Author
Reji Kumar Pillai
President, India Smart Grid Forum
Chairman, Global Smart Energy Federation

Co-Author
Rohan Banik
Research Associate
India Smart Grid Forum

"Quantum to Cloud: The New Compute Stack"

For most of the past decade, the headline questions in computing were about single breakthroughs, a faster chip here or a larger cloud there. That framing is fading - the frontier now sits in how different kinds of compute fit together. In 2026, the computing landscape has entered the "Quantum to Cloud" era, representing a massive shift from traditional IT architectures to a unified, heterogeneous compute stack. Enterprises have moved past the theoretical concept of standalone quantum supremacy and are now focusing on practical hybrid performance. In this modern environment, quantum processing units, artificial intelligence accelerators, high-performance computing clusters, and global cloud infrastructure operate together seamlessly. This convergence allows organisations to tackle computationally intensive bottlenecks by utilising classical machines for orchestration and data preparation while leveraging quantum processors for targeted, highly complex calculations.

Cloud platforms play a foundational role in this transformation by democratising access through quantum-as-a-service model. By abstracting the steep costs and extreme cryogenic complexities of on-premises quantum hardware, hyperscalers and specialised cloud platforms enable businesses to run their quantum computing on a flexible, pay-as-you-go basis. The scale of surrounding investment shows how seriously the market takes the shift. Gartner expects worldwide IT spending to reach 6.31 trillion dollars in 2026, up 13.5% on the year, with data centre systems growing fastest as organisations race to add capacity.[1] Spending on artificial intelligence alone is forecast at 2.52 trillion dollars for 2026.[2] Quantum is the youngest layer in this picture, yet the one drawing the most striking curve. Private investment in quantum technology start-ups reached 12.6 billion dollars in 2025, more than six times the previous year, and 90% of it flowed to quantum computing.[3] As the global market projects exponential growth and investments, the quantum-to-cloud ecosystem is fundamentally redefining the future of digital transformation.

---

[1] https://www.gartner.com/en/newsroom/press-releases/2026-04-22-gartner-forecasts-worldwide-it-spending-to-grow-13-point-5-percent-in-2026-totaling-6-point-31-trillion-dollars

[2] https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026

[3] https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/mckinsey-quantum-technology-monitor-2026-a-commercial-tipping-point

The new compute stack Architecture

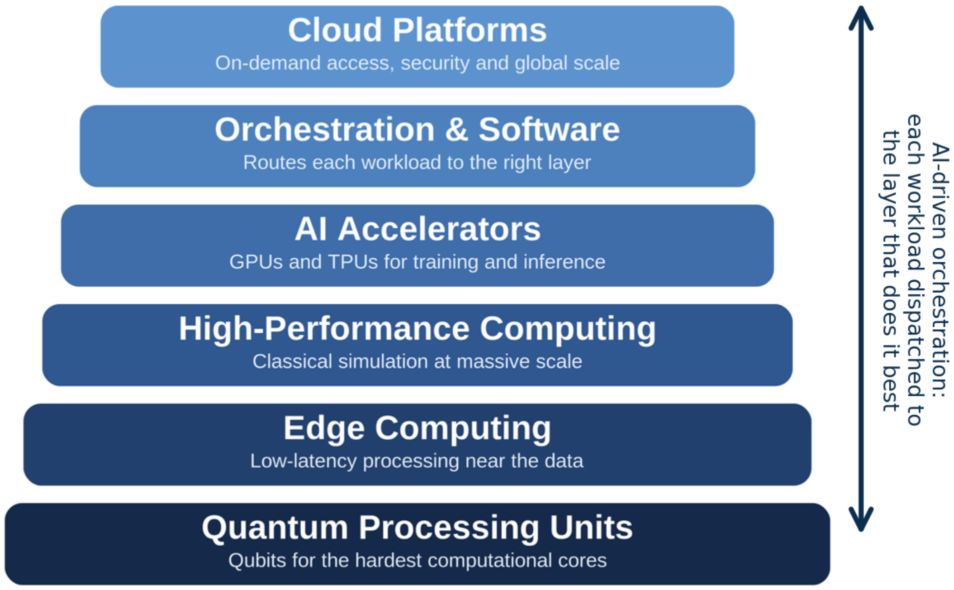

The stack is a set of cooperating layers rather than a sequential order. At the top sits the cloud platform, the layer providing on-demand access, billing and security. The orchestration software and compilers decide what runs where and translate a problem into instructions each machine can execute. The AI accelerators are the workhorses handling learning and inference, while high-performance computing clusters carry large-scale simulation. Closer to the physical world sit edge devices, capturing and processing data where it is generated. At the foundation are the quantum processing units, built on competing designs; these layers interact with one another. A real workload is rarely purely quantum or purely classical. It is usually both, with a scheduler sending each part of a problem to whichever layer solves it best.

Convergence of quantum, cloud and AI are reshaping enterprise computing

This convergence shifts enterprise computing from pure classical systems to a highly integrated hybrid model. AI acts as the orchestration layer, using specialised frameworks and high-speed interconnects to process error syndrome data in real time and tune quantum hardware autonomously. Cloud computing provides the scalable infrastructure to manage these workflows, allowing organisations to automatically route computational tasks to the most appropriate processor (whether that is a CPU, GPU, or QPU) without needing to own or maintain the physical quantum hardware themselves. IBM's new "quantum-centric supercomputing" blueprint envisions QPUs operating alongside GPUs and CPUs on-premises and in the cloud to solve challenges no single approach can tackle[4]. Likewise, AMD argues that "the future of quantum computing will not be built on quantum processors alone; its foundation

[4] https://newsroom.ibm.com/2026-03-12-ibm-releases-a-new-blueprint-for-quantum-centric-supercomputing

will be the convergence of quantum computing, high-performance computing and artificial intelligence[5].

Challenges: Adopting a hybrid "quantum-to-cloud" compute stack

The first challenge is **technical**; it includes high hardware error rates, the extreme fragility of physical qubits, and the need for complex, ultra-low-latency error correction protocols. For instance, an interconnect called NVQLink aims to link quantum processors and GPUs with latency under four microseconds and has already been adopted by seventeen quantum hardware builders and nine United States national laboratories[6].

The second challenge is **people**. Survey work in the State of Quantum 2026 report found that although 89 percent of organisations now claim hands-on quantum activity, only 3 percent have reached deployment at scale, and skills shortages rank as the single biggest barrier.[7] Building genuine quantum expertise typically takes two to five years, so the constraint cannot be solved quickly by hiring. Deployment models are unsettled too. In the same survey, 46 percent of organisations expected on-premises or hybrid quantum infrastructure to feature in their access strategy within three years, a sign that firms want control and integration rather than only a remote endpoint.[8]

The third challenge is **security**. Once quantum machines can break today's public-key cryptography, data stolen now could be decrypted later, a threat already prompting migration to new standards. NIST finalised its first post-quantum cryptography standards in 2024 and expects the most common current algorithms to be phased out by 2030 and disallowed after 2035.[9]

For most enterprises the real challenge is less about buying a quantum computer than about integrating, staffing and securing a stack that now carries a quantum edge.

Accelerating commercialization and accessibility of quantum computing

A quantum computer is costly to build and demands cryogenics, shielding and specialist care, so owning one makes little sense for most organisations, eliminating the multi-million-dollar capital expenditure; cloud services remove that barrier. Amazon Braket[10] offers pay-as-you-go access to superconducting, trapped-ion and neutral-atom machines from several vendors, while Microsoft Azure Quantum and IBM Quantum provide their own hosted routes. This quantum-as-a-service model is growing quickly, rising from an estimated 4.48 billion dollars in 2025 to 26.18 billion dollars in 2030[11].

Through service-based models, they provide vendor-agnostic access to diverse quantum architectures. They lower the technical barrier to entry by integrating generative AI and classical simulators, allowing developers to test, refine and re-run simulations and algorithms across similar challenges before running them on expensive, live quantum processors. For instance, McKinsey notes that private

---

5 https://www.amd.com/en/blogs/2026/amd-advances-the-hybrid-future-of-quantum-computing.html
6 https://nvidianews.nvidia.com/news/nvidia-nvqlink-quantum-gpu-computing
7 https://iqm.tech/reports/IQM-State-of-Quantum-2026.pdf
8 https://iqm.tech/reports/IQM-State-of-Quantum-2026.pdf
9 https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards
10 https://aws.amazon.com/braket/
11 https://www.thebusinessresearchcompany.com/report/quantum-computing-as-a-service-qcaas-global-market-report

companies increasingly reach quantum through cloud providers and that hardware makers such as IBM, IonQ and Quantinuum now sell access through platforms rather than only through bespoke partnerships.[12] The supplier market is maturing alongside that access. By packaging hardware, compilers and orchestration into one repeatable offering, the cloud shortens the path from experiment to production and lets enterprises test ideas without a capital outlay.

Which industries benefit the most

Not every industry gains equally, and the early winners share a trait - they all face challenges whose difficulty explodes as they grow. The electricity grid is the largest machine humanity operates, and we are rebuilding it while it runs. Millions of rooftop solar plants, batteries and electric vehicles are turning a one-way network into a two-way marketplace. Deciding which generators to run, how to reroute power after a fault, when to charge a million electric vehicles – the number of possible combinations in these problems grows faster than any classical computer can follow. Even a one-percent improvement in dispatch efficiency is worth billions of dollars and millions of tonnes of carbon. For instance, Oak Ridge National Laboratory is working with IonQ on power-grid optimisation advancements using hybrid quantum approaches.

Generation scheduling
Unit commitment across
thousands of plants

Network reconfiguration
Faster restoration
after faults

P2P energy trading
Clearing decentralized
markets & microgrids

EV charging orchestration
Coordinating millions
of vehicles

Renewables forecasting
Quantum ML for wind
& solar output

Quantum-safe security
PQC migration & QKD
for critical infrastructure

The pharmaceutical and chemical industries benefit immensely through molecular simulations that cut drug development cycles and identify novel materials without relying entirely on physical lab trials. The financial services sector uses this compute stack for complex portfolio optimisation and high-dimensional risk modeling that traditional machine learning frequently misses. For instance, HSBC worked with IBM; that improved its prediction of whether a corporate bond trade would fill by up to 34 percent against classical techniques, using real market data at production scale[13].

---

[12] https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/mckinsey-quantum-technology-monitor-2026-a-commercial-tipping-point

[13] https://www.hsbc.com/news-and-views/news/media-releases/2025/hsbc-demonstrates-worlds-first-known-quantum-enabled-algorithmic-trading-with-ibm

Way Forward

The next generation will be defined by rapid advancements in quantum error correction and the virtualisation of physical qubits to create highly stable logical qubits. Google's Willow was the first chip to push below the error-correction threshold, the point at which adding qubits starts to reduce errors rather than add them.[14] IBM plans to prototype a real-time decoder in 2026 on the road to its fault-tolerant Starling system, targeted for 2029 with around 200 logical qubits.[15] Secure quantum networking and post-quantum cryptography will also become foundational, driven by innovations like satellite-based quantum key distribution and high-speed terrestrial communication networks. Finally, AI-driven infrastructure software will autonomously calibrate quantum hardware, drastically reducing classical communication overhead and eliminating the need for manual tuning by human physicists. The through-line is integration. The winners of the Quantum-to-Cloud era will be the organisations that treat computing as one adaptable fabric and learn to route each problem to the layer that solves it best.

[14] https://blog.google/innovation-and-ai/technology/research/quantum-echoes-willow-verifiable-quantum-advantage/
[15] https://www.ibm.com/quantum/blog/large-scale-ftqc