



Empowering
Europe's Data
Future



HPC-CLOUD AND QUANTUM COMPUTING: STATE OF THE ART AND INNOVATION ROADMAP

White Paper

16/12/2025



Funded by
the European Union

PARTNERS



HPC-CLOUD AND QUANTUM COMPUTING: STATE OF THE ART AND INNOVATION ROADMAP

WHITE PAPER

Work package	WP Number 8
Task	Task Number 8.3
Due date	31/12/2025
Submission date	19/12/2025
Deliverable lead	NCSR&D
Version	0.1
Authors	Panagiotis Krokidas, Christoforos Rekatsinas, Periklis Terlixidis, Georgios Giannakopoulos, Konstantinos Rallis, Panagiotis Dimitrakis "NCSR Demokritos" and Nikolaos Melissourgos "HPE Hellas S.A."
Reviewers	
Abstract	<p>This white paper presents an overview of the current state of the art in high-performance computing (HPC) and its convergence with cloud technologies, with a strategic focus on innovation management and exploitation. It outlines recent advances in cloud-based HPC services, hybrid architectures, and federated systems that integrate edge, cloud, and HPC resources. Emerging paradigms such as federated learning, AI-driven optimization, and sustainable computing are analyzed for their transformative potential. Special emphasis is placed on the NOUS project, which exemplifies a holistic approach to federated HPC-cloud services, combining technical innovation with robust exploitation strategies. NOUS goes one step beyond and explores the integration of quantum computing in handling data existing in the cloud. As of late 2025, the synergy between Cloud Computing and Quantum Computing (QC) has matured into a functional "Quantum-as-a-Service" (QaaS) model. While physical quantum hardware remains too fragile for on-premise deployment, cloud providers have democratized access to the "Quantum Stack." This report highlights this progress, the issues and real future applications. NOUS addresses European priorities for digital sovereignty and data</p>

interoperability, supporting scalable, privacy-preserving, and AI-enabled HPC workflows. The paper concludes with a roadmap that positions NOUS as a reference architecture and innovation catalyst for Europe's distributed computing ecosystem.

Keywords	High Performance Computing; HPC; Quantum Computer; hybrid computing
----------	---

Document Revision History

Version	Date	Description of change	List of contributor(s)
0.1	19/12/2025	Final Review	Konstantinos Rallis, Panagiotis Dimitrakis

DISCLAIMER

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Health and Digital Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

COPYRIGHT NOTICE

© NOUS Consortium, 2024

This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both. Reproduction is authorised provided the source is acknowledged.

EXECUTIVE SUMMARY

High-performance computing (HPC) and cloud computing are converging, enabling seamless access to computing resources across on-premises and cloud environments. This convergence allows for hybrid usage models, leveraging cloud resources for peak demand and specialized hardware. The integration of HPC and cloud technologies is accelerating innovation, making advanced computing more accessible and adaptable.

The current generation of quantum computers, characterized by a low number of noisy qubits, is unsuitable for real applications but useful for exploring small-scale applications and testing algorithms. The next generation of quantum computers will feature fault-tolerant architectures with a sufficient number of qubits, enabling error-free large-scale quantum computing. European quantum computers are being integrated with high-performance computing (HPC) to enhance capabilities in scientific research, innovation, and industrial applications.

The NOUS project, a European initiative, aims to empower Europe's data-driven future by integrating data spaces, high-performance computing, and edge computing. It focuses on developing advanced cloud technologies for scientists and professionals, enabling efficient data sharing through secure data infrastructures. NOUS's innovation is validated through use cases in automotive, energy, crisis management, and science, leveraging Common European Data Spaces for cross-organization data sharing and analytics.

The convergence of HPC, cloud, and quantum computing is creating a distributed and flexible computational environment. Continued research, infrastructure projects, and supportive policies are needed to nurture this convergence and ensure the next generation of scientific discovery and industrial progress.

TABLE OF CONTENTS

1	HPC-CLOUD CONVERGENCE: TRANSFORMING HIGH-PERFORMANCE COMPUTING	7
2	HPC-CLOUD INFRASTRUCTURE AND SERVICES	7
3	EDGE-TO-CLOUD-TO-HPC: THE DISTRIBUTED COMPUTING CONTINUUM	8
4	FEDERATED LEARNING IN HPC-CLOUD ENVIRONMENTS	8
5	EXASCALE HPC	9
6	HPC-CLOUD AND QUANTUM SYNERGY	9
6.1	Cloud data treatment by quantum computer	11
7	NOUS PROJECT INNOVATION AND EXPLOITATION STRATEGY	12
8	CONCLUSION AND ROADMAP	14

1 HPC-CLOUD CONVERGENCE: TRANSFORMING HIGH-PERFORMANCE COMPUTING

High-performance computing and cloud computing are converging to form a seamless continuum of compute resources. Traditional boundaries between on-premises supercomputers and cloud platforms have begun to blur as workflows adopt cloud-native approaches – for example, containerization and orchestration – across both environments [1].

Scientists now run massive simulations and data analytics interchangeably on university supercomputers and commercial clouds, effectively intersecting these once-distinct worlds [2]. This convergence enables hybrid usage models: organizations can keep most workloads on local HPC clusters while “cloud bursting” to public clouds for peak demand or specialized hardware (e.g. GPUs) [3]. European initiatives are actively driving HPC-cloud convergence; for instance, PRACE partners are developing reference architectures to integrate cloud flexibility with HPC infrastructure [4]. Likewise, EU-funded projects such as LEXIS [5] have built platforms uniting HPC, cloud and big-data technologies, demonstrating unified workflows across these domains [6].

The result is an ecosystem where HPC-grade performance meets cloud agility, allowing researchers and industry to scale computing resources on-demand and access specialized architectures without the traditional barriers [7]. This trend is accelerating innovation by making advanced computing more accessible and adaptable than ever before.

2 HPC-CLOUD INFRASTRUCTURE AND SERVICES

The ecosystem of HPC-cloud infrastructure blends public and private platforms that deliver HPC-as-a-Service with global reach. Major cloud providers have introduced specialized HPC offerings – for example, AWS provides an integrated suite for quickly building HPC clusters on cloud with virtually unlimited capacity and high-speed interconnects, and Google Cloud’s HPC solution offers pre-configured high-performance VMs plus AI/ML integration for data analysis [8], [9].

These services grant researchers and engineers on-demand access to supercomputing-class resources, overcoming the limitations of fixed on-premise capacity. In parallel, new players are innovating with sovereign and sustainable HPC cloud platforms. A notable example is Europe’s Qarnot, which deploys distributed HPC clusters and recycles their waste heat, targeting industries like aerospace, automotive, energy, and AI [10]. Backed by European innovation funds, Qarnot’s cloud is designed to be fully sovereign – allowing European users to innovate without reliance on non-EU providers – aligning with the EU’s digital autonomy goals. On the public infrastructure side, initiatives such as the European Open Science Cloud (EOSC) are federating national HPC centers and cloud resources to provide a common platform for research computing and data services [11].

This federated model, together with emerging cloud marketplaces and HPC resource federations, enables small and medium enterprises (SMEs) and scientific teams to access HPC capabilities on a pay-per-use basis. The net effect is a more agile, service-oriented HPC infrastructure that combines the scalability and flexibility of cloud with the power of traditional supercomputers – a key state-of-the-art enabler for data-intensive science and engineering [12].

3 EDGE-TO-CLOUD-TO-HPC: THE DISTRIBUTED COMPUTING CONTINUUM

Computing now spans an edge-to-cloud-to-HPC continuum, where data and workloads flow from IoT sensors at the network edge to cloud data centers and onward to HPC facilities. In this hyperconnected continuum, billions of devices generate streams of data, initiating AI and data processing pipelines in situ at the edge [13]. Early data filtering and analytics occur on edge or cloud nodes, and the most demanding computations (e.g. training a 100-trillion-parameter AI model or running high-fidelity simulations) execute on large HPC systems that are often linked with cloud infrastructure conferences.computer.org.

Crucially, the historically disparate software environments across edge, cloud, and HPC are converging: workflows are increasingly packaged into containerized services that can be deployed uniformly, with Kubernetes and similar tools managing resources across all layers [14]. This cloud-native paradigm, coupled with high-speed networks, is eroding the technical barriers between local cluster environments and geographically distributed cloud/HPC resources. The availability of diverse hardware (from edge accelerators to GPU- or FPGA-enabled cloud instances and supercomputers) creates a heterogeneous playground for complex workflows. The European research agenda explicitly highlights the need to optimize AI and data applications across this continuum – for example, Horizon Europe calls for techniques to “optimise training times, model accuracy and data management” along the edge-cloud-HPC spectrum while considering metrics like memory usage, energy efficiency and latency [15]. This requires new virtualization and orchestration solutions that can seamlessly integrate heterogeneous architectures and ensure end-to-end performance transparency.

Globally, efforts are underway to develop such distributed computing ecosystems: the NSF-funded Chameleon testbed in the US, for instance, provides a configurable environment spanning edge devices to cloud nodes for large-scale experiments in continuum computing. As the edge-to-cloud-to-HPC continuum matures, it paves the way to novel real-time applications (e.g. AI-driven smart cities, autonomous vehicles, digital twins) while also introducing challenges in programming models, scheduling, and data management across decentralized infrastructure [16].

4 FEDERATED LEARNING IN HPC-CLOUD ENVIRONMENTS

Federated learning (FL) has emerged as a cutting-edge approach to AI that naturally leverages the distributed HPC-cloud continuum. In FL, multiple institutions or devices collaboratively train a machine learning model on their local data, without sharing the raw data – addressing privacy, security and data sovereignty concerns [17]. This approach is inherently complex: it requires federating heterogeneous infrastructures and coordinating iterative training rounds where model updates (not data) are exchanged.

HPC and cloud platforms are increasingly key enablers of federated learning at scale. For example, recent work demonstrated a hybrid FL workflow spanning supercomputers and cloud: two independent HPC clusters each trained a neural network on private datasets (e.g. medical or scientific data silos), while a cloud-based aggregator node periodically collected and averaged their model parameters [18]. This cross-site training pipeline – orchestrated with common workflow descriptions and tools – yielded a unified global model without any institution’s data ever leaving its local HPC center. Such experiments prove the feasibility of running FL across high-performance facilities, effectively creating a federated supercomputing network for AI [19].

They also highlight the need for advanced orchestration: to handle scheduling, placement, and fault tolerance across distributed AI training tasks. Recognizing these needs, the EU's research priorities call for "decentralised and federated computing continuum tools and mechanisms to enable distributed AI architectures," including orchestration techniques that span from edge devices to HPC, with built-in attention to data security and privacy [20].

Moving forward, federated learning stands to benefit from HPC acceleration (for heavy model computations) and cloud elasticity (for scaling coordinator services), combining the strengths of both. This will unlock new scenarios such as cross-hospital disease prediction models, multi-company industry 4.0 analytics, or climate models built from geographically dispersed data – all trained collaboratively without violating data governance rules. In sum, FL exemplifies the convergence of AI, distributed computing, and HPC-cloud infrastructure in the service of both innovation and compliance.

5 EXASCALE HPC

A new era is rising with the deployment and refinement of exascale supercomputers, capable of performing a quintillion calculations per second. Europe is underscoring its leadership in this domain with JUPITER [21], hosted at the Forschungszentrum Jülich in Germany and procured through the EuroHPC Joint Undertaking (EuroHPC JU) [22] – which recently entered public-operation status. JUPITER uses a modular architecture capable of delivering over 1 exaFLOP/s (i.e., more than 10^{18} double-precision calculations per second) and is designed with energy efficiency and European technological sovereignty in mind.

Beyond raw performance, JUPITER is engineered for real-world impact: it supports large-scale scientific simulations (such as climate modelling and astrophysics), AI-model training, and industrial applications in sectors like energy, materials science and health [23].

The arrival of exascale systems like JUPITER represents a significant leap: they are designed to deliver unprecedented processing power, memory bandwidth and connectivity, enabling new classes of application that were previously infeasible. With Europe now fielding such infrastructure, the integration of exascale HPC into the broader HPC-cloud ecosystem paves the way for workflows that span cloud, edge and supercomputing facilities, giving enterprises and researchers access to game-changing performance at scale.

6 HPC-CLOUD AND QUANTUM SYNERGY

The present technology of a Quantum Processing Unit (QPU) is limited by a low number of qubits (tens to a few hundred), which are prone to noise and thus the current quantum computer generation is characterised as Noisy Intermediate-Scale Quantum (NISQ). A qubit operating in a noisy environment loses its state easily, leading to computational errors. This generation of QPU is not suitable for real applications. Nevertheless, it is suitable to provide quantum computers to explore small-size applications, test new quantum algorithms, and develop techniques to mitigate the noise effects on qubits until error-free large-scale quantum computers (EFLS-QC) are achieved. The intermediate step towards EFLS-QC is the next generation of QPU, where fault-tolerant QPUs with a sufficient number of qubits will appear. The qubit devices comprising a QPU are called *physical qubits*, and there are several technologies to fabricate them, such as superconductors, semiconductors, spintronics, and trapped ions. An error-free qubit is called a *logical qubit*. At present, all physical qubits are prone to errors, having error rates between 0.1% and 1%.

Nevertheless, recently, a new approach has been demonstrated named *fault-tolerant QC*, where a significant number of physical qubits are combined in order to form a logical qubit having an error rate in the order of 10^{-6} (1 per million). For the sake of comparison, the error rate in a classical computer is 10^{-18} . Specifically, the information of a logical qubit is encoded across many physical qubits to make it more resilient against errors. IBM Quantum encodes 12 logical qubits into 288 physical qubits [24]. Logical qubits require many physical qubits and quantum gate operations to achieve and maintain their state. Translating gate operations on the abstract logical qubit into operations on the underlying physical qubits incurs costs, or resource overheads, of additional qubits and gates. Quantum error correction (QEC) aims to reduce these overheads in qubits and time to achieve error-free operation utilizing surface and topological codes. Current codes for gates with error rates of 0.1% still have high overheads (ca. many thousand iterations) to create a logical qubit. In the short term, approaches for quantum error mitigation (QEM) and may use QEC to lower, but not eliminate, errors as error rates fall. QEM aims to reduce the effective error rate of quantum calculations to support simple computations or extend the coherence of imperfect qubits for short algorithms in non-gate-based quantum approaches. Typically, this is done in two steps. First, by analyzing how noise affects the output of quantum programs and the formulation of a noise model. Second, classical computing and our noise model are applied to recover what a noise-free result would look like.

Table I. European Quantum Computers

QC name (Country)	Provider	No. Qubits / Technology	HPC integration
JADE (DE)	PASQAL	100 / neutral atoms	<u>JURECA DC</u> (23.5 PFlops)
Ruby (FR)	PASQAL	100 / neutral atoms	<u>Juliot-Curie</u> (22 PFLOPS)
Lucy (FR)	QUANDELA / ATTOCUDE	12 / photonic	<u>Juliot-Curie</u> (22 PFLOPS)
Piast-Q (PL)	AQT	20 / trapped ions	<u>ALTAIR</u> (5.9 PFLOPS)
VLQ (CZ)	IQM	24 / superconducting	<u>KAROLINA</u> (12.9 PFLOPS)
Euro-Q-Exa (DE)	IQM	54 / superconducting	<u>SuperMUC-NG</u> (26.9 PFLOPS)
EuroQCS-Italy (IT)	PASQAL	140 / neutral atoms	<u>LEONARDO</u> (315.74 PFLOPS)
EuroQCS-Spain (ES)	Qilimanjaro	10-25 / superconducting	<u>MARENOSTRUM 5</u> (314 PFLOPS)

A promising architecture with valuable outcomes is the combination of HPC with QC, the *hybrid computing*. The state-of-the-art in Hybrid HPC and Quantum Computers (QC) is defined by a shift from purely conceptual models to the physical deployment and software orchestration of integrated "*Quantum-Centric Supercomputing*" (QCSC) environments. With the present status of various QPU technologies are not suitable for real applications yet. Thus, qubit error mitigation through fault tolerant architectures is necessary making integration of HPC-QC more necessary than ever. For this reason, the European Quantum Computers are connected with HPCs as shown in the table I. Two HPC-QC integration approaches are adopted. In the first approach, the core mission is to treat the QPU as a specialised accelerator within the classical supercomputing workflow, much like a GPU. The second approach relies on the execution of algorithms by the QPU, and the HPC is used to run fault-tolerant procedures. Such a Quantum Centric Computing system requires a new design of the software stack as well as dedicated hardware improving communication speed and latency between HPC and QC [25], [26]. Very recently, Nvidia presented the NVQLink [27] chip that sets the foundation for uncovering the breakthroughs in control, calibration, quantum error correction, and hybrid application development needed to run useful quantum applications.

As coined above, European sovereignty in computing is realised through the EuroHPC JU consortium. Since 2021, the EU and the Consortium [28] discerned the potential of quantum

computing and at the time accepted that quantum computing cannot realise this potential without being integrated with HPC. Under the prism of superiority, all European HPC centres are connected with QC based on QPUs of different technologies, all made in Europe [29]. Such an initiative allows for the acquisition of significant knowledge on HPC-QC integration and the finding of solutions to many software and hardware issues. EuroHPC quantum computers are designed to complement Europe's existing supercomputing infrastructure, significantly enhancing capabilities in scientific research, innovation, and industrial applications. In the long term, the convergence of AI and quantum computing will be realised. This is mainly motivated and foreseen by the significant advantage of QC to improve the time and energy consumption compared to stand-alone HPC systems [30].

The Data Cloud-QC synergy is primarily boosted by potential real-world applications. The current technology is enabling early use cases and catalyzing research across several domains:

A. Scientific Research and Simulation in Materials Science & Drug Discovery:

Quantum simulations of molecular structures and interactions can outperform classical approaches for complex systems. Pharmaceutical firms reportedly use QaaS for pilot drug simulations, potentially accelerating R&D cycles.

B. Optimization Problems in Logistics & Supply Chains:

Optimization tasks that involve large combinatorial search spaces — such as route planning, inventory management, and production schedules — are prime candidates for quantum-assisted solutions. Hybrid quantum-classical methods are actively explored in research and pilot projects.

C. Data Analysis and Machine Learning:

Early research argues quantum accelerators could significantly speed up certain high-dimensional data analyses, such as feature mapping and simulation tasks — although mostly as experimental proofs of concept.

D. Security and Cryptography and Quantum-Resistant Security:

Organizations are concurrently adopting post-quantum cryptography to mitigate the risk that future quantum computers might break widely used encryption schemes (e.g., RSA). Some cloud providers have begun integrating these services into broader security offerings.

E. Finance and Risk Modeling – Portfolio Optimization:

Financial institutions are exploring quantum models for portfolio risk, derivatives pricing, and asset allocation. Though no live production implementations with clear performance gains are widely documented as of late 2025 by Google, research and pilot efforts continue.

6.1 CLOUD DATA TREATMENT BY QUANTUM COMPUTER

Cloud platforms now serve as the primary gateway for enterprises to run hybrid workloads, where classical cloud servers handle data preparation and the Quantum Processing Unit (QPU) executes complex computations such as optimization. However, critical issues and technical bottlenecks exist. The most important of them are summarized here after.

A. The Data Loading Problem (Input/Output)

Quantum computers cannot "read" classical cloud databases directly. Data must be encoded into quantum states (qubits), a process known as State Preparation. For large datasets, the time required to upload data into a QPU often exceeds the "coherence time" (stability) of the qubits. This makes "Big Data" analysis on quantum computers currently inefficient. The preparation of data by an HPC is required. For the management of hybrid workflows, seamless orchestration between classical processors and quantum accelerators remains an architectural and software challenge. Thus a significant software maturity is needed. Although, the Quantum programming frameworks (e.g., Qiskit, Cirq) are evolving, the ecosystem is less mature than classical cloud SDKs, with steep learning curves for developers.

B. "Harvest Now, Decrypt Later" (HNDL) Attacks

Cloud data protected by current RSA/ECC encryption is vulnerable to being intercepted and stored by malicious actors today, to be decrypted once fault-tolerant quantum computers arrive. Cloud-hosted quantum workloads raise unique privacy concerns because users' proprietary data may be transmitted to third-party platforms. Ensuring secure execution and compliance with data regulations (GDPR, HIPAA, etc.) requires new frameworks and encryption standards, prompting a parallel push for quantum-resistant cryptographic methods and secure cloud protocols. In 2025, the UN and NIST have urged organizations to migrate to Post-Quantum Cryptography (PQC) to protect long-life data.

C. Operational Noise (Decoherence)

In the NISQ era of quantum computers, maintaining the "Quantum State" requires special conditions e.g., near absolute zero temperatures. In a cloud environment, any latency or noise during the classical-to-quantum data handoff can cause decoherence, leading to computational errors.

D. Economic and Human Capital issues

Accessing quantum hardware through the cloud is significantly cheaper than owning hardware, but costs can still be high for sustained use—especially in exploratory, iterative experimentation. Skilled personnel with expertise in quantum mechanics, quantum algorithms, and hybrid architectures remain scarce, slowing enterprise adoption.

The above list of issues and challenges make clear the need of improvement of HPC-QC architectures and the adoption of fault-tolerant quantum computers, as described previously. Continued progress in error correction, qubit coherence, and scalable architectures will be pivotal to transitioning from research experimentation to commercial impact.

In July 2025, the Quantum European Strategy [31] was announced. Quantum computing and simulation are one of the main pillars the EU decided to support. Recognizing the strong relationship between HPC and QC technologies, the early stages implementation quantum strategy was amended in the EuroHPC JU regulation [32].

7 NOUS PROJECT INNOVATION AND EXPLOITATION STRATEGY

The NOUS project is a European initiative at the forefront of HPC-cloud convergence, aiming to empower Europe's data-driven future. NOUS focuses on envisioning and deploying advanced cloud services that integrate data spaces, high-performance computing, and edge computing [33].

In practice, this means developing state-of-the-art cloud technologies and making them widely accessible for scientists and professionals who need to run complex, compute-intensive algorithms on demand. A core objective is to enable efficient and trustworthy data sharing through the creation of secure, sovereign data infrastructures, aligning with Europe's values on data privacy and security. Rather than exclusively favoring domestic solutions or barring foreign cloud providers, NOUS adopts a forward-thinking strategy: it prioritizes scientific excellence and the growth of high-performance digital services within Europe.

This is significant given that an estimated 75% of the cloud market is currently controlled by non-European providers. By fostering home-grown alternatives and innovations, NOUS seeks to strengthen Europe's position in the global cloud and HPC arena without isolating it – a balanced approach that encourages open collaboration while addressing strategic autonomy.

NOUS's innovation is being validated through a diverse set of use cases that address specific industry and research challenges:

- Automotive: Improving the perception capabilities of connected and autonomous vehicles by processing data from roadside and in-vehicle cameras in real time.
- Energy: Managing the lifecycle of energy data for power grids and using predictive analytics to optimize electricity pricing in energy markets.
- Crisis Management: Enhancing data lifecycle management for emergency and disaster response, enabling faster and more informed decision-making during crises.
- Science: Handling and analyzing complex datasets from various scientific fields (e.g. genomics, climate science), by providing researchers with on-demand HPC/cloud resources coupled with shared data repositories.

By developing solutions in these real-world scenarios, NOUS drives innovation in an industrially relevant environment, ensuring that its technologies are practical and readily adoptable. The exploitation strategy of NOUS is tightly aligned with broader European digital initiatives. In particular, NOUS will leverage the emerging network of Common European Data Spaces – domain-specific data ecosystems being promoted by the EU – to foster cross-organization data sharing and analytics services.

By connecting to these Data Spaces (which emphasize sovereign data exchange via standards like Gaia-X), the project ensures that its cloud platform can interoperate in a trusted way with other infrastructure and adhere to EU regulations on data protection. The knowledge-sharing mechanisms and tools developed by NOUS (for example, frameworks for secure multi-party computation or federated data queries) will be disseminated to communities and could become blueprints for future services in Europe's digital marketplace.

Furthermore, NOUS benefits from synergies with the EuroHPC Joint Undertaking and Horizon Europe programs, positioning its outcomes to feed into Europe's HPC and cloud roadmap. By the end of the project, the goal is to not only demonstrate technical innovations, but also to deliver a model for sustainable exploitation – via commercial cloud platforms or public services – that can continuously support European enterprises and researchers in harnessing HPC-cloud capabilities.

Moreover, the migration of QCs toward integrated HPC will require many advances in micro- and macro-architecture that address concerns in the difference in infrastructure and performance. In addition, new software components should be developed. It is the scope of NOUS to demonstrate

interface between Data Cloud and QC to treat data for ML as well as to develop hybrid computing application.

In essence, NOUS serves as both a pioneer and an exemplar in Europe's strategy to achieve technological sovereignty and excellence in cloud-era high performance computing.

8 CONCLUSION AND ROADMAP

HPC-cloud integration has become a cornerstone of the modern digital infrastructure, and its importance will only grow in the coming years. The state-of-the-art review shows a clear trajectory: increasingly, HPC, cloud, and edge resources operate not as isolated silos but as parts of a unified computing continuum. The roadmap ahead involves deepening this convergence through technical, organizational, and policy innovations. On the technical front, continued adoption of open standards and interoperable frameworks is paramount. European efforts like Gaia-X are setting federated infrastructure requirements to ensure transparency, interoperability and data sovereignty across cloud and HPC services [34].

In parallel, complementary models such as the International Data Spaces architecture provide standardized connectors for sovereign data exchange directly between organizations, avoiding centralized data silos [35]. Together, these initiatives point toward a future in which a researcher or company can seamlessly deploy a workload that harnesses edge devices for data collection, cloud instances for elasticity, and HPC supernodes for heavy computation – all under a governance model that guarantees trust and compliance. Achieving this vision will require maturing federation services (identity and trust management, catalogues, compliance frameworks) to operate across diverse ecosystems, as well as strong interoperability consortia. Industry groups like the Open Cloud Consortium are already contributing by developing reference implementations and benchmarks to eliminate vendor lock-in and enable multi-cloud/HPC workflows [36].

From an infrastructure perspective, the roadmap entails strategic investments and partnerships. Europe's plan to build a network of world-class supercomputers under EuroHPC is well underway – with machines like LUMI and Leonardo already in the global top 10, and the first European exascale system (Jupiter) being constructed in Germany. These HPC systems are increasingly being designed with cloud interoperability in mind, featuring flexible access modes (such as interactive cloud-like access or elastic burst capacity) alongside traditional batch scheduling.

Moreover, Europe is pursuing technological sovereignty in hardware: the European Processor Initiative (EPI) and projects like EUPEX are developing energy-efficient processors and modular supercomputing architectures to power next-generation exascale machines. Incorporating these novel European components into HPC-cloud environments will be a key milestone in reducing dependency on foreign technology and ensuring long-term resilience [37]. Globally, the private sector will also play a significant role – major cloud providers are expected to further enhance their HPC offerings (e.g. dedicated HPC cloud regions, ultra-fast interconnects, HPC container services), and cross-border collaborations could establish federated clouds that share HPC resources for grand challenges (such as pandemic response or climate modeling).

In terms of emerging practices, we can expect to see more widespread use of hybrid and multi-cloud HPC workflows. Cloud bursting is likely to evolve into intelligent workload orchestration that automatically schedules parts of jobs to the most suitable resources (edge, cloud, or HPC) based on real-time availability, cost, and energy considerations. Data locality and federated data access will be addressed through the proliferation of data spaces and advanced data management tools (for example, high-speed federated filesystems and data lakes accessible across sites). Security

and sovereignty will remain at the forefront: initiatives like the European Alliance for Industrial Data and Cloud are poised to increase the share of trustworthy EU-based cloud/HPC providers and establish common standards for security and portability.

Additionally, new paradigms like quantum computing might be gradually integrated into the cloud-HPC mix, offered via cloud interfaces side by side with classical HPC – early examples of this are already visible as cloud platforms provide access to prototype quantum processors for experimentation [38]. From the above analysis, it becomes clear that the Cloud-QC roadmap is strongly related to the roadmap of the QC providers (IBM, Google etc), which comprises advancements in hardware and software. In this context below is the presented consolidated roadmap based on the latest 2025 industry projections and provider-specific goals.

Phase 1: Current status

In this phase, the goal is not to beat classical computers in speed, but to prove that quantum systems can run complex circuits more accurately than classical simulations.

- Infrastructure: access to "noisy" (nisq) devices with ca. 150 qubits.
- Cloud integration: launch of quantum+hpc (high-performance computing) toolsets
- Milestone: IBM *Heron* processor, Google *Willow* [] and Quantinuum *Helios* system demonstrating 99.9% gate fidelity, allowing for "useful" but small-scale scientific simulations.

Phase 2: Quantum advantage & error correction (2026–2028)

This is the era where quantum computers begin to solve specific commercial problems cheaper or more efficiently than any classical alternative.

- Hardware: Shift toward logical qubits.
- Data readiness: Organizations are expected to move from "experimentation" to "production-like" environments using *quantum-ready* pipelines (*DevOps.com*).
- Key use cases: Early advantages in quantum machine learning, molecular discovery for pharmaceuticals, and complex option pricing in finance.

Phase 3: Fault-tolerant scale (2029–2030 and beyond)

The long-term goal is a quantum computer capable of running millions of gates without errors.

- The Quantum Internet: IBM and Cisco announced (nov 2025) a roadmap to entangle multiple quantum computers across data centers by 2030. In Europe, Quantum Industry Association (QIA) has a proven track record in delivering a pioneering quantum internet technology and will continue to push all key technologies towards building a European quantum internet. As a European platform, the ambition now is to expand our scientific leadership into European leadership in quantum internet innovation. We do this by advancing the European quantum internet ecosystem connecting world-class research and cutting-edge technology development to hardware and software industry all along the value chain. This allows "distributed quantum computing," where data can be processed across a network of QPUs.
- Infrastructure: Systems scaling to >100.000 physical qubits and ca. 2.000 logical qubits.

- Security: Full migration of cloud data to *post-quantum cryptography (PQC)* standards to protect against "harvest now, decrypt later" threats.

In conclusion, the convergence of HPC and cloud computing and quantum computing is fostering an environment where computational power is more distributed, flexible, and ubiquitous. The innovation roadmap calls for nurturing this convergence through continued research (in areas like distributed orchestration, programming models for the continuum, and federated learning), strategic infrastructure projects, and supportive policies. By following this roadmap – investing in exascale systems and sustainable data centers, developing common frameworks like Gaia-X for federated operation, and promoting collaboration between public supercomputing facilities and cloud innovators – Europe and the global community can ensure that the next generation of scientific discovery and industrial progress is powered by an open, powerful, and secure HPC-cloud ecosystem. The journey toward a truly integrated Edge-to-Cloud-to-HPC continuum is underway, and its realization will underpin advancements in AI, personalized medicine, climate modeling, and beyond, translating technological convergence into societal impact.

REFERENCES

- [1] [Container orchestration on HPC systems through Kubernetes](#)
- [2] [Leveraging High-Performance Computing and Cloud Computing](#)
- [3] ['Cloud bursting'](#)
- [4] [PRACE eu](#)
- [5] <https://lexis-project.eu/web/>
- [6] [Technologies and Applications for Big Data Value](#)
- [7] [Cloud Computing and High Performance Computing \(HPC\) Advances](#)
- [8] [Amazon EC2 Hpc7g Instances](#)
- [9] [Google Cloud - HPC](#)
- [10] [HPC cloud provider Qarnot gets funding from European Innovation Council](#)
- [11] [NOUS Project](#)
- [12] [UberCloud Marketplace for HPC as a Service](#)
- [13] [Building a Cloud-Edge Computing Continuum to Address AI Processing Demands](#)
- [14] [HPC Alongside User-space Kubernetes](#)
- [15] [Empowering AI/generative AI along the Cognitive Computing continuum](#)
- [16] [Welcome to Chameleon](#)
- [17] [Federated Learning: A Privacy-Preserving Approach](#)
- [18] [Federated Learning meets HPC and Cloud](#)
- [19] https://link.springer.com/chapter/10.1007/978-3-031-34167-0_39
- [20] [Research & Innovation roadmap for the cognitive continuum](#)
- [21] <https://www.fz-juelich.de/en/jupiter>
- [22] https://www.eurohpc-ju.europa.eu/index_en
- [23] <https://blogs.nvidia.com/blog/jupiter-exascale-supercomputer-live/>
- [24] <https://www.ibm.com/quantum/blog/large-scale-ftqc>
- [25] K. Rallis *et al.*, "Hardware-level Interfaces for Hybrid Quantum-Classical Computing Systems," *2025 6th International Conference in Electronic Engineering & Information Technology (EEITE)*, Chania, Greece, 2025, pp. 1-6, doi: 10.1109/EEITE65381.2025.11166221
- [26] Thomas Beck, *et al.*, "Integrating quantum computing resources into scientific HPC ecosystems", *Future Generation Computer Systems* 161 (2024) 11-25, <https://doi.org/10.1016/j.future.2024.06.058>
- [27] <https://nvidianews.nvidia.com/news/nvidia-nvqlink-quantum-gpu-computing>
- [28] M. P. Johansson, E. Krishnasamy, N. Meyer, and C. Piechurski, "Quantum computing – a European perspective", 2021. PRACE paper, <https://zenodo.org/record/5547408#.YfO0UmAxnkI>
- [29] https://www.eurohpc-ju.europa.eu/eurohpc-quantum-computers/our-quantum-computers_en
- [30] Vanessa García Pineda, *et al.*, "Integrating artificial intelligence and quantum computing: A systematic literature review of features and applications", *International Journal of Cognitive Computing in Engineering*, 7 (2026), pp.26-39; <https://doi.org/10.1016/j.ijcce.2025.08.002>
- [31] <https://digital-strategy.ec.europa.eu/en/library/quantum-europe-strategy>
- [32] https://www.eurohpc-ju.europa.eu/european-commission-proposes-amendment-eurohpc-regulation-support-gigafactories-and-include-quantum-2025-07-16_en
- [33] [nous-project.eu](#)
- [34] [GAI-X](#)
- [35] [Communications of the ACM](#)
- [36] [Rolling plan for ICT standardisation](#)
- [37] [Europe's exascale supercomputer in its starting blocks](#)
- [38] [Unlock Quantum Power from Your Browser](#)
- [39] [Our Quantum Echoes algorithm is a big step toward real-world applications for quantum computing](#)