HPC Application Support service.

The scope of this document covers all support activities coordinated under EPICURE, from initial user requests to completed projects and subsequent follow-up interactions. It analyses recurring themes across EPICURE partners, scientific domains, and user communities to identify common challenges, effective practices, and areas warranting more development at the European level.

The analysis in this document is based on both quantitative and qualitative evidence. Quantitative data comprises support statistics from participating EPICURE partners and aggregated records from the EPICURE support manager. These records include indicators such as the total number of support requests, conversion rates from requests to approved projects, project breakdown across grant types (Benchmark, Development, Regular, Extreme, AI), approval timelines, project durations, estimated and actual person-months of effort, and the number of users served. Additional indicators include the proportion of single site versus multi-site collaborations, the distribution of projects across EPICURE partner sites, and the effort contributions from each partner organisation.

Qualitative data includes user input collected through direct communication and structured surveys. This feedback reflects users' perceptions of support quality, including ratings of helpfulness, staff availability, and likelihood to recommend the service. It also documents whether users achieved objectives such as code porting, performance improvements, intra-node and inter-node optimisation, or GPU-based acceleration, and provides free-text descriptions of effective practices, challenges, and additional learning or time savings. Testimonials and success stories, drawn from feedback entries with dissemination text, show how users feel about the support they received and how useful it was. Internal input from coordination calls, technical meetings, and comments on support and help requests helps us identify process bottlenecks, coordination issues, and aspects needing refinement.

This deliverable is intended for EuroHPC JU, the participating centres, and other stakeholders involved in guiding the future of European HPC application support services. The goal of the document is to turn operational experience into useful suggestions for future program calls, service design, and policy development.

Co-funded by the European Union
EuroHPC Joint Undertaking

1.2 Relation with other EPICURE outputs

Within the wider set of EPICURE outputs, this deliverable serves as a synthesis document that draws upon detailed technical descriptions of individual support actions presented in other project deliverables. Its primary focus is to identify cross-cutting patterns, formulate broad conclusions, and develop recommendations based on the collective experience of application support throughout the project, rather than providing case-level specifics.

The EPICURE support portal serves as the primary repository for recording projects and support actions. The information in this document comes from data stored in this system. This deliverable shows how support activities, together with insights, changed over the course of the EPICURE project, which lasted from Month 1 to Month 24.

1.3 Structure of the document

This deliverable is organised as follows. Section 2 establishes the evidence base for the lessons learned, describes the data sources and collection methods, and provides a quantitative overview of support activity across the project. Section 3 outlines the insights gained, encompassing user requirements and expectations, service and process deficiencies, technical and tooling obstacles, inter-partner collaboration, and particular issues for SMEs and industry stakeholders. Section 4 translates these lessons into structured feedback and recommendations addressed to EuroHPC JU, encompassing service and process improvements, tools and shared infrastructure, support for new user communities, and guidance for future programmes and calls. Finally, Section 5 draws conclusions, demonstrates fulfilment of Milestone MS12, and provides an outlook on how the knowledge acquired can inform the continued development of the European HPC Application Support service.

Co-funded by the European Union
EuroHPC Joint Undertaking

2. Evidence

2.1 Data sources

The insights presented in this deliverable are derived from four primary data types collected throughout the EPICURE project. These data types combine quantitative indicators and qualitative data points to provide a comprehensive perspective on application support activities within the EuroHPC JU ecosystem.

EPICURE Support Manager (ESM) provides the most structured data source. It serves as a centralised registry for all support projects and actions, recording detailed metadata for each activity. For every project, the ESM captures the associated EuroHPC JU supercomputer, scientific domain, user category and affiliation, support level (L2 or L3), grant type (Benchmark, Development, Regular, Extreme, or AI), and status. It also records estimated and actual effort in person-months, project timelines, and the links between support requests, projects, help requests, and user feedback. This structured data enables analysis of workload patterns across EuroHPC JU systems, resource distribution among partners, and the progression of support actions from initial request through to completion.

User feedback is the second data source, providing essential information on the service's quality, value, and effectiveness. This feedback was collected through standardized surveys administered at the end of the project, assessing the service's usefulness, staff accessibility, likelihood of recommending the service, and the extent to which users achieved their objectives. Free-text sections enabled users to describe successful approaches, difficulties experienced, additional learning, and anticipated time savings. Additional qualitative data were obtained through direct email communication and observations during meetings, training sessions, and workshops.

Contributions from internal practitioners within the support staff enhance the research by providing operational information not available in user-facing statistics. This feedback was collected during regular coordination calls between EPICURE partners, technical meetings targeting distinct support challenges and inter-partner dependencies, and opportunities for procedural improvement. Practitioner perspectives are vital to grasping the organizational and coordination aspects of providing support at the European level.

Testimonials and success stories represent the fourth form of evidence, providing qualitative documentation regarding user experiences. These testimonials showcase users' perceived value of the support service, including improved performance, successful code modifications, and the development of new search capabilities. They also highlight gaps and unmet expectations and serve as a foundation for the outreach and community development initiatives discussed later in this document.

Co-funded by the European Union
EuroHPC Joint Undertaking

2.2 Data collection and processing

Throughout the project, we adopted a continuous approach to data collection. The ESM functioned as our central repository: when users requested EPICURE support through EuroHPC JU, the EuroHPC JU forwarded a link to a simple submission form, which the project PI completed to create a support request in the system. As decisions were made—whether to approve, decline, or request further information—those updates were recorded as well. Once a project was underway, the team kept track of effort spent, status changes, and notable milestones like completing a port or running performance benchmarks. To ensure data quality, WP2 and WP3 leaders periodically reviewed project records for consistency, and a final consolidation was performed at the end of the project to make sure everything was as accurate and complete as possible.

Each EPICURE partner entered data through the ESM's web interface, using standardised forms to record project details, link activities to specific EuroHPC JU supercomputers, and track person-months of effort by work package. Staff at each organisation were responsible for keeping their own project records current. User feedback came through a different channel: the ESM automatically issued a survey to users when their projects were marked as complete, asking them to rate aspects like helpfulness and staff availability, and to describe in their own words what went well, what did not, and what they learned. Using the same survey across all partners made it straightforward to aggregate responses and draw Europe-wide conclusions.

To allow for meaningful comparisons across partners and sites, we agreed on to harmonise how we categorised the data. Scientific domains were grouped under a shared taxonomy—life sciences, materials science, engineering, climate and weather, fundamental physics, data analytics, and so on. User affiliations were tagged by institution type (academic, research institute, SME, or industry) and country. We distinguished between L2 support, which covers standard assistance from centre staff, and L3 support, which involves deeper, more specialised engagement. Grant types followed EuroHPC JU's established categories: Benchmark, Development, Regular, Extreme, and AI. These classifications are used purely for statistical and reporting purposes.

Before running any analysis, we put the consolidated dataset through a series of quality checks. We looked for and removed duplicate entries—often caused by administrative corrections or simple data entry mistakes—and ran consistency checks to make sure related records lined up properly. For instance, we verified that every project was correctly linked to its originating support request, that the effort totals partners reported matched the project-level estimates, and that status changes made logical sense (a project cannot be marked complete before it has been approved). Where data were missing, we followed up with the responsible partners when we could; otherwise, we documented the gaps and excluded those records from the relevant parts of the analysis. The result is a dataset we are confident provides a solid basis for the findings that follow.

Co-funded by the European Union
EuroHPC Joint Undertaking

2.3 Role of the EPICURE Support Manager

The EPICURE Support Manager (ESM) plays a crucial role in the EPICURE project, serving as the centralised registry for all support projects and actions. It captures detailed metadata for each activity, enabling analysis of workload patterns across EuroHPC JU systems, resource distribution among partners, and the progression of support actions from initial request through to completion.

For each project and support action, the ESM stores comprehensive information including the target supercomputing centre, scientific domain, user affiliation type (academic, SME, industry, or government), a short description of the work, support level (L2 or L3), and current status. This structured approach provides a unified view of support activities across all EPICURE partners, making it possible to compare performance and identify best practices.

The ESM has supported EPICURE operations in several keyways:

- **Monitoring of ongoing support work**: Project managers can track the status of all active projects, identify bottlenecks, and ensure timely delivery of support services.
- **Pattern identification**: Aggregated data reveals which scientific domains receive the most support, which user groups are underserved, and how demand varies across sites.
- **Data extraction for reporting**: The analysis presented in Section 2.4 was extracted directly from the ESM, demonstrating its value as a reporting tool.

The ESM represents a key outcome of EPICURE, enabling for the first time a European-level perspective on HPC application support that spans multiple centres and countries.

2.4 Quantitative overview of support activity

This section presents a comprehensive set of key performance indicators derived directly from the EPICURE Support Manager (ESM) database. The analysis covers the operational period from January 2025 onwards, when the ESM became the primary tool for tracking all EPICURE support activities. All figures presented here are calculated from the data models described in Section 2.2, ensuring full traceability and reproducibility of the reported metrics.

Co-funded by the European Union
EuroHPC Joint Undertaking

2.4.1 Support Request Pipeline

Since the ESM became operational on 10 January 2025, EPICURE has received a total of **91 support requests** from researchers and institutions across Europe. This represents a steady flow of demand for HPC application support services, averaging approximately 7-8 new requests per month during the operational period.

The breakdown of request outcomes demonstrates the effectiveness of the EPICURE support model:

Status	Count	Percentage
Approved	82	90.1%
Rejected	6	6.6%
Archived	2	2.2%
Pending	1	1.1%

The **90.1% approval rate** is notably high and reflects several factors: the clear eligibility criteria communicated through EuroHPC JU channels, the quality of applications received from the research community, and the collaborative approach taken during the review process. Rejected requests (6.6%) typically fell outside the scope of EPICURE's mandate—for instance, requests for pure system administration rather than application-level support—or lacked sufficient technical detail for partners to provide meaningful assistance. The small number of archived requests (2.2%) represents cases where applicants withdrew or became unresponsive during the clarification phase.

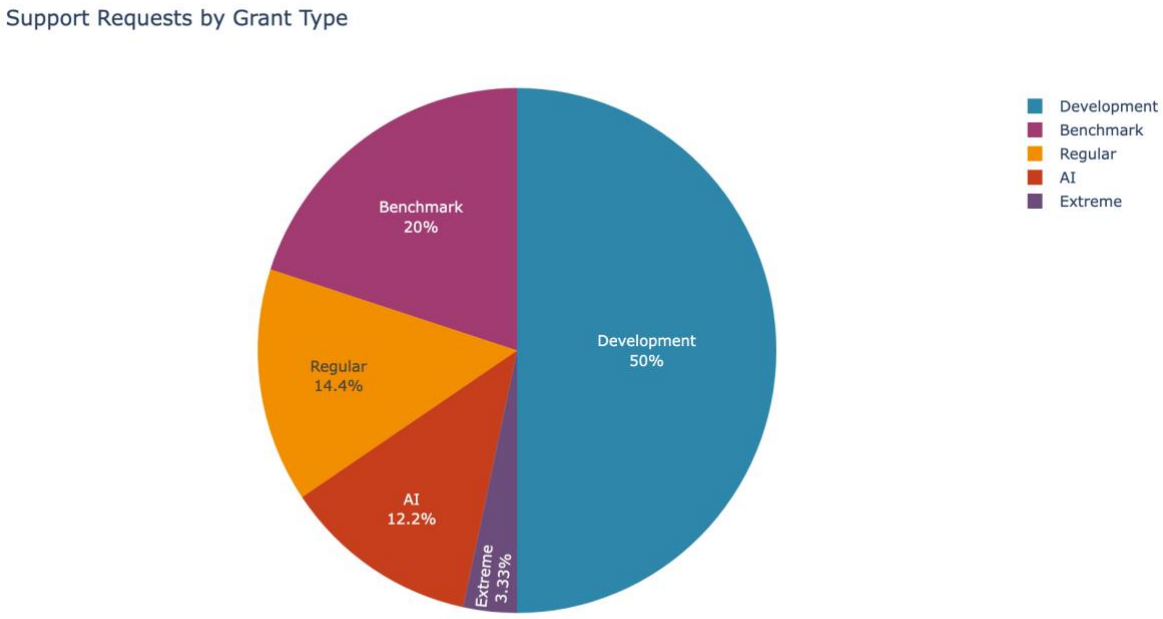
The distribution of requests by EuroHPC JU grant type provides insight into the nature of support being sought:

Grant Type	Count	Percentage	Description
Development	45	49.5%	Code development and optimization
Benchmark	18	19.8%	Performance testing and validation
Regular	13	14.3%	Standard HPC projects
AI	11	12.1%	Machine learning and AI workloads
Extreme	3	3.3%	Large-scale, resource-intensive projects

The dominance of **Development grants (49.5%)** reveals that code optimization and porting remain the primary needs of European HPC users. Researchers frequently require assistance adapting their

Co-funded by the European Union
EuroHPC Joint Undertaking

applications to new architectures, particularly the GPU-accelerated systems that now characterise the EuroHPC JU infrastructure. The significant share of **AI grants (12.1%)** reflects the rapid growth of machine learning workloads in scientific computing, a trend that has accelerated throughout EPICURE's operational period. **Benchmark grants (19.8%)** indicate ongoing demand for performance validation services, which are essential for researchers needing to demonstrate scalability for grant applications or publications. The relative rarity of **Extreme grants (3.3%)** is consistent with expectations: these projects require exceptional resources and represent only the most computationally demanding research endeavours in Europe.



Support Requests by Grant Type

Figure: Distribution of support requests across EuroHPC JU grant types, showing the predominance of Development-focused support needs.

2.4.2 Project Portfolio Overview

Following approval, support requests become active projects in the ESM system. The current portfolio comprises 81 projects representing a substantial investment in European research capacity:

Co-funded by the European Union
EuroHPC Joint Undertaking

Metric	Value
Total Projects (ESM-era)	81
Total Estimated Person-Months	160.50
Average PMs per Project	1.98
Median PMs per Project	1.00
Total Users Involved	192
Average Users per Project	2.37

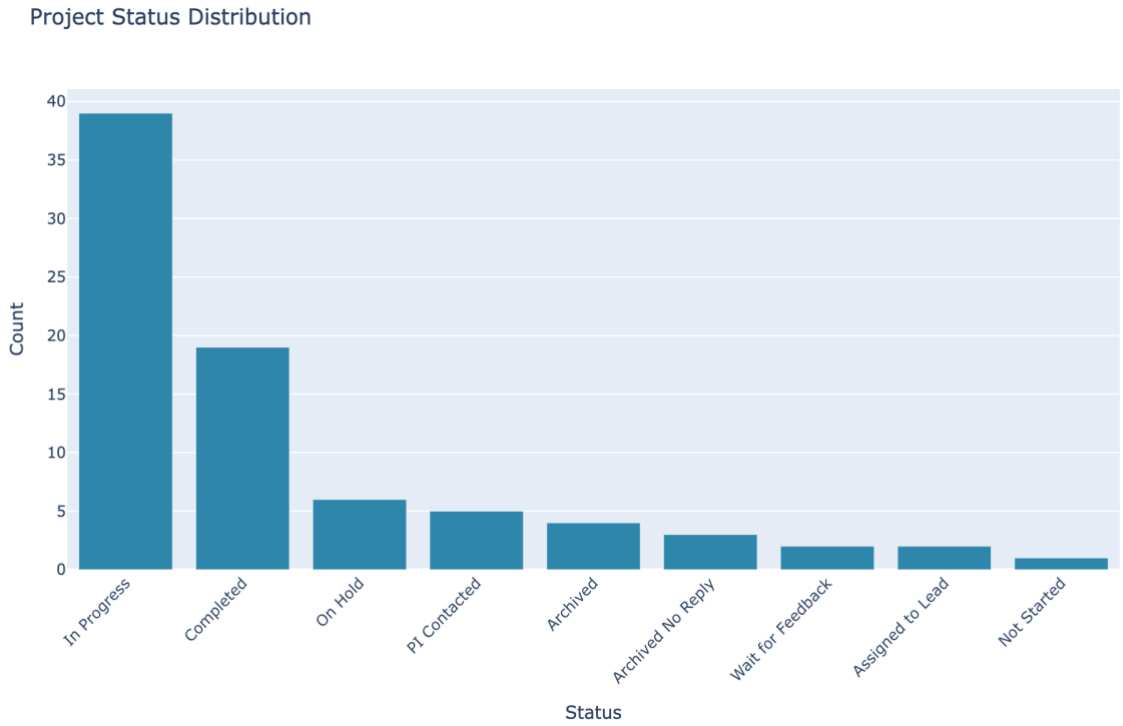
The total of **160.5 person-months** of estimated effort demonstrates EPICURE's significant contribution to enabling European research. With an average of approximately 2 PMs per project, support engagements are substantial enough to deliver meaningful technical improvements while remaining focused and efficient. The 192 users directly involved in these projects represent a growing community of researchers who have benefited from expert HPC application support.

The project lifecycle tracking reveals a healthy distribution across workflow stages:

Status	Count	Percentage	Description
In Progress	39	48.1%	Active support engagement
Completed	19	23.5%	Successfully delivered
On Hold	6	7.4%	Temporarily paused
PI Contacted	5	6.2%	Initial contact phase
Archived	4	4.9%	Closed without completion
Archived No Reply	3	3.7%	User became unresponsive
Wait for Feedback	2	2.5%	Awaiting user input
Assigned to Lead	2	2.5%	Pending partner assignment
Not Started	1	1.2%	Awaiting kickoff

Nearly half of all projects (48.1%) are **In Progress**, indicating a healthy active workload and sustained engagement between EPICURE partners and research teams. The **23.5% completion rate** demonstrates that projects move through the pipeline successfully, with 19 engagements already delivering their intended outcomes. The relatively small proportion of projects stuck in early phases (PI Contacted, Assigned to Lead, Not Started: combined 9.9%) suggests efficient onboarding processes. The archived categories (combined 8.6%) represent expected attrition—some users change priorities, lose funding, or become unavailable during long-running projects.

Co-funded by the European Union
EuroHPC Joint Undertaking



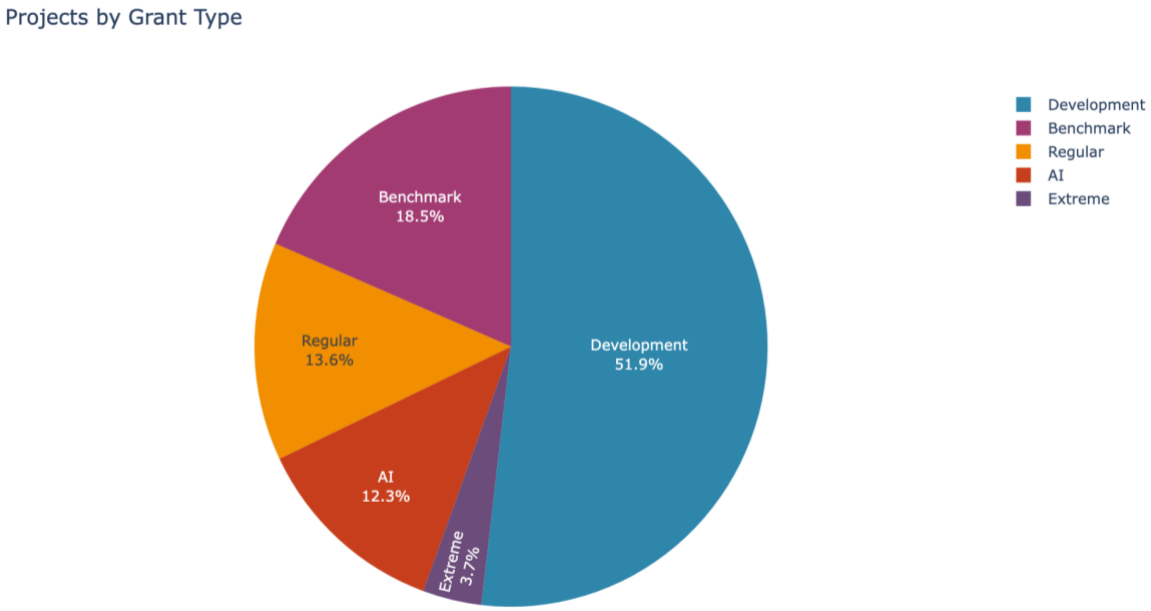
Project Status Distribution

Figure: Status distribution across the project portfolio, demonstrating active engagement and successful delivery.

The grant type distribution at the project level closely mirrors the support request distribution, confirming that approval rates are consistent across grant categories:

Grant Type	Projects	Percentage
Development	42	51.9%
Benchmark	15	18.5%
Regular	11	13.6%
AI	10	12.3%
Extreme	3	3.7%

Co-funded by the European Union
EuroHPC Joint Undertaking



Projects by Grant Type

Figure: Project distribution by grant type, showing consistent patterns with support request intake.

2.4.3 Supercomputer Site Distribution

One of EPICURE's primary objectives is to facilitate access to EuroHPC JU supercomputing resources across Europe. The distribution of projects across the eight participating sites demonstrates broad geographical coverage while reflecting the relative scale and capabilities of each system:

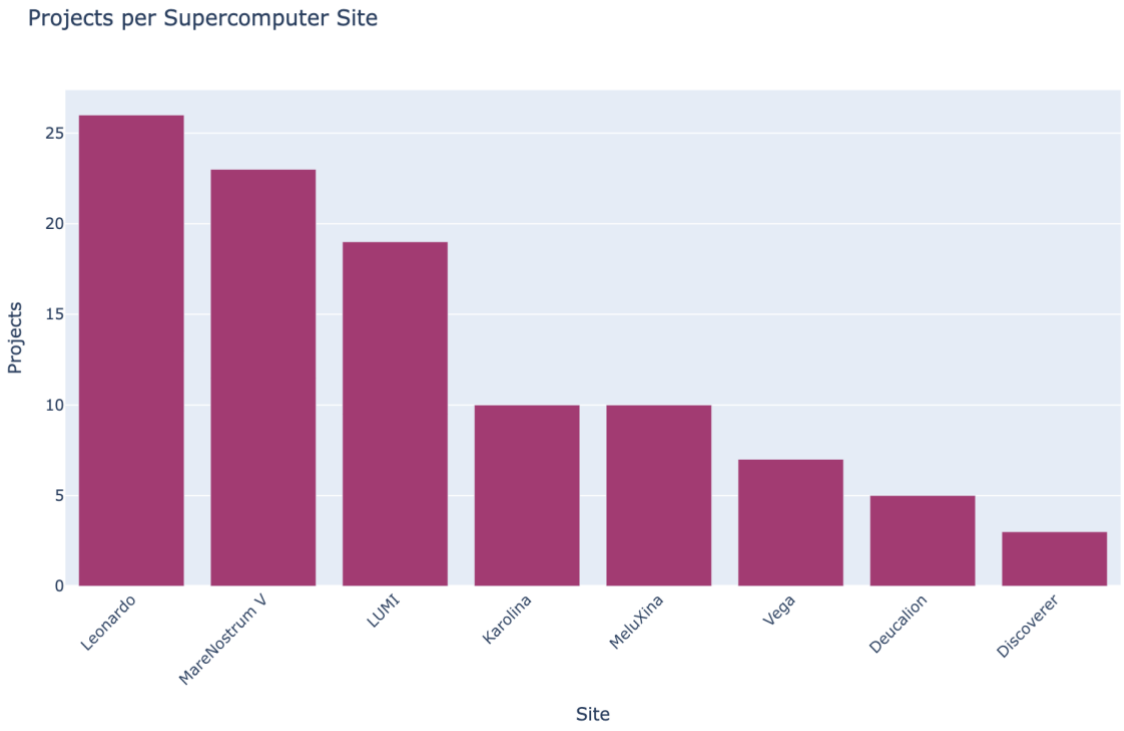
Supercomputer	Country	Projects	Percentage
Leonardo	Italy	26	25.2%
MareNostrum 5	Spain	23	22.3%
LUMI	Finland	19	18.4%
Karolina	Czech Republic	10	9.7%
MeluXina	Luxembourg	10	9.7%
Vega	Slovenia	7	6.8%
Deucalion	Portugal	5	4.9%

Co-funded by the European Union
EuroHPC Joint Undertaking

Discoverer	Bulgaria	3	2.9%

Leonardo (CINECA, Italy) leads with 26 projects, reflecting both its status as one of Europe's most powerful supercomputers and CINECA's extensive experience in HPC application support. **MareNostrum 5** (BSC, Spain) follows closely with 23 projects, benefiting from BSC's strong connections to the European research community and its expertise in performance optimization. **LUMI** (CSC, Finland), Europe's most powerful supercomputer, hosted 19 projects with a particular strength in GPU-accelerated workloads given its AMD MI250X architecture.

The mid-tier sites—**Karolina** and **MeluXina**—each supported 10 projects, demonstrating that EPICURE successfully directs users to appropriate resources rather than concentrating all activity on the largest systems. The smaller sites (**Vega**, **Deucalion**, **Discoverer**) complete the European coverage, ensuring that researchers across all participating countries have access to local expertise and resources.

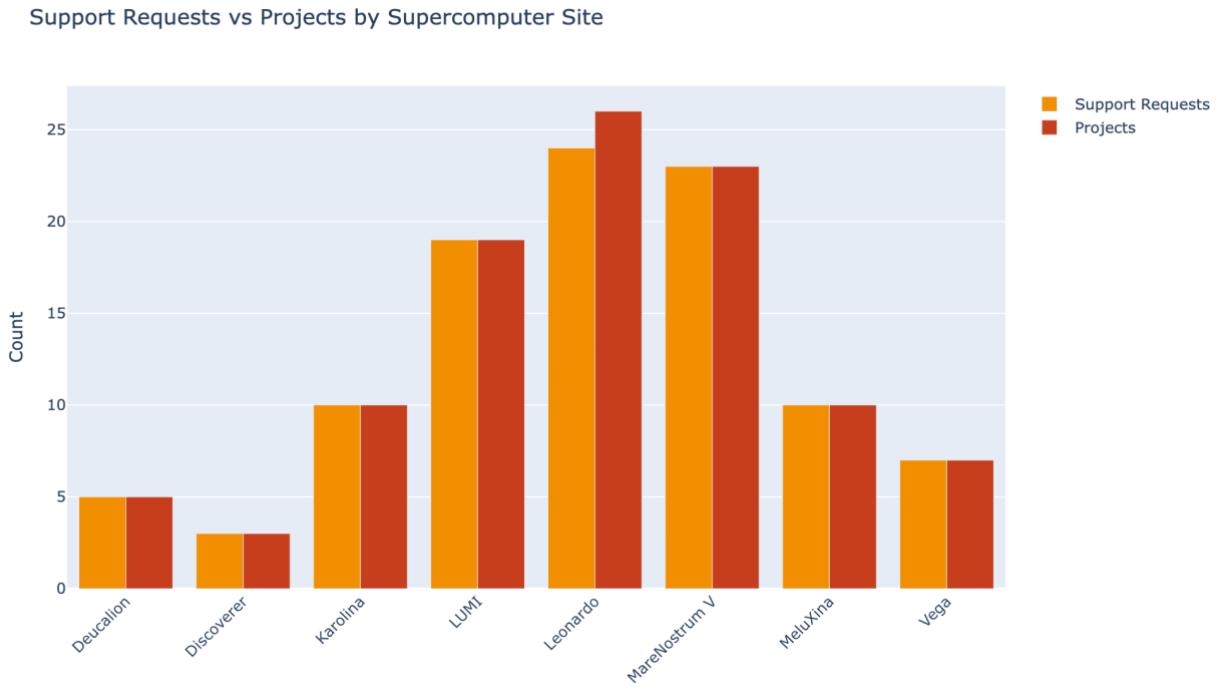


Projects per Site

Figure: Project distribution across EuroHPC JU supercomputing sites, demonstrating broad European coverage.

Co-funded by the European Union
EuroHPC Joint Undertaking

The comparison between support requests and resulting projects per site reveals consistent conversion rates across all centres. This consistency indicates that the approval process operates uniformly regardless of the target supercomputer, and that no site experiences systematic barriers to project activation.

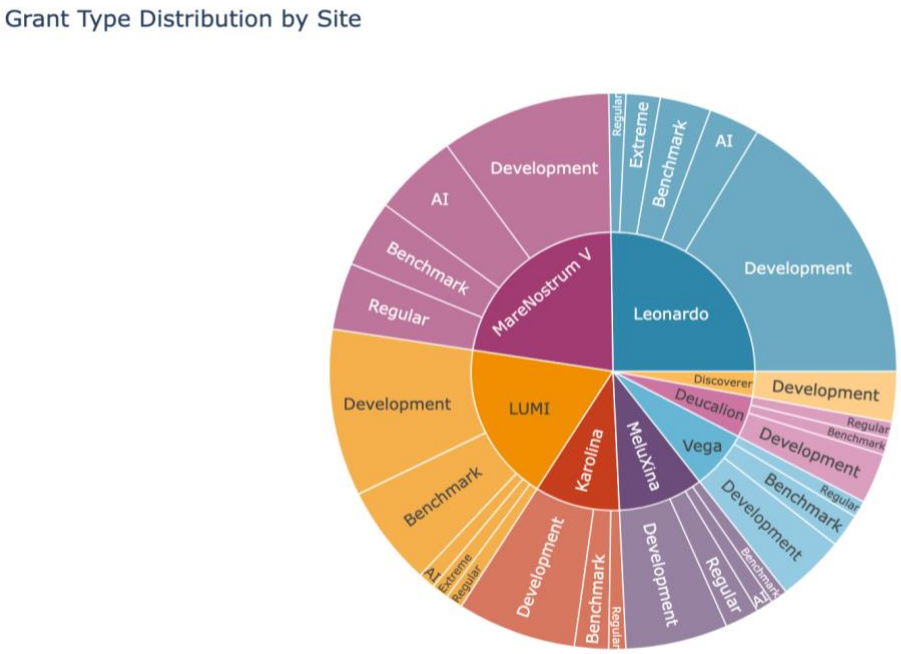


Support Requests vs Projects by Site

Figure: Comparison of support requests received versus projects activated at each site, showing consistent conversion rates. (Note: Leonard has an additional project because it was added to the project after the user request)

The grant type distribution varies somewhat by site, reflecting the different strengths and user communities associated with each supercomputer. The sunburst visualisation below illustrates how different grant types are distributed across sites:

Co-funded by the European Union
EuroHPC Joint Undertaking



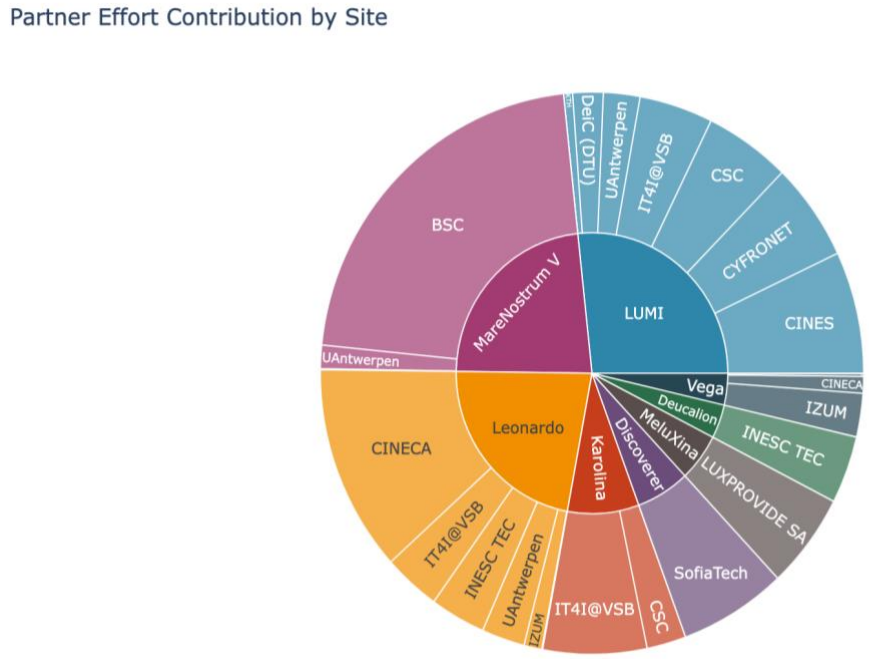
Grant Type Distribution by Site

Figure: Sunburst chart showing the hierarchical distribution of grant types within each supercomputer site.

2.4.4 Partner Effort Distribution

The relationship between partners and sites is visualised in the sunburst chart below, showing how partner effort flows to different supercomputing resources:

Co-funded by the European Union
EuroHPC Joint Undertaking

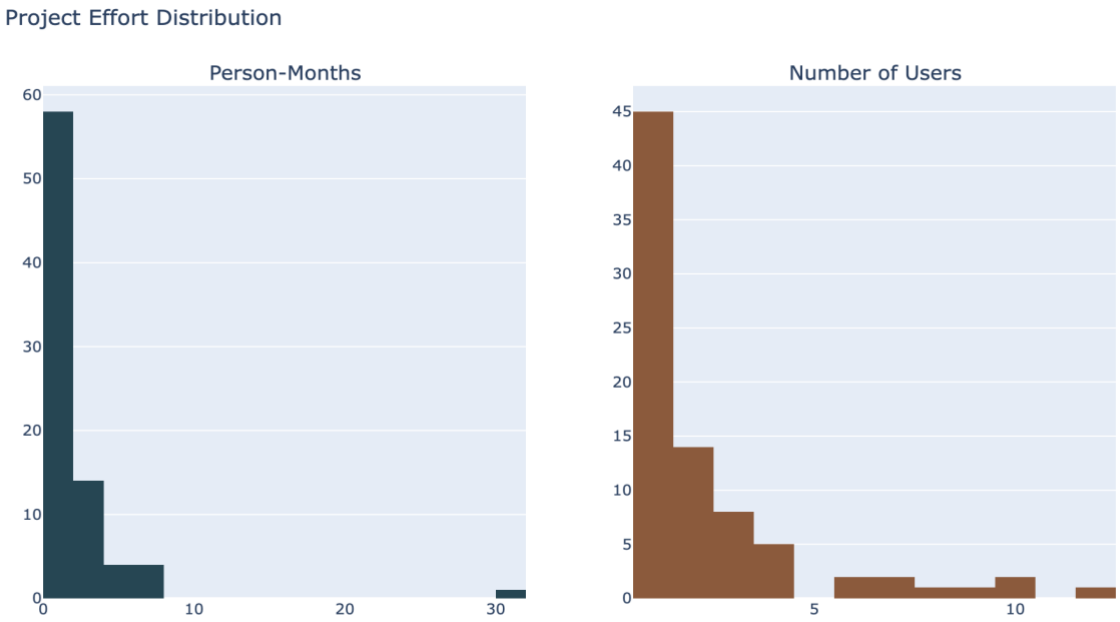


Partner Effort by Site

Figure: Sunburst chart showing partner effort distribution across supercomputer sites, revealing the collaborative support network.

The effort distribution per project provides additional insight into resource allocation:

Co-funded by the European Union
EuroHPC Joint Undertaking



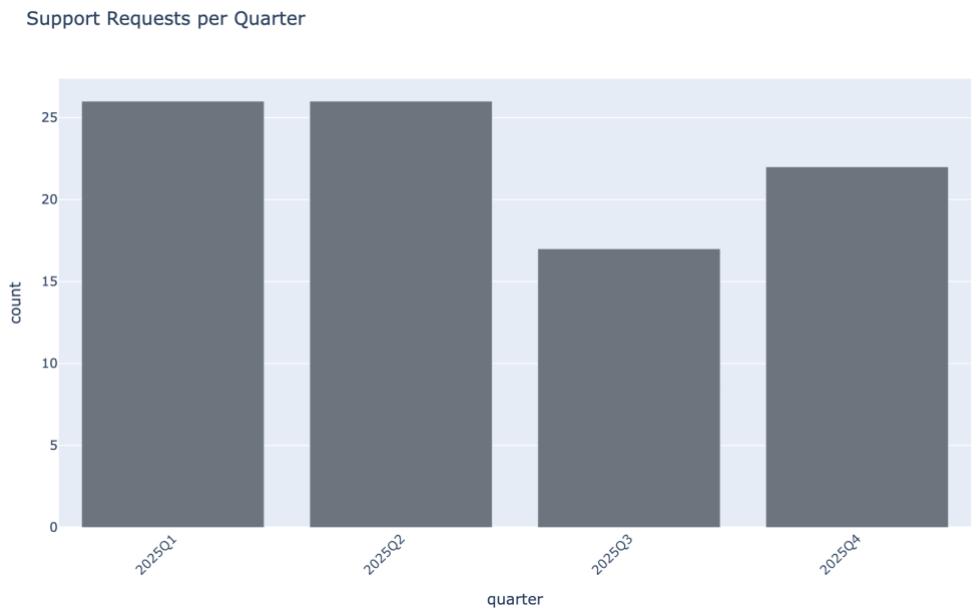
Project Effort Distribution

Figure: Histograms showing the distribution of estimated person-months and number of users per project.

2.4.5 Temporal Trends

The quarterly analysis of support requests reveals the intake pattern since the ESM became operational:

Co-funded by the European Union
EuroHPC Joint Undertaking



Support Requests per Quarter

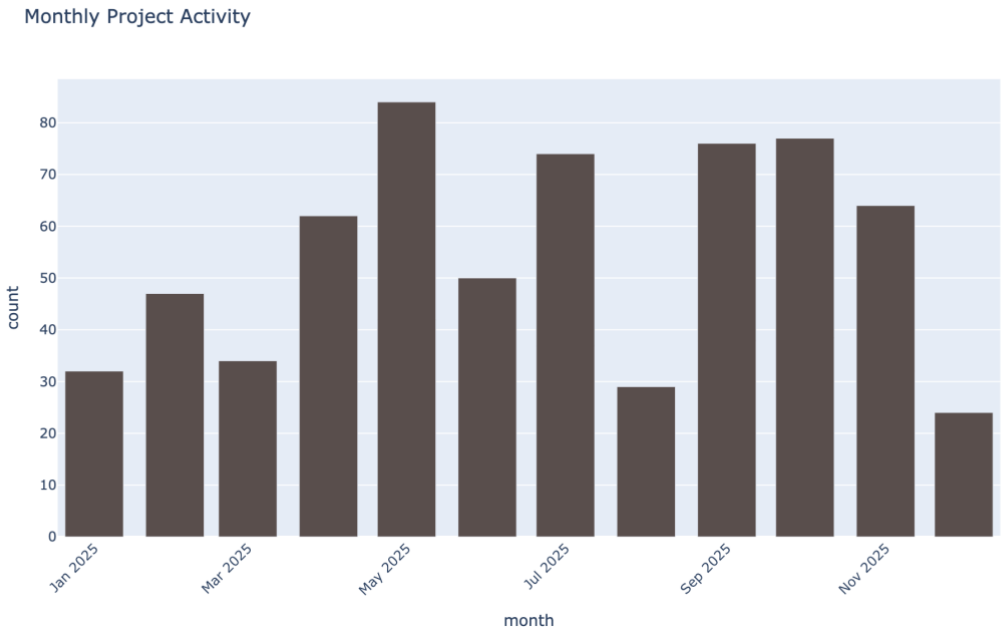
Figure: Support requests received per quarter, showing the growth and stabilisation of EPICURE's intake.

Project activity tracking provides visibility into the workflow patterns across the portfolio:

Activity Type	Count
Project Submitted	81
Project Approved	81
Project status changed to In Progress	68
Project status changed to PI Contacted	38
Meeting with the PI	16
Project status changed to Assigned to Lead	16
Feedback Request Sent	15
Feedback Received	13
Project status changed to Wait for Feedback	12
Project status changed to Completed	11

Co-funded by the European Union
EuroHPC Joint Undertaking

The activity log shows that all 81 projects were approved (100% conversion from submission), with 68 progressing to active work status. The 16 recorded meetings with PIs represent formal kick-off or progress discussions, while the remaining contact may have been via email or informal channels.



Monthly Project Activity

Figure: Monthly project activity showing the temporal distribution of support engagements.

2.4.6 Process Efficiency: Approval Times

A critical performance metric for any support service is the time required to move from initial request to active project. EPICURE's approval process operates efficiently:

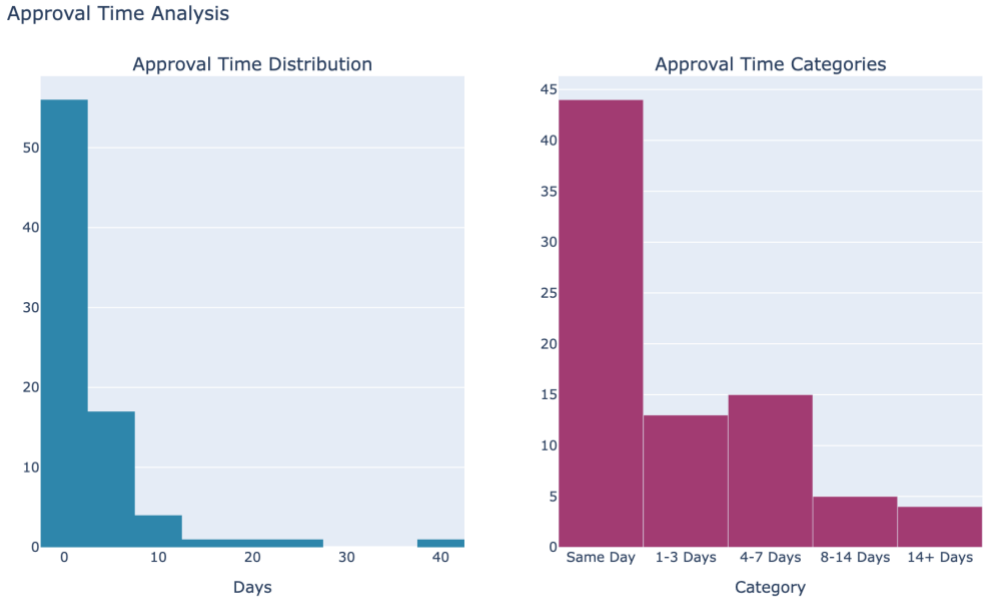
Metric	Days
Average Approval Time	3.1
Median Approval Time	1.0
Minimum	(same day) 0.0
Maximum	38.0

Co-funded by the European Union
EuroHPC Joint Undertaking

The **median approval time of 1.0 day** demonstrates that most requests are processed rapidly, often within 24 hours of submission. The average is somewhat higher (3.1 days) due to outliers requiring additional review or clarification. The distribution of approval times falls into the following categories:

- **Same Day (0-1 days)**: 41% of projects
- **1-3 Days**: 25% of projects
- **4-7 Days**: 15% of projects
- **8-14 Days**: 10% of projects
- **14+ Days**: 9% of projects

The 66% of projects approved within three days indicates a streamlined workflow. Cases taking longer typically involve complex multi-site requests, ambiguous scope definitions requiring clarification, or periods of high demand across the partnership.



Approval Time Analysis

Figure: Left - histogram showing the distribution of approval times in days. Right - categorical breakdown of approval time ranges.

Co-funded by the European Union
EuroHPC Joint Undertaking

2.4.7 User Feedback and Satisfaction

User satisfaction is measured through post-project surveys issued automatically when projects are marked as complete. With a response rate of **86.7%** (13 of 15 completed projects with surveys sent), the feedback data provides reliable insight into service quality:

Metric	Value
Total Feedback Records	15
Submitted Responses	13
Response Rate	86.7%

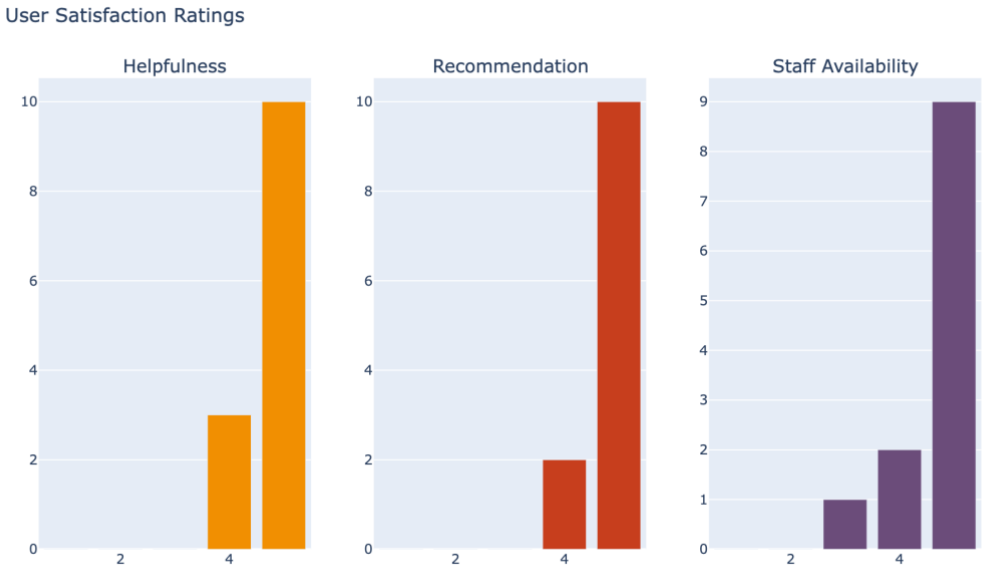
The three core satisfaction metrics, measured on a 1-5 scale where 5 represents the highest rating, show consistently excellent results:

Rating Type	Average	Median
Helpfulness	4.77	5.0
Recommendation	4.83	5.0
Staff Availability	4.67	5.0

Helpfulness (4.77/5) indicates that users find the technical support valuable and effective in addressing their needs. **Recommendation (4.83/5)** suggests strong user loyalty—respondents would recommend EPICURE services to colleagues. **Staff Availability (4.67/5)** confirms that support staff are accessible when needed, an important factor for time-sensitive research projects.

The median of 5.0 across all categories indicates that most users give the highest possible rating, with only occasional lower scores pulling down the average slightly. This consistency suggests reliable service quality rather than variable experiences.

Co-funded by the European Union
EuroHPC Joint Undertaking



Satisfaction Ratings

Figure: Bar charts showing the distribution of user ratings across helpfulness, recommendation, and staff availability metrics.

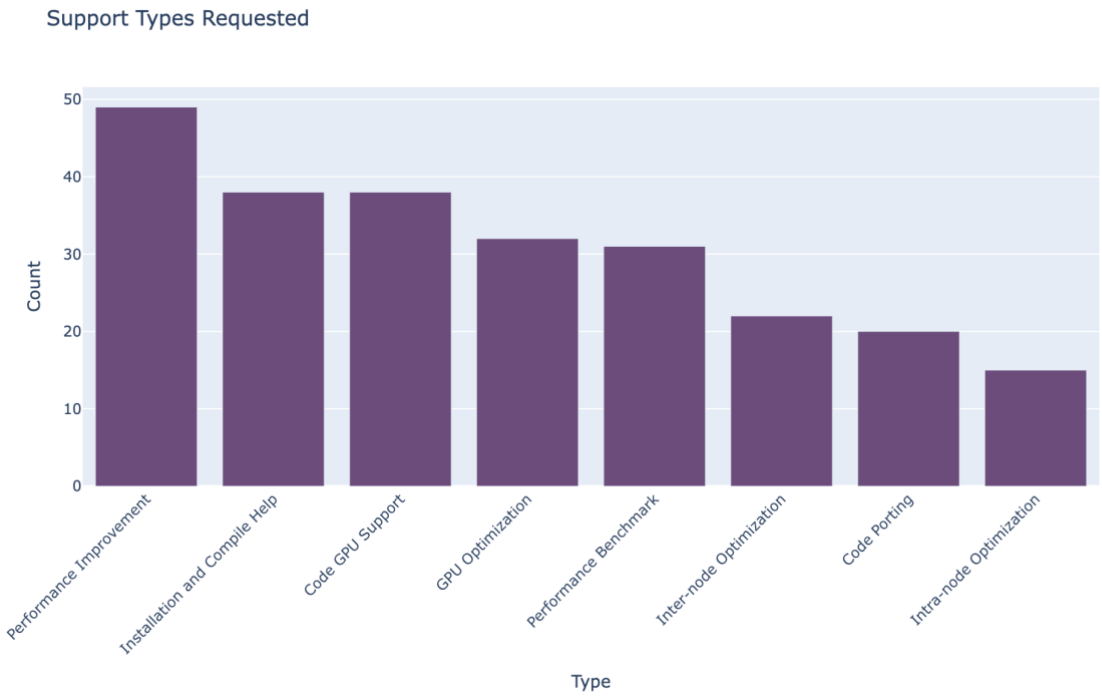
2.4.8 Support Types Requested

The types of support requested reveal the technical priorities and challenges facing European HPC users:

Support Type	Count
Performance Improvement	49
Installation and Compile Help	38
Code GPU Support	38
GPU Optimization	32
Performance Benchmark	31
Inter-node Optimization	22
Code Porting	20
Intra-node Optimization	15

Co-funded by the European Union
EuroHPC Joint Undertaking

Performance Improvement leads with 49 requests, reflecting the universal need to make codes run faster and scale better. The high demand for **GPU-related support** (Code GPU Support: 38, GPU Optimization: 32) confirms the ongoing transition to GPU-accelerated computing across scientific domains. **Installation and Compile Help** (38 requests) addresses the practical challenges of getting complex scientific software to build and run correctly on unfamiliar systems. **Benchmarking** (31 requests) remains essential for validating performance claims and supporting grant applications.



Support Types Requested

Figure: Bar chart showing the frequency of different support types, highlighting the emphasis on performance and GPU-related assistance.

2.4.9 Scientific Domains and User Affiliations

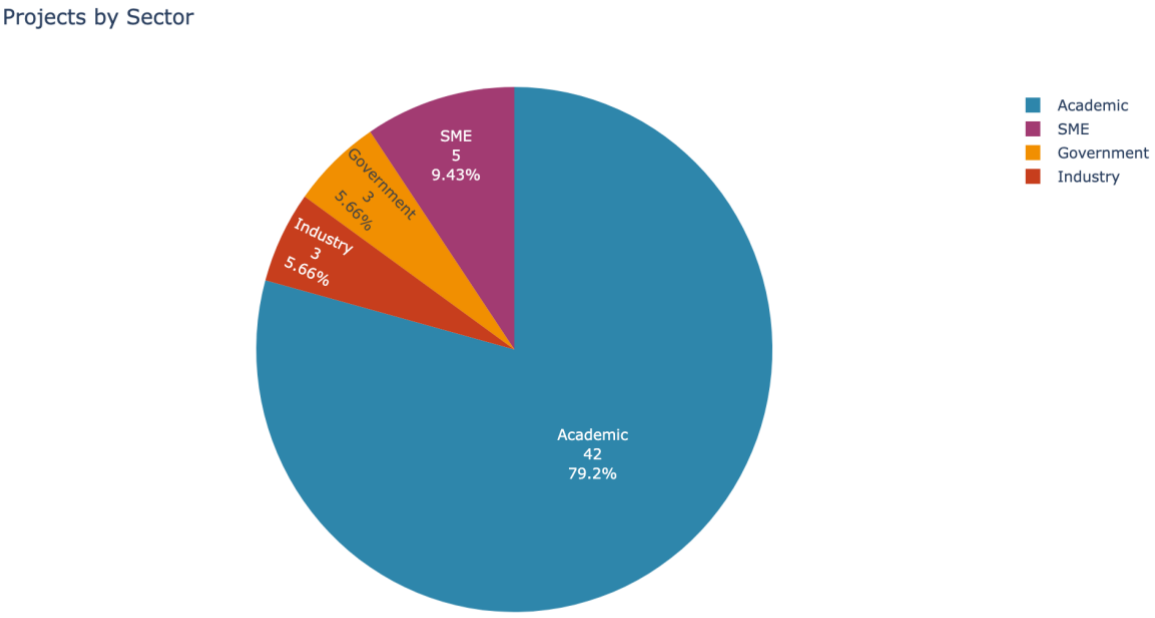
EPICURE serves a diverse user community spanning multiple scientific disciplines and institution types.

User Affiliation by Sector:

Co-funded by the European Union
EuroHPC Joint Undertaking

Sector	Count	Percentage
Academic	42	79.2%
SME	5	9.4%
Government	3	5.7%
Industry	3	5.7%

The predominance of **academic users (79.2%)** is expected given EuroHPC JU's focus on research computing. However, the presence of **SMEs (9.4%)** and **industry (5.7%)** demonstrates that EPICURE also supports the broader innovation ecosystem. **Government users** (5.7%) include research agencies and public sector computing initiatives.



Projects by Sector

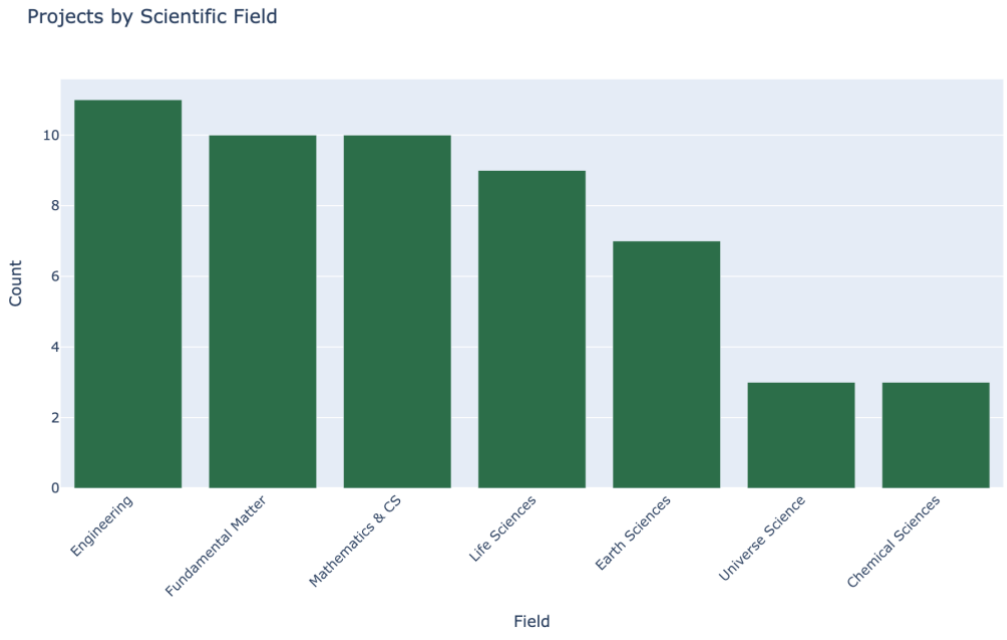
Figure: Pie chart showing user affiliation by sector, demonstrating EPICURE's primarily academic user base with growing industry engagement.

Scientific Field Distribution:

Co-funded by the European Union
EuroHPC Joint Undertaking

Field	Count
Engineering	11
Fundamental Matter	10
Mathematics & CS	10
Life Sciences	9
Earth Sciences	7
Universe Science	3
Chemical Sciences	3

The field distribution shows broad coverage across scientific domains. **Engineering** leads with 11 projects encompassing CFD, structural analysis, and industrial simulations. **Fundamental Matter** (particle physics, condensed matter) and **Mathematics & CS** (algorithms, numerical methods, AI) each contribute 10 projects. **Life Sciences** (9 projects) includes molecular dynamics, genomics, and drug discovery. **Earth Sciences** (7 projects) covers climate modelling, weather prediction, and geophysics.



Projects by Scientific Field

Figure: Bar chart showing project distribution across scientific fields, demonstrating EPICURE's multi-disciplinary reach.

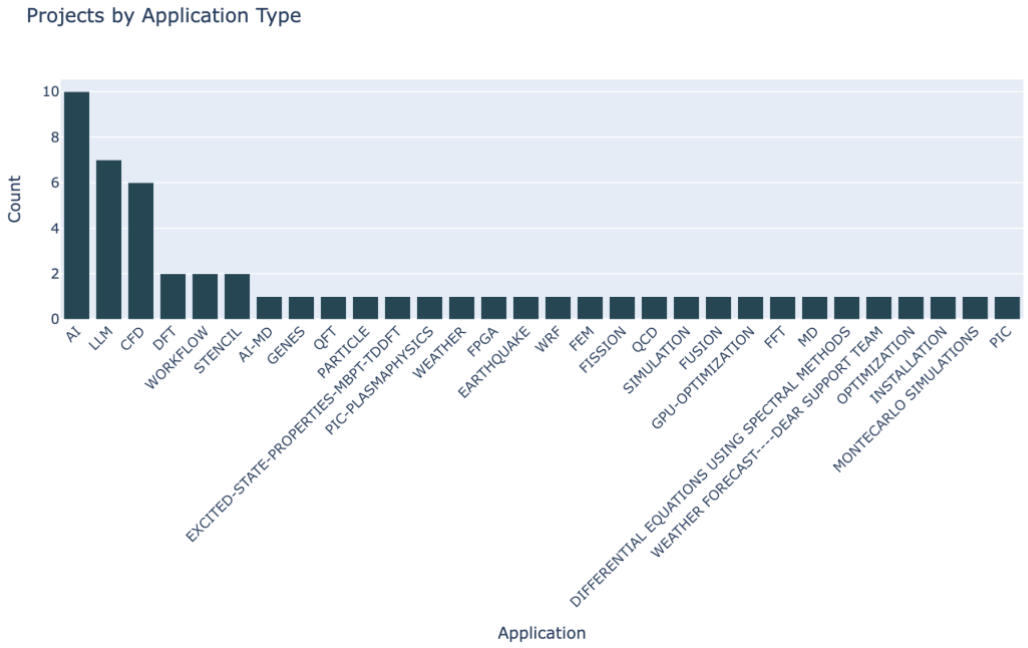
Co-funded by the European Union
EuroHPC Joint Undertaking

Application Types:

The variety of application types supported reflects the breadth of computational methods in modern research:

Application Type	Count
AI	10
LLM	7
CFD	6
DFT	2
WORKFLOW	2
STENCIL	2
Various domain-specific	25+

The strong representation of **AI (10)** and **LLM (7)** applications confirms the growing importance of machine learning in HPC. Traditional simulation codes (CFD, DFT, molecular dynamics) remain well-represented, as do specialised applications for fusion, particle physics, weather forecasting, and quantum chemistry.



Projects by Application Type

Co-funded by the European Union
EuroHPC Joint Undertaking

Figure: Bar chart showing the distribution of application types, highlighting the diversity of computational methods supported.

2.4.10 Key Performance Indicators Summary

The following table consolidates the primary KPIs derived from the ESM data, providing a dashboard view of EPICURE's operational performance:

KPI	Value	Interpretation
Total Support Requests	91	Strong demand for services
Approval Rate	90.1%	High-quality applications
Active Projects	81	Substantial active portfolio
Total Person-Months Allocated	160.50	Significant resource investment
Total Users Supported	192	Growing user community
Median Approval Time	1.0 days	Efficient processing
Completion Rate	23.5%	Successful delivery
Feedback Response Rate	86.7%	Engaged user community
Average Helpfulness Rating	4.77/5	Excellent service quality
Average Recommendation Rating	4.83/5	Strong user loyalty
Sites Covered	8	Broad European reach
Partners Contributing	14	Collaborative network

These indicators collectively demonstrate that EPICURE has established an efficient, well-received support infrastructure that successfully connects European researchers with EuroHPC JU resources. The high approval rates, fast processing times, and excellent user feedback confirm that the operational model is working effectively, while the breadth of scientific domains and institution types shows that EPICURE serves a diverse and growing community.

Co-funded by the European Union
EuroHPC Joint Undertaking

3. Lessons Learned

Section 2 provided quantitative evidence and user feedback, which this section uses to distil the important lessons learned from EPICURE's support operations. The results are categorised into five main sections. Section 3.1 delves into user experience and satisfaction, exploring users' perceptions of the service, their most valued features, and areas where their expectations were not fully met. Using information gathered throughout a project's lifetime, Section 3.2 examines ways to improve process efficiency and identifies opportunities by locating bottlenecks and potential service gaps. The technological environment and the problems it present are described in Section 3.3. This section covers topics such as the variety of programming languages, parallel frameworks, and support types used in the project portfolio, as well as any recurring technical obstacles. Reflecting on cross-centre collaboration, Section 3.4 examines how EPICURE's dispersed approach has enabled partner institutions to share knowledge and collaborate. In conclusion, Section 3.5 offers insights into the involvement of industry and SMEs, highlighting the unique difficulties and potential benefits users face in the private sector.

Keep in mind that the overall number of projects may exceed the counts given above for programming languages, frameworks, or domains combined, simply because many projects employ multiple technologies at once.

3.1 User experience and satisfaction

Analysis of user feedback collected through the EPICURE Support Manager indicates consistent patterns in how users perceive and value application support services. Of the 15 projects for which feedback was solicited, 13 submitted responses (86.7% response rate), providing a substantial basis for perceiving user experience. Satisfaction ratings are consistently high: staff availability received an average score of 4.67 out of 5, helpfulness 4.77, and likelihood to recommend 4.83, with median scores of 5.0 for all three metrics. Project outcome analysis verifies effectiveness, with 92.3% of projects meeting their goals either fully or in part.

Support requests and user needs. The main reasons users seek EPICURE assistance are performance improvement and code optimisation. Performance enhancement was the most common goal, pursued by 46.2% of projects. Code porting (38.5%) and GPU-based acceleration (30.8%) were the next most frequent requests. Intra-node and inter-node optimisation were recognised as specialised requirements, reported by 23.1% and 15.4% of projects respectively.

Responsiveness and technical expertise. The most frequently cited positive aspect is the responsiveness of support teams combined with deep technical expertise. Users described the support as *"highly effective in multiple aspects,"* with *"technical expertise and responsiveness"* that proved *"invaluable, supplying clear guidance on distributed training, memory optimisation, and*

Co-funded by the European Union
EuroHPC Joint Undertaking

debugging complex GPU-related issues." One user specifically noted *"fast feedback, in-depth discussions of custom issues,"* while another observed that *"the support team demonstrated professionalism and rapidly adjusted to the project needs."*

Practical workflow impact. Users consistently report that support delivered concrete enhancements to their research workflows. In one case, a project reported that *"performance of the code was boosted by a factor of five or six."* Another user described receiving *"customised scripts along with a modified container"* that *"greatly simplified multi-node, multi-GPU execution."* The provision of ready-to-use workflow templates proved notably useful: one user reported that after receiving an example workflow for distributed training, *"in approximately one working day, the scripts were ready for distributed training."* Users explicitly acknowledged receiving capabilities beyond what they could develop independently, with one stating that the supplied script *"is something I'd never do myself."*

Knowledge transfer and skill building. Aside from immediate fixes, users valued the systematic approach to problem diagnosis and the educational dimension of support interactions. One user appreciated that the support team used *"a methodical, step-by-step approach to problem solving"* that *"helped stabilise training workflows and improve overall reliability."* Another highlighted that the *"expert helps with diagnosis and setting tools, scripts and optimised inputs, as well as data collection"* and that the *"IT expert teaches how to deal with problems and where to look for solutions."* This knowledge-transfer dimension—building user capability rather than merely resolving tickets—was explicitly valued for its contribution to longer-term research independence. Users view EPICURE as a partner in improving research skills, not merely a service provider.

Service continuity. Several users highlighted the continuum across different support levels. One user specifically praised *"the transition between an initial support and the extended support,"* indicating that the escalation pathway from standard centre support to specialised application support functions works effectively.

Points for development. While overall satisfaction is high, the feedback reveals areas where user expectations were not fully met, providing meaningful insights for service improvement. The most notable problem relates to engagement timing: one user stated, *"only a very short period for interaction"* because *"the first contact was made quite late in the project period,"* suggesting earlier engagement has the potential to boost outcomes.

A small number of users reported disappointment when performance outcomes fell short of expectations. One user candidly reported that *"the gain is much less than we had hoped and anticipated"* and described the outcome as *"a disappointment for sure."* Another noted that while multi-node training was technically possible, *"the performance effectiveness was very low, with only a marginal increase in training speed despite a large increase in resources,"* reflecting inherent constraints in existing distributed computing technologies rather than service quality. These experiences highlight the importance of realistic expectation-setting at the project outset, particularly for complex optimisation challenges where outcomes depend heavily on code quality and algorithmic characteristics.

Co-funded by the European Union
EuroHPC Joint Undertaking

Evolving expectations. User feedback indicates that expectations are evolving. Users want more advanced distributed computing optimisation, long-term collaborative relationships rather than just project-specific help, detailed documentation and learning materials, and expert knowledge in advanced GPU programming and the optimisation of large systems. The challenges identified—scaling optimisation, project timing, and advanced technical support—point to areas for future service development.

3.2 Process efficiency and improvement opportunities

An examination of the EPICURE workflow has shown a well-oiled machine with room for development. There were 91 support requests total; 82 (90.1% approval rate), 6 (6.6% rejection rate), and 2 (2.2% archival) were denied. The low number of rejections suggests effective pre-screening or a clear explanation of eligibility criteria, while the high approval rate indicates that most requests align with EPICURE's capabilities and scope.

The approval process is efficient. With an average of 3.1 days from submitting a support request to creating a project, the time is a mere 1 day. Within 3 days, 66% of requests are approved, and 41% receive the green light on the spot. Some extreme cases, meanwhile, take 38 days or more, which is probably indicative of how complicated they are and how much more review or collaboration with different partners is required.

An examination of the project lifecycle reveals a robust portfolio. Strong throughput is observed across the 81 ongoing projects, with 39 (48.1%) underway and 19 (23.5%) completed. Among those in a less stable position, 6 (7.5%) are on hold, 5 (6.2%) are in the PI contacted stage, and 2 (2.5%) are waiting for feedback. One of the fundamental challenges of research support services is dealing with main investigators who become unresponsive; in three cases (3.8%), no response was received.

Effective follow-through is demonstrated by help request management. Eight out of nine logged assistance requests have been closed, leaving just one active; this indicates that follow-up support issues are promptly addressed.

Thirteen replies out of fifteen requests for feedback indicate a high level of participation (86.7% response rate). With such a high rate, we can be sure that the data on user satisfaction is representative.

While the overall process performs well, several areas warrant attention:

- **Outlier approval times**: A small number of requests take significantly longer than the median, potentially creating frustration for applicants. Earlier identification of complex cases could help set expectations.

Co-funded by the European Union
EuroHPC Joint Undertaking

- **Archived without reply**: The 3.8% of projects archived due to unresponsive PIs represents lost effort. Earlier and more frequent check-ins might reduce this rate.
- **Transitional bottlenecks**: Projects in “Wait for Feedback” or “PI Contacted” states may benefit from clearer timelines or escalation procedures to maintain momentum.

Project analysis also reveals potential service gaps that EPICURE could address:

- **Long-term engagement model**: The current model concludes support when the computing allocation ends, regardless of whether all objectives have been achieved. This creates a hard boundary that may leave users without continued guidance precisely when they have gained enough understanding to make further progress.
- **Standardised tooling across sites**: Different module systems (Lmod configurations), profilers, and software stacks across the 8 EuroHPC JU sites create friction for both users and support staff. Initiatives like the EESSI project are working to address this by developing a common software stack that would ease multi-site projects.
- **Container and cloud-native workflow support**: AI/ML users increasingly rely on containerised deployments and cloud-native tools (e.g., WandB, Ray) that conflict with traditional HPC security and software distribution models. Emerging initiatives such as AI Factories are expected to provide better support for these workflows.
- **Expectation management for complex optimisation**: Some users report disappointment when performance gains fall short of expectations. In many cases, achieving better performance requires code refactoring or selecting appropriate resources rather than simply adding more compute capacity—the assumption that more resources automatically translate into faster performance is often incorrect. Addressing these issues typically requires longer interaction and deeper immersion with the project teams than current allocations allow. Improved scoping processes that set realistic targets based on code quality and problem characteristics could enhance satisfaction.

3.3 Technical landscape and challenges

The EPICURE project portfolio reveals a technically diverse landscape that places significant demands on support teams. Analysis of programming languages shows Python as the dominant language (used in 67 projects), followed by C++ (44 projects), Fortran (40 projects), and CUDA (21 projects). This mix reflects the dual nature of EPICURE’s user base: AI/ML practitioners working primarily in Python alongside traditional scientific computing users maintaining legacy Fortran and C++ codebases.

There is a wide variety of approaches to parallel computing. Most projects require both GPU acceleration and distributed computing, with CUDA leading with 66 projects and MPI following closely with 65. The growing demand for portable GPU programming across vendor architectures is

Co-funded by the European Union
EuroHPC Joint Undertaking

evidenced by the increasing number of projects using OpenMP (29), OpenACC (15), and upcoming technologies such as HIP (10), SYCL (9), and ROCm (7).

Distinct patterns are revealed by analysing language-framework pairings. Given the prevalence of GPU-accelerated AI/ML workloads, the most popular Python and CUDA combination is 35 projects. Many projects still use traditional HPC combinations, such as C++ with CUDA (23 projects), Fortran with MPI (21 projects), or C++ with MPI. Hybrid parallelisation using MPI and OpenMP or OpenACC is common in Fortran/C++ programs (12 projects). Support staff need to be well-versed in a range of programming paradigms to accommodate diverse technical stacks.

Support request analysis confirms that performance-related work dominates the workload. Performance improvement is the most requested support type (50 requests), followed by installation and compilation help (38), GPU support (38), GPU optimisation (32), and performance benchmarking (31). Inter-node optimisation (22 requests) and intra-node optimisation (15 requests) represent more specialised scaling work.

The projects vary greatly in nature and difficulty, but can be categorised as follows:

- **Porting work:** CPU code to GPU, or code to different machines (different GPU/CPU), compilers, or libraries
- **Scaling work:** applications across nodes or cores
- **Performance work:** improving code efficiency
- **Onboarding work:** guidance on HPC usage

Technical domains span AI/ML (21 projects combining AI and LLM categories, typically involving PyTorch training) and traditional scientific applications across engineering (14), life sciences (10), fundamental matter physics (10), mathematics and computer science (10), and earth sciences (9).

Key technical lessons emerge from the portfolio. Despite announced vendor compatibility between compilers (especially for Fortran source code) and libraries (e.g., MPI implementations), moving existing code to another software environment or processor is not trivial. Several project members highlight the usefulness of a standardised approach. This is further compounded by the different module systems (Lmod) used across HPC centres. The diversity of custom code sometimes exposes errors in the HPC centre's software configuration—a pattern that can be interpreted as both a problem and an opportunity, as it surfaces latent issues.

Scaling and performance tuning are activities of unpredictable difficulty. At best, parameter adjustment can suffice to solve performance problems. However, sometimes source code modifications are required, such as changing libraries, moving from runtime to compile-time operations, fixing bugs, improving code quality, or even selecting better algorithms.

AI and ML projects differ significantly from scientific applications. Often, little programming effort is required, but rather high-level Python bug fixing or changes to data preparation approaches. There is very little custom code related to AI training. However, diagnosing root causes for performance or

Co-funded by the European Union
EuroHPC Joint Undertaking

scalability issues can be difficult, as it involves communication analysis and GPU memory usage investigation. Such troubleshooting requires a good understanding of HPC architecture, and the profilers used at each centre (which vary across centres). A documented best-practice approach could benefit future project members.

The project reports on the EPICURE website contain valuable technical information. Exploiting this information (making it searchable and extractable) would encourage project members to report more technical details. From a broader perspective, the complexity of current software development—with its many dependencies, wrappers, and Python JITs—exposes hidden costs. Among programmer communities, those in scientific computing are best placed to adopt simpler, more specific codes with fewer dependencies. EPICURE could provide programming guidelines to raise awareness of the benefits of simplicity.

3.4 Cross-centre collaboration

EPICURE's distributed model spans 14 partner institutions across 8 EuroHPC JU supercomputing sites. This geographic distribution presents coordination challenges but has also enabled unique opportunities for cross-centre collaboration and knowledge sharing.

By establishing a common view across all partners, the EPICURE Support Management (ESM) system has built a solid foundation for collaboration. To ensure all team members use the same data source, ESM consolidates project details, effort tracking, and communication logs. When a project calls for knowledge from many centres, partners may easily access the project history and past interactions, thanks to this transparency, which has proved essential for coordinating support among institutions.

Partners have been able to work more easily together on technical difficulties and solutions thanks to the biweekly sessions. Regular meetings provide an opportunity to share ideas, tackle challenging projects, and identify portfolio trends. These conversations allow partners to develop joint methods rather than tackling the same technical problems separately when they arise in other areas, such as the ongoing problems with the AI training pipeline or GPU porting described in Section 3.3. A culture of cooperative problem-solving has emerged during the meetings, which goes beyond the confines of individual projects.

Faster, more targeted communication is now possible thanks to direct messaging within teams. Direct communication channels allow partners to quickly seek advice from colleagues with specialised knowledge, ask for guidance on unusual technical issues, or facilitate transfers between sites, while ESM serves as the formal documentation of project interactions. For complex or time-sensitive issues that require quick discussion, this informal communication layer has proven a great asset, improving structured ESM operations.

Co-funded by the European Union
EuroHPC Joint Undertaking

Work is distributed among locations in accordance with project needs and partner capabilities. Most proposals are from Leonardo (26, or 25.5% of the total) and MareNostrum 5 (22, or 21.6% of the total), with LUMI coming in third with 19 projects (18.6%). Karolina, MeluXina, Vega, Deucalion, and Discoverer are some of the minor platforms that improve regional coverage and provide specialised features. By distributing tasks in this way, EPICURE ensures that partners can keep their workloads manageable and that projects are assigned to those with the best talents and resources.

Inter-centre collaboration has been a prominent aspect of EPICURE. Among the 81 projects in the portfolio, 16 (20%) involve multiple partners, and 8 (10%) span multiple supercomputing locations. Certain projects have utilised assistance from as many as four partners across all eight EuroHPC JU locations, showcasing the full capabilities of the distributed paradigm. These alliances facilitate knowledge exchange among centres and guarantee that initiatives leverage the consortium's collective expertise.

3.5 Industry and SME engagement

The summary is based on a selection of EPICURE projects: those submitted by SMEs or private entities, completed, in progress, or on-hold (because a project might have started and later been placed on-hold), and for which some effort was reported (to filter out projects without any activity). Overall, only 8 projects matched those criteria.

Although there were few projects in this category, the following lessons can be learned.

Most projects (6 of 8) related to AI, specifically to AI training (or fine-tuning). The main AI training activity focused on the execution of AI training pipelines in the HPC environment, with scalability as the objective. These efforts proved very similar yet challenging: the training pipelines need to be adapted to the Slurm environment and to the internet connectivity restrictions typically in place in HPC. The latter is a problem because some common tools for AI training require a connection to public internet servers (WandB). Also, a recurrent need was the execution of distributed inference using vLLM. This was often accomplished using Ray, but not without tedious setup. This is surprising given the high-level of activity around LLM inference (of large models, such as 405B parameters).

The HPC's Slurm and software distribution mechanisms are often at odds with the containerised deployment of AI software, which appears fragile. Only 2 such projects required just training, with the PI then fully autonomous. Several project teams felt the need to document their efforts, which is a great initiative. A quick win would be collecting these guides and distributing them for future AI projects. This effort would be quite efficient as the software stack required for training is standard (PyTorch, with extensions such as Torchrun, Hugging Face Accelerate).

The two remaining, non-AI, projects were scientific applications that required either installing (compiling) specific software libraries or porting existing code to the GPU. It should be noted that

Co-funded by the European Union
EuroHPC Joint Undertaking

these projects, while officially from private entities, had strong ties to the academic community, typically as spin-offs from research departments. These projects were a better match for the HPC environment and for the EPICURE team members' skill set.

It is also worth noting that the private entities were often SMEs, even startups. Some of the SMEs (2) had even disappeared during the year.

Most PIs were very satisfied with the assistance provided, with one project leading to another. However, the PI's responsiveness varied greatly, with 3 projects losing involvement during the project. This is to be expected of young, small companies that need to pivot early. The AI projects can be seen as representative of the AI business landscape's immaturity.

Co-funded by the European Union
EuroHPC Joint Undertaking

4. Feedback and Recommendations to EuroHPC JU

European researchers cannot fully utilise EuroHPC JU resources without specialised application support services, as shown by the EPICURE project. Not as a continuation of any project, but as an acknowledgement that such services should be set up as a permanent, accessible offering for the wider European research community, this section offers policy suggestions for the EuroHPC JU to consider as it shapes the future of European HPC application support.

4.1 Sustaining application support services

The European HPC ecosystem would benefit greatly from application support services that are both permanent and widely available. European HPC users require continuous access to professional support, in line with demand patterns that span academics, SMEs, and industrial users across a variety of scientific fields. Maintaining application support for the largest potential audience should be a priority for future program design.

There should be support services that can handle various interaction models. Shorter, milestone-based engagements may be more beneficial for SMEs and startups, but deeper, longer-term collaboration is usually required for complex optimisation work. More effective results across user communities might be achieved with more leeway in interaction time and structure.

4.2 Strategic infrastructure investments

The European HPC ecosystem might benefit from several infrastructure developments.

It should remain a top priority to harmonise software environments. With its recent incorporation into the EuroHPC Federation Platform, the EESSI project has taken a giant leap forward. If EuroHPC JU continued to invest in common software stacks and portable environments, users would find it easier to transfer between facilities.

The benefits of implementing a knowledge management strategy at the European level would be substantial. Now, each centre is developing its own best practices and troubleshooting expertise. All European HPC users would benefit from a concerted effort to gather and distribute this expertise across the ecosystem.

Co-funded by the European Union
EuroHPC Joint Undertaking

Modernising legacy code should be given dedicated assistance. It would be beneficial to migrate to new standards (Fortran 2018+, C++17/20) or reimplement in more maintainable languages, as a large portion of the scientific computing base in Europe is based on outdated Fortran and C codebases. Code portability, maintainability, and the capacity to take advantage of new hardware features would all be enhanced by such a modernisation.

The AI Factories project offers an opportunity to resolve the conflict between conventional HPC security models and AI operations native to the cloud. The development of user services for AI Factory should be informed by operational experiences learned from application support efforts. The entrance hurdles for this quickly growing user base might be reduced with investment in standardised onboarding services for AI workloads.

4.3 European HPC community coordination

Centres of Excellence, AI Factories, infrastructure projects, national competence centres, and application support teams are now all part of the EuroHPC JU ecosystem. Various coordination methods are now in place. A network of 33 National Competence Centres (NCCs) has been set up across Europe as part of the EuroCC project. These NCCs provide HPC services tailored to each country's needs. A single gateway to European supercomputing resources, the HPC in Europe Portal was developed by CASTIEL, which also acts as a coordinator and support activity supporting contact between NCCs. EVITA, the EuroHPC Virtual Training Academy, is developing a European-wide training framework.

Coordination across these efforts' respective domains is highly beneficial. On the other hand, there is still a lack of coordination across the larger ecosystem linking NCCs, AI Factories, Centres of Excellence, and application support teams. To enhance current frameworks, a forum uniting all projects and actions supported by EuroHPC JU could facilitate cross-domain discussion of shared technical problems, the exchange of solutions and best practices, and the discovery of gaps between current efforts.

The CASTIEL coordination model and the HPC in Europe Portal infrastructure could be used to create such a forum, encompassing all initiatives sponsored by EuroHPC JU. Connecting projects that are running in parallel is more important than duplicating existing coordination.

To better facilitate benchmarking, best practice identification, and evidence-based policy formation at the European level, standardised metrics reporting across application support activities is necessary.

Co-funded by the European Union
EuroHPC Joint Undertaking

5. Conclusions

5.1 Summary of key lessons

The EPICURE project has generated key lessons about the provision of application support services at the European level. User satisfaction ratings demonstrate that the service model works: users appreciate rapid responses, technical expertise, and the depth of engagement that specialised support teams can provide. At the same time, the experience has revealed important lessons for the future.

Users approach EuroHPC JU resources with a range of necessities and varying levels of HPC expertise. Performance optimisation remains the dominant requirement, but the growing importance of AI workloads—particularly among SME users—demands new competencies and different engagement models. Non-expert users gain considerable advantage from guided onboarding and realistic expectation-setting at the outset of projects.

Opportunities and challenges coexist in the technical landscape. The variety of scientific areas, programming languages, and parallel frameworks that EPICURE supports highlights the extensive skill set required. There is a great chance to fortify Europe's scientific computing foundation through legacy code modernization. Complex projects benefit greatly from cross-centre collaboration; there is room for improvement in the coordination procedures.

SME and industry involvement require specialised strategies that account for the intrinsic volatility in the startup ecosystem and the specific constraints these users face regarding IP governance and access pathways.

The most important lessons for the evolution of European HPC application support are:

1. Application support services deliver considerable value and should be established as a permanent component of the EuroHPC JU ecosystem
2. Early engagement and realistic expectation-setting improve project outcomes
3. Flexible engagement models are needed to serve assorted user groups efficiently
4. Cross-centre collaboration improves capability but requires active coordination
5. AI workloads present special challenges that warrant dedicated onboarding resources
6. Legacy code modernisation offers major rewards for European scientific computing
7. Coordination across EuroHPC JU initiatives (NCCs, CoEs, AI Factories, application support) remains an untapped opportunity

Co-funded by the European Union
EuroHPC Joint Undertaking

5.2 Fulfilment of Milestone

Compiling lessons learnt from EPICURE's support activities and converting them into structured feedback and recommendations for EuroHPC JU is how this deliverable fulfils Milestone MS12.

The quantitative data is presented in Section 2, which covers information gathered by the EPICURE Support Manager since January 2025, including support requests, project portfolio characteristics, user feedback, and operational metrics. Part 3 examines the information to draw conclusions on the following five areas: user experience and satisfaction, process efficiency and improvement potential, technical landscape and challenges, cross-centre collaboration, and industry and SME participation.

Section 4 translates these lessons into policy recommendations for EuroHPC JU, organized around sustaining application support services, strategic infrastructure investments, broadening participation, and European HPC community coordination.

MS12 Requirement	Deliverable Coverage
Compile lessons learned from user requests	Sections 2 and 3
Compile lessons learned from support actions	Sections 2 and 3
Provide feedback to EuroHPC JU	Section 4
Provide recommendations to EuroHPC JU	Section 4

5.3 Outlook

This deliverable contains lessons that will be useful long after EPICURE has ended. They lay the groundwork for future successor projects, centre operations, or new program initiatives to build upon when it comes to supporting HPC applications in Europe.

With these takeaways, EPICURE partners can better their local services and help shape a unified European offer. Both individual centre strategies and collective European methods can be informed by the patterns observed, which revolve around user demands, technical challenges, and effective collaboration. These patterns apply across sites.

Co-funded by the European Union
EuroHPC Joint Undertaking

Future program design for EuroHPC JU can be informed by the suggestions made in Section 4. The demand patterns and user comments reported here provide credence to the case for application support services that are both permanent and widely accessible. By enhancing existing frameworks such as EuroCC, CASTIEL, and EVITA, the proposed coordination mechanisms would strengthen the entire European HPC ecosystem.

The dedicated application support greatly increases the value that European researchers get from EuroHPC JU resources, as shown by EPICURE. The lessons learned from that experience are included in this deliverable, which will aid in developing a more integrated and effective support ecosystem, ultimately benefiting long-term research in Europe.

Co-funded by the European Union
EuroHPC Joint Undertaking