

Academic, public administration and industrial end-user requirements in high-performance computing in the European Union

Final report



EuroHPC
Joint Undertaking

The User Study (“EuroHPC/2023/034/OP - Study on HPC academic, public administration and industrial end-user requirements in the European Union”) is funded by the EuroHPC Joint Undertaking under contract number EuroHPC/SC/03/2023/ADMIN. The contract was carried out by Blue Ocean Sustainable Solutions LDA, Rua do Norte, N16, 2655-321 Ericeira, Portugal.

Disclaimer:

The information and views set out in this report are proposed by the author(s) of the study (Blue Ocean Sustainable Solutions LDA) and do not necessarily reflect the official position of the EuroHPC JU. EuroHPC JU does not guarantee the accuracy of the data included in this report. EuroHPC JU, nor any person acting on their behalf, may not be held responsible for the use which may be made of the information contained therein. This report does not involve EuroHPC JU in liability of any kind.

The information contained in this document represents the research of Blue Ocean Sustainable Solutions LDA as of the date published. Blue Ocean Sustainable Solutions LDA does not guarantee that any information contained herein is error-free, or fully up to date. Blue Ocean Sustainable Solutions LDA makes no warranties, express, implied, or statutory, by publishing this document. The information and guidance provided in this document should not be considered legal advice.

This document is not intended to replace professional legal advice.

Manuscript completed in February, 2025

Luxembourg: Publications Office of the European Union, 2026

© European High Performance Computing Joint Undertaking, 2026



Unless otherwise noted, the reuse of this document is authorised under the Creative Commons Attribution 4.0 International (CC BY 4.0) licence (<https://creativecommons.org/licenses/by/4.0/>). This means that reuse is allowed, provided appropriate credit is given and any changes are indicated.

For any use or reproduction of elements that are not owned by EuroHPC Joint Undertaking, permission may need to be sought directly from the respective rightholders. The European High Performance Computing Joint Undertaking does not own the copyright in relation to the following elements:

- cover photo, source: © Sashkin/stock.adobe.com
- page 96, photo, source: © Sashkin/stock.adobe.com

PRINT
PDF

ISBN 978-92-95228-10-8
ISBN 978-92-95228-09-2

doi:10.2793/4886019
doi:10.2793/4929863

LL-01-25-003-EN-C
LL-01-25-003-EN-N

Academic, public administration and industrial end-user requirements in high-performance computing in the European Union

Final report

Prepared by:



For:



EuroHPC
Joint Undertaking

Document information	
Client	EuroHPC Joint Undertaking
Document type	Report
Document internal reference	BO-PRO-23-020-FinalStudy
Version	1
Recommended citation	Almeida, D., Gomes, T., Ramos, C., Kozubek, T., Martinovic, J. et al. (2024), Study on high-performance computing academic, public administration and industrial end-user requirements in the European Union, Euro HPC Report [reference], blueOASIS and IT4Innovations National Supercomputing Centre

Contents

Abbreviations	6
Glossary	7
01 Introduction	8
02 Methodology for the study	12
03 Geographic coverage and high-performance computing resources around the world	17
04 High-performance computing case studies	34
05 High-performance computing technologies, ecosystems and trends	47
06 Legal aspects	63
07 Financial aspects	67
08 Environmental aspects	70
09 Socioeconomic benefits	73
10 Final remarks: insights, challenges and future directions	81
Endnotes	85
Appendices	91

Figures

Figure 1.	The participating countries in the June 2024 TOP500 list and their share of systems.....	19
Figure 2.	Geographical distribution of the stakeholders interviewed.....	24
Figure 3.	Characterisation of the participants based on their role within organisations, funding source and area.....	25
Figure 4.	Geographical distribution of the stakeholders from the web-based questionnaire.....	28
Figure 5.	Characterisation of the participants based on demographics, stakeholder group, sector, type of organisation, funding source and revenue.....	29
Figure 6.	Representativeness of different sectors according to the type of organisation.....	31
Figure 7.	Responses from HPC providers based on the type of funding their organisation receives.....	32
Figure 8.	Responses of users and providers regarding which sectors would benefit from specialised HPC services.....	34
Figure 9.	Opinions on whether services should be specialised, per stakeholder group.....	40
Figure 10.	Marketing strategies used to attract new users to HPC, from the HPC providers' perspective.....	46
Figure 11.	HPC environments.....	50
Figure 12.	Summary of the HPC resources used and acquired by the users.....	51
Figure 13.	Different types of interfaces provided by the HPC providers.....	54
Figure 14.	Emerging technologies that could make HPC more accessible.....	58
Figure 15.	Types of hardware expected by participants to enhance competitiveness in future HPC challenges, by organisation type.....	59
Figure 16.	Responses given by the developers on which sectors will benefit from quantum computing.....	60
Figure 17.	Examples given by the developers of how quantum computing might benefit different sectors.....	61
Figure 18.	Number of responses given by the developers on whether exascale computing will be widely used in the future, by organisation type.....	63
Figure 19.	Legal obligations to consider when submitting data or workloads, by stakeholder group.....	66
Figure 20.	Pricing models for using HPC services.....	68
Figure 21.	Financial income from providing HPC services, for public and private organisations.....	70
Figure 22.	The three most energy-efficient supercomputers.....	72
Figure 23.	Significant achievements made possible by HPC applications.....	76

Tables

Table 1.	Topics addressed during the stakeholder consultations	15
Table 2.	Stakeholder engagement during the interview period	16
Table 3.	Summary of the considerations expressed by users from different organisations regarding access to HPC resources	42
Table 4	Reasons cited by users for having or not having their own HPC cluster	52
Table 5.	Summary of users' technical considerations	54
Table 6.	Opinions about quantum computing	60
Table 7.	Summary of the legal requirements for accessing HPC resources as stated by users	66
Table 8.	Financial return on investment from HPC, by select industry	75
Table 9.	Strategies for improving access to HPC resources in regions with limited availability	80
Table 10.	Categories definition for the study analysis	112

Abbreviations

AI	artificial intelligence
AMD	Advanced Micro Devices
API	application programming interface
ASIC	application-specific integrated circuit
AUP	acceptable use policy
AWS	Amazon Web Services
BSC	Barcelona Supercomputing Center
COE	centre of excellence
CPU	central processing unit
EU	European Union
EuroHPC Joint Undertaking	European High-Performance Computing Joint Undertaking
flops	floating-point operations per second
FPGA	field programmable gate array
FTP	File Transfer Protocol
GDPR	General Data Protection Regulation
GPU	graphical processing unit
GUI	graphical user interface
HBM	high-bandwidth memory
HPC	high-performance computing
HPCaaS	high-performance computing as a service
IP	intellectual property
IT	information technology
ICT	information and communication technology
LLNL	Lawrence Livermore National Laboratory
LLM	large language model
MPI	message-passing interface
MW	megawatt
NGO	non-governmental organisation
RAM	random-access memory
RDP	Remote Desktop Protocol
R & D	research and development
SFTP	Secure File Transfer Protocol
SLA	service-level agreement
SME	small to medium-sized enterprise
SSH	Secure Shell (Protocol)
TACC	Texas Advanced Computing Center
VPN	virtual private network

Glossary

Cloud-based HPC

Cloud-based HPC utilises cloud computing resources for intensive computational tasks, offering scalability and flexibility without the need for on-premises infrastructure.

Grid computing

Distributed computing model that aggregates computational resources from multiple locations to solve large-scale problems, often used for scientific research, engineering simulations and data analysis.

High-performance edge computing

Extension of HPC capabilities to edge devices or locations, allowing for real-time processing of data generated at or near the point of data collection, reducing latency and enabling time-sensitive applications.

HPC as a service

HPC as a service is a cloud computing model that provides on-demand access to HPC resources, such as clusters of servers and specialised hardware accelerators, over the internet.

Hybrid HPC

Utilisation of both on-premises and cloud-based HPC resources in a combined infrastructure, allowing organisations to dynamically allocate workloads based on demand, cost and performance requirements.

On-premises HPC facilities

HPC infrastructure located within an organisation's premises, typically comprising interconnected servers or nodes optimised for parallel processing.

Quantum computing

Computing paradigm that applies the principles of quantum mechanics to perform computations using quantum bits, offering the potential for exponential speedup over classical computing for certain types of problems, such as cryptography, optimisation and simulation.

Shared HPC clusters

HPC clusters that are shared among multiple users or organisations, typically managed by a centralised entity such as a centre of excellence to maximise resource utilisation and efficiency.

01

Introduction

1.1 Project background

The present study is the final deliverable of the work contracted by the European High-Performance Computing Joint Undertaking (EuroHPC Joint Undertaking) under the call for tender with reference EuroHPC/SC/03/2023/ADMIN. The EuroHPC Joint Undertaking is a public-private partnership in the high-performance computing (HPC) sector established under Article 187 of the Treaty on the Functioning of the European Union by Council Regulation (EU) [2021/11732](#). As provided in Article 3.1 of the regulation, the mission of the joint undertaking will be to:

- develop, deploy, extend and maintain in the European Union a world-leading federated, secure and hyper-connected supercomputing, quantum computing, service and data infrastructure ecosystem;
- support the development and uptake of demand-oriented and user-driven innovative and competitive supercomputing systems based on a supply chain that will ensure components, technologies and knowledge limiting the risk of disruptions and the development of a wide range of applications optimised for these systems;
- widen the use of that supercomputing infrastructure to a large number of public and private users and support the twin transition and the development of key skills for European science and industry.

1.2 Scope of work

The main goal of the study is to understand what specific technical, legal, financial, commercial and socioeconomic considerations shape academic, public administration and industrial end-user requirements and how these requirements influence the business decisions of HPC users to access HPC solutions and services. The analysis considers the requirements in different industrial and scientific sectors, such as computer science, climate science, healthcare, government and defence, information technology (IT) and telecoms, physics, mobility, energy and maritime engineering.

In terms of geographic areas, the study investigated the state-of-the-art HPC systems, solutions and services currently available to end users in Europe and non-European countries such as Australia, Canada and the United States.

1.3 Executive summary

This report evaluates the current state of HPC resources, aiming to assess their availability, understand decision-making factors and identify challenges users face in accessing HPC services. The findings reveal both significant advancements and notable gaps within the HPC landscape.

HPC Environments	Challenges	Solutions
<p>HPC plays a crucial role in both academic and industrial sectors</p> <p>IT, telecoms, smart cities and robotics are increasingly adopting HPC alongside traditional fields such as weather forecasting, healthcare and engineering</p> <p>On-premises facilities, physical infrastructures and centre of excellence (COE) clusters are most widely used</p> <p>Cloud-based HPC services are an attractive option owing to scalability and flexibility</p> <p>Growing interest in quantum computing and hybrid HPC environments that combine on-site and cloud resources.</p>	<p>Concerns about data transfer, security and storage</p> <p>Budgetary constraints impact users despite the provision of free public access to academic and research and development (R & D) users</p> <p>A need to bridge gaps in knowledge and expertise and keep pace with rapid advancements</p> <p>Strict legal and security measures required for those in sensitive sectors such as healthcare and biosciences</p> <p>Sustainability concerns require alignment of software practices with environmental objectives.</p>	<p>Developing collaborative partnerships with industry to increase dialogue and access pathways</p> <p>Tailoring solutions for industry users including extension of competencies training, reduction of technical barriers and knowledge exchange</p> <p>Addressing financial shortfalls by creating flexible payment models for small to medium-sized enterprises (SMEs) and other smaller organisations and institutions</p> <p>Improved and secure technical infrastructure that addresses the concerns of private users.</p>

HPC environments and utilisation.

HPC is crucial across academic and industrial sectors, with traditional fields such as weather forecasting, healthcare and engineering continuing to leverage its capabilities. Emerging sectors like IT, smart cities and robotics are increasingly adopting HPC to drive innovation. On-premises facilities and COE clusters remain the most widely utilised resources, although cloud-based services are gaining traction due to their flexibility and scalability. The interest in hybrid environments and quantum computing highlights a shift towards more advanced computational frameworks. Users express concerns regarding computational power, network connectivity and data security, indicating a need for robust infrastructure to support intensive workloads.

Financial considerations.

Financial considerations are a main challenge for organisations that use HPC services to support their activities and find it difficult to establish or maintain their own in-house HPC infrastructure. Most users allocate less than 10 % of their budgets to private HPC resources, with a primary focus on hardware expenditures. Public providers offer free access for academic users, while private users access HPC services that employ various pricing models, reflecting the economic constraints they face when choosing an HPC provider. Notably, many public HPC service providers report revenues of less than EUR 0.5 million, highlighting a need for sustainable funding models to enhance better service provision.

Security and legal considerations.

Legal and security issues remain paramount, particularly in sensitive fields like healthcare. Users prioritise data privacy, compliance and security measures, while HPC providers focus on protecting sensitive data from breaches. The necessity for service-level agreements (SLAs) and non-disclosure agreements between HPC providers and users underscores the complexity of managing legal obligations in HPC environments.

Environmental considerations.

There is a growing emphasis on sustainability within the HPC community. While users recognise the importance of energy efficiency, it often takes a back seat to immediate computational needs. HPC providers are implementing energy-efficient technologies and exploring renewable energy sources to reduce their carbon footprint. However, a gap in communication between users and providers exists, indicating a need for collaborative efforts to align operational practices with sustainability goals.

Socioeconomic benefits and accessibility challenges.

The integration of HPC has significantly influenced the socioeconomic landscape, fostering advancements in healthcare, climate science and mobility, among others. However, challenges remain, particularly in improving access for regions and subject areas with limited availability of HPC resources, such as developing countries. Barriers such as network latency and funding constraints highlight the need for strategic investments in infrastructure and training. Stakeholders agree on the importance of public investment and sustainable funding models to enhance HPC accessibility and maximise its socioeconomic impact.

Conclusion.

As HPC continues to evolve, its ability to drive job creation, innovation and economic growth will depend on addressing accessibility challenges and fostering collaborative partnerships between public and private sector entities. This report underscores HPC's critical role in R & D while highlighting the need for enhanced communication and cooperation between users and providers. This report does not review the impact artificial intelligence (AI) will have on the sector but stipulates that ongoing research is necessary to keep abreast of this rapidly advancing technological landscape.

1.4 Document structure

Introduction

This section will provide an overview of the project, including its objectives and scope.

Methodology for the study

This section will offer an explanation of the methodologies implemented in the literature review and stakeholder consultations, including an analysis of the interviews and the web-based questionnaire.

Geographic coverage and resources around the world

In this section, the focus will be on the global impact of HPC. It will cover the TOP500 list, highlighting the world's most powerful supercomputers and their performance metrics. The discussion will include HPC contributions from European countries, along with China, the United States and other countries, emphasising their major facilities and technological advancements.

High-performance computing case studies

Three case studies with methodologies similar to those of this report will be presented. The results from the stakeholder consultations will be discussed and aligned with the findings of the case studies.

High-performance computing technologies, ecosystems and trends

This section will explore key aspects of HPC, including the anatomy of HPC systems, relevant software and various HPC environments. It will also highlight emerging technologies like exascale and quantum computing that are shaping the future of HPC.

Legal aspects

This section will examine the legal and ethical issues in HPC and AI, including data protection and compliance with the General Data Protection Regulation (GDPR).

Financial aspects

This section will offer a review of the financial models for HPC, including pay-per-use, subscription, auction-based and advertising-based pricing. It will highlight how these models address diverse user needs and support sustainability, with academic institutions often providing resources through grants.

Environmental aspects

This section will cover the sustainability aspect of HPC, focusing on energy efficiency and environmental impacts. It will examine the challenges of high energy consumption in HPC systems and highlight the Green500 list's role in promoting energy-efficient supercomputers.

Socioeconomic benefits

The role of HPC in advancing social and economic development will be discussed. The section will highlight HPC's impact on decision-making, innovation and various sectors, including healthcare and climate science, along with its economic benefits and job creation potential.

Final remarks: insights, challenges and future directions

A summary of the main results from the stakeholder consultations – including insights from the dedicated interviews and the web-based questionnaire – will be presented.

02

Methodology for the study

2.1 Literature review

To accomplish the objectives of this study, a literature review of the topics related to the scope of work was done, including:

- HPC services and resources available worldwide;
- user experience in accessing HPC clusters;
- financial trends of the global HPC market;
- legislation and regulations that impact HPC systems;
- sustainability considerations of the functioning of HPC systems.

A full review of scientific articles, peer-reviewed papers, conference proceedings, scientific presentations and informative bulletins was performed using Google Scholar, peer-reviewed academic journals (e.g. IEEE Xplore), industry press (e.g. HPCwire, Data Centre Dynamics and vendor reports), government reports and research organisation publications to gather the information needed for the literature review.

The main keywords used were: *HPC technologies, supercomputing architectures, supercomputing emerging technologies: HPC + [type of environment] (e.g. HPC + cloud-services, HPC + on-premises facilities), trends in HPC, investment in HPC, HPC in science, HPC geographic coverage, HPC regulations, HPC policies, energy use, sustainable data centres.*

A total of 90 documents were reviewed and analysed and subsequently stored in the Zotero platform, facilitating easy access and accurate referencing throughout this report (see Endnotes). Moreover, the scientific articles were sourced from highly cited journals with significant impact factors, covering the period from 2010 to the present. Industry press provided up-to-date studies, statistics and vendor reports.

2.2 Stakeholder consultations

Three stakeholder groups were created for the purpose of the study.

HPC providers. Entities, both public and private, that offer HPC resources and associated services. These include HPC centres, cloud computing services and other organisations that provide access to supercomputing capabilities.

Developers. Individuals or organisations that design, create and maintain software frameworks, tools and libraries specifically for HPC applications. These developers focus on improving the efficiency, scalability and usability of HPC systems.

Users. Individuals or teams from various organisations, including universities, research institutes, large companies, SMEs, start-ups and public organisations, all of which utilise supercomputing resources for their computational needs.

The objectives of the stakeholder consultations were:

- to assess the HPC resources available to users;
- to understand the factors that influence the decision-making process that impacts the provision of HPC resources;
- to identify challenges in accessing HPC resources and services.

Stakeholder consultations were performed using two main methods: direct interviews and a web-based survey. Both were conducted in English, targeting stakeholders from multiple continents or regions.

Based on the literature review, several key topics and gaps were identified related to the current state of the art in HPC and future trends to be addressed during the interviews and web-based questionnaire.

Questions for the stakeholder consultations were developed in accordance with the topics outlined in **Table 1**.

Table 1. Topics addressed during the stakeholder consultations

Area	Topics
General information and organisational context	<p>Location of the HPC entity and target countries</p> <p>Main funding source (public, private, collaborative)</p> <p>Sector of the organisation</p> <p>Ownership of an HPC system and rationale</p> <p>Industry focus and target fields of knowledge</p> <p>Specific industries where HPC services could bring significant benefits</p>
Decision-making and prioritisation factors	<p>Key factors in selecting an HPC service (knowledge area, user community, ease of scaling, resources, security)</p> <p>Influence of legal, financial and technical factors</p> <p>Preferred types of contracts and factors influencing contract choice</p> <p>Importance of specific resources (computational power, storage, performance)</p> <p>Marketing strategies that attract organisations to HPC providers</p>
Legal, compliance and security considerations	<p>Legal requirements for submitting workloads</p> <p>Relevance of legal considerations across providers/services/countries</p> <p>Agreements related to software tools and access rights</p> <p>Data protection strategies (malware protection, software updates)</p> <p>Access methods to HPC services (physical tokens, two-factor authentication, public key)</p>
Financial considerations and contractual aspects	<p>Budget allocation for HPC services</p> <p>Financial aspects prioritised by the organisation (computational power, data storage, hardware)</p> <p>Types of contracts signed (metered usage, fixed price, limited/unlimited usage)</p> <p>Research and commercial contracts</p>
Technical requirements and infrastructure	<p>Technical factors in prioritising HPC usage (software stack, workflows, security features)</p> <p>Type of HPC systems used and future preferences (central processing unit (CPU), graphical processing unit (GPU), field programmable gate array (FPGA), application-specific integrated circuit (ASIC))</p> <p>Homogeneous versus heterogeneous systems</p> <p>Storage configurations (local, parallel, object storage)</p> <p>HPC system interfaces used (command-line interface, graphical user interface (GUI), virtual private network (VPN))</p> <p>Programming environments in different fields (data science, natural science, engineering, creative software)</p>
Socioeconomic impact and collaboration	<p>Impact of HPC services on local and regional economic growth (job creation, business development)</p> <p>Collaboration with government agencies (societal challenges, public policy)</p> <p>Access to HPC systems in regions with limited infrastructure and challenges</p> <p>Engagement with the local community to broaden access to HPC benefits</p>

2.2.1 Dedicated interviews

Dedicated interviews were conducted to gather different perspectives regarding the HPC landscape. Interview questions covered current and future trends related to HPC, technical perspectives and the socioeconomic impact of HPC technology.

Interviewees were selected through online searches and from the existing network of the project team and the EuroHPC Joint Undertaking. Interviews were conducted online, via Microsoft Teams, to facilitate participation from stakeholders in different geographical locations. All interviewees participated voluntarily, agreed that their feedback could be included in this report, and were not remunerated for their time. A guideline was developed as a baseline for the interviews (see Appendix 1) to guide the conversation. Several topics were addressed through open-ended questions, allowing for a broad discussion of various topics, including technical, financial, commercial and socioeconomic aspects of the HPC field. Interviews were conducted between 7 March and 19 July 2024, and each session lasted between 30 minutes and one hour. The target groups, summarised in Table 2, mainly consisted of representatives from organisations that provide or use HPC resources, along with associations that work closely with industries. These stakeholders were then organised into the three HPC categories defined previously (**users**, **developers** and **providers**).

Table 2. Stakeholder engagement during the interview period

Stakeholders	Engagement approach
Academia researchers	Research team/ EuroHPC network, conferences, contacts from the interviews/ survey
Commercial users/ providers	Online search, research team network, contacts from the interviews/survey
Industry users	Online search, research team/EuroHPC network, platforms, contacts from the interviews/survey
HPC providers	Online search, research team network, conferences, contacts from the interviews/ survey
International associations	Online search, research team/EuroHPC network, contacts from the interviews/ survey

2.2.2 Web-based questionnaire

A web-based questionnaire was conducted, aimed at reaching a broader audience and capturing the perspective of the different stakeholders in a structured format. The questionnaire was created based on the literature review and insights from the interviews. The platform [Typeform](#) was used for distribution. The survey included different types of questions: multiple choice, ranking questions and open-ended questions with short answers. The list of questions can be found in Appendix 2. In the questionnaire, participants were given the option to share their contact details if they were willing to participate in dedicated follow-up interviews, which are described in Section 2.2.1.

The web-based questionnaire was conducted between 20 April and 19 July 2024 and was distributed through several methods to reach a wider community across continents: emailing organisations that are part of the existing network, posting on social media, reaching out to international/regional organisations (e.g. the [Portuguese Foundation for Science and Technology](#), the [Big Data Value Association](#), the [Digital Research Alliance of Canada](#)).

Questionnaire analysis

Close-ended questions were analysed using quantitative statistics (e.g. percentage of answers with the same result), and the main results were compared according to several categories (e.g. gender, education level, stakeholder type, organisation type and level of expertise (see Appendix 3)).

The responses to the open-ended questions were analysed by sorting all responses, organising the main ideas and trends and extracting information relevant to the study's objectives.

03

Geographic coverage and high-performance computing resources around the world


HPC resources have significant contributions from various regions around the world in several areas of knowledge, and are provided globally through national supercomputing centres, academic institutions and cloud-based platforms. Notable examples include the Oak Ridge National Laboratory in the United States, the RIKEN research institute in Japan and the National Supercomputing Centre in China. In Europe, the EuroHPC initiative combines resources from several countries to support advanced systems like LUMI in Finland, MareNostrum in Spain, Leonardo in Italy and others. Additionally, cloud services such as Amazon Web Services (AWS), Google Cloud and Microsoft Azure offer on-demand HPC access, enabling scalable computing solutions.

The [TOP500](#) list was first published in 1993 and is the semi-annual ranking of the world's 500 most powerful computer systems. It provides a standardised ranking of the world's most powerful supercomputers, facilitating performance comparison, promoting technological development and raising awareness of the critical role these systems play in addressing complex scientific and industrial challenges.

The ranking is determined by the performance of these systems on the Linpack benchmark, which assesses their capacity to solve complex sets of linear equations.

This benchmark offers a consistent metric for measuring and comparing the computational strength of supercomputers. The latest version of the TOP500 list was released in June 2024 and features HPC systems from around 40 countries from several continents and regions. The continents with the most systems listed are North America (39 %), Europe (31 %) and Asia (29 %). Of the five countries with the most systems ranked, the United States occupies the first place with 34.2 % of systems, followed by China, Germany, Japan and France (see [Figure 1 \[1\]](#)).

Figure 1. The participating countries in the [June 2024 TOP500 list](#) and their share of systems



The following trends can be inferred from the TOP500 supercomputer rankings.

- **Increased focus on exascale computing.** The emergence of exascale systems, like Frontier and Aurora, indicates a significant leap in computational capabilities, enabling more complex simulations and data processing across various scientific fields, such as climate modelling and AI.
- **Diversity in supercomputer architectures.** There is a trend towards heterogeneous architectures that combine CPUs and GPUs, optimising performance and energy efficiency. This shift allows for better utilisation of resources to handle diverse workloads.
- **Growing importance of energy efficiency.** Growing attention is placed on energy-efficient designs, seen through an overlapping focus on the [Green500](#) list. This reflects the need for sustainable supercomputing solutions in an energy-conscious world.
- **Rise of cloud-based HPC solutions.** Cloud computing is playing a more significant role in the HPC landscape, with providers like Microsoft Azure and AWS making HPC resources more accessible to a broader range of users.
- **Global collaboration in HPC development.** Initiatives like the EuroHPC illustrate the collaborative efforts among countries to enhance their HPC capabilities, emphasising the importance of shared resources and knowledge. LUMI supercomputers and MareNostrum 5 are good examples of this joint effort, enabling Europe to compete globally in HPC capabilities and research.
- **Expanding application domains.** The applications of supercomputers are broadening beyond traditional fields, with growing use in AI, machine learning and biomedical research, showcasing the versatility of these systems.

Europe

In Europe, HPC resources are widely provided by a robust infrastructure of supercomputing centres, national data centres, research institutions and collaborative initiatives such as the EuroHPC Joint Undertaking, making these resources available to scientific researchers, industry users and government agencies.

Several EuroHPC Joint Undertaking supercomputers have been highly ranked on the TOP500 list, with LUMI (Finland), Leonardo (Italy) and MareNostrum5 (Spain) all appearing within the top 10, given their computational power. Alps (Switzerland), a non EuroHPC Joint Undertaking system, also ranks within the top 10.

Other major European supercomputing centres are also well known for their HPC resources, such as the Jülich Supercomputing Centre and the Leibniz Supercomputing Centre in Germany, and facilities like the Très Grand Centre de Calcul (managed by GENCI) and the Institut du Développement et des Ressources en Informatique Scientifique in France. In Eastern Europe, the Poznań Supercomputing and Networking Centre and the Interdisciplinary Centre for Mathematical and Computational Modelling at the University of Warsaw in Poland, along with the IT4Innovations National Supercomputing Centre in Ostrava in Czechia, are also important data centres. European HPC resources have traditionally been used in the academic sector, but recently SMEs and the industry sector have also started to utilise this technology.

In Europe, the EuroHPC Joint Undertaking plays a crucial role in enhancing access to HPC resources. It was established in 2018 with the objective of advancing HPC research towards exascale computing. It has a budget of approximately EUR 7 billion for 2021–2027, funded by the European Commission, EU Member States, associated states and private-

sector entities. The EU's multiannual financial framework contributes EUR 3 billion: EUR 1.9 billion from the digital Europe programme, EUR 900 million from Horizon Europe and EUR 200 million from the Connecting Europe Facility II. This EU funding is matched by contributions from the participating countries and supplemented by an additional EUR 900 million from private-sector entities [2].

Access to HPC systems in Europe, including those provided through EuroHPC and national initiatives, is facilitated through different calls, such as regular access, extreme scale access and development access, among others. EuroHPC systems offer public access to European scientific, industrial and public-sector users, with details published on the [EuroHPC website](#). National HPC initiatives provide access through centres such as GENCI in France (Jean Zay and Irene Joliot-Curie supercomputers) or the Jülich Supercomputing Centre in Germany (JUWELS and JURECA supercomputers). Additionally, the EuroCC national HPC competence centres facilitate local access and expertise for researchers and SMEs in each Member State, with calls aligned with national priorities. These initiatives collectively support Europe's leadership in cutting-edge research and technological innovation.

The **United Kingdom** is also making advances in HPC infrastructures. ARCHER2, the UK's national supercomputing service, is hosted by the University of Edinburgh. Another notable system is Dawn, hosted by the University of Cambridge, which provides HPC resources for scientific research and industry collaboration. These systems receive funding through UK research and innovation, which supports HPC development across its councils, such as the Engineering and Physical Sciences Research Council, the Science and Technology Facilities Council and the Biotechnology and Biological

Sciences Research Council, ensuring sustained investment in infrastructure, computational resources and training. Similarly to other countries, initiatives like the Hartree Centre foster industry engagement, enabling businesses – particularly SMEs – to harness HPC for innovation in AI, data analytics and simulation technologies.

Since 2024, the United Kingdom has been associated with Horizon Europe, granting UK researchers access to future EuroHPC supercomputers. The UK government will provide match funding for EuroHPC grants, enabling businesses, researchers and scientists to bid for advanced computing resources across Europe. This investment is part of the UK's GBP 1.5 billion (EUR 1.75 billion) plan to enhance its supercomputing capabilities, which includes the launch of the Isambard-AI and Dawn supercomputers, along with the creation of a new national exascale facility in Edinburgh. The United Kingdom also participates in Horizon Europe, a EUR 93 billion (GBP 80 billion) programme, offering UK businesses grants worth an average of EUR 510 000 (GBP 450 000). Additionally, the United Kingdom has joined the Chips Joint Undertaking, securing EUR 1.2 billion (GBP 1.1 billion) for semiconductor research [3].

United States

The United States has a long history of providing HPC resources, primarily at the national level, starting with the establishment of some of the first HPC facilities in the 1940s. The Lawrence Livermore National Laboratory (LLNL), founded in 1952, played a key role in this development, primarily focusing on the research and development of nuclear weapons. Another well-known HPC organisation is the Oak Ridge National Laboratory, founded in 1943, which hosts several HPC systems, including the Frontier and Summit supercomputers. Five American HPC systems

are among the top 10 on the TOP500 list. The Frontier supercomputer from the Oak Ridge National Laboratory has held first place for several years, followed by the Aurora supercomputer from the Argonne National Laboratory in second place. Large private American companies are also known for their HPC resources, with Amazon and Microsoft being principals in this field and providing commercial cloud computing services to users. Notably, the Eagle supercomputer from Microsoft Azure currently occupies third place on the TOP500 list. Advanced Micro Devices (AMD), NVIDIA and Intel are American companies that provide GPU technology to HPC facilities worldwide, significantly contributing to their high computing power.

In the United States, financing for major HPC supercomputers comes from federal funding, grants and private partnerships. The Department of Energy, one of the major public funding institutions, manages several key supercomputing facilities and received USD 8.8 billion (EUR 8.27 billion) from the 2024 US budget for emerging technologies, including HPC. The National Science Foundation also provides significant support, distributing nearly USD 165 million (EUR 155 million) in grants, and planning a USD 11.3 billion (EUR 10.63 billion) budget for new facilities like the Leadership-Class Computing Facility. The 2024 US budget proposed USD 25 billion (EUR 23.5 billion) for emerging technologies, with a substantial portion allocated to supercomputing [4].

China

China submitted its last entry to the TOP500 in 2016: the Sunway TaihuLight supercomputer, developed by the National Research Centre of Parallel Computer Engineering and Technology. Since then, no new machines have been featured in the list. This does not mean that there have been no new HPC systems. On the contrary, an exascale machine, Sunway Oceanlite, was partially presented in the Gordon Bell finalist talk at the SC23 conference [5]. It is built using more than 100 000 SW26010Pro heterogeneous many-core processors also developed within the country. Recently, the Tianhe-3 exascale supercomputer was launched, and a Sugon exascale system was built [6]. China finances its supercomputing advancements through major government investments and dedicated funds, such as the CNY 344 billion (EUR 44.72 billion) semiconductor fund. The government focuses on developing domestic technology and chips, supported by policies that enhance technological independence and position China as a global supercomputing leader.

Japan

Japan's TOP500 landscape is dominated by Fugaku, developed by RIKEN and Fujitsu, which held the first place for two consecutive years (2020 and 2021). Fugaku utilises ARM-based processors, making it highly versatile in applications like drug discovery, weather forecasting and AI. In 2024, Japan significantly boosted its investment in HPC to enhance its technological capabilities and independence. The Japanese government is allocating approximately JPY 72.5 billion (EUR 470 million) to develop an AI supercomputer, aiming to reduce reliance on US technology.

Additionally, Japan's Ministry of Economy, Trade and Industry announced plans to deploy a supercomputer dedicated to generative AI, with a peak performance of 54.34 petaflops, at the National Institute of Advanced Industrial Science and Technology. Japan is also set to develop a supercomputer for quantum computing research incorporating NVIDIA H100 GPUs [7].

Japan funds its supercomputing projects through substantial government investments, primarily from the Ministry of Economy, Trade and Industry. Significant subsidies and cost-sharing with private companies aim to boost domestic AI and supercomputing capabilities, reduce reliance on foreign technology and support local innovation.

Other regions

HPC resources are expanding across other regions too. **India** is making advancements in HPC, with a focus on both hardware development and applying HPC to sectors like AI, climate research and healthcare. The AIRAWAT PSAI system reflects India's ambition to become a leader in AI, which is a priority in the country's broader digital strategy [1]. AI research, along with quantum computing and big data, is being heavily funded through initiatives like the National Supercomputing Mission. Under the mission, India aims to build a network of supercomputing facilities across the country, connecting research institutions and universities. Additionally, there is a growing focus on indigenously developed supercomputers, with the goal of reducing dependency on foreign technology, as seen in the PARAM series of supercomputers, to benefit areas such as space exploration and semiconductor manufacturing [8].

Singapore

Singapore's HPC ecosystem is led by the National Supercomputing Centre, which manages the country's premier systems, including the ASPIRE 2A supercomputer. Academic institutions, government agencies and industries can access the centre's systems through proposal-based allocations. The ASTAR Computational Resource Centre complements this by providing specialised HPC services for government research projects under the Agency for Science, Technology and Research. Singapore's focus on fostering industrial innovation is evident through collaborations that leverage HPC for sectors like fintech, manufacturing and urban sustainability [9].

Taiwan

Taiwan has rapidly advanced its HPC capabilities, driven by its strategic focus on AI, semiconductors and biotechnology. The National Center for High-Performance Computing, under the National Applied Research Laboratories, operates Taiwan's most powerful supercomputers, including the Taiwania series, such as Taiwania 4. These systems are employed in areas ranging from semiconductor development and AI research to disaster prevention and precision medicine. Academic and industry users can apply for access through the centre's application process. The Industrial Technology Research Institute further supports industry-focused HPC applications, fostering collaboration between researchers and businesses to enhance Taiwan's competitiveness in high-tech manufacturing and innovation [10].

South Korea

South Korea's HPC infrastructure is among the most advanced in Asia, supporting fields such as AI, climate modelling and material science. The Korea Institute of Science and Technology Information operates Nurion, one of the region's fastest supercomputers, while the Korea Supercomputing Centre collaborates on projects like energy-efficient computing and autonomous vehicles. The government is developing its seventh supercomputer, targeting exaflop performance by 2025 to accelerate AI tasks. The SSC-21, the country's most powerful supercomputer, operates at 25.18 petaflops and currently ranks 20th on the TOP500 list.

As part of its national HPC plan, the Ministry of Science and ICT (information and communication technology) will provide open-source software for computational fluid dynamics and structural analysis. It will also offer HPC software as a service to support start-ups with computational modelling and consulting. Tailored resources will be provided to industries like aerospace, autonomous driving and nuclear fusion. The government aims to expand KREONET's⁽¹⁾ data traffic to 1 335 petabytes by 2025, positioning South Korea to close the gap with global supercomputing leaders and drive innovation in AI, aerospace and other sectors [11].

1 [KREONET](#) is a high-performance backbone network designed to support the research and education community in Korea. By connecting 18 regional centers across the country, it offers transmission speeds ranging from 200 gigabits per second to 1.2 terabits per second. By linking with domestic and international backbone networks, internet exchange nodes and national data exchange nodes, KREONET provides research network services and internet services to over 200 member institutions.

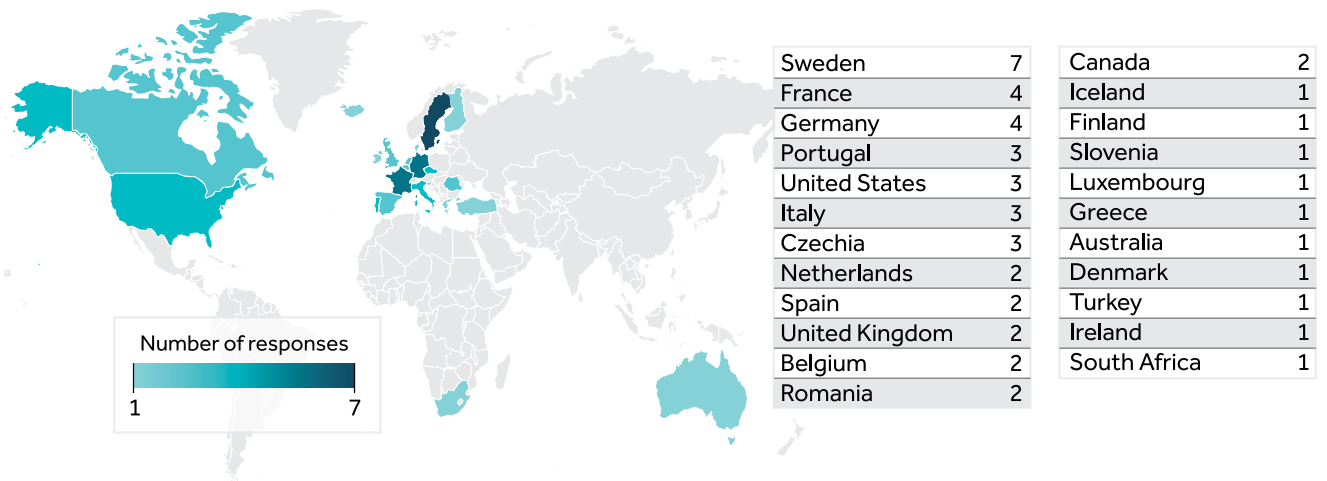
Global financial trends

In 2023, global investment in HPC stood at around EUR 35 billion (approximately USD 37 billion), experiencing inconsistent growth due to various market forces. On-premises/shared cluster server sales declined slightly by 2.7 %, influenced by supply chain challenges and slower-than-anticipated adoption of large systems such as exascale systems. However, the cloud HPC market saw significant growth, reflecting the ongoing shift towards AI-driven workloads that blend HPC with AI technologies. The market for HPC storage also remained robust, estimated at EUR 5.9 billion in 2023 and expected to increase to EUR 9.5 billion by 2028 [12].

AI’s influence on the HPC market has become transformative, as AI-augmented HPC is creating a profound paradigm shift. In response to this change, Hyperion Research recently updated its market sizing to reflect the rapid growth in AI-centric servers, which are being increasingly integrated into both traditional HPC workflows and enterprise data centres for AI applications. In 2023, the inclusion of AI-centric servers was projected to boost the overall HPC server market growth rate to 36.7 %, including an additional EUR 12.8 billion (USD 13.6 billion) in revenue predicted by 2028 [13].

As AI models advance and large foundation models – such as large language models (LLMs) – continue to merge with HPC, new types of buyers and suppliers are entering the market, significantly broadening its scope. The rising adoption of GPU- and accelerator-based systems for AI-related tasks has been a major driver behind the sector’s growth. Hyperion Research notes that this influx of AI into HPC is reshaping both the technology landscape and the market dynamics, signalling a new era of accelerated HPC adoption and investment. The growth forecast for 2024 was positive with server sales expected to rebound to EUR 15.2 billion, while the HPC–AI hybrid model is likely to dominate future investment trends. This underscores the increasing importance of AI in reshaping the HPC sector globally [12].

Figure 2. Geographical distribution of the stakeholders interviewed



3.1 Stakeholder consultations

To gain a comprehensive understanding of how HPC resources are provided globally, and which sectors benefit most, stakeholders from around the world were consulted through dedicated interviews and a web-based questionnaire.

3.1.1 Dedicated interviews

Fifty-one interviews were conducted in 24 countries, as shown in Figure 2. Forty-three participants were from Europe, and eight were from non-European countries (Australia, Canada, South Africa, Turkey and the United States).


Of the participants interviewed, 46 % were users representing universities, research institutes, large companies, SMEs and start-ups; 36 % were providers of HPC services representing national HPC centres, private HPC providers and national agencies; 16 % were developers representing start-ups, large companies and universities; and 2 % were non-users representing start-ups.

Participants came from a range of sectors: 52 % represented organisations that provide services across multiple sectors or worked closely with users from different sectors, and 48 % represented specific sectors, including IT and telecoms, computer science, applied mathematics, biosciences, earth sciences, weather and climate science, physics, renewable energy, biomolecular science, maritime engineering, aerospace engineering and AI.


Participants had extensive experience in the HPC field, with 72 % having worked in this area for over five years. Additionally, participants held various roles within their organisations, including head of department/team manager (41.2 %), researcher (21.6 %), director or chief executive officer (19.6 %), senior researcher/professor (17.6 %).

Figure 3. Characterisation of the participants based on their role within organisations, funding source and area

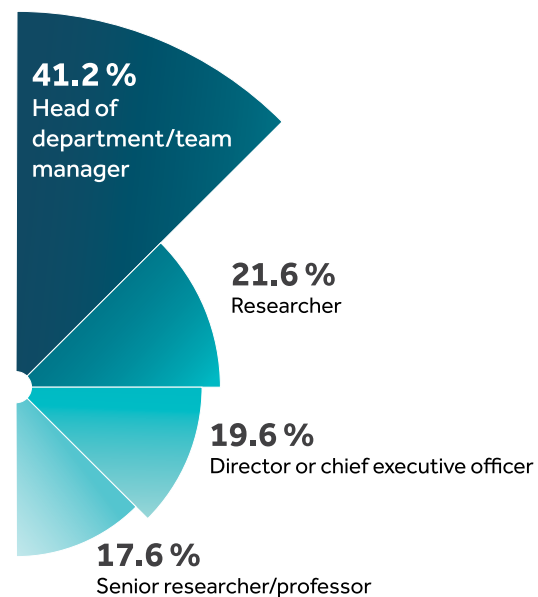
Area



Funding source



Organisational role classification



Participants in the dedicated interviews predominantly came from academia (62.7 %), followed by those from industry (37.25 %), and a smaller proportion (3.95 %) coming from organisations involved in both areas. In terms of funding sources, most participants were from public institutions (58.8 %), followed by the private sector (35.3 %), and 5.9 % were from organisations that received both public and private funding (see **Figure 3**).

All stakeholders provided an overview of how HPC resources are structured in their respective countries. HPC systems are organised and articulated through a variety of collaborative initiatives at both the local and regional levels, enhancing resource sharing, research capabilities and innovation for users from academia and industry. National supercomputing centres, primarily funded by public government sources, facilitate access to HPC resources for local researchers, institutions and external users, thereby enhancing computational research and collaboration. Such centres are widely established across Europe, and in Canada, South Africa and the United States. There are data centres that provide HPC resources across continents and regions, but users primarily tend to utilise resources from their own country first, then their continent.

Europe

HPC resources in Europe are provided through a variety of national and regional centres, supporting both academic and industrial research. Several European supercomputing facilities were involved in the interviews. In Spain, the Barcelona Supercomputing Center (BSC) is a leading hub for HPC, attracting users from across the country and Europe. As highlighted in the interviews, the majority of users are Spanish researchers working in fields such as healthcare, government and defence, energy, mobility, biosciences, physics, computer science and civil

and environmental engineering. A significant proportion of these users come from the Spanish Supercomputing Network, which is managed by the BSC and comprises 14 institutions, and from the EuroHPC Joint Undertaking. According to the interviews, the BSC is a preferred choice for many Spanish researchers due to its substantial computational capacity. GPUs, which account for 80–90 % of the centre's computing system, offer enhanced scalability, reduced energy consumption and superior performance for parallel applications, making them a more sustainable and efficient choice for HPC.

The Jülich Supercomputing Centre, located in Germany, is a leading facility for HPC, with users coming from across Europe. It is primarily funded through public resources provided by the German government and the European Commission. Its users represent a wide range of sectors, including energy, material science, biosciences, aerospace, environmental engineering, astronomy, earth sciences, computer science and applied mathematics. The centre places significant emphasis on GPU-based systems, which are highly valued by its users. Many users have the expertise to optimise their codes to maximise the efficiency of these systems, reflecting the strong interest in GPU capabilities. This demand is further supported by the centre's immense computational capacity and its first-place ranking on the Green500 list, recognising it as a global leader in energy-efficient supercomputing.

In Sweden, the Tetralith supercomputer, managed by the National Supercomputer Centre, primarily serves academic institutions, with 90 % of its users coming from Swedish universities. The centre, funded by the Swedish government, plans to expand its access to include private companies and public agencies, marking a shift in its user base. In France, CRIANN, located in Normandy, was described

in the interviews as a unique, independent computing centre that serves local universities and industries. It focuses on fields like fluid mechanics, chemistry and physics, with a particular emphasis on material science at the microscopic level. Meanwhile, in Italy, CINECA, the national supercomputing centre, is funded equally by the Italian Ministry of Research and European projects. As mentioned in the interviews, while its primary users are from academia, there is a growing number of industry users, particularly in fields like engineering and AI, highlighting the expanding role of HPC in both academic research and industry across Europe.

North America

The HPC landscape in the United States plays a vital role in advancing the nation's scientific and technological efforts. Supported by a vast network of national laboratories, research institutions and universities, it receives substantial funding from federal agencies such as the National Science Foundation, the Department of Energy and the National Institutes of Health. Large companies in the oil, gas and energy sectors frequently rely on HPC resources. The Texas Advanced Computing Center (TACC) exemplifies a leading HPC centre, providing essential computational resources for both academic and industrial research. While most users are based in the United States, TACC also serves some international users, including Europeans. TACC is primarily publicly funded, with around 80 % of its budget sourced from the National Science Foundation, while the remainder comes from government agencies, the National Institutes of Health, industry partners and the University of Texas. To ensure effective access to its HPC systems, TACC offers training for industry partners, as these systems are not as readily available as commercial platforms like Amazon. TACC supports a diverse array of sectors,

including healthcare, government, IT, energy (both renewable and non-renewable), logistics and biosciences, with approximately 30 000 users having direct or indirect access to its resources.

In Canada, HPC resources are primarily provided through a national system managed by the Digital Research Alliance of Canada, which coordinates and funds HPC infrastructure across the country. Regional organisations such as ACENet, which is responsible for Atlantic Canada, offer support, consulting and training services to users in their respective regions. In the interviews, it was mentioned that large HPC systems are located in Quebec, Ontario and British Columbia, while users from across Canada access these national resources.

The main areas benefiting from these HPC resources include oceanography, physics, chemistry, computer science, biology and engineering, along with digital humanities and social sciences. Most of the funding is public, with 60 % provided by the federal government through the Digital Research Alliance of Canada. ACENet secures the remaining 40 % from provincial governments and their innovation agencies. Additionally, ACENet receives some industry funding for special projects, particularly for training initiatives. The organisation also provides services to industry clients on a not-for-profit basis, which makes it somewhat unique, as the national alliance does not yet have a formal industry mandate.

3.1.2 Web-based questionnaire

From the stakeholder consultations conducted through the web-based questionnaire, 189 answers were collected and analysed.

Figure 4 and Figure 5 provide a geographical distribution of the stakeholders and a characterisation based on demographics, funding source, sector and revenue.

- Most of the participants were men (81 %), 14.3 % were women and 4.7 % preferred not to say.
- Most participants were between 31 and 50 years old.
- In terms of education level achieved, most had a doctorate (84.1 %) or master's degree (13.8 %).
- 61.4 % were users, 27 % were developers, 9 % were providers and 2.6 % had never used HPC before.
- 84 % of the participants represented non-profit organisations, such as academic institutions, governmental organisations, legal associations, non-governmental organisations (NGOs) or research institutes, and 16 % represented for-profit organisations, such as SMEs, start-ups and large companies.

- Most of the participants represented academic institutions (53.4 %) and research institutes (24.3 %), followed by governmental organisations (6.9 %), SMEs (5.3 %), start-ups (4.8 %), large companies (3.7 %) and others (1.6 %).
- Participants from large companies were only developers and users, while participants from academic institutions, research institutes, SMEs, start-ups and governmental organisations were represented by users, developers and providers.
- Most of the participants had more than five years of experience in the HPC field (70.4 %), 22.8 % had one to five years of experience, 4.2 % had less than one year of experience and 2.6 % had never worked with HPC.

Organisations from a total of 28 sectors were represented, with the highest participation from computer science (21.2 %), physics (16.9 %) and government and defence (6.9 %). Aerospace engineering, biosciences, maritime engineering, IT and telecoms were equally represented, with 11 answers each. The complete list of sectors is presented in Figure 6.

Figure 4. Geographical distribution of the stakeholders from the web-based questionnaire

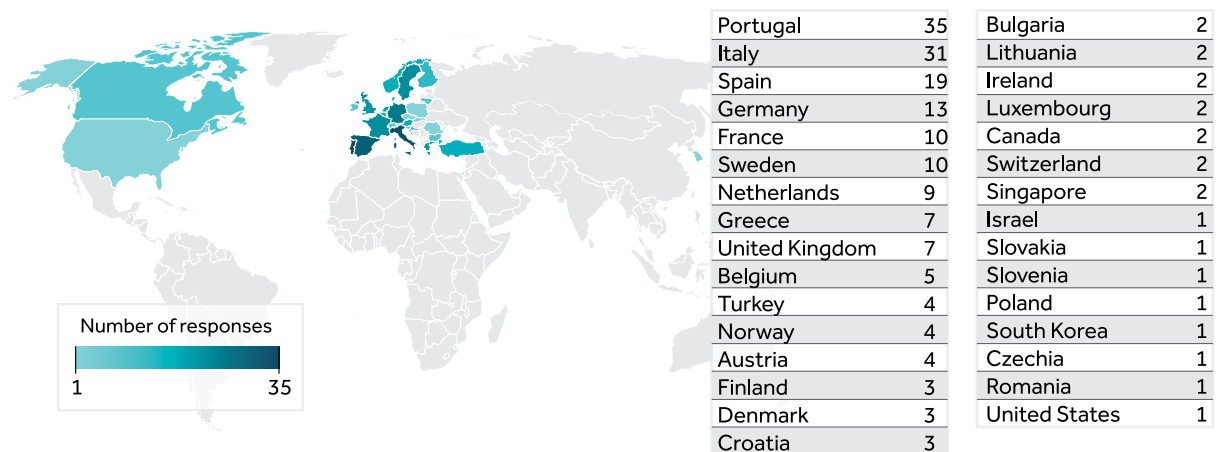
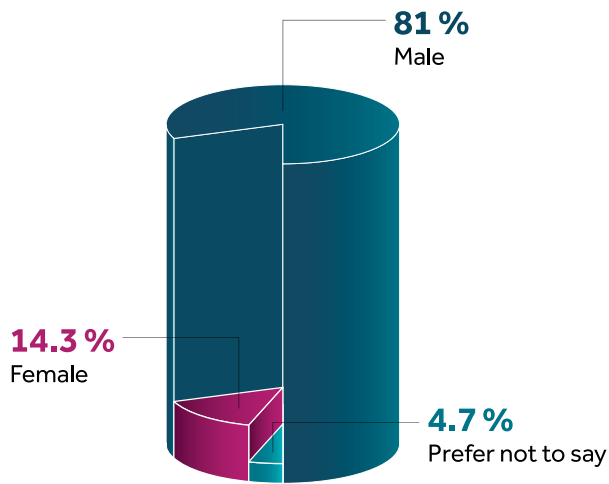



Figure 5. Characterisation of the participants based on demographics, stakeholder group, sector, type of organisation, funding source and revenue

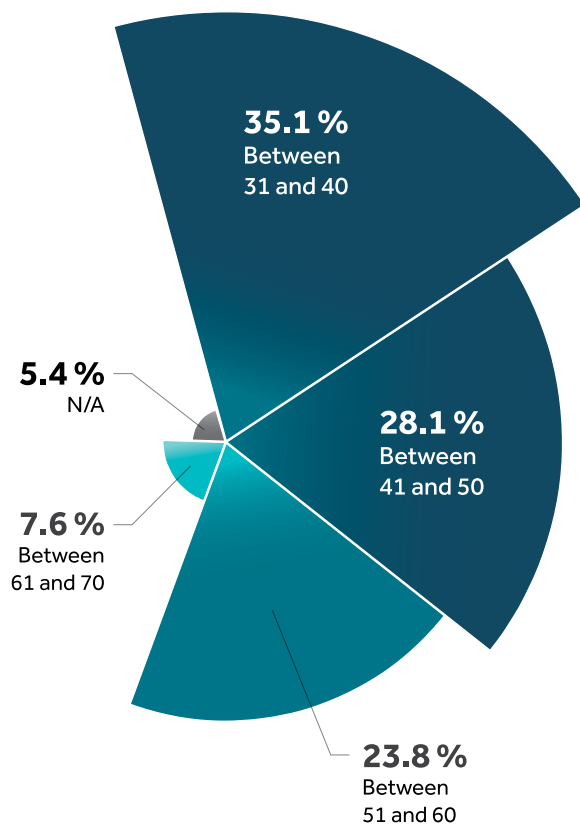
Gender



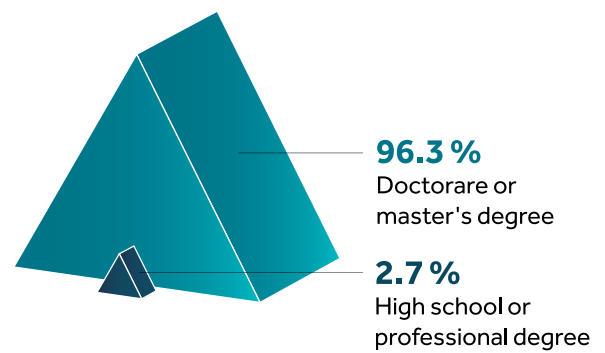
Geographic coverage



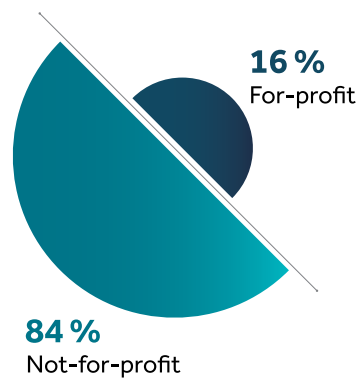
Age



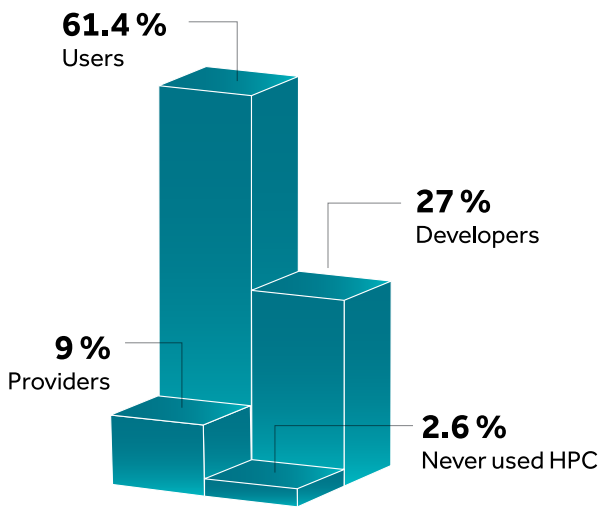
Education level



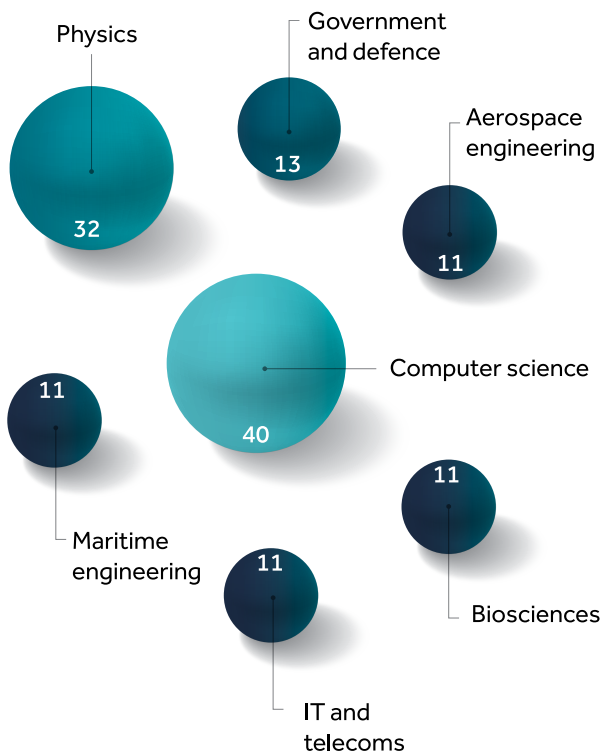
Revenue



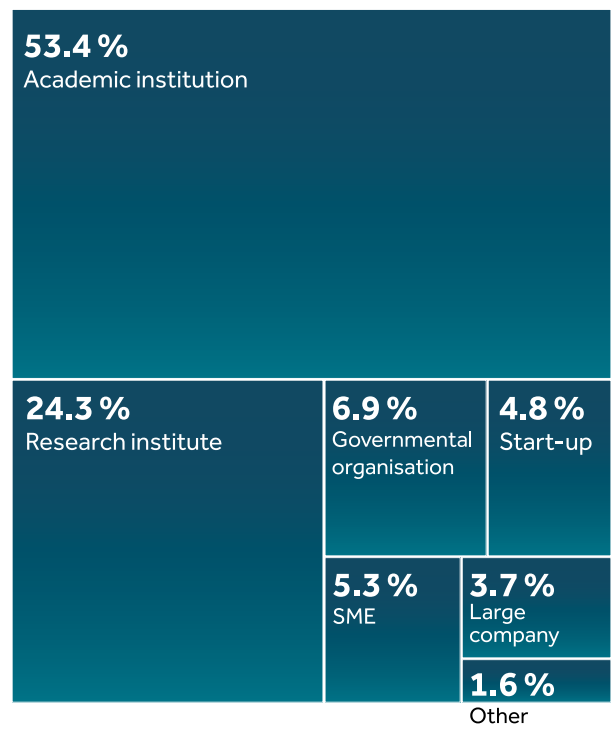
Stakeholders group



Sector



Type of organisation



Funding source




Figure 6. Representativeness of different sectors according to the type of organisation

Sector	Type of organisation									
	Academic institution	Governmental organisation	Large company	NGO	No organisation	Research Institute	SME	Start-up	Legal Association	Total
Academic/research	0	0	0	0	0	1	0	0	0	1
Aerospace engineering	9	0	1	0	0	0	1	0	0	11
Applied mathematics	6	0	0	0	0	1	0	0	0	7
Astronomy	3	0	0	0	0	2	0	0	0	5
Biosciences	8	0	0	0	0	2	0	1	0	11
Chemistry	4	0	0	0	0	0	0	0	0	4
Civil engineering	1	0	0	0	0	0	0	0	0	1
Computer science	21	5	2	0	1	6	5	0	0	40
Earth sciences	2	1	0	0	0	4	0	0	0	7
Education	1	0	0	0	0	0	0	0	0	1
Energy (renewable)	0	0	0	0	0	3	0	1	0	4
Energy (non-renewable)	0	0	0	0	0	0	0	1	0	1
Environmental engineering	0	0	0	1	0	0	0	0	0	1
Epistemology	0	0	0	0	0	1	0	0	0	1
Government and defence	5	5	0	0	0	3	0	0	0	13
HPC	2	2	0	0	0	0	0	0	0	4
Higher education	1	0	0	0	0	0	0	0	0	1
IT and telecoms	0	0	2	0	0	3	2	3	1	11
Information, materials and energy	0	0	0	0	0	1	0	0	0	1
Life sciences	1	0	0	0	0	0	0	0	0	1
Manufacturing	0	0	0	0	0	1	0	0	0	1
Maritime engineering	4	0	1	0	0	5	0	1	0	11
Material science	6	0	0	0	0	2	1	0	0	9
Media and entertainment	0	0	0	0	0	0	0	1	0	1
Nanosciences	1	0	0	0	0	0	0	0	0	1
Other engineering	0	0	0	0	0	0	1	0	0	1
Physics	19	0	1	0	0	11	0	1	0	32
University	7	0	0	0	0	0	0	0	0	7
Total of answers	101	13	7	1	1	46	10	9	1	189
Total of sectors represented	18	4	5	1	1	15	5	7	1	

A total of 18 sectors were represented by participants from academic Institutions, 15 sectors by research institute participants and 17 sectors by industry participants (large companies, SMEs and start-ups). Among the sectors with the most representatives, computer science (21.8 %) and physics (18.8 %) were the most common for academic institutions and research institutes, while IT and telecoms (30 %) and computer science (50 %) were the leading sectors for industry participants (large companies, SMEs and start-ups).

The higher percentages of participation from computer science, IT and telecoms can be explained by their strong reliance on innovation and development through HPC. Computer science is a sector valuable for both academia and

industry, enabling advancements in theoretical exploration and emerging technologies like AI, quantum computing and cybersecurity. In academia, it drives interdisciplinary research, influencing fields such as biology and social sciences, while also developing a skilled workforce through education in programming and computational thinking.

In industry, computer science enhances efficiency through automation, data-driven decision-making and the creation of innovative products and services. HPC is particularly critical in telecommunications, supporting large-scale data processing, network optimisation and the roll-out of technologies like 5G and AI. The collaboration between academic research and industry application fosters continuous technological progress across sectors.

Users

This group was composed of 116 participants, of which 79 % were men, 16 % were women and 5 % preferred not to say. In terms of education level, most users held academic degrees, with 85 % holding a doctoral degree, 12 % holding a master's degree and the remaining participants holding either a bachelor's degree or a professional degree. Additionally, 72 % of users worked in academic institutions and research institutes, while the rest came from governmental organisations, large companies, SMEs and start-ups.

Most users (66 %) said they use HPC daily, 23 % use it once a week and the remaining users access it sporadically throughout the year.

Developers

A total of 51 developers participated in the web-based questionnaire, where 65 % were men and 14 % were women. Most of them (76 %) had been working with HPC for more than five years and held a doctorate degree (80.4 %).

In terms of employers, 41 % of developers

came from academic institutions and 31 % were affiliated with research institutes. The remaining developers were spread across governmental organisations, large corporations, SMEs and start-ups.

Providers

A total of 17 responses were received from providers, of which 82 % were men and 18 % were women. Age distribution was as follows: 17 % were between 31 and 40 years old; 41 % between 41 and 50; 35 % between 51 and 60; and 7 % between 61 and 70.

Additionally, 88 % of providers had worked in the HPC field for over five years, while the remainder had at least one year of experience. Regarding education, 76 % of providers held a doctorate degree, 12 % held a master's degree and the remaining held a bachelor's degree. The majority (71 %) of the organisations were fully publicly funded, a small percentage were fully privately funded (12 %) and 17 % were funded by both public and private sources (see Figure 7).

Figure 7. Responses from HPC providers based on the type of funding their organisation receives



HPC resources in each sector

Participants were also asked which sectors utilise HPC resources. In the interviews, HPC providers stated that they provide HPC resources to a broad range of sectors. Key sectors mentioned include weather forecasting, material science, healthcare, mobility and engineering, as these have long relied on HPC for complex simulations and predictions. However, as computational capabilities advance, other areas of significant interest are emerging. HPC providers also mentioned sectors such as IT and telecoms, smart cities, government applications, agriculture, tourism and robotics as growing users of HPC. For example, robotics is gaining attention for its potential to develop systems that can mirror human movement. These emerging areas are increasingly reliant on HPC due to the need for real-time processing of vast amounts of data. HPC enables faster network communications and data analytics in IT and telecoms, supports urban planning and resource management in smart cities and assists governments with complex simulations and policy development. In agriculture, it improves crop forecasting and precision farming, while in tourism, it enhances customer behaviour analysis and virtual simulations.

In the web-based questionnaire, participants were asked whether they believe certain industries or sectors could benefit from specialised HPC services. A majority of 79 % indicated that specialised HPC services would be beneficial across various sectors, while 21 % disagreed. Among those who responded positively, the specific sectors they believed would benefit varied. Notably, 60 % of these respondents identified material science and biosciences as areas that could gain from HPC, whereas fewer than 50 % saw potential benefits for sectors such as mobility and logistics. The same questions were asked to providers. 94 %

of HPC providers expressed the belief that HPC specialised services will bring significant benefits to specific sectors, particularly healthcare and material science (see **Figure 8**).

Moreover, HPC providers highlighted several ways in which specialised HPC systems would benefit various sectors, particularly in areas like AI and speeding up R & D efforts.

Providers catering to multiple scientific domains see the need for different specialised systems to accelerate specific applications and offer tailored solutions to users' needs. Some providers noted that specialised HPC systems could enhance processing for sensitive data, optimise hardware for specific tasks, and provide dedicated machines for hybrid HPC and AI, which could lead to training LLMs more effectively. In life sciences, the availability of HPC clusters for specific use cases could significantly accelerate research. Additionally, some sectors would benefit from digital twins and high-throughput computing, while others could reduce the number of personnel required to maintain these systems by optimising their hardware and related services.

Conclusion

- The growing interest in specialised HPC services across both traditional and emerging sectors underscores its potential to drive innovation and efficiency, particularly in areas like healthcare, material science and AI.
- While the majority of participants recognise the benefits of specialised HPC, there is variability in how different sectors perceive its value.
- HPC services and systems should be tailored to meet the unique needs of various industries.
- HPC providers need to assess what specific sectors – such as mobility and logistics – might require in terms of additional outreach or development to realise HPC's advantages fully.

Figure 8. Responses of users and providers regarding which sectors would benefit from specialised HPC services

	 USERS	 PROVIDERS
More than 60 %	Material science Biosciences Aerospace engineering Healthcare Astronomy Applied mathematics Earth sciences Computer science Energy (renewable) Physics	Material science Biosciences Healthcare Physics Earth sciences
Between 50 % and 60 %	Government and defence It and telecom Banking and finance Civil engineering Environmental engineering Energy (non-renewable)	Energy (renewable) Agriculture Aerospace engineering Civil engineering Environmental engineering Astronomy Computer science
Less than 50%	Mobility Logistics Retail and consumer goods Agriculture Media and entertainment	Government and defence Applied mathematics Banking and finance Retail and consumer goods Media and entertainment Energy (not renewable) It and telecom Mobility Logistics

04

High-performance computing case studies

In this section, we will present three case studies that were selected due to their relevance in illustrating user experiences with HPC systems, employing methodologies like those implemented in this study.

4.1 Case study 1

Evaluation of criteria for the implementation of HPC in Danube-region countries (2020)

The increasing demand for processing large-scale scientific and industrial data has prompted a collaborative exchange between scientists and data analysis experts. This collaboration is crucial for generating insights and solutions to address specific challenges. HPC providers highly value such exchanges, recognising the significance of partnering with industries for the future development of their organisations. They believe that by working closely with businesses they can effectively assist enterprises in meeting their requirements through HPC.

Intermediaries in the HPC sector that offer consulting or advisory services play a vital role in bridging the gap between HPC service users and centres. This is particularly crucial as numerous companies and public research centres lack the technical expertise required for the effective utilisation of HPC. In Europe, independent software vendors are notably prevalent and successfully collaborate with research institutions. However, their business expansion is hindered by challenges in securing financing. Traditional commercial banks typically focus on financing private and commercially oriented HPC centres.

The significant size of the financial risk brings forth related business risk. This means that, beyond allocating resources for the analysis required in selecting an HPC architecture, it is also essential to earmark resources for the architecture itself. This includes adhering to the standards necessary for the architecture to function optimally. Moreover, a dedicated team of experts is indispensable to support the implementation of the HPC architecture and ensure positive business outcomes, particularly when it comes to complex problems or projects. Creating space for future projects also hinges on this support. Various risks can lead to the failure of implementing the HPC architecture, often occurring when specific requirements or needs are not met.

The process of selecting an architecture typically involves multiple steps. Initially, there's a critical assessment of the current state to determine if architecture changes are necessary and, if so, the reasons and conditions for such changes. Subsequently, plans, budgets and requirements are analysed to assess the need for a new solution. Considering this information and any constraints, an exploration of current technologies and their capabilities follows.

The next crucial step is deciding on the most suitable solution for the desired changes. This entire process culminates in decisions that serve as the foundation for negotiations and the creation of contracts for procuring and implementing the desired architecture.

In this process, it is imperative to gather as much information as possible before making decisions. Establishing an HPC architecture within an organisation requires meticulous attention to each criterion and its domain to avoid making flawed decisions that may hinder the desired outcomes. Failure to achieve the desired results can compromise the implementation and application of the desired architecture. Therefore, it is essential to implement a multidimensional database model and analytical procedures to facilitate decision simulations based on specific criteria.

The objective of this paper was to evaluate which criteria several countries in the Danube region are using to implement HPC infrastructure. The study aimed to gain valuable insights from HPC experts regarding the development of HPC in specific countries, covering a range of critical aspects such as infrastructure, skills, education, funding, awareness, collaboration and legislative support. The degree of cooperation between science and public authorities, as it relates to HPC, was the most important criterion, followed by the availability of skilled human resources and competitive public funding (e.g. direct public funding, grants, awards, baseline funding), and the least important criterion was the availability of private funding for R & D related to HPC.

In conclusion, the development of HPC infrastructure in the Danube region relies heavily on strong collaboration between science and public authorities, supported by skilled human resources and public funding. While private funding for HPC-related R & D is less critical, addressing financial and technical challenges through careful planning and cooperation is essential. This ensures the effective implementation of HPC systems and fosters sustainable growth in the sector.

4.2 Case study 2

Leveraging HPC and cloud computing with unified big-data workflows: the LEXIS project (2021)

This paper discusses the Horizon 2020 LEXIS project, which integrates big data, HPC and cloud computing to streamline automated data analysis and simulation workflows across European supercomputing centres. The project's goal is to merge scientific HPC capabilities with the big data tools utilised by industry and SMEs, breaking down barriers to accessibility. One of the principal results of the project is the development of an orchestration system that automates complex workflows, optimising the execution of tasks across both HPC and cloud resources. The system streamlines task management and integrates diverse computing resources.

Additionally, the distributed data infrastructure of LEXIS supports large-scale, data-intensive workflows by providing unified access to distributed storage systems. It utilises the integrated rule-oriented data system ⁽²⁾ to manage data access and metadata, ensuring reliable and secure data handling. Security measures include a role-based access control model and Keycloak ⁽³⁾ for authentication, which together provide secure and efficient access to distributed resources.

The platform was tested through various pilot programmes, covering fields such as aeronautics, earthquake and tsunami simulations and weather and climate modelling. These programmes demonstrated the platform's versatility and its ability to handle diverse data-heavy tasks effectively.

In conclusion, the LEXIS project successfully integrates HPC, cloud and big data technologies into a cohesive platform that makes HPC more accessible. By lowering the technical barriers and facilitating collaboration between scientific and industrial communities, LEXIS enhances data analytics and simulation capabilities, supporting innovation across multiple sectors.

2 The [integrated rule-oriented data system](#) is an open-source data management software used by research, commercial and governmental organisations worldwide.

3 [Keycloak](#) is a project that provides authentication, user federation, authorisation and more for modern applications and services.

4.3 Case study 3

Reinventing HPC: challenges and opportunities (2022)

This study identifies significant challenges and gaps in the current state of the HPC field. One major issue is the slowing of Moore's Law and Dennard scaling, which traditionally drove advancements in semiconductor technology by shrinking transistors and increasing performance while reducing costs. As these effects wane, the ability to achieve performance improvements through smaller and faster chips alone is diminishing. This has created a pressing need for innovative approaches in semiconductor design. Additionally, the cost of semiconductor fabrication has risen dramatically. As manufacturing technologies advance to the nanometre scale, building state-of-the-art fabrication facilities now requires billions of dollars. This economic burden makes it increasingly difficult to maintain the pace of technological advancements that were once driven by incremental improvements in chip technology. The high costs of developing and deploying new HPC systems pose another challenge. The shift from traditional computing companies to cloud and smartphone giants has further complicated the landscape. These new industry leaders, such as Amazon, Google and Apple, now wield significant financial influence, surpassing that of traditional hardware manufacturers like Intel and AMD. This shift means that HPC, which was once a primary driver of innovation, is now secondary to the needs of large cloud providers that set the direction for future technology. In addition to these technological and economic challenges, the industry is facing a broader shift in application domains. HPC's traditional focus on complex optimisation and scientific simulations is expanding into new areas such as AI and machine learning. This expansion requires adapting HPC systems to handle a wider range of applications, further complicating the design and development of new systems. Finally, the rise of cloud computing has reshaped the HPC supply chain and financial dynamics. Cloud service providers have far greater financial leverage compared to government and academic projects, which are increasingly outpaced by the scale of commercial cloud investments. This shift has led to new models of collaboration and funding, emphasising the need for government and academic stakeholders to adapt their strategies to remain relevant in an evolving landscape. Overall, the future of HPC is at a crossroads, facing significant challenges in technology, economics and application scope. Addressing these challenges requires innovative solutions, strategic collaborations and a reassessment of traditional approaches to computing.

4.4 Insights from the case studies and stakeholder perspectives

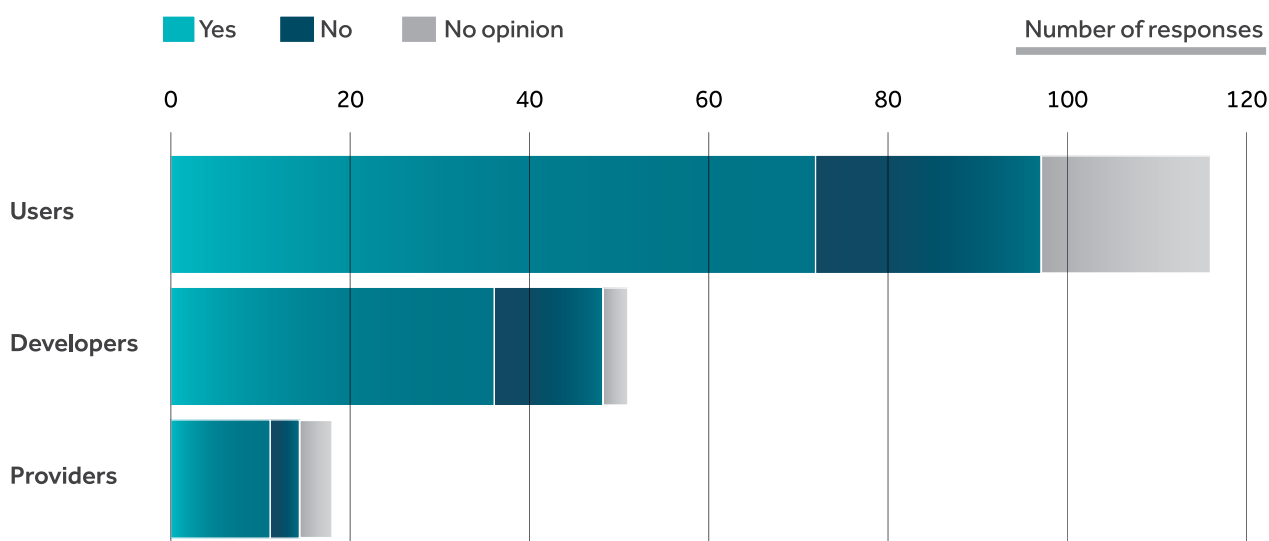
The debate between generalisation and specialisation in HPC systems is highlighted by the findings from the three case studies and the stakeholder consultations.

Case 1 emphasises the need for collaboration between public authorities and the scientific community in the Danube region to build a versatile HPC infrastructure that supports diverse research needs. Stakeholders echoed this sentiment, expressing a preference for general-purpose HPC systems that can adapt to rapid technological changes and serve a broader user base.

In contrast, **Case 2** showcases the benefits of specialised HPC systems designed for specific applications like AI and big data. This case illustrates how specialisation can enhance performance and scalability, aligning with user feedback that recognises the advantages of tailored systems for data-intensive tasks. However, stakeholders also raised concerns about the challenges of accessing these specialised resources, such as job queuing and data management issues, which can hinder adoption.

Case 3 reveals the significant economic and technological challenges facing HPC, including the slowing of Moore's Law and rising semiconductor costs. These issues resonate with stakeholders, many of whom cited budget constraints and a lack of expertise as barriers to accessing HPC resources. Despite recognising the potential benefits, particularly from cloud solutions like AWS, users emphasised the need for improved support and clearer pathways to HPC access. Overall, the stakeholder consultations underscored the importance of addressing these concerns while balancing the benefits of specialised and general-purpose HPC systems to foster innovation and collaboration across sectors.

Figure 9. Opinions on whether services should be specialised, per stakeholder group



4.5 Stakeholder consultations

Generalisation versus specialisation

In the web-based questionnaire, most of the participants agreed that **HPC systems or services optimised for a specific application (e.g. AI) can provide a competitive advantage** compared to traditional HPC services that address a diverse range of applications (see Figure 9).

Where do stakeholders agree?

Performance improvement. Optimised hardware can lead to faster execution of tasks, particularly for applications like AI and machine learning, which benefit from specialised components such as GPUs.

Computational and energy efficiency. Tailored HPC systems can achieve better efficiency in terms of computing power and energy consumption, which is crucial for high-performance applications. Some applications require specific algorithms or tools that perform better on optimised hardware, enhancing the overall efficiency and effectiveness of the computations. Optimised systems can provide better energy efficiency, which is particularly important for large-scale computations and environmentally conscious operations.

Cost-effectiveness. Specialised HPC systems can be more cost-effective by providing better performance for specific tasks, reducing the overall computation time and resource usage.

Resource optimisation. Dedicated resources for specific applications can prevent bottlenecks and ensure optimal use of the available hardware, leading to better overall system performance.

Scalability and flexibility. Optimised systems can be scaled more effectively to meet the specific demands of certain applications, such as large-scale AI training, which requires significant computational power.

Domain-specific benefits. Different applications have varying requirements, and an HPC system optimised for a particular domain (e.g. AI, climate modelling, molecular dynamics) can significantly improve performance and outcomes in that area.

Enhanced user experience. Users can benefit from systems that are easier to use and tailored to their specific needs, leading to more straightforward and efficient workflows.

Competitive advantage. Organisations using specialised HPC systems can achieve competitive advantages by reaching solutions faster and more efficiently, enabling quicker innovation and development.

Where do stakeholders not agree?

Lack of versatility. General-purpose HPC systems are more versatile and capable of addressing a wider range of applications and more complex problems, which is crucial for supporting diverse research needs.

Rapid technological advancements. The fast pace of technological advancements can render specialised systems obsolete quickly. General-purpose systems can adapt more easily to new developments.

Cost and resource allocation. Specialised systems are costly to develop and maintain. Splitting resources across many specialised systems can be inefficient and expensive compared to using versatile, general-purpose systems.

Broad user base. Large facilities and institutions need to cater to a broad range of users and applications. Specialised systems could limit access and usability for non-targeted applications, reducing the overall effectiveness.

Scientific inquiry and flexibility. Scientific research often requires exploring unknowns and unplanned areas. General-purpose systems provide the necessary flexibility to support diverse and evolving research requirements. Highly specialised clusters can narrow down research perspectives and potentially stifle innovation by focusing too much on specific applications.

Future-proofing. Trends and specific application requirements can change over time. A general-purpose system is more future-proof and can adapt to new demands as they arise.

Equity and accessibility. Ensuring broad accessibility to HPC resources for various research fields is important. Specialised systems might prioritise certain fields over others, leading to inequitable access. Also, optimising for one application can suggest fewer resources and optimisation for other applications, which can be detrimental to fields that do not fit within the specialised focus.

Complexity and usability. General-purpose systems tend to be easier to use and manage. Specialised systems can be highly restrictive and complex, potentially limiting their usefulness for a wider audience.

Users' concerns and preferences when accessing HPC resources

During the stakeholder consultations, users' concerns and preferences regarding HPC resources use were discussed, focusing on aspects such as resource availability, system performance and reliability, data storage and management, security, support and usability, along with access and policy considerations. During the interviews conducted, many users shared that their preferred resources are on-premises HPC facilities and shared clusters. Typically, national or European HPC centres primarily serve researchers from academia, with a smaller proportion of users from industrial sectors.

Based on the insights from the interviews, the decision-making process for selecting an HPC centre is multifaceted. Users weigh various factors with differing levels of importance. For some, particularly those from outside the academic sector, the appeal of private cloud providers lies in their ease of use, quick access and user-friendly interfaces, even though these services tend to be more expensive than publicly funded shared clusters. In contrast, others prioritise data storage and security, particularly in sensitive sectors like healthcare.

Users from the healthcare sector specifically highlighted the critical importance of data transfer and storage locations when handling public health data, underscoring the need for robust privacy and security measures in their computing choices. In addition to the available HPC resources, users shared insights on other factors to consider when accessing different HPC systems, which are summarised in Table 3.

Table 3. Summary of the considerations expressed by users from different organisations regarding access to HPC resources

Users' considerations in accessing HPC resources	
Academic users	<ul style="list-style-type: none"> • Experienced HPC users require adequate resources to handle their workloads. • Adequate random-access memory (RAM) is necessary to meet the demands of complex workloads. • Efficient internode communication is essential for running complex tasks. • The granted access time to the HPC system is sometimes insufficient to complete tasks according to the research project's duration. • User support provided by the HPC provider is crucial for users.
Industry users	<ul style="list-style-type: none"> • Industry users prefer easy access to HPC systems, prioritising user-friendly interfaces for seamless navigation and quick setup that allows them to start using resources without lengthy installation or configuration processes. • They need certainty and reliability in accessing HPC systems (e.g. guaranteed access, predictable costs and clarity on prioritisation of their workloads). • There is a shared interest in user-friendly interfaces between new users and industry users that usually do not have a background in using or programming HPC systems. • Industry users seek guarantees that they have a dependable and well-managed containerised environment. • They prioritise scalable resources from HPC providers to adapt quickly to changing workloads and project demands. This flexibility enables efficient management of computational needs while minimising costs. Ultimately, scalable solutions are vital for maintaining competitiveness in a fast-paced environment, along with the ability to easily and cost-effectively transfer workloads across different HPC systems. • Some sectors, such as healthcare, require strict legal and security measures to ensure data privacy, including transparency of the data's location within the HPC system, which allows users to know where their data is stored and ensures compliance with regulatory requirements. This transparency is crucial for verifying that sensitive information is isolated from other users and protected against unauthorised access. Additionally, users seek isolation from other users to further safeguard their data integrity and confidentiality. • Industry users invest in in-house HPC systems for enhanced security, greater control over data privacy and the ability to adapt to changing project demands. In-house systems offer operational flexibility, allowing companies to maintain control over resources during fluctuating workloads.
SME and start-up users	<ul style="list-style-type: none"> • SME and start-up users experience financial and technical constraints to building their own HPC infrastructure. • They require greater unification of access processes and authentication mechanisms across different HPC clusters.

Conclusion

The debate between generalisation and specialisation in HPC systems reflects the diverse needs and priorities of stakeholders.

The insights gained from the three case studies and the stakeholder consultations highlight the importance of collaborative efforts, addressing user concerns and adapting to the evolving technological landscape.

By balancing the strengths of both specialised and general-purpose systems, stakeholders can ensure that HPC remains accessible, efficient and capable of driving innovation across diverse fields. Moving forward, fostering a supportive ecosystem will be essential for maximising the potential of HPC technologies.

Providers from national, public and private HPC/data centres across Europe (Czechia, Finland, France, Germany, Iceland, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Slovenia, Spain and Sweden), North America (Canada and the United States), Africa (South Africa) and Oceania (Australia) have shared their perspectives on how their resources are offered to users. Most of these providers (89 %) are from the academic sector. As providers, the participants were able to identify users' concerns when accessing the systems.

Regarding resource availability, providers reported users' concerns linked to CPU and GPU access, scalability, job queuing and software licence, as described below.

- **CPU and GPU access.** Users are often concerned about whether they can access the computational resources they need, especially CPUs and GPUs. For example, users might worry about whether they can secure the required number of CPUs or whether they can get access to powerful GPUs for parallel processing.
- **Scalability.** Some users require large-scale resources, such as multiple nodes or even thousands of nodes, for massive simulations, and they may worry about the system's ability to scale effectively.
- **Job queueing.** The efficiency of the job queueing system is a concern, particularly in environments where thousands of jobs might be queued, leading to potential delays.
- **Software licences.** Most HPC providers offer a range of software licences to support various tasks. However, in cases where a provider does not have the specific software needed, users are often allowed to bring their own licences.

Regarding system performance and reliability, two main concerns were highlighted: (a) the file system responsiveness, particularly in shared environments where many users are active simultaneously; and (b) the queuing system overhead caused by complex queuing systems and network bottlenecks.

The main concerns related to data storage and management are linked to data loss, where users are concerned about the safety and persistence of their data, especially in environments with scratch spaces that have automated deletion policies. Losing valuable data due to these automated policies is a major concern, underscoring the need for reliable archival storage that can keep data safe for the long term, and systems that offer professionally managed storage solutions to ensure data remains accessible for years.

On the security side, providers confirm that users' concerns mainly relate to the data security of their data on HPC systems, which is amplified when sensitive or proprietary data is involved and there is a possibility of disruptions caused by security breaches such as hacking, which can result in system downtimes or loss of access to resources.

With regard to support and usability, providers noted that it is critical for users to know the quality and availability of technical support that could range from basic troubleshooting to advanced assistance with tuning and optimising. Support helps ease access to the systems, particularly when getting their code up and running on the HPC system. For example, complex systems can require significant effort to adapt code, which can be a barrier for some users if the right support is missing. Finally, regarding access and policy, providers commented that obtaining access to HPC resources often requires applying through competitive processes.

This is particularly true for large or national systems, where users need to navigate bureaucratic procedures, which can be difficult if the processes are cumbersome or unclear.

Another aspect to consider is the contractual agreements made between providers and users. Usually, researchers from universities can access HPC resources for free from public HPC providers, such as national HPC centres or universities hosting clusters. For users paying for HPC services, whether from public or private providers, key factors include the terms of the contracts, the pricing models and the ability to flexibly scale resources as required. The pricing model often mentioned in the answers was the 'fixed price per core hour' model.

The results from the web-based questionnaire were similar to the feedback provided during the dedicated interviews. Participants referred to the 'abundance of computational power' as the first major concern and 'sufficient RAM to accommodate intensive workloads' as the second major concern when accessing and using HPC systems. Opinions diverged for the third major concern: users and developers mentioned 'reliable and high-speed network and connectivity', while providers referred to 'high-speed file system/data storage'.

Findings from both the web-based questionnaire and the interviews revealed a shared understanding between users and providers of the most critical resources for HPC systems. The two groups agreed that 'abundance of computational power' and 'sufficient RAM' are the main priorities when accessing and using HPC systems. These core concerns reflect the need for systems capable of handling complex, resource-intensive tasks efficiently. This alignment suggests that providers are aware of, and responsive to, the fundamental requirements of their users, particularly in academic and research contexts.

However, differences emerged when ranking the third most important concern. Users prioritised 'reliable and high-speed network and connectivity', reflecting their focus on seamless data transfer and access, which are crucial for their day-to-day interactions with the system. In contrast, providers ranked 'high-speed file system/data storage' higher, emphasising the importance of file system performance in maintaining overall system efficiency, especially in environments with multiple active users. This divergence highlights the distinct role each group plays: users are concerned with their immediate experience, while providers focus on ensuring the infrastructure's long-term reliability.

System performance issues were also discussed in the interviews, with concerns about queuing system overhead raised by both groups. Users found job scheduling delays and network bottlenecks frustrating, while providers acknowledged the complexity of managing these resources efficiently. Providers highlighted the need for responsive file systems in shared environments, recognising that slow file systems can degrade the performance for all users. Addressing these concerns will be key to optimising both user experience and system efficiency in HPC environments.

Perspectives on the uptake of HPC technologies by new users

During the stakeholder consultations, some questions focused on how HPC resources could become more accessible and extended to new users. Stakeholders with experience in the HPC field provided their perspectives alongside those who had not previously worked with HPC.

In the dedicated interviews, the following recommendations were made.

- 1. Simplification of the user experience.** Users and providers emphasised the importance of simplifying the user experience from both hardware and software perspectives. A key suggestion was the development of GUIs to make HPC systems more intuitive, particularly for users that are less familiar with command-line operations. A simplified interface would reduce technical barriers, allowing new users from various disciplines to leverage HPC without requiring extensive technical knowledge or training. This measure has already been implemented in some HPC systems, and one HPC provider mentioned that it has been successful among users.
- 2. API protocols for accessing HPC centres.** Stakeholders recommended the creation of API protocols to facilitate easier access to HPC resources. Standardised APIs would enable users to integrate HPC more seamlessly into their existing workflows, improving efficiency and encouraging innovation in areas where computational resources are crucial, such as scientific research and data analysis.
- 3. Awareness and education initiatives.** Raising awareness of HPC is considered vital to broadening its accessibility. Stakeholders advocated for hosting events, workshops and summer schools aimed at students and researchers. Such initiatives would not only build technical expertise but also enhance understanding of the practical applications of HPC. This educational outreach is key to nurturing the next generation of HPC users and bridging the knowledge gap between experienced professionals and emerging talent.

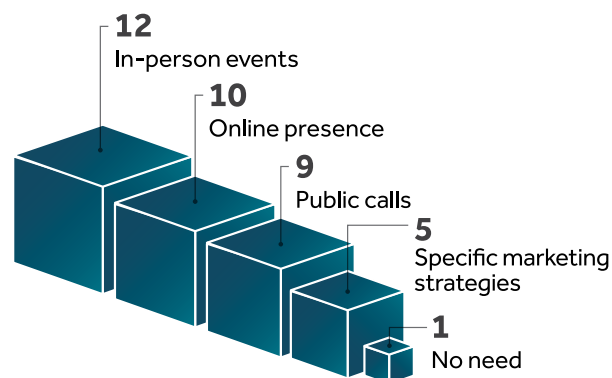
4. Clear and non-sales-oriented information delivery. A common concern among stakeholders was the need for clear and accessible information about HPC systems, presented in a more informative and less sales-driven manner. IT managers and decision-makers often struggle to filter out excessive promotional material from vendors. Providing straightforward, technical explanations and unbiased guidance would help organisations make informed decisions about adopting HPC technology without being overwhelmed by marketing messages.

5. Showcasing success stories and use cases. To further demonstrate the value of HPC, stakeholders recommended highlighting real-world success stories and evidence-based use cases. By showcasing examples from different sectors, potential new users – particularly in industries that have yet to fully adopt HPC – would gain a better understanding of its practical benefits. Demonstrating tangible outcomes such as increased productivity, faster processing or significant discoveries would encourage more organisations to invest in HPC.

These recommendations indicate that enhancing the accessibility of HPC technology requires not only technical improvements but also greater educational outreach and transparent communication. Stakeholders from various backgrounds agree that making HPC more user-friendly, increasing awareness and providing clear, practical information are essential steps in driving its wider adoption across academia and industry.

More specifically, in the web-based questionnaire, HPC providers employ diverse strategies to attract new users, including in-person events, online presence and participation in public calls (see **Figure 10**).

Figure 10. Marketing strategies used to attract new users to HPC, from the HPC providers' perspective



When it comes to selecting a specific HPC provider, users' responses reveal a range of factors influencing the decision. A significant portion of users, 41 %, reported that they were not influenced by specific marketing strategies, while 40 % noted that public calls played a crucial role in their choice. Online presence also emerged as a key factor, with 23 % of users finding it appealing. In-person events were important for 15 % of users, and 13 % were drawn to sophisticated user environment initiatives, such as user-friendly interfaces or advanced programming options. Only a small percentage, 6 %, were influenced by specific marketing strategies.

Perspectives of participants who had never worked with HPC

Of the five participants who had never worked with HPC, two were from start-ups, and the remaining represented a large company, an academic institution and a research institute.

Respondents who had never worked with HPC provided the following reasons for not utilising HPC resources.

- Budget limitations, data transfer and latency concerns, and a lack of experienced human resources were common barriers mentioned by academic institutions, research institutes, start-ups and large companies.
- Bureaucracy and lack of expertise in submitting proposals were also cited, particularly by start-ups and large companies. Respondents found it challenging to fill out HPC applications correctly due to insufficient technical knowledge, and they felt there was a lack of support when applying.
- Those unfamiliar with HPC systems noted that cloud-based alternatives, such as AWS, were quicker and easier to use, with fewer application hurdles.

Among the participants who expressed the belief that HPC could benefit their organisations (large companies, academic institutions and start-ups), the following important considerations were identified.

- Availability of in-house expertise, awareness of HPC services and relevance of HPC to their work were seen as critical factors. In addition, securing adequate funding and proper planning were deemed essential.
- Infrastructure, data management, security and software or toolchain compatibility were also highlighted as important by respondents across different sectors, including academic institutions, research institutes, start-ups and large companies.

Although they could see the potential benefits of using HPC, all participants agreed that the main constraints preventing them from setting up their own HPC systems were budget limitations and a lack of skilled personnel. The following questions explore how these barriers to HPC adoption can be addressed.

- **Simplifying the application process.** How can the administrative complexities of the HPC application process be reduced, and what forms of support could be offered to facilitate easier access?
- **Bridging the expertise gap.** What initiatives, such as training programmes or workshops, could be introduced to equip organisations with the necessary technical knowledge, particularly for those with limited experience?
- **Addressing budget constraints.** Could the introduction of more flexible pricing models, such as pay-as-you-go schemes, make HPC more financially accessible to smaller organisations and institutions?
- **Improving technical infrastructure.** What measures could be taken to improve data transfer capabilities and reduce latency, in order to address the technical challenges faced by potential HPC users?

05

High-performance computing technologies, ecosystems and trends

HPC technologies involve both computing hardware and a highly optimised software stack (in terms of both system software and parallel applications) that aim at superior performance and scalability to solve complex problems or process large datasets. At the same time, the HPC environment refers to the specific setup or infrastructure designed to handle HPC tasks. This section describes the current state of the art in HPC technologies, environments and trends, with the intent of supporting the concepts mentioned in the survey.

HPC hardware technologies

CPUs are undoubtedly still very important workhorses for HPC. Currently, there are several key players delivering CPUs for the HPC market: Intel and AMD are the most important vendors for x86-based CPUs worldwide [1]. However, ARM-based CPUs are also gaining momentum worldwide and deliver production-ready HPC chips [14] and chips for cloud computing [15]. Europe invests in HPC CPU development in two directions: (1) ARM-based processors [16] and (2) RISC-V-based processors. Additionally, EuroHPC has recently launched a RISC-V-based CPU development initiative focused on HPC [17].

GPUs are currently the most important accelerators for general-purpose accelerated computing in HPC and, as such, are present in almost every pre-exascale and exascale supercomputer. NVIDIA plays a dominant role in this area with its accelerators, and also provides a rich software stack.

The current trend is to develop unified and coherent memory between CPUs and GPUs in a single compute node. AMD, NVIDIA and Intel provide this feature at the system software level, but it is important to mention the AMD APU series [18], which is a chiplet-based system on a chip that contains both CPU cores and a GPU accelerator, along with shared high-bandwidth memory (HBM). Recently, the AI community and industry have become heavy users of GPU-accelerated supercomputers. Therefore, there is a trend in the architecture design of the chips to devote more silicon to special compute units, which, among other benefits, accelerate AI workloads [19].

There are other types of computing devices optimised for specific tasks that are being used in HPC applications, such as remote sensing [20]. These include FPGAs, the circuitry of which can be reprogrammed, unlike ASICs, which are optimised to perform a specific task or set of tasks with high efficiency, speed and lower power consumption [21] [22].

Another notable aspect of HPC systems is storage [23], [24]. The most important part of parallel computing is scratch storage, which is short-term storage specifically designed for HPC to handle highly concurrent accesses, exhibit low latency and deliver very high throughput. Scratch storage is usually implemented using parallel file systems such as Lustre or GPFS [23], [25] for shared storage or node-local storage, if data is private for individual tasks. In recent years, burst buffer technology has also been adopted in HPC [26], [27].

Currently, the most widely used network technologies in HPC are InfiniBand and Ethernet-based Slingshot [28]. However, there is a European interconnect, BXI, developed by Eviden [29], and a Japanese interconnect, TOFU, developed by Fujitsu for RIKEN's machine (K-computer and Fugaku) [30], [31]. China also develops interconnect technologies for their home-grown machines [32].

HPC software technologies

Over the last few years, Linux has indisputably become the leading operating system for HPC, as can be perceived from the TOP500 list. Preferred distributions are Red Hat Enterprise Linux, SUSE Linux Enterprise Server and some popular free versions, such as Rocky Linux. In order to interact with the operating system and its applications, a user interface is necessary. Remote terminal access using the Secure Shell (SSH) Protocol is still the most popular interface for advanced users. Another option is based on GUIs. Popular options include the Remote Desktop Protocol (RDP) or X11. Some HPC providers provide a web

interface that allows users to transfer data to the target cluster, define high-level workflows and execute them on distributed HPC resources, even across multiple clusters at the same time (e.g. Open OnDemand).

For the overall system, the job manager attempts to maximise the usage of its resources and contributes to its stability. There are several legacy job schedulers from different vendors: PBS Pro, Sun Grid Engine, Torque/Moab or IBM LSF. Currently, Slurm is one of the most widely used job schedulers in Europe and worldwide, being both free and open source. The more recent scheduler developed by LLNL for exascale machines is FLUX, which will be installed and used on the exascale machine El Capitan at LLNL [33]. FLUX addresses the key challenges: (a) throughput (being able to schedule a massive number of jobs); (b) co-scheduling (complex coupling of jobs); (c) job communication and coordination; and (d) portability [34]. The European representative in this area is the OAR scheduler [35], developed by Inria, but it has been used only on several smaller systems.

Finally, there are several programming models that unify parallel application development and improve the portability of parallel applications. They are divided into the following categories:

- distributed memory – the most common is message-passing interface (MPI) library;
- shared memory – OpenMP, Pthreads, TBB and OmpSs are typical representatives;
- distributed shared memory – PGAS implementations like UPC [36], Chapel [37], SHMEM [38] [39], Coarray Fortran [40].

One can also use a combination of distributed and shared memory models, such as MPI+OpenMP. Other important combinations are between MPIs and programming models for GPU accelerators (CUDA, HIP, SYCL, OpenACC, OpenMP, Kokkos, ALPAKA, etc.). These are mostly developed or integrated into open-source projects by individual vendors, which results in limited portability. Herten [41] looks at the compatibility between hardware devices and the C/C++ and Fortran languages. The new wave of programming models and runtimes focuses on portability among different accelerators (Kokkos/Raja) or task-based processing (Legion, PaRSEC) [42]. MPI implementations (OpenMPI, MPICH, MVAPICH, Intel MPI, etc.) are distributed as libraries. Others are supported by the compilers as, in many cases, they are extensions of standard programming languages like C, C++, Fortran and others.


HPC technology environments

HPC technology environments (also known as HPC environments) include the following (see Figure 11).

On-premises facilities, storage and network resources located in the same physical location. This requires significant upfront investment but offers more control over resources, which can be used internally or offered to the public under various funding schemes. Resources are often shared among multiple users.

Cloud-based services, mostly offered as high-performance computing as a service (HPCaaS) or through cloud computing on-premises supercomputers [43]. HPC resources are offered on demand for a fee and are typically managed by large technology companies like Amazon, Microsoft and Google. This is highly scalable and customisable but often more expensive and less efficient than traditional on-premises HPC [44].

Figure 11. HPC environments



Hybrid HPC merges both worlds and consists of solutions offered by cloud service providers as a middle ground to entities that already have on-premises infrastructure but might need to expand some part of the computation (e.g. requesting GPU resources) or data storage, utilising the cloud for higher availability. No practical examples of this paradigm are known to the wider HPC community.

It is quite common for organisations – such as COEs – to have **shared clusters** in order to maximise their usage [44]. Moreover, grid computing offloads tasks from busy machines to idle ones in the same grid or network, further maximising resource utilisation [45]. Distributed computing, conversely, uses multiple machines in the same network to perform the same task.

HPC edge computing brings high-performance processing closer to data sources, reducing latency and costs, and enhancing privacy and scalability [46]. This is beneficial for applications like autonomous vehicles, healthcare diagnostics and smart city management systems. A limitation is available performance, which is very low compared to on-site facilities HPC.

Figure 12. Summary of the HPC resources used and acquired by the users

HPC environments/services used		Costly to acquire/operate	
1	On-premises HPC clusters	1	Emerging technologies (e.g. quantum computing)
2	Shared HPC clusters	2	On-premises HPC clusters
3	Cloud-based HPC	3	High-performance edge computing

5.1 Stakeholder consultations

HPC environments

In the web-based questionnaire, users reported using multiple HPC environments and services to meet their computing needs, with some options proving to be more financially challenging than others. On-premises HPC facilities and COE clusters are the most used, with nearly 69 % of users relying on them. However, they also ranked as the most financially challenging, as noted by 33.6 % of users. Emerging technologies, such as quantum computing, were identified as the most financially demanding, with 34.5 % of users citing them as difficult to acquire. Shared HPC clusters and cloud-based HPC services are also widely utilised, with 31 % and 28.4 % of users, respectively. Despite their popularity, they present fewer financial challenges – only 6 % of users found shared clusters and HPCaaS financially challenging, while 11.2 % expressed concerns about cloud-based HPC. Hybrid HPC environments are used by 24.1 % of users, but only 6.9 % find them financially difficult to manage. High-performance edge computing, used by 9.5 % of users, presents financial challenges for 12.9 % of users. A significant portion of users, 38.8 %, expressed uncertainty about which HPC environments are the most financially challenging, indicating a potential gap in understanding or mentioning a dissimilar experience. Grid computing, used by 12.9 % of users, was considered financially challenging by only 3.4 %.

Overall, as summarised in **Figure 12**, while a variety of HPC environments are employed, the financial viability of emerging technologies and on-premises facilities pose significant challenges to many users.

Comparing HPC providers with owning a cluster

In the web-based questionnaire, users were specifically asked whether they have their own cluster. An organisation having its own cluster refers to a dedicated network of interconnected computers or servers that work together as a single system to perform HPC tasks. This setup allows the organisation to efficiently process large datasets, run complex simulations and manage computationally intensive applications, providing advantages such as enhanced performance, scalability, control and cost-effectiveness.

71 % of users have their own HPC cluster, while 29 % do not. Users with their own clusters have outlined several key advantages, whereas those without have shared their reasons for not having them.

Table 4 summarises their answers.

Table 4. Reasons cited by users for having or not having their own HPC cluster

Do you have your own cluster?	
<p>Yes - Why?</p> <p>Independence and control. The ability to develop and run code independently, without depending on external facilities, along with complete control over software, data and system access, and customisation tailored to specific applications and needs.</p> <p>Accessibility and flexibility. Immediate access to computational power ensures that users face no delays waiting for resources or responses. This setup offers exceptional flexibility, with no time limits on usage, making it easier to test, debug and prototype effectively.</p> <p>Cost and resource management. Potential for cost savings by avoiding external service fees and efficient management allows for better planning and oversight of HPC workloads and resources.</p> <p>Data management and compliance. Enhanced control over data privacy and compliance with regulations like the GDPR or direct management of data storage and post-processing.</p> <p>Support and development. Immediate, tailored support for specific needs and facilitates code development, prototyping and training with fewer constraints.</p>	<p>No - Why?</p> <p>Insufficient initial investment. The primary concern is the substantial initial cost required for acquiring and setting up an HPC cluster. This includes not only the hardware itself but also the associated infrastructure and installation costs.</p> <p>Lack of specific expertise. Many respondents pointed to a shortage of specialised knowledge needed for the installation, configuration and ongoing maintenance of HPC systems. This includes the technical skills required to manage and operate such complex systems effectively.</p> <p>Difficulty in predicting future computational needs. There is considerable difficulty in forecasting future computational demands, which makes it challenging to determine the appropriate scale and configuration of an HPC system.</p> <p>Space and power requirements. The significant physical space and power requirements necessary to house and operate an HPC cluster pose logistical challenges. Ensuring adequate facilities to support such infrastructure is a substantial barrier.</p> <p>Integration with existing systems. Integrating an HPC cluster with current IT infrastructure and workflows can be complex and problematic, potentially disrupting existing processes and systems.</p> <p>Other factors with less importance</p> <p>Technological advancements. The rapid pace of technological change in HPC hardware and software makes it difficult to keep up with the latest innovations and maintain a competitive edge.</p> <p>Risk of obsolescence. The potential for hardware and software to become outdated quickly raises concerns about the long-term viability of investing in an in-house HPC system.</p> <p>Security and compliance issues. Ensuring that an in-house HPC system meets security and compliance standards adds another layer of complexity.</p>

Most users use several HPC clusters for several reasons. 50 % only use one cluster, 22 % use more than four HPC clusters and the remaining use two or three clusters. Most users employ several HPC services from external HPC providers.

When asked about their preferred method of accessing HPC resources, 18 % of users favoured a private cluster (HPC environment that is dedicated solely to a single user or organisation), while 23 % preferred a shared cluster (HPC environment where multiple users or organisations share the same resources). Nearly 27 % opted for HPCaaS, and the remaining 32 % indicated no particular preference. Regarding the choice between publicly and privately funded HPC systems, nearly 70 % of users preferred public systems, only 4 % favoured private ones and about 26 % expressed no specific preference. This could reflect trust in public resources or satisfaction with the services they provide.

Multiple factors drive the use of more than one HPC system, with nearly 40 % of users citing the need to manage HPC-system-specific projects and workloads as the primary reason. Capacity planning and scalability are also significant considerations, influencing over 25 % of users. Cost optimisation and budget allocation account for about 15.5 % of responses, while 11.2 % of users rely on multiple systems to access specialised expertise and services. Geographic distribution is a factor for 10.3 % of users, and compliance and data sovereignty concerns affect 4.3 %. Lastly, a small percentage, 1.7 %, use multiple systems to mitigate vendor lock-in.

Services provided

In terms of HPC services, all HPC providers offer hardware infrastructures and software tools and applications. Sixteen provide data management and storage, along with technical support and consulting. Fifteen include training services and 10 offer other services, such as cloud computing. Among the types of HPC services, on-premises HPC facilities and COE clusters are the most prevalent, provided by 16 stakeholders, followed by cloud-based HPC (8), emerging technologies like quantum computing (7) and hybrid HPC (6). Grid computing and high-performance edge computing are the least common, each offered by only one stakeholder.

In terms of technical storage aspects, the types of storage most provided to users are parallel file systems (e.g. scratch) (17) and local storage on compute nodes (11). These are followed by object storage (e.g. S3, Swift) (10) and file storage shared across multiple HPC systems (e.g. project) (10). Among these options, parallel file systems (e.g. scratch) are the most preferred by users (14), with file storage shared across multiple HPC systems (e.g. project) being the next most preferred (5). The interfaces provided to users are delivered through various approaches, which are summarised in **Figure 13**.

The information collected in the dedicated interviews also supports some findings from the web-based questionnaire, as summarised in **Table 5**.

Figure 13. Different types of interfaces provided by the HPC providers

Command-line interface

Batch job submission (e.g., Slurm, PBS/Torque)	17
SSH Protocol	15
File transfer and management (e.g. SCP, SFTP)	13
Containerisation interfaces (e.g. Docker, Singularity)	13
Environment and software management (e.g. Module System)	12
Data management tools	11
Other	2

Data access/transfer

FTP/ SFTP	10
SCP	10
Data transfer nodes	8
GridFTP	4
Globus	4
Other	3

Graphical user interface

Visualisation software (e.g. ParaView)	8
Job management tools (e.g. SlurmView)	7
Cluster monitoring dashboards (e.g. Ganglia)	5
File management tools (e.g. Cyberduck)	5
Remote desktop access (e.g. RDP, Teamviewer)	5
Workflow management platforms (e.g. Pegasus)	1
None	3
Other	5

Network access

SSH Protocol	16
VPN	7
Web portal login	5
RDP	3
APIs	1
Other	1

Table 5. Summary of users’ technical considerations

Technical considerations
CPU and GPU are the most common architectures. User preference depends on the specific applications and simulations.
InfiniBand is widely used and is a good solution for internode communication.
Interfaces available are mainly SSH and batch job submission for on-premises HPC systems.
Slurm is universally used as a workload manager.
Local storage on compute nodes and parallel file systems are most desirable for users.
Common programming languages mentioned: Python (ease to use), C++ (higher performance), Fortran.
Physical tokens, two-factor authentication and public key usage are widely implemented by and acceptable for users.

5.2 Emerging high-performance computing technologies

Considering that HPC technological development occurs at several levels in both software and hardware, only some key areas at the frontier of what defines HPC will be covered in these subsections: exascale and quantum computing. While exascale is a type of computing that is already a reality, quantum computing is still limited to small research systems with few applications. Nevertheless, both represent technologies that are pushing the performance frontier and even the definition of HPC.

Exascale computing

The computational power of an HPC system is reported in floating-point operations per second (flops). In the case of the TOP500, the machines are sorted based on performance measured using the Linpack benchmark.

Over the last decade, there has been a race to be the first to create a system capable of exascale computing (i.e. of performing at least 10^{18} calculations per second in double precision). This poses several technological challenges, including the necessity of porting existing codes to handle heterogeneous hardware and limitations in memory and data transfer speeds and energy consumption [47].

In May 2022, the supercomputer Frontier, installed at the Oak Ridge National Laboratory, United States, became the first computer to surpass that milestone and, as of August 2024, it still holds first place on the TOP500 list with a maximum performance of 1.2 exaflops. The second US exascale system is the Intel-based Aurora, featuring Intel Xeon Max CPUs with HBM and Intel Data Center Max GPU, and is installed at the Argonne National Laboratory [48]. The third exascale system will be El Capitan at LLNL. This system is built on AMD MI300A accelerated processing units, which integrate CPU cores and GPUs into a single package with shared HBM. El Capitan is anticipated to exceed two exaflops in double precision [48].

In Europe, the HPC centre selected to host the first exascale machine Jupiter is located at the Jülich Supercomputing Centre in Germany [49]. Jupiter has two main partitions. A cluster partition will be built on a European Rhea1 processor. The Booster partition will feature roughly 6 000 compute nodes to achieve the compute performance of one exaflop. Each node has four NVIDIA GH200 superchips (i.e. four GPUs, each closely attached to a partner CPU). The second European exascale machine – Alice Recoque – will be based in France [50]. Other exascale-grade machines are also installed in China as outlined below.

Quantum computing

Quantum computing, while not yet ready for practical applications, has seen consistent development by technological giants like Google and IBM. This emerging technology promises a paradigm shift from classical computing, which relies on bits, to a system based on quantum bits that exploit quantum mechanics. This shift could revolutionise fields such as cryptography and cybersecurity, drug discovery and material science, financial modelling, supply chain and logistics, AI and machine learning, weather forecasting and climate research, energy and quantum chemistry by solving problems deemed intractable for classical systems [51].

Despite the potential, scaling and maintaining stability in quantum systems remain significant challenges. These issues have prevented developers from achieving 'quantum supremacy', a term used within the community to describe a quantum system capable of solving problems that classical computers cannot handle effectively [52] with many extensions and improvements in subsequent years. The first 10 chapters cover the circuit model and the main quantum algorithms (Deutsch-Jozsa, Simon, Shor, Hidden Subgroup Problem, Grover, quantum walks, Hamiltonian simulation and HHL). Additionally, the establishment of benchmarks for quantum computing and the practical relevance of solved problems are still areas of contention.

Recent developments in hybrid classical-quantum computing offer a promising pathway forward. These systems combine classical computing architectures with quantum processors, acting as co-processors or accelerators to tackle new classes of problems [53].

For instance, hybrid systems are being explored for more efficient drug discovery [54], complex financial modelling [55] and climate simulation, where quantum algorithms accelerate specific computations that pose bottlenecks in classical processes.

Looking ahead, organisations like ETP4HPC [56] predict that quantum computing will play a pivotal role in the HPC community, particularly as heterogeneous architectures become more prevalent. Bartsch et al. outline the anticipated requirements and implications of quantum computing integration, emphasising the transformative potential of hybrid systems in extending the capabilities of current technologies.

As the field progresses, the integration of quantum and classical systems is expected to unlock unprecedented computational capabilities, paving the way for breakthroughs in multiple domains. The journey towards practical quantum computing continues to be a compelling story of technological advancement, underscored by both challenges and immense opportunities [57].

5.2.1 Stakeholder consultations

In the dedicated interviews, stakeholders were asked about the future of HPC technologies. From the feedback provided, the following three topics were most prominently mentioned.

- **Hybrid architectures.** With the rise of AI and machine learning, there is a growing interest in hybrid HPC architectures that incorporate CPUs, GPUs and AI accelerators, among others. Users are concerned about the compatibility of their workloads with these evolving technologies and whether they will need to invest in new skills or code adaptations. Moreover, hybrid architectures have the potential to maximise hardware performance. While GPUs provide massive parallelism at a lower cost, they are difficult to code for.
- **AI integration.** As AI becomes more prevalent, users are interested in how HPC systems will support AI workloads and whether this will affect the traditional scientific computing models they rely on. This tendency was mentioned by interviewees from Europe, Canada and the United States, where government organisations are investing directly in AI computing.
- **Quantum computing.** Quantum computing is viewed by some as having significant potential, though its adoption is likely to be selective. Despite its substantial computational power, it is anticipated that not all users will incorporate it into their operations. However, certain industries, such as finance, are already investing in quantum computing and exploring its capabilities for specialised applications.

In the web-based questionnaire, emerging technologies were extensively explored. Looking at the growth of HPC environments, most of the participants agreed that hybrid HPC (88 participants) and cloud-based (82) will grow faster, followed by on-premises HPC facilities and shared clusters (60). The main

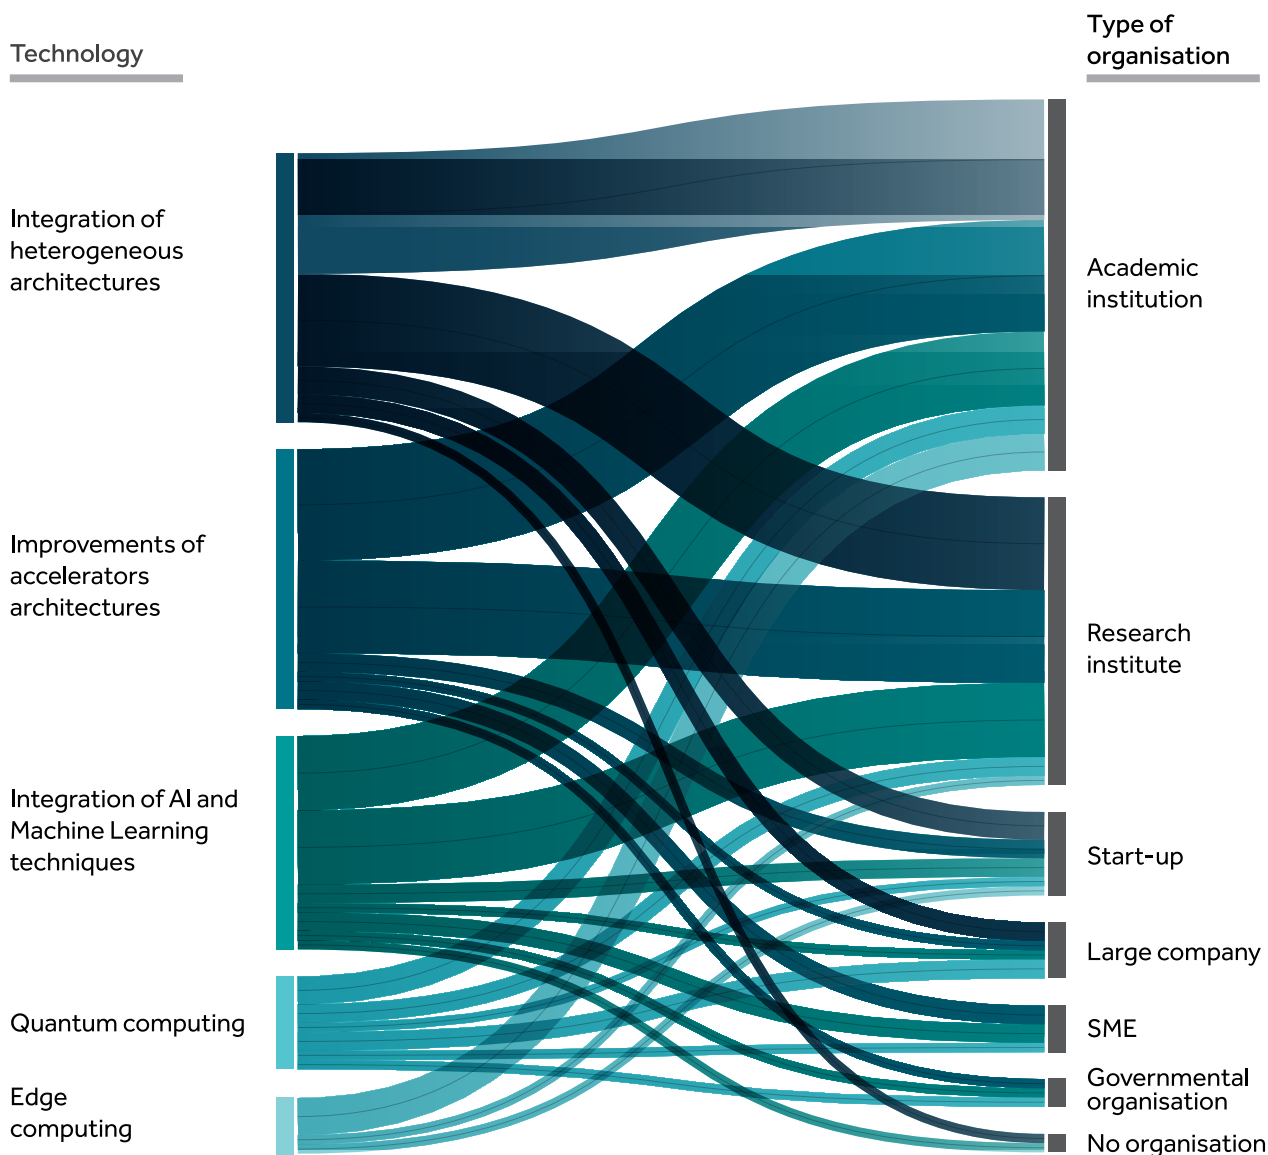
factors affecting these views are cost and accessibility, scalability and versatility, along with emerging technologies and future trends, with a special focus on quantum computing. However, there are slight differences in justifications:

- **users** place significant emphasis on the scalability and practicality of cloud-based HPC, especially for AI and data analysis;
- **developers** lean towards the value and flexibility of hybrid solutions, with a forward-looking view on quantum computing;
- **providers** value the balance and optimisation provided by hybrid HPC, highlighting the necessity of on-site HPC facilities for high-performance research tasks and the promising future of quantum computing.

Developers were asked in the questionnaire to share their opinions on the future of emerging technologies and on the status HPC system hardware. More than 50 % of developers consider that 'integration of heterogeneous architectures' (56.9 %; 29 participants) and 'improvements of accelerator architectures' (e.g. specialised accelerators for specific tasks) (54.9 %; 28 participants) are potential emerging technologies that might contribute to increasing HPC accessibility for different types of organisations (see **Figure 14**).

According to the developers, several types of hardware will bring more competitiveness in addressing HPC future challenges by increasing performance and energy efficiency. The most valued hardware were GPUs (88 %; 45 participants) and specific AI accelerators (66.6 %; 34 participants) (see **Figure 15**).

Figure 14. Emerging technologies that could make HPC more accessible




There was no consensus on future trends in ARM and x86 CPU architectures. ARM led in popularity with 39 % of respondents (20) favouring it, followed by those supporting both architectures equally at 35 % (18). The study showed that x86 is favoured by 12 % (6) of participants, while 8 % (4) prefer other architectures and 6 % (3) believe none of the mentioned architectures will dominate.

The participants that prefer ARM architecture gave the following reasons:

- it achieves high performance and maintains energy efficiency;
- it is seen as the most flexible and affordable in the market;
- many vendors in the AI market favour ARM for its potential to generate higher profits;
- it offers a more interesting instruction set architecture compared to its competitors.

Figure 15. Types of hardware expected by participants to enhance competitiveness in future HPC challenges, by organisation type



The participants that favour both architectures stated the following:


- ARM is growing strongly and garnering increasing interest, but x86 continues to dominate in the CPU/HPC space for now;
- each architecture has specific strengths – ARM offers better performance per watt, while x86 continues to excel in traditional CPU/HPC applications;
- ARM's efficiency comes with the need for more developer effort to optimise performance, whereas x86's established ecosystem provides a robust foundation.

RISC-V is expected to play an increasingly important role in HPC in the future, once any licensing issues are removed. Its importance will grow with the development of the software stack.

Table 6. Opinions about quantum computing

Will quantum computing shape future trends?	
<p>Yes - Why?</p> <p>Advancing integration. Quantum computing is rapidly evolving and may soon integrate with HPC systems, transforming computational science and industry workflows.</p> <p>Breakthrough potential. It holds promise for breakthroughs in drug discovery and encrypted communications, offering advanced calculations and energy efficiency.</p> <p>Development progress. Significant strides are being made in quantum-compatible algorithms and compilers, with a focus on quantum and hybrid architectures.</p> <p>Future impact. Quantum computing could become a key component of HPC systems, enhancing computational power and revolutionising problem-solving.</p>	<p>No - Why?</p> <p>Current status and timeline. Quantum computing is still in its early stages and remains largely experimental, making it unlikely to have a meaningful impact in the near to medium term, with practical applications expected to be several years away.</p> <p>Limited applicability. It is mainly useful for specific problems like Fourier transforms, and its applications are currently limited to a narrow range of issues.</p> <p>Future uncertainty. The future of quantum computing remains uncertain, with predictions suggesting it could take 10 to 15 years before it becomes a viable HPC resource.</p> <p>Current perception. Quantum computing is seen as a speculative and experimental technology with substantial potential, but it is not yet a realistic or practical option, similar to the status of fusion energy.</p>

Figure 16. Responses given by the developers on which sectors will benefit from quantum computing

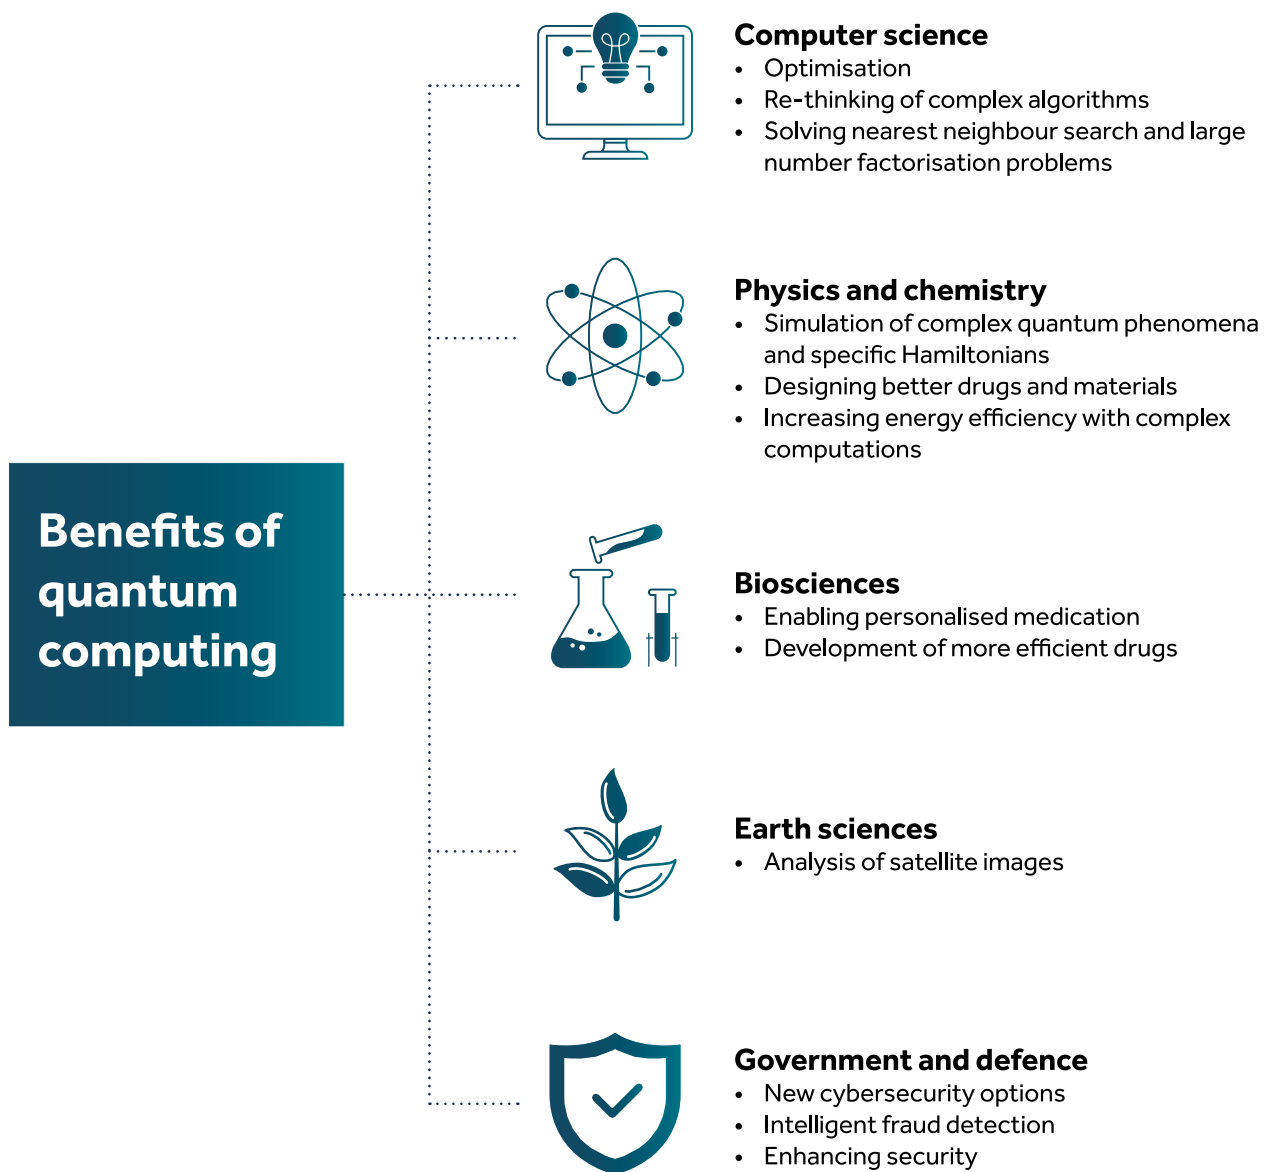


Quantum computing

Regarding the future of quantum computing, developers have divided opinions. 51 % of developers (26) do not believe this technology will shape future trends, while 49 % (25) believe it will. The reasons given for each perspective are summarised in Table 6.

Despite this division, most developers agree that quantum computing might benefit several sectors (88 %; 45 participants), including physics (26), material science (20) and banking and finance (20), as shown in Figure 16. Moreover, developers have enumerated some specific benefits of this technology for various sectors (see Figure 17).

Figure 17. Examples given by the developers of how quantum computing might benefit different sectors



Exascale computing

In the case of exascale computing, 68.1 % of developers (35) consider that it will be widely used in the future, whereas 31.3 % (16) do not agree with this presumption (see **Figure 18**). Developers pointed out several advantages of this technology:

- the possibility to perform billions of computations per second using a combination of CPUs and GPUs, significantly enhancing data processing and analysis capabilities;
- high performance for the incoming demand for data processing and complex analysis;
- the ability to address cutting-edge scientific problems and support new computational paradigms like digital twins, AI, advanced models and extensive datasets related to climate change and renewable energy;
- with the integration of more heterogeneous architectures such as GPUs, power consumption may be further reduced;
- growing funding for exascale HPC centres, with widespread adoption in Europe and the United States.

Developers that are sceptical about the rise of exascale computing offered the following justifications for their doubts.

- Exascale computing is primarily needed for specialised applications and is not widely required by the general population.
- Access to exascale computing is restricted and software development struggles to keep up with rapidly evolving supercomputer architectures.
- Many applications do not effectively parallelise or utilise the full potential of exascale systems, leading to underutilisation for most current workloads.
- Exascale computing may be premature for many industrial applications, which often require less computing power than exascale systems provide.

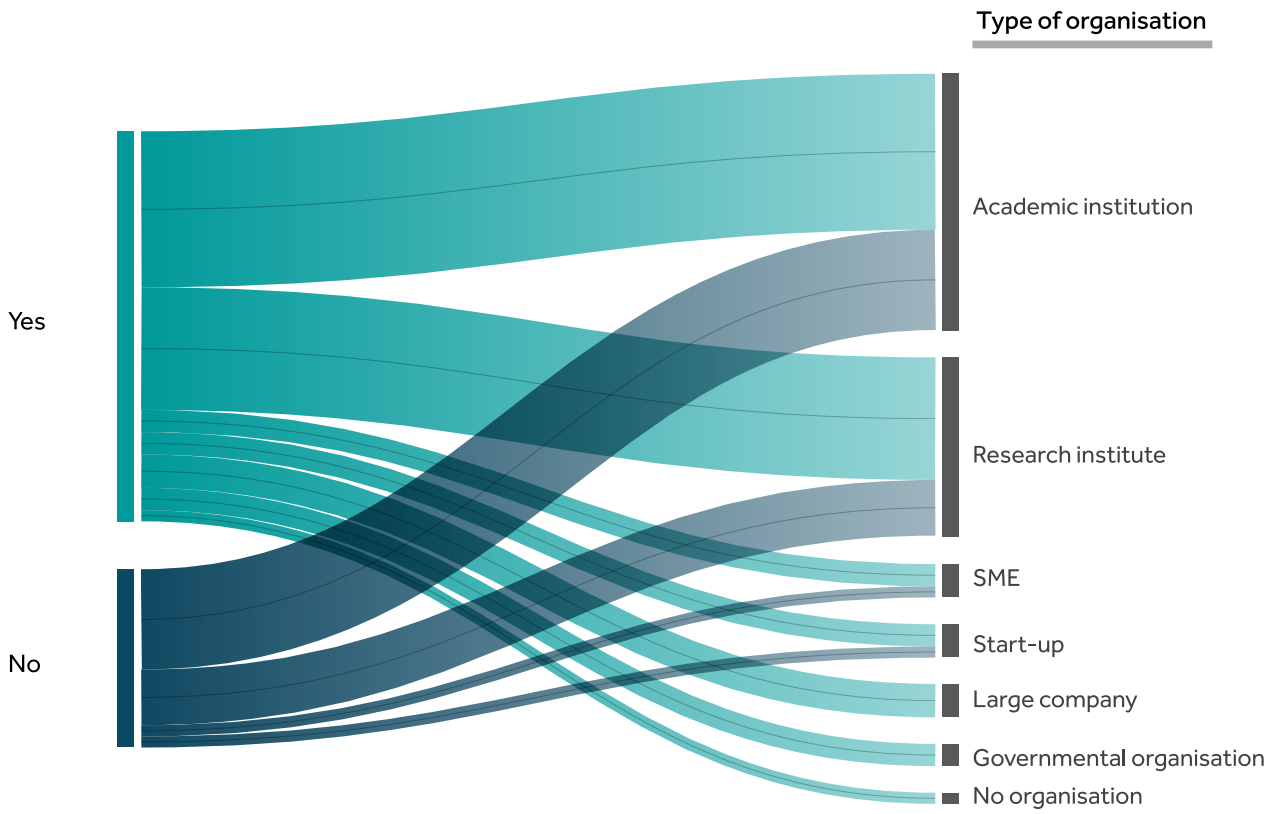
- While exascale systems offer enormous computing power, they are primarily useful for reducing operational expenses and might not be necessary for all applications. Most current HPC users are not prepared for exascale jobs, and many applications do not yet require or effectively use exascale computing.

More than 50 % of the developers believe that several sectors will benefit from exascale computing, such as energy (renewable), physics, aerospace engineering, earth sciences, biosciences, material science and astronomy. Meanwhile, only 10 % of the developers believe media and entertainment, agriculture, logistics, retail and consumer goods will benefit from this advanced technology.

Developers were also asked which aspects should be addressed to increase the accessibility of exascale computing, for software and hardware architectures. In terms of software technology, the highest scoring measures were 'portability to heterogeneous hardware architectures' (70 %; 36 participants), 'parallelisation in heterogeneous hardware architectures' (63 %; 32 participants) and 'energy efficiency' (59 %; 30 participants). In terms of hardware technology, the most voted measures were 'better hardware interconnection technology (higher bandwidth and lower latency)' (59 %; 30 participants) and 'improvements in RAM (volatile) technology' (49 %; 25 participants).

In conclusion, emerging technologies are expected to benefit various sectors, with exascale computing showing more immediate promise than quantum computing. The current focus is on improving computational efficiency to meet the demands of complex tasks, with the integration of heterogeneous architectures or the adoption of ARM or x86 architectures seen as viable solutions.

Figure 18. Number of responses given by the developers on whether exascale computing will be widely used in the future, by organisation type



06

Legal aspects

HPC poses several legal challenges that intersect with ethical considerations, data protection and equitable access.

Europe

Key issues include the handling of large datasets, particularly ones that contain personal data, which must comply with data protection laws like the GDPR in the European Union. Ethical frameworks are essential for managing the privacy of individuals, especially in sensitive applications like health data. Legal frameworks are also relevant to managing disparities in access to HPC resources, which can exacerbate digital divides and impact equity in technological advancements.

An acceptable use policy (AUP) in the context of HPC defines the rules and guidelines governing the ethical and legal use of HPC resources. It specifies who can access the system, outlines prohibited activities such as illegal or disruptive actions and sets standards for data management, user responsibilities and system monitoring.

The AUP ensures that HPC resources are used appropriately and fairly, maintaining system integrity and security while supporting legitimate research and educational objectives [58]. Many HPC systems, such as those in universities, national centres and cloud-based services, implement an AUP to govern user behaviour.

The basis of this document typically includes legal compliance, ethical use and resource management to ensure responsible access to computing power.

Compliance with the AUP is required from all users, including researchers, students, external collaborators and administrative staff, ensuring that everyone understands their responsibilities and the implications of their actions while using HPC resources.

Moreover, the legal frameworks governing collaboration and data sharing in HPC environments must address ownership, usage rights and compliance with international regulations. These frameworks ensure that data sharing and processing do not infringe on privacy rights or intellectual property (IP) laws. Overall, robust legal and ethical guidelines are crucial for navigating the complexities of HPC, ensuring data security and fostering fair access to computational resources [59].

The GDPR is crucial for HPC in Europe, focusing on protecting individuals' data in both private and public sectors. It enhances individual control over personal data, modernises and harmonises regulations, reduces business administrative burdens and builds consumer trust.

The GDPR also introduces several new rights for individuals, such as data portability and the right to erasure [\[60\]](#).

Security in HPC is complex due to its unique architecture and multi-user environments. Traditional security measures can impair HPC performance, so specialised solutions like the Science DMZ are used to balance security and efficiency. The Science DMZ is a network meant to optimise high-speed data transfers for scientific research while maintaining security by isolating data-intensive operations from the rest of the institutional network, using specialised security measures that prevent performance bottlenecks without compromising data protection [\[61\]](#) [\[62\]](#).

As HPC moves to the cloud, protecting sensitive and strategic data becomes crucial, with confidential computing and trusted execution environments playing key roles. Despite improvements from companies like Intel, AMD and ARM, integrating trusted execution environments with cloud HPC requires further research. Advances in machine learning and containerisation also show promise for enhancing HPC security [\[63\]](#), [\[64\]](#).

With the advancement of AI applications and their impact on society, in March 2024 the EU adopted the AI Act (Regulation (EU) [2024/1689](#)), which regulates AI usage within Member States. The objective is to ensure AI is safe, ethical and respects fundamental rights. The framework introduces a risk-based approach, categorising AI into four risk levels – unacceptable, high, limited and minimal – each with corresponding regulatory requirements. High-risk AI applications, like those in healthcare and law enforcement, will face stringent controls, while lower-risk systems will have fewer obligations [\[65\]](#).

To ensure responsible use of technology, integrating adequate legal and ethical frameworks into HPC and AI is crucial.

Regulations like the AI Act and the GDPR safeguard data, protect privacy and promote ethical standards while fostering innovation. As AI and HPC evolve, these guidelines will be essential for ensuring safe, equitable and ethical advancements.

United States

The legal requirements governing HPC in the United States encompass various critical areas, including funding, export controls, IP and national security [\[66\]](#). The [CHIPS and Science Act](#) of 2022 allocates substantial federal funding to strengthen semiconductor manufacturing and support HPC research, enhancing US competitiveness in advanced computing technologies [\[67\]](#). HPC technologies are subject to strict export control regulations under the [Export Administration Regulations](#). The Export Control Classification Number system determines the classification of items for export based on sensitivity [\[68\]](#). Technologies with military applications are further regulated under the [International Traffic in Arms Regulations](#).

IP protections, including patents and copyrights, safeguard innovations in HPC, encouraging R & D. Additionally, compliance with federal regulations is essential for entities involved in HPC, particularly concerning financial accountability for federally funded projects. Data security is also fundamental, with the Federal Information Security Modernization Act [\[69\]](#) mandating the protection of information systems. Lastly, HPC facilities must adhere to environmental regulations related to energy consumption and waste management. This complex legal framework is designed to foster innovation while ensuring the protection of sensitive technologies and upholding national security.

6.1 Stakeholder consultations

The legal considerations that users must address when accessing HPC systems were consistently noted by all stakeholders during the dedicated interviews, as summarised in **Table 7**.

The dedicated interviews revealed that most academic users were not fully aware of the legal framework and compliance requirements when accessing HPC providers. This could be because these users come from academia, where specialised departments typically handle such issues.

In the web-based questionnaire, participants were asked to indicate which legal protocols must be followed when accessing HPC resources. Developers and users agreed that 'data privacy and protection regulations' are the most important aspects to consider when submitting data or workloads to an HPC system, while most of the providers selected 'acceptable use policy' as the main aspect. Most of the participants ranked 'intellectual property rights' in second place, regardless of their group, likely due to its strong connection with software licensing.




Table 7. Summary of the legal requirements for accessing HPC resources as stated by users

Legal considerations
Legal considerations such as SLAs and AUPs were frequently mentioned by participants from the different types of organisations. Participants representing sectors such as biosciences and healthcare, emphasised the need to apply strict measures, including IP rights, to safeguard their interests and ensure compliance with legal and regulatory standards.

Additionally, users ranked 'acceptable use policy' and providers ranked 'ethical considerations' in second place. In third place, users ranked 'confidentiality and non-disclosure agreements', while developers and providers selected 'compliance with specific regulations, laws or institutional policies'. Providers also ranked 'data privacy' in third place (see **Figure 19**).

When asked about the HPC environments and types of services that represent the greatest legal challenges, most participants mentioned 'cloud-based HPC'. Most users and developers admitted that they 'do not know' as their second choice, while providers mentioned 'hybrid HPC'. In third place, most users mentioned 'shared HPC clusters', and providers gave an equal number of responses for 'do not know'. These results could once again be explained by data security considerations.

Figure 19. Legal obligations to consider when submitting data or workloads, by stakeholder group

 Users	 Developers	 Providers
1 Data privacy	1 Data privacy	1 Acceptable use
2 IP rights and acceptable use	2 IP rights	2 IP rights and ethical aspects
3 Confidentiality	3 Regulatory compliance	3 Data privacy and regulatory compliance

The findings from the stakeholder consultations revealed a pressing need for clarity and communication regarding the legal considerations in accessing HPC resources. Users from sectors such as biosciences and healthcare emphasised the importance of SLAs and AUPs, reflecting their concerns about compliance with IP rights and regulatory standards. The distinction between the priorities of users and providers (users focus on data privacy and protection regulations, while providers emphasise AUPs) suggests an opportunity for enhanced dialogue and understanding between the two groups. The legal challenges associated with cloud-based and hybrid HPC environments are also notable, as many stakeholders admitted a lack of understanding in these areas. This indicates a need for providers to improve user education and support, ensuring that all parties are aligned in their understanding of the legal frameworks governing data handling. Addressing these gaps is crucial for fostering trust and promoting a secure HPC ecosystem. To effectively address the need for training on the legal considerations in accessing HPC resources, a collaborative approach involving several key entities is essential. HPC providers should lead the development of tailored training programmes that cover SLAs, AUPs and compliance with data privacy regulations. Industry associations and professional bodies can support these efforts by creating standardised materials, and academic institutions can integrate legal training into their HPC curricula. Additionally, engaging legal experts to conduct specialised sessions will enhance users' understanding of the complexities surrounding cloud-based and hybrid HPC environments. This combined effort will ensure users are well-informed and foster a secure, trustworthy HPC ecosystem.

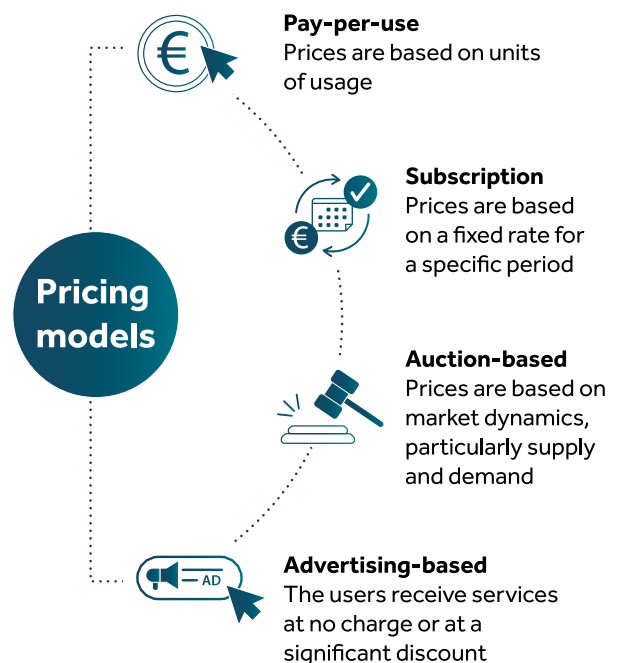
07

Financial aspects

Users can freely access public HPC resources through open-access grant competitions or calls based on a peer review process [70]. On the other hand, utilising HPC resources for commercial purposes may involve various financial models. A wide range of users – from small research teams to large corporations – have the option to pay HPC providers directly [71]. Key models include (1) pay-per-use, where customers pay only for the resources they use; (2) subscription, which involves a fixed rate for a set period; (3) auction-based pricing, where prices adjust based on supply and demand; and (4) advertising-based pricing, where services are offered for free or at a discount with revenue from advertising [72] cloud computing has revolutionized the information and communication technology (ICT) (see Figure 20).

The global HPC financial landscape combines public funding, private investment and collaborative efforts to support research, innovation and economic growth.

Figure 20. Pricing models for using HPC services



7.1 Stakeholder consultations

Financial considerations in accessing HPC resources were addressed in the web-based questionnaire and in the dedicated interviews. HPC providers offer a range of access options, with many public organisations providing free contracts primarily to academic researchers and R & D users, typically following an application review. In addition to these free options, providers also offer various pricing models, including variable pricing, metered usage and fixed pricing with limited or unlimited usage. While these models are well-documented for public providers, there is limited data available on the practices within the private sector.

In the web-based questionnaire, providers were specifically asked to provide information on the annual budget for and financial investment in HPC resources within their organisation. The allocation of these budgets to HPC services varies significantly among organisations. The majority, 61.2 % of organisations, stated they allocate between 0 % and 10 % of their budget to HPC resources, indicating a relatively modest investment. 17.2 % of organisations dedicate between 11 % and 30 % of their budget, 5.2 % allocate between 31 % and 50 %, 3.4 % invest between 51 % and 70 %, and a minimal 0.9 % allocate more than 90 % of their budget to HPC services. Additionally, 12.1 % of organisations did not provide information on their budget allocation. This distribution highlights a tendency towards conservative investment in HPC services, with only a few organisations committing a more significant portion of their budget.

Participants ranked the following concerns

related to budget allocation necessary to maintain HPC services and ensure optimal performance for users, highlighting several key areas of expenditure that receive the highest financial resources within HPC costs.

Top-ranked concerns (rank 1)

- 'Operational costs' were ranked first by seven respondents, making up 43.75 % of the total responses.
- 'Infrastructure investments' and 'hardware maintenance' received four responses each, accounting for 25 % each.

Second-ranked concerns (rank 2)

- 'Infrastructure investments' led with six responses (37.5 %).
- 'Operational costs' followed closely with five responses (31.25 %).

Third-ranked concerns (rank 3)

- 'Hardware maintenance' was ranked third by five respondents (31.25 %).
- 'Operational costs' came next with four responses (25 %).

Middle-ranked concerns (ranks 4–6)

- 'Human resources' received three responses (18.75 %) at rank 4 and four responses at rank 5 (25 %).
- 'Software resources' received six responses at both ranks 5 and 6, representing 37.5 % each.
- 'Security' emerged as the most pressing issue at rank 6, with six responses (37.5 %).

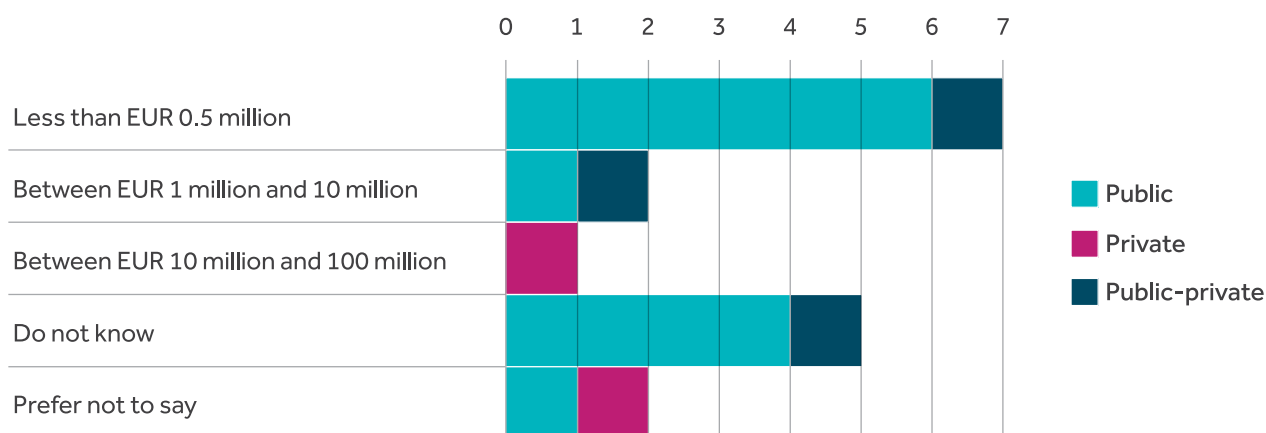
Lower-ranked concerns (ranks 7–9)

- 'Educational initiatives' were highly ranked at lower levels, with seven responses (43.75 %) at rank 7 and seven responses (43.75 %) at rank 8.
- 'Financial planning' increased in importance at ranks 8 and 9, receiving eight responses (50 %) at rank 8 and seven responses (43.75 %) at rank 9.
- 'Environmental considerations' dominated rank 9 with nine responses, accounting for 56.25 % of the total.

In terms of financial aspects, providers had the option to communicate their current financial income from providing HPC services. Some of them did not know their financial income related to these services, and most admitted it was less than EUR 0.5 million per year (see **Figure 21**).

In summary, financial investment in HPC services tends to be modest, with most provider organisations allocating less than 10% of their budget. Key expenses include hardware, operating costs and energy consumption, and providers generally report low annual income from HPC services, highlighting a conservative approach to financial commitment in this area.

Figure 21. Financial income from providing HPC services, for public and private organisations



08

Environmental aspects

HPC and AI are driving groundbreaking discoveries, free from the constraints of traditional theoretical models. In the era of big data [73], it is estimated that up to 80 % of big data contains spatial information with locational components attached [74]. HPC refers to using supercomputers to tackle computational problems that are either too complex for standard computers or would take too long to solve. It leverages supercomputing power and extensive parallel processing techniques to address complex issues through computer modelling, simulation and data analysis. Unlike a typical desktop computer, which usually has a single CPU, an HPC system comprises a network of nodes, each containing one or more processing chips and memory [75].

Energy efficiency

HPC applications demand extensive computational power, often reaching petaflop levels, and recently also exaflop according to the TOP500 list of the most powerful systems [1]. A significant challenge in modern data centres housing these high-performance facilities is the substantial energy consumption required for their operations [76]. To fully harness the potential of HPC systems, significant efforts are being dedicated to advanced programming and code optimisation techniques. However, these performance-driven initiatives often overlook energy efficiency considerations [77].

Reducing energy consumption and developing sustainable computing infrastructure have now become primary goals for the HPC community [20]. The motivation driving these efforts is in facing several challenges to build even more powerful systems. A powerful HPC cluster requires a large source of power, which has a major impact on the location of the centres. The leading five HPC systems, according to the TOP500 list, currently consume from 7.1 to 29.9 megawatts (MW). Operating a system with such high power consumption leads to energy costs of several million euro per year. Moreover, these systems show a power consumption variability, which must be handled at both the hardware and software levels. The large gap between power consumption in idle and in overloaded states may influence the stability of the power grid. For example, the Cori supercomputer infrastructure (rated fifth on the TOP500 list when it started operation in 2016) was built to handle 5.7 MW power consumption. However, its peak consumption is 4.6 MW, base consumption is 3.9 MW and idle-state consumption is 2 MW [78].

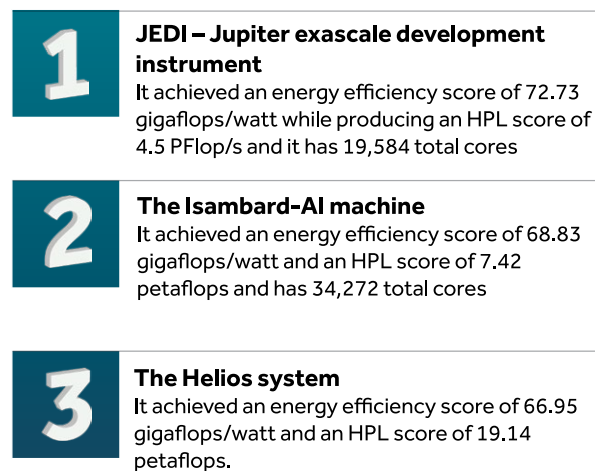
The Green500 list ranks the world's most energy-efficient supercomputers, focusing on measuring and comparing their performance per watt. This is crucial as energy consumption and efficiency have become significant considerations in the design and operation of HPC systems. By recognising and ranking supercomputers that achieve high performance with lower energy consumption, the Green500 list promotes the development and adoption of energy-efficient computing technologies and encourages the industry to prioritise sustainable computing practices. Energy efficiency is measured through a combination of power consumption data and performance benchmarks, such as the Linpack benchmark. Systems that rank on the Green500 list often feature advanced cooling technologies, energy-efficient processors and innovative architectural designs. The primary metric used to present energy efficiency is 'performance per watt', typically measured in flops per watt. This metric provides a clear indication of how efficiently a supercomputer can perform computations relative to the power it consumes. The [latest version](#) of the Green500 list is from June 2024, where the top three places are held by European supercomputers: JEDI (Germany), the Isambard-AI machine (United Kingdom) and the Helios system (Poland) (see **Figure 22**).

Greenhouse gas emissions

It is projected that, by 2030, data centres and HPC systems could contribute up to 8 % of global greenhouse gas emissions if no action is taken. Just a single exascale system may produce approximately 150 000 metric tons of carbon dioxide annually [\[78\]](#). There is increasing pressure to use renewable energy sources to power large IT servers, which is not an easy task given that these systems require a powerful and stable power source.

Figure 22. The three most energy-efficient supercomputers

Source: Green500 list, June 2024)



Consequently, major technology companies have made substantial investments in offsetting carbon emissions from their data centres, and there is a growing research focus on enhancing the carbon efficiency of large-scale computing systems [\[79\]](#) which requires urgent intervention. In this work, we present a comprehensive analysis of the carbon footprint of highperformance computing (HPC). Developing energy-aware solutions is crucial and can be achieved through various approaches at the levels of infrastructure (i.e. cooling), computing hardware, system software or applications. It would also involve optimising energy supply and generating strategies to manage heat dissipation and reduce carbon dioxide emissions. These efforts are essential for enhancing the sustainability of computational systems [\[77\]](#).

Resources

Producing HPC hardware involves intricate industrial processes and relies on valuable resources, such as rare-earth elements, extracted through extensive mining. This extraction has profound ecological and health impacts, releasing toxic and radioactive materials that contaminate the surrounding

land and communities, so addressing these issues is critical to mitigating the environmental impact of HPC technology [80]. An ongoing challenge in environmentally sustainable computational research is accurately assessing the environmental impact of software and algorithms, which is often underestimated. However, there has been growing awareness of this issue across various computational communities. Initiatives like the Green Software Foundation have emerged to prioritise sustainability in software development practices [81].

8.1 Stakeholder consultations

Environmental considerations were a concern mainly for HPC providers. Most users in the dedicated interviews acknowledged that environmental considerations are important for the future, although they view them as a secondary concern. Nevertheless, some users suggested that continuously optimising code to reduce computational power consumption would be a valuable approach to improving sustainability.

Both private and public data centres expressed a concern for enhancing the sustainability of data centres and making them green data centres. Several measures were mentioned to improve energy efficiency and sustainable practices:

- optimising energy use in supercomputing operations by using energy-efficient hardware and cooling systems to reduce power consumption;
- developing hybrid architectures, such as merging CPUs with GPUs, to further reduce energy consumption;
- implementing heat reuse in data centres, where the excess heat generated is utilised to warm local homes;
- partnering with vendors that provide environmental certificates for their products and have system retirement plans to ensure sustainable practices throughout the data centre life cycle;

- utilising renewable energy sources to power facilities and in turn lower carbon emissions;
- utilising efficient hot-water-cooling systems and direct liquid cooling, which require significantly less energy than conventional air cooling, while allowing the generated heat to be reused downstream.

Similar ideas emerged from the web-based questionnaire. Among the initiatives undertaken, energy-efficient cooling systems are the most common, with 14 providers having adopted this approach. Additionally, 12 providers have taken steps to reduce energy consumption, while 10 have started using renewable energy sources. Seven providers are reusing generated heat, and four are conducting carbon footprint assessments and utilising carbon credits. Other measures include managing energy surpluses and planning for system retirement (three providers), along with increasing awareness of material sourcing and implementing waste reduction and recycling programs (two providers).

In conclusion, the stakeholder consultations revealed an important convergence between HPC advancements and environmental sustainability. While providers are actively implementing measures to enhance energy efficiency and reduce emissions, users are still prioritising computational performance. Bridging this gap through collaborative efforts and a shared commitment to sustainable practices will be essential for the future of HPC, ensuring that it remains both innovative and environmentally responsible.

09

Socioeconomic benefits

Digital technologies have significantly shaped social and economic development, particularly through the widespread availability of online services. HPC provides essential tools for real-time socioeconomic data collection, empowering AI-driven decision-making.

By using HPC, policymakers can gather and analyse data promptly, leading to more precise and timely decisions [82] we investigate the digital transformation of socio-economic development through the quality-of-life vision, highlighting the importance of sustainable and inclusive strategies that prioritize the well-being of individuals and communities.

Digital technologies have transformed socio-economic development worldwide. Assessing the quality of life in the digital economy is essential to ensure that everyone benefits from this transformation. Artificial intelligence (AI). HPC plays a crucial role for both private and public sectors. Private companies use it to drive innovation and maintain market leadership, while public institutions leverage it to improve citizens' quality of life and tackle social issues [83]. HPC systems have the capability of handling large datasets that benefits various fields, including healthcare, weather forecasting, renewable energy and cybersecurity [84].

Expanding access to HPC beyond traditional scientific fields and into areas like social media analytics, behavioural finance and smart cities provides organisations with a competitive edge, drives innovation and delivers significant socioeconomic benefits. This broader application enhances decision-making and boosts competitiveness across industries [85]. However, making HPC more accessible might not be an easy task. Samuel J. et al. [85] proposed three strategies to achieve this:

- availability involves expanding infrastructure to ensure HPC resources are widely distributed;
- accessibility simplifies user connections through intuitive interfaces and cloud services;
- usability enhances user experience with effective training and user-friendly tools.

These combined efforts aim to make HPC more widely available, easier to use and more impactful across various fields.

The HPC sector has also contributed to job creation across various industries, particularly in research, engineering and IT services [86].

The demand for specialised skills in HPC is growing, requiring expertise in areas such as parallel programming, data analysis and system management [87].

Hyperion Research reports that supercomputers have generated products worth EUR 92 trillion (over USD 100 trillion) in the automotive, aerospace and pharmaceutical industries over the past 25 years (89). Hyperion also estimates that Linux-based supercomputers alone have generated more than EUR 2.76 trillion (USD 3 trillion) in economic value. In a study of 175 industrial firms, Hyperion found an average return of EUR 415 (USD 452) for every dollar invested in HPC (see Table 8). In finance, life sciences, manufacturing and transportation, the return was EUR 463 (USD 504) in sales revenue and EUR 35 (USD 38) in profits or cost savings per euro or dollar invested. These HPC-supported projects created around 2 335 new jobs in the companies studied [88].

Table 8 presents the financial return on investment from HPC in different industries.

Socioeconomic contexts for HPC use include:

- climate simulations to provide early warnings for storms and develop strategies to mitigate or adapt to climate change;
- managing transportation systems to optimise traffic flow;
- addressing potential safety issues;
- automating of freight delivery;
- detecting suspected criminals in real time;
- counter-terrorism efforts;
- detecting infectious diseases;
- analysing virus genomic sequences.

For example, hospitals in Germany utilise HPC to predict which expectant mothers may need caesarean sections, aiming to prevent last-minute, riskier childbirth decisions. These applications demonstrate HPC’s utility in areas such as climate science, transportation, security and healthcare [89]. Thus, HPC serves as a cost-effective tool despite the initial investment, significantly speeding up the R & D process [86].

Another application of HPC technology is training AI models, due to its ability to handle the massive computational demands of modern AI algorithms. As AI tools and techniques advance, including LLMs and sophisticated machine learning algorithms, HPC provides the necessary power and efficiency to handle these demands. Examples of LLMs include ChatGPT-4 and DALL-E 2. Beyond these, AI is also making significant impacts in other areas. In healthcare, for instance, AlphaFold predicts protein folding to advance drug discovery. In mobility, AI is driving progress in self-driving cars [90]. The recent increase of interest in AI tools by organisations is notable and, according to a recent study from the International Monetary Fund, 40 % of the global employment is exposed to AI, with potential benefits and drawbacks for economies [91].

Table 8. Financial return on investment from HPC, by select industry [88]


Industry	Average revenue per HPC (dollar invested)	Average profit or cost savings per HPC (dollar invested)
Defense	75.00	18.80
Financial	641.70	47.40
Insurance	175.70	280.00
Life Sciences	205.60	40.90
Manufacturing	216.50	28.40
Oil and Gas	416.00	53.70
Telecommunications	210.70	30.40
Transportation	1,804.30	15.60
TOTAL	452.10	37.60

Overall, supercomputing has become a critical tool for supporting better decision-making [84]. It plays a crucial role in both enhancing the EU's industrial capabilities and improving its citizens' well-being. Investments in HPC yield high returns at the macroeconomic level, enabling leading companies and countries in this field to excel in science and economics. By processing massive amounts of data and facilitating simulations, HPC drives advancements in technology, economics and society. It boosts competitiveness in key industries, fostering economic growth and positioning itself as an evolving multidisciplinary paradigm with profound implications for today's society [83], [85], [89].

9.1 Stakeholder consultations

Findings from both the interviews and the web-based questionnaire reveal important trends regarding user engagement with HPC services. From the HPC providers interviewed, several mentioned that they serve between 300 and 4 000 users and expect the numbers to increase in the future. The web-based questionnaire indicated that some providers support around 100 users, while others cater to as many as 6 000 or even 12 000 users. When asked about the average number of new users per year, providers' responses varied significantly, ranging from as few as 25 to as many as 500 new users annually. Common averages reported include 50, 60, 100, 300 and between 200 and 500 new users per year. Looking at the past, the number of users ceasing HPC services also varied considerably. Reported figures include 50, 25 and 10 users, with one provider reporting no cessations. In some cases, user cessation was minimal, particularly in the academic segment where accounts remain active until users leave the university. Industrial users were noted as more critical, with some ceasing services due to changing needs or reconsideration under current access policies. Common reasons for cessation included users moving to other universities, scaling up to national resources, changes in research focus and retirement.

Figure 23. Significant achievements made possible by HPC applications

 Users	 Developers	 Providers
1 Scientific research achievements	1 Scientific research achievements	1 Scientific research achievements
2 Innovative approaches	2 Technological advancements	2 Technological advancements and innovative approaches
3 Technological advancements	3 Innovative approaches	3 Collaborations with industries

Perspectives on the contribution of HPC to addressing societal challenges

HPC technology has made significant contributions to research and industry, enabling achievements that are only possible through HPC applications. From the interviews with users and both national and private data centres from all the continents, it was noted that they usually collaborate with government agencies, for example those working in the weather and healthcare sectors. The following examples related to research topics highlight the main benefits of using HPC.

- Accurate weather predictions.
- Support for research into COVID-19 investigation and treatment.
- Enhancement of early detection and treatment of cancers (e.g. colorectal and pancreatic cancer) through collaboration with local hospitals and analysis of public health data.
- Simulation of the formation of the universe after the Big Bang to gain an understanding of the statistical behaviour of the cosmos – including the formation of supermassive black holes – and analyse the behaviour of atoms within materials at an atomic scale, made possible by the computational power of HPC.
- Discovery of a new superconductor and prediction of a new thermoelectric material, both of which were confirmed through experiments and made possible by HPC, with the latter even leading to a material patent.

- HPC simulations play a crucial role in advancing nuclear fusion research by enabling scientists to model the extreme conditions required for fusion, test different approaches and optimise methods for achieving stable and sustainable nuclear fusion, which could lead to breakthroughs in clean energy.
- HPC simulations help identify potential drug candidates, particularly by predicting the effects of mutations in proteins. This approach has been instrumental in improving clinical trials by helping determine which patients should be included or excluded based on how likely a specific drug is to work, especially in cases where protein mutations affect drug efficacy.
- Accurate simulation of lava flow to predict movement in volcanic areas and optimise the management of local resources during volcanic activities.

For all stakeholder groups from the web-based questionnaire, the most significant results achieved through the utilisation of HPC systems were linked to scientific research. The second and third options with more selections among users and developers were linked to technological advancements and the development of innovative approaches. Providers had a similar number of responses to the options on technological advancements, innovative approaches and collaborations with industries (see **Figure 23**).

Some of the mentioned benefits were linked to energy efficiency, support for industries with a focus on real-world challenges, improvements in data processing and the possibility to establish contracts with relevant business actors.

Users enumerated the following benefits of HPC simulations that directly support their work and help companies gain advantage in their respective sectors.

- 1. Research and development advancements**
 - Enabling cutting-edge research in fields like drug development, material science and astrophysics.
 - Facilitating large-scale simulations and advanced modelling for scientific breakthroughs.
- 2. Technological innovation**
 - Supporting the development of AI models and deep learning algorithms.
 - Optimising technologies such as hydrogen injectors and renewable energy systems.
- 3. Industrial and commercial applications**
 - Improving product development, cost reduction and efficiency in sectors like aerodynamics and fluid mechanics.
 - Providing SMEs and industries with essential computational resources to accelerate R & D.
- 4. Efficiency and productivity**
 - Reducing computational time, enabling faster data analysis and research progress.
 - Increasing productivity by handling complex simulations unfeasible for standard systems.
- 5. Environmental and cost benefits**
 - Aiding in reducing emissions and improving the reliability of renewable energy technologies.
 - Offering cost savings by optimising processes and reducing the need for physical prototypes.

9.1.1 Strategies to improve access to high-performance computing in regions with limited availability

Regions with limited availability of HPC are areas where access to advanced computing resources is constrained due to factors such as geographic isolation, economic limitations, insufficient institutional support and inadequate awareness of or training on HPC technology. These challenges hinder the ability of local users to leverage HPC for research, innovation and economic development.

In the web-based questionnaire, over half of the users, 53 %, stated that they do not access HPC resources outside their continent. Many organisations find that HPC facilities are readily available within their own countries or regions, which eliminates the need to seek out international options. Local HPC centres and national HPC centres often provide tailored support services, ensuring that users receive assistance relevant to their specific needs. Additionally, these facilities typically offer sufficient capacity and advanced capabilities to handle computational tasks effectively. Utilising local HPC resources ensures minimal latency and optimal performance, which are critical for data-intensive applications. These advantages, combined with the cost-effectiveness of local services – such as reduced data transfer fees and more competitive pricing – makes domestic HPC options more appealing. As a result, users are more inclined to rely on the resources available within their own regions rather than exploring options located outside their continent.

Those that do seek out external options face several significant challenges:

- 'network latency' and 'bandwidth issues' are major hurdles, affecting 17.2 % of respondents;
- 'funding constraints' and 'difficulties with data transfer and storage' pose challenges for 15.5 % of users;
- 'limited remote technical support' is a concern for 7.8 % of users;
- 'software licences that are restricted to specific systems' (5.2 %);
- 'data security' (6.0 %) and 'regulatory compliance' (6.0 %);
- 'limited local technical support impacts' affects 4.3 % of respondents;
- 'geographical barriers' and 'cultural aspects' affect a smaller percentage of users, at 1.7 % and 0.9 %, respectively.

Participants were asked to rank the relevance of different strategies for improving access to HPC in regions with limited availability, from 'not at all' to 'to a great extent'. All stakeholder groups agreed that public investment in HPC infrastructures, education and training for the community, along with developing sustainable funding models for long-term availability, operation and growth of HPC resources/ services, could be good strategies. Unlike developers, users and providers agreed that investment in remote collaboration tools will not contribute to improving accessibility (see **Table 9**).

The results indicate a strong belief among users (67.2 %) in the importance of public investment in HPC infrastructures, reflecting a reliance on government support. In contrast, only 19.8 % of users recognise significant private funding, suggesting a potential funding gap that could hinder the growth of HPC resources. While 37.1 % of users see substantial availability of cloud-based HPC services, this optimism is not shared by developers (21.5 %) and providers (17.6 %), who are more aware of the challenges associated with cloud solutions. This discrepancy highlights the need for improved communication regarding the capabilities and limitations of these services.

Additionally, the importance of education and training is underscored, with 50.8 % of users and 52.9 % of developers recognising its critical role in enhancing HPC accessibility. The strong support for developing sustainable funding models (53.4 % of users) suggests that collaboration between the public and private sectors, along with targeted educational initiatives, is essential for advancing HPC capabilities and ensuring that organisations can fully leverage these technologies. Overall, there is clear demand for increased public investment and strategic partnerships to address challenges in the HPC landscape.

Table 9. Strategies for improving access to HPC resources in regions with limited availability

	Scale	Users (116)	Developers (51)	Providers (17)
Public investment in HPC infrastructures	To a great extent	67.2 %	58.8 %	76.5 %
	Somewhat	25 %	27.4 %	17.6 %
	Very little/not at all	4.3 %	5.9 %	29.4 %
Private investment in HPC infrastructures	To a great extent	19.8 %	17.6 %	17.6 %
	Somewhat	47.4 %	39.2 %	47.1 %
	Very little/not at all	26.7 %	29.4 %	29.4 %
Availability of cloud-based HPC services	To a great extent	37.1 %	21.5 %	17.6 %
	Somewhat	38.8 %	43.1 %	58.8 %
	Very little/not at all	18.1 %	21.5 %	23.5 %
Providing HPCaaS to access HPC infrastructures remotely in other regions	To a great extent	43.9 %	37.3 %	35.3 %
	Somewhat	35.3 %	19.6 %	52.9 %
	Very little/not at all	15.5 %	21.5 %	11.8 %
Public-private partnerships	To a great extent	19.8 %	45.1 %	17.6 %
	Somewhat	39.6 %	37.3 %	35.3 %
	Very little/not at all	32.7 %	33.3 %	29.4 %
Regional HPC consortia	To a great extent	35.3 %	49 %	29.4 %
	Somewhat	43.1 %	39.2 %	52.9 %
	Very little/not at all	14.6 %	5.9 %	17.6 %
Open access policies to make research findings, data and HPC resources freely available to promote transparency, collaboration and knowledge in the community	To a great extent	36.2 %	45.1 %	23.5 %
	Somewhat	30.2 %	23.5 %	47.1 %
	Very little/not at all	25.8 %	25.5 %	17.6 %
Education and training for the community	To a great extent	50.8 %	52.9 %	70.5 %
	Somewhat	31.9 %	29.4 %	23.5 %
	Very little/not at all	13.8 %	31.4 %	6 %
Investment in remote collaboration tools (e.g. remote access solutions)	To a great extent	28.4 %	33.3 %	11.7 %
	Somewhat	28.4 %	31.4 %	35.3 %
	Very little/not at all	36.2 %	23.5 %	47.1 %
Developing sustainable funding models for long-term availability, operation and growth of HPC resources/services	To a great extent	53.4 %	39.2 %	70.5 %
	Somewhat	28.4 %	35.3 %	23.5 %
	Very little/not at all	13.8 %	9.8 %	0

Local and regional contributions of HPC provider organisations

HPC technology contributes to the local and regional economy in several significant ways. Some providers directly employ over 100 people in their computing centres, actively hire local staff and collaborate with regional partners such as public organisations and industry.

Through their research activities and the support they provide to researchers, provider organisations play a key role in economic growth and business development. They also support local businesses and academic institutes by providing valuable research output. Additionally, the organisations offer extensive training to high-qualified professionals and provide HPC experience to students and researchers, enhancing their future job prospects. They also organise and participate in local community events. While some responses suggest a minimal impact or uncertainty regarding the extent of these contributions, the overall involvement in job creation, economic development and skill-building is notable.

In conclusion, the widespread adoption of HPC technology has led to significant advancements across various sectors, demonstrating its vital role in driving research and innovation. While providers are increasingly catering to large user bases, challenges remain in user retention and accessibility, particularly in regions with limited resources. The contributions of HPC extend beyond computational power and include fostering economic growth, supporting local communities and enhancing educational opportunities. As stakeholders recognise the importance of sustainable practices and collaborative investments, the focus on improving accessibility and maximising the societal benefits of HPC will be crucial for future developments in this field.

10

Final remarks: insights, challenges and future directions

10.1 High-performance computing environments

HPC plays a crucial role in both academic and industrial sectors. Traditional fields such as weather forecasting, healthcare and engineering have long utilised HPC for complex simulations and predictions. As computational capabilities continue to advance, emerging sectors like IT, telecoms, smart cities and robotics are increasingly adopting HPC to drive innovation and research.

On-premises HPC facilities and COE clusters are the most widely used and the main service delivered by HPC providers, reflecting their continued importance in HPC tasks. Cloud-based HPC services are also favoured due to their scalability and flexibility, addressing a broad spectrum of computational needs. Additionally, all stakeholders indicated growing interest in advanced technologies, such as quantum computing and hybrid HPC environments that combine access to on-premises or shared COE facilities and cloud resources.

The research confirms that on-premises and shared-cluster HPC facilities are a staple for many users.

Cloud-based services are also prevalent, though opinions on private cloud solutions vary. Some users appreciate the convenience and accessibility offered by private clouds, while others are cautious due to concerns about data transfer and storage security.

The perspectives gathered from the interviews and the survey highlight users' concerns regarding HPC resources, such as insufficient computational power or RAM to accommodate intensive workloads. In the survey, opinions diverged on the third major concern: while users and developers prioritised reliable, high-speed network connectivity, providers emphasised the importance of a high-speed file system and data storage. Specifically, in the interviews, providers noted the need for secure, long-term data storage, particularly in environments with automated deletion policies.

Strategic investments in infrastructure and training will help democratise access to HPC resources, allowing a wider range of users to benefit. Furthermore, fostering collaborative partnerships between users and providers can ensure that both parties effectively address existing challenges, paving the way for innovation and sustainable practices within the HPC community.

10.2 Financial aspects

Some companies still face financial and technical challenges in building or having their own HPC infrastructure. Most users stated that organisations allocate less than 10% of the budget to HPC resources, and, when investing, the main expenditure is linked to hardware resources.

There are also differences in how providers offer HPC services. Usually, public organisations offer free HPC access to academic and R&D users based on a competitive application process. On the other hand, when the services are paid, the pricing models include variable pricing, metered usage and fixed pricing with limited or unlimited usage. Among the providers that stated their financial income from HPC services, most reported an income of less than EUR 0.5 million.

10.3 Security and legal aspects

Based on the interviews and the web-based survey, the main legal concerns are linked to data privacy and protection, legal considerations, IP rights, compliance, confidentiality and security. Users and developers prioritise strict legal and security measures, especially in sensitive sectors like healthcare and biosciences. These include data location transparency and isolation from other users. From the providers' perspective, the main concerns are linked to securing sensitive and proprietary data, particularly against hacking and security breaches.

Users and developers highlighted the importance of SLAs and AUPs to ensure data protection, compliance and clear legal responsibilities, and the need for non-disclosure agreements and confidentiality measures to protect sensitive and proprietary information. Additionally, IP rights are also understood as critical for safeguarding proprietary data and ensuring compliance with legal and regulatory standards, particularly in sensitive industries. Providers still recognise the importance of adhering to sector-specific regulations and laws.

10.4 Environmental considerations

The stakeholder consultations indicated a growing recognition of environmental considerations within the HPC community, particularly among HPC providers. While users acknowledge the importance of sustainability, many view it as a secondary concern compared to immediate computational needs. This divergence in prioritisation suggests a potential gap in communication and shared goals between users and providers. Users are increasingly aware that optimising code to reduce computational power consumption could play a significant role in sustainability, signalling a desire for collaborative efforts to align software practices with environmental objectives.

Providers outlined various strategies to enhance energy efficiency and promote sustainable practices in data centres. Initiatives such as employing energy-efficient hardware, developing hybrid architectures that combine CPUs and GPUs and implementing heat reuse systems exemplify proactive measures taken to minimise energy consumption. These approaches not only address the operational costs associated with high energy use but also respond to the growing pressure to reduce the carbon footprint of HPC systems. The integration of renewable energy sources into data centre operations reflects a commitment to long-term sustainability goals, demonstrating that the HPC sector is adapting to emerging environmental challenges.

The findings from the web-based questionnaire reinforce the notion that the HPC community is taking substantial steps towards achieving greener practices.

With many providers adopting energy-efficient cooling systems and assessing their carbon footprint, there is clear momentum towards creating environmentally friendly HPC infrastructures. This is particularly relevant in the context of the projected contribution of data centres to global greenhouse gas emissions. As stakeholders continue to prioritise sustainability, the HPC community has an opportunity to lead in developing technologies that not only drive scientific discovery but also support broader environmental goals, ensuring that HPC remains relevant in an increasingly eco-conscious world.

10.5 Socioeconomic benefits

The integration of HPC into various sectors has transformed not only how organisations operate but also the broader socioeconomic landscape. As revealed during the stakeholder consultations, HPC data centres now serve a diverse and expanding user base, with some providers supporting thousands of users. This growth underscores HPC's critical role in facilitating real-time data analysis and AI-driven decision-making, which are essential for effective policy formulation and innovation. Users from both private and public sectors highlight HPC's contributions to significant advancements in fields such as healthcare, climate science and transportation, showcasing its versatility and impact.

However, challenges persist, particularly in enhancing accessibility for regions with limited access to HPC resources.

The barriers identified, including network latency and funding constraints, echo the broader need for strategic investments in infrastructure and training to democratise access to HPC resources. Stakeholders emphasised the importance of public investment and sustainable funding models as key strategies for expanding HPC's reach and impact. The stakeholder consultations highlighted a collective acknowledgement that overcoming these obstacles will not only improve user retention but also amplify HPC's socioeconomic benefits. As HPC continues to evolve, its potential to drive job creation, innovation and economic growth will be contingent upon addressing these accessibility challenges and fostering collaborative partnerships across sectors.

In summary, as HPC continues to evolve, its potential to drive job creation, innovation and economic growth will hinge on addressing accessibility challenges and fostering collaborative partnerships across sectors. The positive feedback from stakeholders reinforces HPC's significant contributions to research, economic growth and technological advancement, making it a vital tool for fostering collaboration and unlocking new possibilities for a more connected and informed society.

10.6 Recommendations

HPC is essential for addressing the growing computational demands across diverse fields, offering transformative benefits such as accelerating innovation, solving complex challenges and enhancing data analysis capabilities. To further enhance the role of HPC in advancing research, driving innovation and fostering collaboration with industry, several key strategic recommendations are proposed:

- commissioning further research with different academic and industrial sectors to better address current and future gaps;
- providing leadership in access, along with knowledge exchange on key HPC competencies and technological advances in different markets, including AI;
- developing new mechanisms for engaging with industry users as strategic partners in ways that bridge sectoral practices;
- supporting preference for on-site HPC access by addressing financial and legal barriers;
- developing collaborations to drive innovation in sustainability, software development and HPC;
- increasing pathways for public administrations and industry to take up HPC while acknowledging that their needs will differ from those of developers or academia;
- HPC providers assessing what certain sectors – such as health, aerospace, energy, mobility and finance – might require in terms of additional outreach or development to fully realise HPC’s advantages for industry;
- encouraging technological design that supports different levels of user proficiency.

Endnotes


- [1] TOP500, 'June 2024', TOP500 The List website, accessed 18 September 2024, <https://top500.org/>.
- [2] European High-Performance Computing Joint Undertaking, 'Discover EuroHPC JU', European High-Performance Computing Joint Undertaking website, accessed 18 September 2024, https://www.eurohpc-ju.europa.eu/about/discover-eurohpc-ju_en.
- [3] Department for Science, Innovation and Technology and The Rt Hon Michelle Donelan, 'UK joins European super-computing scheme as world-class British researchers now able to bid for funding', GOV.UK. Accessed: Nov. 29, 2024, <https://www.gov.uk/government/news/uk-joins-european-super-computing-scheme-as-world-class-british-researchers-now-able-to-bid-for-funding>.
- [4] The White House, 'Fact sheet: President Biden's 2024 budget invests in American science, technology, and innovation to achieve our nation's greatest aspirations', The White House website, 13 March 2023, <https://bidenwhitehouse.archives.gov/ostp/news-updates/2023/03/13/fy24-budget-fact-sheet-rd-innovation/>.
- [5] Fu, Y., Shen, W., Cui, J., Zheng, Y., Yang, G. et al., 'Toward exascale computation for turbomachinery flows', *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, Article 4, Association for Computing Machinery, New York, 11 November 2023, pp. 1–12, <https://doi.org/10.1145/3581784.3627040>.
- [6] Trader, T., 'Three Chinese exascale systems detailed at SC21: Two operational and one delayed', HPCwire, 24 November 2021, accessed 11 September 2024, <https://www.hpcwire.com/2021/11/24/three-chinese-exascale-systems-detailed-at-sc21-two-operational-and-one-delayed/>.
- [7] Trueman, C., 'Japanese government to fund \$470m AI supercomputer', Data Center Dynamics, 19 April 2024, accessed 15 October 2024, <https://www.datacenterdynamics.com/en/news/japanese-government-to-fund-new-ai-supercomputer/>.
- [8] Trueman, C., 'India unveils three supercomputers to boost HPC ambitions', Data Center Dynamics, 19 April 2024, accessed 17 October 2024, <https://www.datacenterdynamics.com/en/news/india-unveils-three-supercomputers-to-boost-hpc-ambitions/>.
- [9] Telecom Review, 'High-performance computing drives Singapore's digital growth', Telecom Review, 5 June 2023, accessed 29 November 2024, <https://www.telecomreviewasia.com/news/interviews/3324-high-performance-computing-drives-singapore-s-digital-growth/>.
- [10] Peckham, O., 'Nvidia announces four supercomputers, with two in Taiwan', HPCwire, 29 May 2023, accessed 29 November 2024, <https://www.hpcwire.com/2023/05/29/nvidia-announces-four-supercomputers-with-two-in-taiwan/>.
- [11] Lee, Hae-Sung, 'S.Korea to develop supercomputer for industrial use', The Korea Economic Daily, 31 May 2023, accessed 29 November 2024, <https://www.kedglobal.com/artificial-intelligence/newsView/ked202305310010>.
- [12] Russell, J., 'ISC 2024: Hyperion Research predicts HPC market rebound after flat 2023', HPCwire, 13 May 2024, accessed 17 October 2024, <https://www.hpcwire.com/2024/05/13/isc24-hyperion-research-predicts-hpc-market-rebound-after-flat-2023/>.
- [13] Eadline, D., 'Hyperion Research announces a 36.7 % increase in the HPC/AI market size', Alwire, accessed 17 October 2024, <https://www.aiwire.net/2024/08/26/hyperion-research-announces-a-36-7-increase-in-the-hpc-ai-market-size/>.
- [14] Connatser, M., 'Nvidia Grace Superchip loses to Intel Sapphire Rapids in HPC performance benchmarks, but promises greater efficiency', Tom's Hardware, 7 February 2024, accessed 18 September 2024, <https://www.tomshardware.com/pc-components/cpus/nvidia-grace-superchip-loses-to-intel-sapphire-rapids-in-hpc-performance-benchmarks-but-promises-greater-efficiency>.
- [15] Amazon Web Services, 'AWS Graviton performance testing: Tips for independent software vendors', AWS White Paper, 2025, <https://docs.aws.amazon.com/pdfs/whitepapers/latest/aws-graviton-performance-testing/aws-graviton-performance-testing.pdf>.
- [16] European Processor Initiative, 'Successful conclusion of European Processor Initiative Phase One', European Processor Initiative website, 22 December 2021, accessed 18 September 2024, <https://www.european-processor-initiative.eu/successful-conclusion-of-european-processor-initiative-phase-one/>.

- [17] European High-Performance Computing Joint Undertaking, 'New call for developing an HPC ecosystem based on RISC-V', European High-Performance Computing Joint Undertaking website, 1 February 2023, accessed 18 September 2024, https://eurohpc-ju.europa.eu/new-call-developing-hpc-ecosystem-based-risc-v-2023-02-01_en.
- [18] AMD, 'AMD Instinct™ MI300A Accelerators', AMD website, accessed 18 September 2024, <https://www.amd.com/en/products/accelerators/instinct/mi300/mi300a.html>.
- [19] NVIDIA, 'NVIDIA Tensor Cores – Unprecedented acceleration for generative AI', NVIDIA website, accessed 18 September 2024, <https://www.nvidia.com/en-us/data-center/tensor-cores/>.
- [20] Lee, C. A., Gasster, S. D., Plaza, A., Chang, C.-I. and Huang, B., 'Recent developments in high performance computing for remote sensing: A review', *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, Vol. 4, No 3, September 2011, IEEE, 11 August 2011, pp. 508–527, <https://doi.org/10.1109/JSTARS.2011.2162643>.
- [21] Altera, 'FPGAs, SoCs, CPLDs', Altera website, accessed 18 September 2024, <https://www.altera.com/fpga>.
- [22] AMD, 'AMD completes acquisition of Xilinx', AMD website, accessed 18 September 2024, <https://www.amd.com/en/newsroom/press-releases/2022-2-14-amd-completes-acquisition-of-xilinx.html>.
- [23] Brim, M. J., 'A survey of HPC storage systems', Oak Ridge National Laboratory, 5 August 2022, accessed 7 August 2024, <https://extremecomputingtraining.anl.gov/wp-content/uploads/sites/96/2022/11/ATPESC-2022-Track-3-Talk-3-Brim-HPC-Storage-Survey.pdf>.
- [24] Shah, P., Ye, J. and Sun, X.-H., 'Survey the storage systems used in HPC and BDA ecosystems', *Internet of Things and Cloud Computing*, Vol. 10, Issue 1, 19 May 2022, <https://doi.org/10.48550/arXiv.2112.12142>.
- [25] Calucci, P., 'Parallel file systems for HPC – Introduction to Lustre', International School for Advanced Studies, Trieste, November 2008, accessed 8 July 2024, https://people.sissa.it/~calucci/smr1967/lustre/par_fs_slides.pdf.
- [26] Hewlett Packard, 'XC Series DataWarp Installation and Administration Guide CLE7.0UP01 S-2564', Hewlett Packard website, accessed 18 September 2024, https://support.hpe.com/hpesc/public/docDisplay?docId=a00113970en_us&page=DataWarp_Use_Cases.html.
- [27] Atos, 'Extreme data: Storage for high performance computing', Atos website, accessed 18 September 2024, <https://atos.net/en/solutions/high-performance-computing-hpc/extreme-data>.
- [28] Lu, P.-J., Lai, M.-C. and Chang, J.-S., 'A survey of high-performance interconnection networks in high-performance computer systems', *Electronics*, Vol. 11, No 9, April 2022, p. 1369, <https://doi.org/10.3390/electronics11091369>.
- [29] Atos, 'BXI V2 - Shifting to a new scale', Atos website, accessed 16 September 2024, <https://atos.net/en/solutions/high-performance-computing-hpc/bxi-v2>
- [30] Ajima, Y., Kawashima, T., Okamoto, T., Shida, N., Hirai, K. et al., 'The Tofu Interconnect D', 2018 *IEEE International Conference on Cluster Computing (CLUSTER)*, IEEE, Belfast, 10–13 September 2018, pp. 646–654, <https://doi.org/10.1109/CLUSTER.2018.00090>.
- [31] Fujitsu, 'Supercomputer Fugaku', Fujitsu website, accessed 21 September 2024, <https://www.fujitsu.com/global/about/innovation/fugaku/>.
- [32] Feldman, M., 'China fleshes out exascale design for Tianhe-3 supercomputer', The Next Platform, 2 May 2019, accessed 18 September 2024, <https://www.nextplatform.com/2019/05/02/china-fleshes-out-exascale-design-for-tianhe-3/>.
- [33] Lawrence Livermore National Laboratory, 'El Capitan: NNSA's first exascale machine', Lawrence Livermore National Laboratory website, accessed 18 September 2024, <https://asc.llnl.gov/exascale/el-capitan>.
- [34] Ahn, D. H., Bass, N., Chu, A., Garlick, J., Grondona, M. et al., 'Flux: Overcoming scheduling challenges for exascale workflows', *Future Generation Computer Systems*, Vol. 110, September 2020, pp. 202–213, <https://doi.org/10.1016/j.future.2020.04.006>.
- [35] start, 'What is OAR?', start, accessed 18 September 2024, <https://oar.imag.fr/start>.
- [36] Bachan, J., Scott, B. B., Bonachea, D., Corbino, J., Grossman, J. et al., 'UPC++ v1.0 Programmer's Guide, Revision 2023.9.0', Lawrence Berkeley National Laboratory, December 2023, <https://doi.org/10.25344/S4P01J>.

- [37] Chamberlain, B.L., 'Chapel (Cray Inc. HPCS Language)', in: Padua, D. (ed.) *Encyclopedia of Parallel Computing*, Springer, Boston, 2011, pp. 249-256, https://doi.org/10.1007/978-0-387-09766-4_54.
- [38] The Open MPI Community, 'OpenSHMEM', The Open MPI Community, accessed 18 September 2024, <https://docs.open-mpi.org/en/v5.0.x/man-openshmem/man3/OpenSHMEM.3.html>.
- [39] NVIDIA, 'NVIDIA OpenSHMEM Library (NVSHMEM) Documentation', NVIDIA website, accessed 18 September 2024, <https://docs.nvidia.com/nvshmem/api/index.html>.
- [40] Intel, 'Tutorial: Using Coarray Fortran'.
- [41] Herten, A., 'Many cores, many models: GPU programming model vs. vendor compatibility overview', *Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, Association for Computing Machinery, New York, 12 November 2023, pp. 1019-1026, <https://doi.org/10.1145/3624062.3624178>.
- [42] 'The Exascale Computing Project, 'Addressing a national imperative', The Exascale Computing Project website, accessed 18 September 2024, <https://www.exascaleproject.org/research/>.
- [43] García, A. L., 'The evolution of the cloud – The work, progress and outlook of cloud infrastructure', Massachusetts Institute of Technology, February 2015, accessed 8 August 2024, <https://dspace.mit.edu/bitstream/handle/1721.1/100311/932065967-MIT.pdf;sequence=1>.
- [44] Carlyle, A. G., Harrell, S. L. and Smith, P. M., 'Cost-effective HPC: The community or the cloud?', *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, IEEE, Indianapolis, 30 November-3 December 2010, pp. 169-176, <https://doi.org/10.1109/CloudCom.2010.115>.
- [45] IBM, Jacob, B., Brown, M., Fukui, K. and Trivedi, N., '*Introduction to grid computing*', IBM Redbooks, 1st edition, December 2005, <https://www.redbooks.ibm.com/redbooks/pdfs/sg246778.pdf>.
- [46] Conway, S., 'Edge computing and HPC', Hyperion Research, November 2021, accessed 8 August 2024, https://hyperionresearch.com/wp-content/uploads/2021/11/Hyperion-Research-SC21-Market-Update-Briefing_Edge-Computing_Conway.pdf.
- [47] Draeger, E., 'Supercomputing at the exascale and beyond: Future trends and challenges', 2023 IPAM Workshop '*New Mathematics for the Exascale: Applications to Materials Science*', University of California, Los Angeles, 13 March-16 June 2023.
- [48] Thomas, J., 'LLNL's El Capitan debuted on new Top500 list of world's most powerful supercomputers', Lawrence Livermore National Laboratory, 13 May 2024, accessed 18 September 2024, <https://www.llnl.gov/article/51231/llnl-el-capitan-debuted-new-top500-list-worlds-most-powerful-supercomputers>.
- [49] Jülich Supercomputing Centre, 'JUPITER technical overview – A deep dive into JUPITER's building blocks', Jülich Supercomputing Centre website, accessed 18 September 2024, <https://www.fz-juelich.de/en/ias/jsc/jupiter/tech>.
- [50] European High-Performance Computing Joint Undertaking, 'Signature of the hosting agreement for the second European exascale supercomputer, Alice Recoque', European High-Performance Computing Joint Undertaking website, 21 June 2024, accessed 18 September 2024, https://eurohpc-ju.europa.eu/signature-hosting-agreement-second-european-exascale-supercomputer-alice-recoque-2024-06-21_en.
- [51] Bharti, K., Cervera-Lierta, A., Kyaw, T. H., Haug, T., Alperin-Lea, S. et al., 'Noisy intermediate-scale quantum algorithms', *Reviews of Modern Physics*, Vol. 94, Issue 1, 015004, American Physical Society, 15 February 2022, <https://doi.org/10.1103/RevModPhys.94.015004>.
- [52] de Wolf, R., 'Quantum computing: Lecture notes', Cornell University arXiv, 16 January 2023, accessed 9 August 2024, <https://doi.org/10.48550/arXiv.1907.09415>.
- [53] Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S. et al., 'Variational quantum algorithms', *Nature Reviews Physics*, Vol. 3, No 9, pp. 625-644, 12 August 2021, <https://doi.org/10.1038/s42254-021-00348-9>.
- [54] Weidman, J. D., Sajjan, M., Mikolas, C., Stewart, Z. J., Pollanen, J. et al., 'Quantum computing and chemistry', *Cell Reports Physical Science*, Vol. 5, Issue 9, 102105, 18 September 2024, <https://doi.org/10.1016/j.xcrp.2024.102105>.
- [55] Orús, R., Mugel, S. and Lizaso, E., 'Quantum computing for finance: Overview and prospects', *Reviews in Physics*, Vol. 4, 100028, November 2019, <https://doi.org/10.1016/j.revip.2019.100028>.
- [56] Bartsch, V., de Verdière, G. C., Nominé, J.-P., Ottaviani, D., Dragoni, D. et al., 'Quantum for HPC – The impact of quantum computers on HPC applications and the integration of quantum computers in HPC centres', European Technology Platform for High Performance Computing, White Paper, 5 October 2021, https://greencompute.uk/References/ETP4HPC/ETP4HPC_WP_Quantum4HPC_FINAL.pdf.

- [57] Gyongyosi, L. and Imre, S., 'A survey on quantum computing technology', *Computer Science Review*, Vol. 31, February 2019, pp. 51–71, <https://doi.org/10.1016/j.cosrev.2018.11.002>.
- [58] ISMS Manager, 'ID024a Acceptable Use Policy', May 2024, <https://docs.hpc.cam.ac.uk/srcp/isms-docs/aup.html#acceptable-use>.
- [59] Cortés, U., Cortés, A., Garcia-Gasulla, D., Pérez-Arnal, R., Álvarez-Napagao, S. et al., 'The ethical use of high-performance computing and artificial intelligence: Fighting COVID-19 at Barcelona Supercomputing Center', *AI and Ethics*, Vol. 2, May 2022, pp. 325–340, <https://doi.org/10.1007/s43681-021-00056-1>.
- [60] Radley-Gardner, O., Beale, H., Zimmermann, R. and Schulze, R. (eds), *Fundamental Texts on European Private Law*, 2nd edition, Hart Publishing, 2016, <https://doi.org/10.5040/9781782258674>.
- [61] Hawrylak, P. J., Louthan, G., Hale, J. and Papa, M., 'Practical cyber-security solutions for the Science DMZ', *Practice and Experience in Advanced Research Computing 2019: Rise of the Machines (learning)*, Article 50, Association for Computing Machinery, New York, 28 July 2019, pp. 1–6, <https://doi.org/10.1145/3332186.3332213>.
- [62] Peisert, S., 'Security in high-performance computing environments', *Communications of the ACM*, Vol. 60, No 9, 1 September 2017, pp. 72–80, <https://doi.org/10.1145/3096742>.
- [63] Chen, K., 'Confidential high-performance computing in the public cloud', *IEEE Internet Computing*, Vol. 27, No 1, IEEE, 30 January 2023, pp. 24–32, <https://doi.org/10.1109/MIC.2022.3226757>.
- [64] Guo, Y., Chandramouli, R., Wofford, L., Gregg, R., Key, G. et al., 'High-performance computing security: Architecture, threat analysis, and security posture', National Institute of Standards and Technology, Gaithersburg, NIST SP 800-223, February 2024, <https://doi.org/10.6028/NIST.SP.800-223>.
- [65] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) (OJ L, 2024/1689, 12.7.2024, ELI: <http://data.europa.eu/eli/reg/2024/1689/oj>).
- [66] Networking and Information Technology Research and Development (NITRD) Program, 'HPCC/NITRD Program Authorizing Legislation', NITRD website, accessed 23 October 2024, <https://www.nitrd.gov/legislation/>.
- [67] U.S. National Science Foundation, 'CHIPS and Science', U.S. National Science Foundation website, accessed 23 October 2024, <https://new.nsf.gov/chips>.
- [68] U.S. Department of Commerce: Bureau of Industry and Security, 'Export Administration Regulations (EAR)', U.S. Department of Commerce: Bureau of Industry and Security website, accessed 23 October 2024, <https://www.bis.gov/regulations/ear>.
- [69] CMS Information Security and Privacy Group, 'Federal Information Security Modernization Act (FISMA)', CyberGeek, 20 June 2023, accessed 23 October 2024, <https://security.cms.gov/learn/federal-information-security-modernization-act-fisma>.
- [70] European High-Performance Computing Joint Undertaking, 'Supercomputers Access Policy and FAQ', European High-Performance Computing Joint Undertaking website, accessed 18 September 2024, https://eurohpc-ju.europa.eu/access-our-supercomputers/access-policy-and-faq_en.
- [71] Breslow, A. D., Tiwari, A., Schulz, M., Carrington, L., Tang, L. et al., 'Enabling fair pricing on HPC systems with node sharing', *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, Article 37, Association for Computing Machinery, New York, 17 November 2013, pp. 1–12, <https://doi.org/10.1145/2503210.2503256>.
- [72] Wu, D., Terpenney, J. and Gentsch, W., 'Economic benefit analysis of cloud-based design, engineering analysis, and manufacturing', *Journal of Manufacturing Science and Engineering*, Vol. 137, Issue 4, 040903, 1 August 2015, <https://doi.org/10.1115/1.4030306>.
- [73] Kitchin, R., 'Big data, new epistemologies and paradigm shifts', *Big Data & Society*, Vol. 1, No 1, April 2014, <https://doi.org/10.1177/2053951714528481>.
- [74] Liu, P. and Biljecki, F., 'A review of spatially-explicit GeoAI applications in urban geography', *International Journal of Applied Earth Observation and Geoinformation*, Vol. 112, 102936, August 2022, <https://doi.org/10.1016/j.jag.2022.102936>.
- [75] Aghimien, E. I., Aghimien, L. M., Petinrin, O. O. and Aghimien, D. O., 'High-performance computing for computational modelling in built environment-related studies – a scientometric review', *Journal of Engineering, Design and Technology*, Vol. 19, Issue 5, 4 November 2020, pp. 1138–1157, <https://doi.org/10.1108/JEDT-07-2020-0294>.

- [76] Guyon, D., Orgerie, A.-C., Morin, C. and Agarwal, D., 'How much energy can green HPC cloud users save?', *2017 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*, IEEE, St. Petersburg, 6-8 March 2017, pp. 416–420, <https://doi.org/10.1109/PDP.2017.62>.
- [77] Nana, R., Tadonki, C., Dokladal, P. and Mesri, Y., 'Energy concerns with HPC systems and applications', Cornell University arXiv, 31 August 2023, accessed 5 July 2024, <https://doi.org/10.48550/arXiv.2309.08615>.
- [78] Energy Technologies Area: Berkeley Lab, Broughton, J., 'HPC energy efficiency in the exascale era', Energy Technologies Area: Berkeley Lab, 20 November 2019, accessed 18 September 2024, <https://datacenters.lbl.gov/index.php/resources/hpc-energy-efficiency-exascale-era>.
- [79] Li, B., Basu Roy, R., Wang, D., Samsi, S., Gadepally, V. et al., 'Toward sustainable HPC: Carbon footprint estimation and environmental implications of HPC systems', *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, Article 19, Association for Computing Machinery, New York, 11 November 2023, pp. 1–15, <https://doi.org/10.1145/3581784.3607035>.
- [80] de Oliveira Castro, P., 'High performance computing code optimizations: Tuning performance and accuracy', Université Paris-Saclay, 26 October 2022, <https://theses.hal.science/tel-03831483#:~:text=We%20discuss%20the%20design%20of%20verificarlo%2C%20an%20opensource,Functional%20Theory%20quantum%20mechanical%20modeling%2C%20or%20structure%20simulations>.
- [81] Nature Computational Science, 'The carbon footprint of computational research', *Nature Computational Science*, Vol. 3, p. 659, 17 August 2023, <https://doi.org/10.1038/s43588-023-00506-2>.
- [82] Barykin, S. E., Kadyrova, O. V. and Senchin, P. K., 'Digital transformation of socio-economic development by the quality-of-life vision', *Information Sciences Letters*, Vol. 12, No 11, pp. 2859–2868, 1 November 2023, <http://dx.doi.org/10.18576/isl/121102>.
- [83] Lee, K. and Lee, S., 'Knowledge structure of the application of high-performance computing: A co-word analysis', *Sustainability*, Vol. 13, Issue 20, 11249, 12 October 2021, <https://doi.org/10.3390/su132011249>.
- [84] European Investment Bank, *Financing the future of supercomputing: How to increase investments in high performance computing in Europe*, European Investment Bank website, 26 June 2018, https://www.eib.org/files/pj/financing_the_future_of_supercomputing_en.pdf.
- [85] Samuel, J., Brennan-Tonetta, M., Samuel, Y., Subedi, P. and Smith, J., 'Strategies for democratization of supercomputing: Availability, accessibility and usability of high performance computing for education and practice of big data analytics', Cornell University arXiv, 19 April 2021, accessed 5 July 2024, <https://doi.org/10.48550/arXiv.2104.09091>.
- [86] Rončević, B., Coscodaru, R. and Fric, U. (eds), *Go with the flow: High performance computing and innovations in the Danube region*, Vega Press, London, 16 July 2019, <https://revis.openscience.si/lzpisGradiva.php?id=5936&lang=slv>.
- [87] Melicherčík, M., Siládi, V., Svítek, M. and Huraj, L., 'Spreading high performance computing skills with e-learning support', *2018 16th International Conference on Emerging eLearning Technologies and Applications (ICETA)*, IEEE, Stry Smokovec, 15-16 November 2018, pp. 361–366, <https://doi.org/10.1109/ICETA.2018.8572134>.
- [88] Ezell, S., 'A new frontier: Sustaining U.S. high-performance computing leadership in an exascale era', Information Technology & Innovation Foundation, 12 September 2022, <https://itif.org/publications/2022/09/12/high-performance-computing-leadership-in-an-exascale-era/>.
- [89] 'Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions – High-performance computing: Europe's place in a global race, COM(2012) 45 final of 15 February 2012, <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2012:0045:FIN:EN:PDF>.
- [90] Egli, A., 'ChatGPT, GPT-4, and other large language models: The next revolution for clinical microbiology?', *Clinical Infectious Diseases*, Vol. 77, No 9, 11 November 2023, pp. 1322–1328, <https://doi.org/10.1093/cid/ciad407>.
- [91] Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J. et al., 'Gen-AI: Artificial intelligence and the future of work', International Monetary Fund Staff Discussion Notes, 14 January 2024, <https://doi.org/10.5089/9798400262548.006>.



Academic, public administration and
industrial end-user requirements in
high-performance computing in the
European Union

Appendices

Appendices

Appendix 1 – Interview guidelines for different stakeholders

Section	Questions
<p>Section A. General information The objective of this first section is to ask general questions aimed at gaining insight into the context of HPC entity, within the HPC community. The responses will help us better understand the role, mission and services of the HPC entity.</p>	<p>1. What is the main funding source of the HPC entity? Public Private Collaborative initiatives</p> <p>2. What is the target field of knowledge for the services provided by the HPC system? You may choose several options.</p> <ul style="list-style-type: none"> • Healthcare • Government and defence • IT and telecom • Banking and finance • Mobility • Energy (non-renewable) • Energy (renewable) • Logistics • Retail and consumer goods • Agriculture • Media and entertainment • Material science • Biosciences • Aerospace engineering • Civil engineering • Environmental engineering • Astronomy • Physics • Earth sciences • Computer science • Applied mathematics • Other (provide detailed information) <p>3. How do you quantify the reach of your HPC entity’s network, specifically in terms of the number of users and developers involved?</p> <p>4. Are there specific industries or sectors where you believe specialised HPC services could bring significant benefits? If yes, please specify.</p> <p>5. What are the primary concerns of the users, when utilising an HPC service?</p> <ul style="list-style-type: none"> • CPU architecture (e.g. x86 versus ARM) • GPUs • Accelerators/heterogenous processing (e.g. GPU, ASIC, FPGASs, TPUs, PPU, DPUs, SOCs) • Internode communication (e.g. InfiniBand, Ethernet, Intel Omni-Path, Fibre Channel) • Network connectivity (e.g. intranet cloud, internet-based connectivity, extranets, VPNs) • Available software stack • Friendly interface • Workflows • Containerised environment (e.g. Docker, Singularity, Charliecloud, Shifter, Podman) • Calculation time limit (e.g. software based or checkpoint-restart, such as DMTCP) • Data storage • Software licences security measures

Section	Questions
<p>Section B. Commercial perspective In this section of the questionnaire, we pretend to infer how HPC services are overviewed from a commercial perspective. Our questions are focused on understanding the general commercial viewpoint of an HPC system user and comprehending the commercial requisites of the user community.</p>	<p>6. In the context of seeking an HPC service, what factors does the HPC provider consider pivotal for users to prioritise in their decision-making process? <i>You may order the following topics according to their importance, with the most important listed first and the least important topic listed last.</i></p> <ul style="list-style-type: none"> • Area of knowledge of the host entity • User community • Ease of scaling • Resources provided (performance, capability of parallel processing, storage capacity, internal connectivity and networking) • Security features • Contract conditions • Support services • Environmental/sustainability considerations • Others (you can specify) <p>7. What are the marketing strategies currently used to attract new end users to use HPC services (e.g. marketing strategies, online presence, public calls)?</p> <p>8. Does your organisation have its own HPC system? Yes or no?</p> <p>9. Could you give reason for your choice? What are the advantages of having or not having your own HPC system?</p>
<p>Section C. Legal perspective In this section of the questionnaire, we intend to analyse how HPC services are considered from a legal perspective. Our questions are focused on understanding the legal considerations of accessing an HPC system and comprehending the legal considerations and agreements needed between an HPC provider and the user community.</p>	<p>10. When users and developers submit data or workloads to an HPC system, it involves considerations related to data privacy, security and usage rights. What are the legal requirements needed for developers and users? <i>You may give examples for each legal consideration.</i></p> <ul style="list-style-type: none"> • Acceptable use policy • Data privacy and protection regulations • Intellectual property rights • Confidentiality and non-disclosure • Liability and indemnification regulations • Service-level agreements • Compliance with specific regulations, laws or institutional policies <p>11. Are there other agreements that dictate the terms and conditions associated to the software tools provided by the centre? For example, can users access software from other users? Or does the user require some specific measures regarding these topics?</p> <p>12. What types of agreements or contracts are signed with HPC developers and users?</p> <ul style="list-style-type: none"> • Variable price, metered usage • Fixed price, limited usage (e.g. core-hour) • Fixed price, unlimited usage <p>13. Does your entity undertake research contracts (non-monetary)?</p> <p>14. Do your entity undertake commercial contracts (monetary)? Yes or no? If yes, are there any loyalty periods? What are the associated advantages?</p>
<p>Section D. Financial perspective In this section of the questionnaire, our aim is to analyse how HPC services are viewed from a financial perspective. Our questions are focused on understanding the main financial aspects of an HPC provider entity and the financial considerations addressed by users.</p>	<p>15. In the last two years, on average:</p> <ul style="list-style-type: none"> - How many users does the HPC centre currently have? - How many have left and why? - How many new users are expected in the next two years? <p>16. Could you prioritise these topics based on the allocation of the budget to maintain HPC services and ensure optimal performance for users?</p> <ul style="list-style-type: none"> • Operational costs (facility costs, power and cooling costs, backup power) • Hardware maintenance (life cycle planning, regular maintenance and repairs) • Software resources • Infrastructure investments (hardware upgrades, networking infrastructure, storage infrastructure) • Human resources • Security • Educational initiatives (training and outreach programmes) • Financial planning (contingency and emergency funds) • Environmental considerations (sustainability of the HPC centre) <p>17. What type of HPC service do you find more financially challenging to acquire?</p>

Section	Questions
<p>Section E. Technical perspective In this section of the questionnaire, we intend to analyse how HPC services are available to the users. Our questions are focused on understanding the technical overview of an HPC system and inferring what systems are available, for examples architectures and programming languages.</p>	<p>18. What is the general architecture of the HPC system (e.g. CPU, GPU, DPU, VPU, FPGA and ASIC)?</p> <p>19. How do you categorise architecture in terms of homogeneous or heterogeneous system?</p> <p>20. Specific storage configuration in an HPC environment can vary widely based on the specific needs and goals of the computing cluster or centre. What types of storages are available in the HPC centre, which are accessible from the HPC systems? You may select several options.</p> <ul style="list-style-type: none"> • Local storage on compute nodes • Parallel file systems (i.e. scratch) • File storage shared across multiple HPC systems (i.e. project) • Object storage (i.e. S3, Swift) • Others <p>21. What are the interfaces available of the HPC system to the end user community? You may select several options.</p> <ul style="list-style-type: none"> • Command-line interface (SSH Protocol, batch job submission, interactive job sessions) • Graphical user interface (RDP, Teamviewer) • Data access (FTP/ SFTP, data transfer nodes, GridFTP, Globus) • Network access (VPN-web portal login) • Others <p>22. What types of job managers are implemented in the HPC centre for managing computational tasks and resource allocation (e.g. Slurm, PBS)?</p> <p>23. What are the programming environments accessible to the users? You may give several options according to each knowledge area.</p> <ul style="list-style-type: none"> • Data science (e.g. database format such as PostgreSQL, MariaDB, Mongo, other processing software) • Natural science (e.g. COMSOL Multiphysics, TensorFlow, PyTorch, scikit-learn) • Biosciences (e.g. AMBER, GROMACS, Schrodinger Suite) • Engineering (e.g. Openfoam, Ansys, start CCM+) • Scientific programming language (e.g. MATLAB, OCTAVE, Python) • IDE (e.g. Spyder, Visual studio, PyCharm, Eclipse) • Art/creative (e.g. Blender, Photoshop, Cinema 4D, Maya, Houdini, 3ds Max, Unreal) • IT (e.g. JAVA, Puppet) • Other programming languages (e.g. C/C++, HTML) <p>24. How are security measures implemented to safeguard HPC resources, infrastructures and user data (for example, in terms of malware and software security, secure software development, regular software updates)?</p> <p>25. How can developers and users gain access to the HPC services (e.g. physical tokens, two-factor authentication, public key, passwords)?</p>

Section	Questions
<p>Section F. Socioeconomic perspective</p> <p>In this section of the questionnaire, our aim is to analyse the potential influence of HPC services on the socioeconomic landscape of a society. Our questions are focused on understanding how HPC resources may contribute to the local and regional development.</p>	<p>26. How has the presence of the HPC centre contributed to the local and regional economy in terms of job creation, business development and overall economic growth? Do you have an estimate of the total number of jobs created, both directly and indirectly, due to the presence of HPC facilities in the region. Could you give examples?</p> <p>27. How does the HPC centre collaborate with government agencies and institutions to address societal challenges, public policy issues or national priorities?</p> <p>28. How does the HPC centre engage with the local community, and what initiatives are in place to ensure that the benefits of HPC are accessible to a broader audience?</p> <p>29. How does the HPC centre engage with the local community, and what initiatives are in place to ensure that the benefits of HPC are accessible to a broader audience?</p> <p>30. What measures has the HPC entity taken to address environmental concerns? You may provide an example for each topic.</p> <ul style="list-style-type: none"> • Energy consumption • Renewable energy use • Material sourcing awareness • Cooling system • Life cost assessment of an HPC • Energy surplus management • System retirement plans

Appendix 2 – Web-based questionnaire structure

1. Participant profile

1. What is the gender you identify with?

Question type: multiple choice with one answer

Required answer

- A. Man
- B. Woman
- C. Non-binary
- D. Transgender
- E. Prefer not to say
- F. Other

2. What is your age group?

Question type: multiple choice with one answer

Not required answer

- A. < 21 years
- B. 21-30 years
- C. 31-40 years
- D. 41-50 years
- E. 51-60 years
- F. 61-70 years
- G. 70 years
- H. Prefer not to say

3. In which country do you currently live?

Question type: dropdown with one answer

Required answer

- A. All of the countries

4. What is your highest level of education completed?

Question type: multiple choice with one answer

Required answer

- A. High school diploma
- B. Bachelor's degree
- C. Professional degree
- D. Master's degree
- E. Doctorate
- F. Other

2. Organisation profile

1. What is the sector that best represents your organisation?

Question type: multiple choice with one answer

Required answer

- | | |
|------------------------------|------------------------|
| A. Healthcare | R. Astronomy |
| B. Government and defence | S. Physics |
| C. IT and telecom | T. Earth sciences |
| D. Banking and finance | U. Computer science |
| E. Mobility | V. Applied mathematics |
| F. Energy (non-renewable) | W. Other (open answer) |
| G. Energy (renewable) | |
| H. Logistics | |
| I. Retail and consumer goods | |
| J. Agriculture | |
| K. Media and entertainment | |
| L. Material science | |
| M. Biosciences | |
| N. Aerospace engineering | |
| O. Civil engineering | |
| P. Environmental engineering | |
| Q. Maritime engineering | |

2. Which one of the following options best describes the organisation you work for?

Question type: multiple choice with one answer

- A. Required answer
- B. Academic institution
- C. Research institute
- D. Start-up
- E. SME
- F. Large company
- G. Other (open answer)

3. What is the main funding source of your organisation?

Question type: multiple choice with one answer

Not required answer

- A. Public
- B. Private
- C. Both

4. Your organisation is...

Question type: multiple choice with one answer

Required answer

- A. For-profit
- B. Not-for-profit

3. General characterisation as an HPC stakeholder

The following questions will help us to understand your engagement within the HPC community.

1. How long have you been working with HPC?

Question type: multiple choice with one answer

Required answer

- A. Less than 1 year
- B. 1-5 years
- C. More than 5 years
- D. I've never worked with HPC

If the participant answer D on the question 1. The next question is:

1.1 Do you think your organisation could benefit from using HPC resources?

Question type: yes or no question

Required answer

- A. Yes
- B. No

If the participant answer A on the question 1.1. The next question is:

1.1.1. If yes, what are the reasons for your organisation do not procure HPC resources?

Question type: multiple choice with several answers (up to 4)

Required answer

- A. Business objectives and use cases are not related with HPC service
- B. Budget limitations
- C. Data sensitivity and security concerns
- D. Lack of customisation or control from your organisation
- E. Data transfer and latency concerns
- F. Risk aversion
- G. Insufficient expertise human resources to deal with an HPC systems
- H. I'm not aware of HPC
- I. Other (open answer)

After the question 1.1.1 the survey ends for this type of participant.

If the participant answer B on the question 1.1. The next question is:

1.1.2. In your opinion, what are the reasons that HPC may not bring benefits to your organisation?

Question type: multiple choice with several answers (up to 5)

Required answer

- A. Cost of implementing and maintaining HPC infrastructure
- B. Uncertain ROI (Return of investment)
- C. Workloads are not suited to HPC infrastructures
- D. Difficulties with data accessibility and data storage
- E. Data sensitivity and security concerns
- F. I don't know
- G. Other (open answer)

1.1.3. If you consider to use or provide HPC in the future, how relevant are the following aspects?

Question type: matrix (select one answer per topic)

Required answer

	Not important	Moderately important	Very important	No opinion
Understanding the business needs				
Availability of funding (private, public, both)				
Experienced human resources				
Legal support				
Infrastructure and resources				
Planning				
Software and tool chain compatibility				
Data management and security				
Awareness of HPC services				

After the question 1.1.3. the survey ends for this type of participant.

If the participant answer A, B, C on the question 1. The next question is:

2. Do you think an HPC system or service optimised for a specific application (e.g. AI, Big Data) can provide a competitive advantage compared to a traditional HPC service, that addresses a diverse range of applications?

Question type: multiple choice with one answer

Required answer

- A. Yes
- B. No
- C. I don't have an opinion

If the participant answered A, B on the question 2. The next question is:

- 2.1 Could you justify your previous answer?

Question type: open question with short answer

Required answer

After this question the participant goes to 3.

3. What types of HPC environments do you think will grow faster in the future?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. On-premises facilities and COE clusters
- B. Cloud-based HPC
- C. Hybrid HPC
- D. Grid computing
- E. High-performance edge computing
- F. Emerging technologies (e.g. Quantum computing)
- G. Other (open answer)

4. Could you justify your previous answer?

Question type: Open question with short answer.

Not required answer.

5. From your perspective, and considering the following topics, what are the user's concerns when accessing and using HPC systems?

Question type: ranking

Required answer

Please order by preference, where 1 is the most important.

1. Abundance of computational power
2. Sufficient RAM to accommodate intensive workloads
3. Reliable and high-speed network connectivity
4. High-speed file system/data storage
5. Existence of accelerators or heterogenous processing
6. Less strict job manager policies
7. Availability of necessary software environments and libraries
8. Robust security and data privacy policies
9. Effective user support
10. Friendly interface (e.g. Graphic User Interface)
11. Generous data storage policies (e.g. quotas, long-term storage)

4. Section 'Legal considerations'

From the following questions, we would like to understand your perspective related to the legal considerations when accessing HPC systems.

1. In your opinion, what are the legal obligations that developers and users have to consider regarding data privacy, security and usage rights when submitting data or workloads to an HPC system?

Question type: multiple choice with several answers (unlimited)

Required answer

Please select all the considerations you agree with.

- A. Acceptable Use Policy
- B. Data privacy and protection regulations
- C. Intellectual Property Rights
- D. Confidentiality and non-disclosure
- E. Liability and indemnification regulations
- F. Service-level agreements
- G. Compliance with specific regulations, laws or institutional policies
- H. Ethical considerations
- I. Other (open answer)

2. From the following HPC environments and types of services, which ones represent the greatest legal challenges?

Question type: multiple choice with several answer (up to 3)

Required answer

Please, select the options that you agree most.

Click on 'Glossary' for assistance in addressing the question.

- A. On-site HPC facilities and COE clusters
- B. Cloud-based HPC
- C. Hybrid HPC
- D. Shared HPC clusters
- E. HPC as a service
- F. Grid computing
- G. Emerging technologies (e.g. Quantum computing)
- H. High-performance edge computing
- I. I don't know

5. Section 'Socioeconomic considerations'

From the following questions, we would like to understand your perspective about the socioeconomic impact of using HPC systems.

1. What significant results has your organisation achieved through the utilisation of HPC systems?

Question type: multiple choice with several answer (unlimited)

Required answer

Please, select the options that apply to your main achievements.

- | | |
|--|--|
| A. We did not achieve significant results | E. Contribution to government agencies |
| B. Scientific research achievements | F. Innovative approaches |
| C. Technological advancements | G. Competitive advantage |
| D. Collaborations with industries to address real-world challenges | H. Other (open answer) |

2. From your previous answer, could you provide specific examples of significant achievements that your organisation has made through the utilisation of HPC systems?

Question type: open question with short answer

Required answer

3. What strategies do you think could be implemented to improve accessibility to HPC services in regions with limited availability?

Question type: matrix

Not required answer

	Not at all	Very little	Somewha	To a great extent
Public investment in HPC infrastructures				
Private investment in HPC infrastructures				
Availability of cloud-based HPC services				
Providing HPC-as-a-service to access HPC infrastructures remotely in other regions				
Public-private partnerships				
Regional HPC consortia				
Open access policies to make research findings, data and HPC resources freely available to promote transparency, collaboration and knowledge in the community				
Education and training to the community				
Investment in remote collaboration tools (e.g. remote access solutions)				
Developing sustainable funding models for long-term availability, operation and growth of HPC resources/services				

6. Section for each type of stakeholder

1. From the following options which one best describes you as a stakeholder?

User: Entities that utilise HPC resources and infrastructure to perform computationally intensive tasks, simulations, analysis or research in various domains. These users rely on HPC systems to efficiently process large volumes of data, solve complex computational problems and enhance workflows in scientific, engineering and industrial areas.

Provider: Entities that offer HPC as a service or use them to provide data at large-scales, either for a profit or not. These services can encompass a range of offers, including hardware infrastructure (such as supercomputers or cloud-based HPC resources), software tools and applications optimised for parallel computing, data management and storage solutions, technical support, consulting services and access to specialised expertise in HPC-related fields. Example of these stakeholders: HPC centres.

Developer: Entities that are specialised in designing, building, optimising and maintaining software and hardware systems for HPC environments. These developers specialise in developing applications, tools, algorithms and frameworks that utilise parallel processing, distributed computing and advanced hardware architectures to achieve high levels of computational performance, scalability and efficiency. They create the infrastructure that the HPC providers maintain and distribute to HPC users, such as software developers.

Question type: multiple choice with one answer

Required answer

- A. I'm a user
- B. I'm a provider
- C. I'm a developer

If the participant answers A, they go to the 'User' section.

If the participant answers B, they go to the 'Provider' section.

If the participant answers C, they go to the 'Developer' section.

6.1. Section 'User': In this section we would like to understand your perspective as an HPC user...

1. How often do you use HPC services or infrastructures?

Question type: multiple choice with one answer

Required answer

- A. Daily
- B. Once a week
- C. 2-3 times a month
- D. Once a month
- E. Every 2-3 months
- F. 2-3 times a year or less

2. Are there specific industries or sectors where you believe specialised HPC services would bring significant benefits?

Question type: yes or no question

Required answer

- A. Yes
- B. No

If the participant answer A in the question 2. The next question is:

2.1 If 'Yes', which sector or sectors?

Question type: multiple choice with several choice (unlimited)

Required answer

- A. Healthcare
- B. Government and defence
- C. IT and telecom
- D. Banking and finance
- E. Mobility
- F. Energy (non-renewable)
- G. Energy (renewable)
- H. Logistics
- I. Retail and consumer goods
- J. Agriculture
- K. Media and entertainment
- L. Material science
- M. Biosciences
- N. Aerospace engineering
- O. Civil engineering
- P. Environmental engineering
- Q. Maritime engineering
- R. Astronomy
- S. Physics
- 20. Earth sciences
- U. Computer science
- V. Applied mathematics
- W. Other (open answer)

After answered the question 2.1, the participants go to 3.

3. Does your organisation have its own HPC system?

Question type: yes or no question

Required answer

- A. Yes
- B. No

If the participant answer A in the question 3. The next question is:

3.1 What are the main advantages of having your own HPC system?

Question type: open question with short answer

Required answer

If the participant answer B in the question 3. The next question is:

3.2 If no, why? Please, rank by order.

Question type: ranking

Required answer

Please select 1 for the main reason.

1. Insufficient specific expertise to installation, configuration and ongoing maintenance
2. Insufficient Initial investments
3. Difficulty in predicting future computational needs
4. Space and power requirements
5. Integration with existing systems, existing IT infrastructure and workflows
6. Constant technological advancements in the software and hardware innovations entering the market
7. Risk of obsolescence
8. Security and compliance concerns
9. Other

If the participant also consider the option 'Other' in the previous question. The next questions is:

3.2.1. If the option 'Other' in the previous question is not ranked in the position 9, please, specify:

Question type: short answer

Required answer

4. How many HPC systems is your organisation using, at the moment?

Question type: multiple choice with one answer

Required answer

- A. 1
- B. 2
- C. 3
- D. More than 4

If the participant answers B, C, D, E, the next question is:

4.1 Which factors contribute to use more than one HPC system?

Question type: multiple choice with several answer (unlimited)

Required answer

- A. HPC-system specific projects/workloads
- B. Geographic distribution
- C. Capacity planning and scalability
- D. Specialised expertise and services
- E. Cost optimisation and budget allocation
- F. Compliance and data sovereignty
- G. Vendor lock-in mitigation
- H. All above

5. Do you prefer to use a public funded or a private funded HPC system?

Question type: multiple choice with one answer

Required answer

- A. Public
- B. Private
- C. I don't have any preference

6. Do you prefer to access....

Question type: multiple choice with one answer

Required answer

- A. A private cluster
- B. A shared cluster
- C. HPC as a service
- D. I don't have any preference

7. What types of HPC environments and services have your organisation used/acquired?

Click on 'Glossary' for assistance in addressing the question.

Question type: multiple choice with several answers (unlimited)

Required answer

- A. On-premises HPC clusters
- B. Cloud-based HPC
- C. Hybrid HPC
- D. Shared HPC clusters
- E. HPC as a service
- F. Grid computing
- G. Emerging technologies (e.g. quantum computing)
- H. High-performance edge computing
- I. Other (open answer)

8. From the following HPC environments and services, which are the most financial challenging to use/acquire?

Click on 'Glossary' for assistance in addressing the question.

Question type: multiple choice with several answers (up to 3)

Required answer

- A. On-premises HPC clusters
- B. Cloud-based HPC
- C. Hybrid HPC
- D. Shared HPC clusters
- E. HPC as a service
- F. Grid computing
- G. Emerging technologies (e.g. quantum computing)
- H. High-performance edge computing
- I. Other (open answer)

9. What are the most common funding sources to acquire/access HPC services?

Question type: multiple choice with several answers (up to 4)

- A. Required answer
- B. Public funding
- C. Private funding
- D. Industry partnership
- E. Public-private partnerships
- F. Internal resources
- G. Other (e.g. crowdfunding, venture capital)

10. Which factors influence the decision to finalise a specific contract with an HPC provider entity?

Question type: matrix

Required answer

	Definitely not	Probably not	Probably	Very probably	Definitely
Pricing structures (e.g. upfront costs, subscription fees, usage-based charges)					
Performance guarantees offered by the provider (e.g. minimum processing speeds, memory allocation, network bandwidth)					
Flexibility in resource allocation/scalability options given by the HPC provider					
Service-level agreements refers to service quality, uptime commitments, support responsiveness, compensation for service disruptions)					
Stringent security measures (e.g. encryption of data in transit and at rest, access controls, strong network security, RAID data storage redundancy)					
Compliance with industry regulations					
Effective incident response and disaster recovery plans (e.g. advanced intrusion detection systems, regular backups)					
Good support services offered by the HPC provider (e.g. 24/7 technical assistance, proactive monitoring, and regular maintenance)					
Vender reputation and reliability					
Customisation options (e.g. hardware configurations, software environments)					
Clear exit strategy for the user (e.g. contract termination terms, data migration support and vendor lock-in risks)					
Environmental/sustainability considerations of the HPC provider					
Geographical location					

11. Are there other important aspects regarding HPC user choices to be considered?

Question type: open question with short answer

Required answer

12. What marketing strategies are HPC providers applying that make their services appealing enough to you, as an user, to select a specific provider?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. None
- B. Specific marketing strategies
- C. Online presence
- D. In-person events
- E. Public calls
- F. Other (open answer)

13. Which percentage of year budget does your organisation allocate to HPC services?

Question type: multiple choice with one answer

Not required answer

- A. 0-10 %
- B. 11-30 %
- C. 31-50 %
- D. 51-70 %
- E. 71-90 %
- F. 91 %

14. Please select the categories receiving the highest financial resources within all HPC costs.

Question type: multiple choice with several answers (up to 3)

- A. Required answer
- B. Hardware resources
- C. Networking infrastructure
- D. Operating costs
- E. Software licences
- F. Contract fees
- G. Technical support
- H. Data storage expenses
- I. Data transfer expenses
- J. Secure software development
- K. Regular software update
- L. Energy consumption
- M. Training and contracting personnel
- N. Outside consulting services
- O. Other (open answer)

15. How do you classify the access to HPC resources in your region?

Question type: opinion scale

Required answer

0	1	2	3	4	5	6	7	8	9
Inexistent									Excellent

The participant has to select between a scale 0-9.

16. What are the main challenges you face when you need/attempt to access HPC resources outside of your continent?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. I don't access HPC resources outside of my continent
- B. Geographical barriers
- C. Network latency and bandwidth
- D. Limited local technical support
- E. Limited remote technical support
- F. Funding constraints
- G. Data transfer and storage
- H. Data security
- I. Regulatory compliance
- J. Cultural aspects
- K. None
- L. Other (open answer)

This is the end of the survey for users.

6.2. Section 'Provider': In this section we would like to understand your perspectives as an HPC provider....

1. As an HPC provider, do you offer:

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Hardware infrastructures
- B. Software tools and applications
- C. Data management and storage
- D. Technical support and consulting
- E. Training services
- F. Other services (e.g. cloud services)
- G. All above

2. What types of HPC environments does your organisation provide?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. On-premises HPC clusters
- B. Cloud-based HPC
- C. Hybrid HPC
- D. Grid computing
- E. Emerging technologies (e.g. quantum computing)
- F. High-performance edge computing
- G. Other (open answer)

3. What are the sectors represented by your main users?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Healthcare
- B. Government and defence
- C. IT and telecom
- D. Banking and finance
- E. Mobility
- F. Energy (non-renewable)
- G. Energy (renewable)
- H. Logistics
- I. Retail and consumer goods
- J. Agriculture
- K. Media and entertainment
- L. Material science
- M. Biosciences
- N. Aerospace engineering
- O. Civil engineering
- P. Environmental engineering
- Q. Maritime engineering
- R. Astronomy
- S. Physics
- T. Earth sciences
- U. Computer science
- V. Applied mathematics
- W. Other (open answer)

4. How many active users and developers does your organisation have?

Question type: open question with short answer

Required answer

5. How many users stopped using your services in the last two years and why?

If you don't know why, please refer that.

Question type: open question with short answer

Required answer

6. How many new users are expected in the next two years?

Question type: open question with short answer

Required answer

7. What marketing strategies do you use to make your services more attractive to new users?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Specific marketing strategies
- B. Online presence
- C. In-person events
- D. Public calls
- E. Other (open answer)

8. Select the target continent of your services.

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Africa
- B. Asia
- C. Europe
- D. North America
- E. South America
- F. Oceania
- G. All of the above

9. Are there specific industries or sectors where you believe specialised HPC services bring significant benefits?

Question type: yes or no question

Required answer

- A. Yes
- B. No

If the participant answer A in the question 9. It goes to the next question:

9.1 If 'Yes', which sector or sectors?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Healthcare
- B. Government and defence
- C. IT and telecom
- D. Banking and finance
- E. Mobility
- F. Energy (non-renewable)
- G. Energy (renewable)
- H. Logistics
- I. Retail and consumer goods
- J. Agriculture
- K. Media and entertainment
- L. Material science
- M. Biosciences
- N. Aerospace engineering
- O. Civil engineering
- P. Environmental engineering
- Q. Maritime engineering
- R. Astronomy
- S. Physics
- T. Earth sciences
- U. Computer science
- V. Applied mathematics
- W. Other (open answer)

10. What types of agreements or contracts are provided by your organisation?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Variable price, metered usage
- B. Fixed price, limited usage
- C. Fixed price, unlimited usage

11. Does your organisation allow for non-paid research HPC work?

Question type: yes or no question

Required answer

- A. Yes
- B. No

12. Does your organisation allow for commercial (paid) HPC work?

Question type: yes or no question

Required answer

- A. Yes
- B. No

If the participant answer A in the question 12. The next question is:

12.1 If 'Yes', is there any loyalty period?

Question type: yes or no question

Not required answer

- A. Yes
- B. No

If the participant answer A in the question 12.1. The next question is:

12.1.1. If 'Yes', what are the advantages of having a loyalty period?

Question type: short answer

Not required answer

13. From the following contracts, which are the most common in your organisation?

Question type: multiple choice with one answer

Not required answer

- A. Commercial contracts
- B. Research (non-paid) contracts

14. Do you think this will change in the future? Why?

Question type: open question with short answer

Not required answer

15. What is the current financial income of your entity from providing HPC services?

Question type: multiple choice with one answer

Not required answer

- A. Less than 0.5 million €
- B. Between 0.5 million and 1 million €
- C. Between 1 million and 10 million €
- D. Between 10 million and 100 million €
- E. Between 100 million and 400 million €
- F. More than 400 million €
- G. I don't know
- H. Prefer not to say

16. Rank the following topics based on the allocation of budget to maintain HPC services and ensure optimal performance for users?

Please select 1 the topic with higher budget allocation.

Question type: ranking

Not required answer

1. Operational costs (e.g. facility costs, power and cooling costs, backup power)
2. Hardware maintenance (e.g. life cycle planning, regular maintenance and repairs)
3. Software resources
4. Infrastructure investments (e.g. upgrades for network infrastructure or for storage infrastructure)
5. Human resources
6. Security
7. Educational initiatives (e.g. training and outreach programmes)
8. Financial planning (e.g. contingency and emergency funds)
9. Environmental considerations (e.g. sustainability of the HPC centre)

17. Specific storage configurations in an HPC environment can vary widely based on the specific needs and goals of the computing cluster or centre. What types of storage are available in your organisation for the HPC systems provided?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Local storage on compute nodes
- B. Parallel file systems (e.g. scratch)
- C. File storage shared across multiple HPC systems (e.g. project)
- D. Object storage (e.g. S3, Swift)
- E. Other (open answer)

18. From the previous topics, which are the ones predominantly preferred by users?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Local storage on compute nodes
- B. Parallel file systems (e.g. scratch)
- C. File storage shared across multiple HPC systems (e.g. project)
- D. Object storage (e.g. S3, Swift)
- E. Other (open answer)

19. Which factors could influence the choice of the user for a specific storage configuration?

Question type: open question with short answer

Required answer

20. What are the interfaces provided by your entity to users through command-line interface?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. SSH Protocol
- B. Batch job submission (e.g. Slurm, PBS/Torque)
- C. File transfer and management (e.g. SCP, SFTP)
- D. Environment and software management (e.g. Module System)
- E. Containerisation interfaces (e.g. Docker, Singularity)
- F. Data Management Tools
- G. Other (open answer)

21. What are the interfaces provided by your entity to the users through graphical user interfaces?

Question type: multiple choice with several answers (unlimited)

- A. Required answer
- B. Job Management Tools (e.g. SlurmView)
- C. Cluster Monitoring Dashboards (e.g. Ganglia)
- D. File Management Tools (e.g. Cyberduck)
- E. Remote Desktop Access (e.g. RDP, Teamviewer)
- F. Visualisation Software (e.g. ParaView)
- G. Workflow Management Platforms (e.g. Pegasus)
- H. Other

22. What are the interfaces provided by your entity to users for data access/transfer?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. FTP/SFTP
- B. Data transfer nodes
- C. GridFTP
- D. Globus
- E. SCP
- F. Other (open answer)

23. What are the interfaces provided by your entity to users for network access to the HPC systems?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. VPN
- B. Web portal login
- C. APIs
- D. SSH Protocol
- E. RDP
- F. Other

24. How does your organisation contribute to the local and regional economy in terms of job creation, business development and overall economic growth?

Please provide some examples

Question type: open question with long answer

Required answer

25. What measures has your entity taken to address environmental concerns?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Energy reduction
- B. Renewable energy usage
- C. Material sourcing awareness
- D. Energy efficient cooling systems
- E. Waste reduction and recycling programs
- F. Energy surplus management
- G. System retirement plans
- H. Carbon footprint assessments and carbon credits
- I. Other

This is the end of the survey for providers.

6.3. Section 'Developer': In this section we would like to understand your perspective as an HPC developer...

1. From these emerging technologies, what do you believe will mostly contribute to make HPC more accessible to different types of organisations (e.g. large companies, SME, Startups)?

Question type: multiple choice with several answers (unlimited)

Required answer

- A. Quantum computing
- B. Improvements of accelerators architectures (e.g. specialised accelerators for specific tasks)
- C. Integration of heterogeneous architectures
- D. Integration of AI and machine learning techniques
- E. Edge computing
- F. Other (open answer)

2. Do you think exascale computing will be widely available in the near future?

Question type: yes or no question

Required answer

- A. Yes
- B. No

3. What sectors do you think will mostly benefit from exascale computing, in your perspective?

Question type: multiple choice with several answers (unlimited)

Required answer

- | | |
|------------------------------|-------------------------|
| A. Healthcare | Q. Maritime engineering |
| B. Government and defence | R. Astronomy |
| C. IT and telecom | S. Physics |
| D. Banking and finance | T. Earth sciences |
| E. Mobility | U. Computer science |
| F. Energy (non-renewable) | V. Applied mathematics |
| G. Energy (renewable) | W. Other (open answer) |
| H. Logistics | |
| I. Retail and consumer goods | |
| J. Agriculture | |
| K. Media and entertainment | |
| L. Material science | |
| M. Biosciences | |
| N. Aerospace engineering | |
| O. Civil engineering | |
| P. Environmental engineering | |

4. In your opinion, what aspects of software technology development still need to be addressed for the expansion of exascale computing?

Question type: multiple choice with several answers (up to 3)

Required answer

- A. I don't know
- B. Nothing
- C. Parallelisation in heterogenous hardware architectures
- D. Portability to heterogenous hardware architectures
- E. Energy-efficiency (power usage versus computational performance)
- F. Data handling
- G. Fault tolerance and recovery mechanism
- H. Other

5. In your opinion, what aspects of hardware technology still need to be addressed for the expansion of exascale computing?

Question type: multiple choice with several answers (up to 3)

Required answer

- A. I don't know
- B. Nothing
- C. New/improved specialised hardware (e.g. FPGAs, ASICs, etc.)
- D. Improvements in RAM (volatile) technology
- E. Better hardware interconnection technology (higher bandwidth and lower latency)
- F. More effective cooling systems
- G. Robust power delivery systems
- H. Other (open answer)

6. From the following types of hardware, which ones will bring more competitiveness in addressing HPC future challenges (e.g. higher performance, energy efficiency)?

Question type: multiple choice with several answers (up to 3)

Required answer

- A. FPGAs
- B. ASICs
- C. CPUs
- D. GPUs
- E. Specific AI accelerators
- F. Other

7. From your perspective, what CPU architecture will bring more benefits for HPC systems, in the future?

Question type: multiple choice with one answer

Required answer

- A. ARM
- B. x86
- C. Both
- D. None
- E. Other

8. Could you justify your previous answer?

Question type: open question with long answer

Not required answer

9. Do you believe quantum computing will contribute to shape future trends in the HPC community?

Question type: yes or no question

Required answer

- A. Yes
- B. No

10. Could you justify your previous answer?

Question type: open question with long answer

Not required answer

11. What sectors do you think will mostly benefit from quantum computing, in your perspective

Question type: multiple choice with several answers (unlimited)

Required answer

- | | |
|------------------------------|------------------------------|
| A. Healthcare | M. Biosciences |
| B. Government and defence | N. Aerospace engineering |
| C. IT and telecom | O. Civil engineering |
| D. Banking and finance | P. Environmental engineering |
| E. Mobility | Q. Maritime engineering |
| F. Energy (non-renewable) | R. Astronomy |
| G. Energy (renewable) | S. Physics |
| H. Logistics | T. Earth sciences |
| I. Retail and consumer goods | U. Computer science |
| J. Agriculture | V. Applied mathematics |
| K. Media and entertainment | W. Other (open answer) |
| L. Material science | |

This is the end of the survey for developers.

7. Section 'Contact info'

(at the end of the survey all participants will see the following section)

If you are available to be contacted for an interview, please leave here your contact details.

1. First name
2. Last name
3. Email
4. Company

8. Section 'End screen'

Thank you for your contribution!

Appendix 3 – Categories defined in the study analysis

Table 10. Categories definition for the study analysis

Category	Classification
Gender	Female, male, prefer not to say
Education level	High school diploma, Professional degree, Bachelor's degree, Master's degree, Doctorate
Organisation type	Academic institution, research institute, large company, SME, start-up, governmental organisation, NGO, no organisation
Level of expertise working in HPC field	Never worked with HPC, less than one year, one to five years, more than five years
Stakeholder type	User, developer, provider

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct centres. You can find the address of the centre nearest you online (european-union.europa.eu/contact-eu/meet-us_en).

On the phone or in writing

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

- by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls);
- at the following standard number: +32 22999696;
- via the following form: european-union.europa.eu/contact-eu/write-us_en.

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website (european-union.europa.eu).

EU publications

You can view or order EU publications at op.europa.eu/en/publications. Multiple copies of free publications can be obtained by contacting Europe Direct or your local documentation centre (european-union.europa.eu/contact-eu/meet-us_en).

EU law and related documents

For access to legal information from the EU, including all EU law since 1951 in all the official language versions, go to EUR-Lex (eur-lex.europa.eu).

EU open data

The portal data.europa.eu provides access to open datasets from the EU institutions, bodies and agencies. These can be downloaded and reused for free, for both commercial and non-commercial purposes. The portal also provides access to a wealth of datasets from European countries.



EuroHPC
Joint Undertaking

**The European High Performance Computing
Joint Undertaking (EuroHPC JU)**

is a joint initiative between the EU, European
countries and private partners to develop
a world class supercomputing ecosystem



Publications Office
of the European Union
Law | Data | Publications

ISBN 978-92-95228-09-2