TECHNOLOGY REPORT

Why the Future of Computing Is Hybrid

Classical and quantum together: computing, networks, and cryptography in the age of convergence

An analysis of hybrid computing, quantum algorithms, error correction, quantum networking, and post-quantum cryptography

August 2026

Table of Contents

Why the Future of Computing Is Hybrid: Classical and Quantum Together

For decades, quantum computing has been portrayed as a technology destined to *replace* classical computing: machines capable of solving in seconds problems that would take today's most powerful supercomputers centuries. The reality emerging from laboratories and data centers around the world is quite different — and, in some ways, more interesting: the future will not be "quantum instead of classical," but **quantum together with classical**, in increasingly intertwined architectures.

Quantum processors (QPUs) excel at a narrow but valuable set of problems — combinatorial optimization, molecular simulation, certain cryptography and machine learning tasks — thanks to phenomena such as superposition and entanglement. But they remain entirely unsuited to tasks such as data entry, memory management, graphics rendering, or training large language models, where classical CPUs and GPUs remain irreplaceable. This complementarity is what creates the need for a hybrid infrastructure.

> *The future will not be quantum instead of classical, but quantum together with classical: not the replacement of one paradigm by another, but an ever-tighter integration between the two worlds.*

The central issue, however, is not *whether* classical and quantum should collaborate, but *how*. And this is where a crucial distinction emerges between two very different models of cooperation: the **sequential** architecture (or "separate islands") and the **tightly-coupled hybrid** architecture, in which the two types of computation intertwine in real time.

The Sequential Model: Two Computers Passing the Ball

In the simplest model — and historically the first to be implemented in quantum cloud computing services — a classical computer prepares the parameters of a problem, sends them to a quantum computer over a network or cloud service, waits for the QPU to perform the calculation, and receives a result back. It is a linear flow, almost a "request-response" pattern, triggered only for that narrow class of problems where the quantum advantage is clear and demonstrable.

Advantages of this approach

- **Architectural simplicity.** There is no need for a dedicated, ultra-low-latency interconnect between CPU and QPU: an API and a network connection are enough, exactly as with today's public cloud quantum computing services (IBM Quantum, Azure Quantum, Amazon Braket).
- **Isolation and reliability.** If the quantum computer is unavailable or returns an error, the classical system keeps running on its own; the two domains are decoupled and easier to test, update, and secure separately.
- **Accessibility.** Companies can experiment with quantum computing without owning quantum hardware themselves, renting compute time on remote QPUs and integrating it into existing classical applications.
- **Well suited to "closed" problems.** For well-defined tasks with clean inputs and outputs (e.g., a one-off optimization problem), the sequential model is more than sufficient and does not introduce unnecessary complexity.

Disadvantages and limitations

- **Latency and communication overhead.** Every call to the QPU incurs queuing, data transfer, and synchronization time that can far exceed the actual computation time, especially on shared cloud services.
- **Inadequacy for variational algorithms.** Many of the most promising quantum algorithms in the near-to-medium term — such as VQE (Variational Quantum Eigensolver) or QAOA (Quantum Approximate Optimization Algorithm) — require **thousands of fast iterations** between classical and quantum systems to optimize parameters, not a single data exchange. The sequential model, with its inherent latency, makes these cycles enormously slower and more expensive.
- **No real-time error mitigation.** Today's quantum hardware (the so-called NISQ era, *Noisy Intermediate-Scale Quantum*) is noisy and subject to decoherence. Without a tight, fast dialogue with the classical side, it becomes difficult to correct errors while the computation is in progress.
- **Underutilization of the QPU.** Quantum resources, still scarce and expensive, risk sitting idle during preparation and data-transfer phases, reducing return on investment.

The Tightly-Coupled Hybrid Model: Quantum and Classical "Breathing Together"

In the tightly-coupled hybrid architecture, the QPU, CPU, and GPU don't just exchange a final result: they **collaborate in continuous, closely spaced cycles**, often within the same data center or even the same rack, with ultra-high-speed interconnects (such as NVIDIA's NVQLink or IBM's quantum-centric supercomputing architectures). A concrete example: in the SQD (Sample-based Quantum Diagonalization) technique developed by IBM together with AMD and NVIDIA, the QPU runs the quantum circuits while GPUs process the heaviest tensor operations in parallel — a process that has achieved a speedup of roughly 100 times over classical-only methods, with further performance gains of 1.8 to 3 times from the addition of the latest-generation AMD and NVIDIA GPUs.

Advantages of this approach

- **Makes the most of NISQ hardware.** The continuous iterative cycle between the QPU and classical systems enables real-time error mitigation and correction techniques, essential until fully fault-tolerant QPUs exist.
- **Enables new classes of algorithms.** Hybrid variational algorithms become truly practical when the latency across thousands of classical-quantum iterations is minimized, paving the way for applications in computational chemistry, drug discovery, materials science, and large-scale logistics optimization.
- **Maximizes the use of quantum resources.** By integrating the QPU as a specialized accelerator within an HPC cluster (exactly as happens with GPUs today), idle time is reduced and overall system throughput is increased.
- **Lays the groundwork for fault-tolerant computing.** Even once quantum computers become more reliable, real-time error correction will continue to require enormous classical compute power: the tightly-coupled hybrid infrastructure built today will be the foundation for tomorrow's.

Disadvantages and challenges

- **Engineering complexity.** Synchronizing hardware with radically different physical characteristics (cryogenic temperatures for superconducting QPUs, standard operating conditions for CPUs/GPUs) requires extremely sophisticated interconnection, orchestration, and software solutions.

- **High infrastructure costs.** Ultra-low-latency networks, dedicated hardware, and often the physical co-location of QPUs and classical systems are required: a significantly higher investment than simple sequential cloud access.
- **Scarcity of skills.** Designing software for these systems requires professionals who master both quantum computing and classical HPC — a professional profile still rare in the market.
- **Immature ecosystem and standards.** Frameworks such as NVIDIA's CUDA-Q or IBM's quantum-centric architectures are still converging toward common standards; the risk of lock-in to a single hardware or software vendor is real.
- **Limited testbeds.** Research itself notes that, despite solid theoretical foundations, large-scale experimental infrastructure and empirical datasets demonstrating benefits in production — not just in the lab — are still lacking.

Feature	Sequential Model	Tightly-Coupled Hybrid Model
Latency between classical and QPU	High (cloud queue, data transfer)	Minimal (co-location, dedicated interconnects)
Suited to	Isolated problems, a single data exchange	Iterative algorithms (VQE, QAOA), real-time QEC
Infrastructure cost	Low, public cloud access	High, dedicated and co-located hardware
Engineering complexity	Low	High (synchronization, orchestration)
Ecosystem maturity	Established (IBM Quantum, Azure, Braket)	Rapidly evolving (CUDA-Q, NVQLink)

Table 1. Summary comparison between the sequential architecture and the tightly-coupled hybrid architecture.

Hybrid Algorithms in Action: VQE and QAOA

If there is concrete proof that the tightly-coupled hybrid architecture is not an engineering whim but a necessity, it lies in the algorithms that today represent the beating heart of quantum computing in the NISQ era: **VQE** (Variational Quantum Eigensolver) and **QAOA** (Quantum Approximate Optimization Algorithm).

Both share the same underlying architecture, known as the "variational loop": a parameterized quantum circuit is run on the QPU, the measurement result is sent to a classical optimizer (algorithms such as COBYLA, SPSA, or gradient-based methods), which updates the circuit's parameters to move closer to an objective — minimizing an energy, maximizing a cost function — and the cycle repeats. Not once, but **hundreds or thousands of times**, until convergence.

- **VQE** is designed mainly for computational chemistry: it estimates the ground-state energy of a molecule, a calculation crucial to drug discovery, catalyst design, and the development of new materials (for example, more efficient batteries or solar cells).
- **QAOA**, on the other hand, targets combinatorial optimization problems — logistics routing, scheduling, financial portfolio allocation, max-cut-type problems — by alternating layers of quantum operators whose intensity is governed by classically optimized parameters.

The reason these algorithms *require* a tightly-coupled hybrid architecture is clear: today's QPUs are noisy and can only run relatively shallow circuits before decoherence destroys the quantum information. VQE and QAOA were designed precisely to work around this limit, using shallow circuits and offloading the "heavy" part of the optimization to classical computing. But this continuous exchange only works if the latency between one iteration and the next is minimal: with a cloud-based sequential model, every single iteration would incur queuing and data-transfer delays, multiplying a lag that would quickly become unsustainable across thousands of iterations.

Open limitations remain, however: the "barren plateaus" phenomenon (gradients that vanish as the number of qubits grows, making optimization increasingly difficult), sensitivity to hardware noise, and the absence, to date, of an unambiguous demonstration of a clear quantum advantage over the best classical heuristics for most real-world use cases. VQE and QAOA therefore remain fundamental testbeds rather than solutions already mature for industrial scale.

The Crucial Role of Quantum Error Correction

If variational algorithms show why classical and quantum systems must talk to each other *during* the computation, quantum error correction (QEC) shows why they must do so at an almost unimaginable speed. We have already covered this topic in our Deep Dive of May 29, 2026; we revisit it here as a practical example of hybrid integration.

Physical qubits are extremely fragile: a single interaction with the surrounding environment, a vibration, an unwanted photon, is enough to introduce an error that quickly propagates through the computation (decoherence). The conceptual solution, known since the 1990s, is to encode a single, stable and reliable **logical qubit** by spreading the information across many redundant **physical qubits**, using schemes such as the *surface code* or the more recent *qLDPC* codes (quantum Low-Density Parity-Check). The system continuously measures "error syndromes" without directly touching the encoded quantum state (so as not to destroy it), and these syndromes must be interpreted and corrected.

This is where the classical side becomes indispensable: **decoding** — translating the measured syndromes into concrete correction instructions — is a computation that must happen in real time, within the qubits' coherence time, on the order of **microseconds**. A decoder that is too slow means errors accumulate faster than they are corrected, defeating the entire purpose of QEC. This is why dedicated hardware decoders have emerged (often based on FPGAs or specialized chips), physically located close to the cryostat housing the superconducting qubits, capable of decoding in under a microsecond.

Here lies an interesting paradox: the larger and more reliable quantum computers become, **the more — not less — classical computing power will be needed**. As the number of physical qubits required per logical qubit grows, and as systems scale toward fault tolerance, the volume of syndrome data to be decoded in real time grows enormously. It is no coincidence that IBM's roadmaps toward fault-tolerant systems by the end of the decade envision embedding GPUs and classical computing directly within the cryogenic infrastructure of quantum systems, precisely to handle real-time error correction at scale. QEC, in other words, may be the most extreme — and most necessary — example of tightly-coupled hybrid cooperation between the two worlds.

Sequential or Tightly-Coupled Hybrid? A Matter of Problem, Not Dogma

Neither model is "superior" in absolute terms: the choice depends on the type of problem being solved. For isolated tasks with a demonstrated quantum advantage and a single data exchange (for example, factoring a number or a single, well-defined optimization problem), the sequential model remains efficient, economical, and more than adequate. Conversely, whenever a problem requires iteration, progressive learning, or continuous error correction — that is, for the vast majority of practical quantum computing uses in the NISQ era — the tightly-coupled hybrid architecture is set to become the standard. It is no coincidence that industry observers expect hybrid quantum-classical computing, supported by middleware platforms such as CUDA-Q, to become the standard implementation mode in data centers by 2026, with QPUs increasingly integrated as accelerators within high-performance computing clusters.

Who Is Building This Future

The shift toward tightly-coupled hybrid computing is no longer purely academic theory: it is already a strategic priority for the leading players in the technology industry.

- **IBM** is betting on a *quantum-centric supercomputing* architecture, demonstrating together with its partners concrete results such as the already-mentioned 100x speedup achieved by integrating QPUs and GPUs at the Oak Ridge National Laboratory's Frontier supercomputer and RIKEN's Miyabi supercluster in Japan.
- **AMD** has formed a partnership with IBM focused on quantum-centric supercomputing architectures that combine quantum processors, HPC, and AI resources, providing EPYC CPUs, Instinct GPUs, and Versal adaptive SoCs as the classical infrastructure for quantum systems. AMD is also working with organizations such as JPMorganChase and Oak Ridge National Laboratory to explore the integration of quantum systems, artificial intelligence, and high-performance computing.
- **NVIDIA** is emerging as a central hub of this convergence thanks to its CUDA-Q software platform and the NVQLink interconnect, designed to link QPUs and GPUs in real time. **Google Quantum AI** itself uses NVIDIA infrastructure to simulate 40-qubit quantum systems using more than 1,000 H100 GPUs, while the startup **Infleqtion** has announced the integration of its Sqale quantum computers with NVIDIA GPU systems via NVQLink at the Illinois Quantum & Microelectronics Park, to enable a unified, real-time hybrid computing architecture.
- **Microsoft**, with its Azure Quantum platform, continues to position itself as an integration layer between heterogeneous quantum hardware and existing classical cloud workflows, betting in the long run on topological fault-tolerant architectures.
- On the specific front of real-time error correction, the UK's **Riverlane** is today one of the most frequently cited specialists in the field: with its Deltaflow stack and hardware decoders capable of decoding errors in under a microsecond, the company works with several quantum hardware makers to reach the "MegaQuOp" scale (one million error-free quantum operations). Alongside it operate **Quantum Machines** (the OPX1000 control platform) and **Q-CTRL** (the Fire Opal mitigation software), while Google and IBM continue publishing their progress on surface codes and qLDPC codes.

What About the Quantum Internet? Hybrid, Separate-but-Cooperative, or Destined to Prevail?

The same question that applies to computing — replacement, tight integration, or coexistence as separate islands — arises in a mirrored form for networks: will the **quantum internet** replace the classical one, remain

a parallel infrastructure that only cooperates from a distance, or will the two eventually merge into a single hybrid network?

Before answering, it's worth clarifying what "quantum internet" means: not a replacement for the network that carries video, email, and web pages, but an infrastructure designed for tasks the quantum world does better than anything else — distributing **entanglement** between distant nodes, transmitting intrinsically secure cryptographic keys (QKD, Quantum Key Distribution), and, prospectively, linking multiple quantum computers together to form distributed clusters more powerful than a single QPU. It is not meant to replace everyday data traffic, which will remain classical.

The total-replacement scenario is the least likely. Quantum networks do not carry data "in the clear" in the classical sense: they transmit extremely fragile quantum states, subject to decoherence, that cannot be amplified with ordinary optical repeaters (the "no-cloning" theorem prevents copying an unknown quantum state). Building a global quantum network that replaces existing optical fiber would require quantum repeaters, quantum memories, and infrastructure that is still immature; no serious industry observer expects the classical internet to be supplanted.

The most realistic scenario today is a hybrid/cooperative one, built on a single shared physical infrastructure. The latest research shows that quantum and classical signals can coexist on the same existing optical fiber by exploiting different wavelength bands (WDM multiplexing techniques): one experiment demonstrated that quantum channels more than 40 nanometers apart in the O-band can co-propagate alongside high-power classical signals in the C-band while preserving the fidelity of the quantum signal. An even tighter approach is "classical-decisive": a classical signal guides the real-time routing of quantum entanglement across the network, preserving its integrity — a hybrid protocol that enables dynamic, high-fidelity entanglement routing over existing fiber networks, opening a concrete path toward a scalable quantum internet built on top of today's infrastructure, rather than alongside it.

In the medium term, however, a "separate but cooperative" model still prevails. The metropolitan quantum networks actually in operation — in Chattanooga, New York, Beijing-Shanghai — today function as dedicated, distinct infrastructures, designed initially to serve specific use cases (government security, the financial sector, research) rather than mass internet traffic. Cooperation with the classical network happens at the orchestration level — routing, management, software interfaces — rather than through physical sharing of the transmission medium. It is the same pattern already seen in computing: first separate islands that cooperate on request, then, as the technology matures, an increasingly tight integration.

> *The quantum internet will not replace the classical one, but will increasingly overlap with it: first as a separate network that communicates only for orchestration and security, then as a truly hybrid infrastructure on the same fiber.*

Who is working on the quantum internet

- **Aliro Quantum** is today the leading software specialist for orchestrating quantum networks: its *Entanglement-as-a-Service* platform and the Aliro QN router allow enterprise IT departments to manage a quantum network the way they would a classical software-defined network (SDN). A partnership announced with **Cisco** in February 2026 signals that classical networking incumbents are also taking this orchestration layer seriously.
- **Photonic Inc.**, after a quantum teleportation demonstration over 30 km of metropolitan fiber carried out with TELUS, is now shifting toward more commercial execution, working closely with Microsoft on integration with Azure Quantum.

- **Xanadu**, a Canadian maker of photonic quantum processors, has signed an agreement with **TELUS** to build Canada's first sovereign hybrid quantum-classical infrastructure, integrating its processors directly into the national PureFibre network.
- **China Telecom**, in collaboration with the University of Science and Technology of China, operates the Beijing-Shanghai quantum communication backbone, a roughly 2,000 km network, and has launched the Micius satellite for intercontinental quantum key distribution (QKD).
- **EPB** (the electric utility of Chattanooga, Tennessee) operates the first commercial quantum network in the United States and, together with **IonQ**, is building a hub that will unite quantum computing and networking in the same facility.
- **Qunnect** (the GothamQ network in New York), **ID Quantique**, **Qubitekk** (now part of IonQ, listed on the NYSE), and **Arqit** (listed on the Nasdaq) round out the picture of leading commercial players, while government initiatives such as EuroQCI in Europe, AUKUS Pillar 2, and the U.S. National Quantum Initiative continue to represent the bulk of current demand.

Post-Quantum Cryptography: The Classical Defense Against the Quantum Threat

There is one last piece, perhaps the most urgent in practical terms, that completes the picture of cooperation between the two worlds: **post-quantum cryptography** (PQC). Unlike QKD and the quantum internet, which use quantum physical phenomena to protect communications, PQC is an entirely **classical** approach: new mathematical algorithms, executable on today's ordinary computers and chips, but designed to withstand attacks from a future quantum computer.

The threat is concrete and has a name: **Shor's algorithm (1994)**. A sufficiently large, fault-tolerant quantum computer could factor huge prime numbers and extract discrete logarithms in a practical amount of time, effectively demolishing the mathematical foundations of RSA, Diffie-Hellman, and elliptic-curve cryptography (ECC) — the very schemes that today protect the vast majority of internet communications, banking transactions, blockchains, and government data. The moment this becomes possible is informally called "Q-Day," and no one knows for certain when it will arrive — but the risk is not only in the future: under the threat model known as "harvest now, decrypt later," an adversary can intercept and store encrypted traffic today, waiting to decrypt it retroactively once quantum hardware is powerful enough. For data that must remain confidential for decades — state secrets, medical records, intellectual property — or for some cryptocurrency transactions, the threat is therefore already active, even though the computer capable of exploiting it does not yet exist.

The cryptographic community's response did not wait for quantum hardware to arrive. After a selection process lasting nearly a decade, the U.S. NIST has published the first official standards for quantum-resistant cryptography, based mainly on mathematical problems related to **lattices** (lattice-based cryptography), believed to be secure even against an attacker equipped with a quantum computer:

- **FIPS 203 (ML-KEM)** — for key exchange;
- **FIPS 204 (ML-DSA)** — for digital signatures;
- **FIPS 205 (SLH-DSA)** — an alternative, hash-based signature scheme, intended as a safety net should weaknesses emerge in lattice-based schemes;
- in addition, **HQC** was selected by NIST as a fifth public-key encryption standard, and a future **FIPS 206** is intended to offer more compact, bandwidth-optimized signatures.

What makes this topic perfectly consistent with the hybrid thread running through this entire article is how the migration is actually taking place: almost no one is abruptly replacing classical cryptography with post-quantum cryptography. The path chosen by giants such as Google, Cloudflare, and Apple (with iMessage's PQ3 protocol) is **hybrid cryptography**: within the same secure connection, an established classical algorithm (for example, an elliptic curve) is combined with a new-generation post-quantum algorithm, so that communication remains protected even if one of the two schemes should later prove weaker than expected. It is the same engineering caution seen in hybrid computing: the mature tool is not abandoned until the new one has proven fully reliable in the field.

On the regulatory front, pressure is mounting rapidly. In the United States, the NSA's Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) sets precise deadlines for national security systems, with full adoption of quantum-resistant cryptography expected by 2035, while an executive order from June 2026 (EO-14412) mandates an acceleration of migration across the entire U.S. federal government, with binding deadlines for the most critical assets. Several sectors — defense, banks with long-lived secrets, government archives — are expected to migrate well ahead of these general thresholds.

Who is working on post-quantum cryptography

- Among the "pure-play" specialists are **PQShield**, **SandboxAQ** (a Google spin-off), **ISARA** (originally a spin-out of BlackBerry's research lab, now focused on "crypto-agility" tools and cryptographic inventory more than on algorithms themselves), **Crypto4A**, **evolutionQ**, **Post-Quantum**, **QuintessenceLabs**, and **Quantum Xchange**.
- Among the major integrators already deploying PQC at global scale are **Cloudflare**, **AWS**, **Microsoft**, and **Google**, which over the course of 2026 have raised their targets and timelines for making TLS connections "quantum-safe" by default.
- On the public side, NIST leads the standardization process together with CISA and the NSA, while initiatives such as the *Quantum Computing Cybersecurity Preparedness Act* require U.S. federal agencies to inventory their cryptographic systems and plan their migration — an obligation that inevitably flows down to private vendors and contractors as well.

In a sense, post-quantum cryptography is the clearest proof that the coexistence of classical and quantum computing is not just about data-center hardware, but about the entire digital infrastructure we take for granted every day: we don't wait for quantum computing to replace classical computing before protecting ourselves from it — we use classical computing, made smarter, to shield ourselves in advance against a threat that will come from the quantum world itself.

Conclusion

The picture that emerges is clear, for computing as much as for the networks and the security that runs through them: the race is no longer (only) about who builds the QPU with the most qubits or the longest quantum link, but about who can integrate the quantum world with the classical infrastructure that already underpins most of the digital world today — better, faster, and more efficiently.

> *It is in this integration — not in the replacement of one paradigm by another — that the future of computing, of the networks that connect it, and of the security that protects it will be decided.*