Sep 11, 2026 – Sep 12 • Daily Brief

How Fast Are Quantum Computers? A Unified Stopwatch Has Finally Arrived

28 items • 8 sources | PDF | This issue's roundup is free

CONTENTS

Items	28
Sources	8
Reading time	~73 min read

CURRENT SOTA

Physical qubits	**1,180**
2Q fidelity	**99.99%**
Logical qubits	**96 LQ**

Quantum hardware SOTA table →

01 Other

Today's Three Key Developments

1 **A new yardstick for quantum processor speed: one million layers per second becomes the new benchmark** A demonstrated sustained execution rate of "over one million circuit layer operations per second" sets a new benchmark for measuring quantum computer performance; the new CLOPS_h metric ties speed directly

to hardware limits [8]—until now the industry compared fidelity, not throughput.

1 **Pasqal and LG CNS sign three-year MoU to build a hybrid computing data center** The two parties will build a scalable hybrid AI and quantum computing platform over a three-year term to power next-generation AI infrastructure [12]—the neutral-atom approach gains a data center entry point via a major Korean systems integrator.

1 **Qoro Quantum and the UK's STFC Hartree Centre to build a quantum-supercomputing demonstration platform** Qoro's Divi SDK and Maestro simulation engine will be integrated into the Hartree Centre's high-performance computing clusters via the open-source QRMI interface [5]—an enterprise-grade hybrid orchestration testbed for UK industry and public-sector research.

---

02 Hardware

Hardware Frontier

Neutral Atoms • Quantum Simulation

A Harvard team used a cold-atom quantum simulator to observe a gradual crossover to a "pseudogap metal" in the Hubbard model (the foundational model describing how electrons interact on a lattice), at temperatures several times lower than previously achieved[14].

Technical significance: The pseudogap is the hardest piece of the high-temperature superconductivity phase diagram to explain, and until now could only be inferred indirectly in real materials; cold-atom simulators turn it into a directly observable object under tunable parameters, and the several-fold reduction in temperature is precisely what crosses that observation threshold—the pseudogap regime of the Hubbard model requires temperatures beyond the reach of earlier experiments.

Landscape impact: This kind of simulation does not follow the universal quantum computing route and does not rely on error correction; it is a textbook case of "special-purpose machines deliver results first." It is an immediate endorsement of the scientific value of neutral-atom platforms, but it runs on a separate track from that route's weaknesses on fault tolerance (cycle repetition rates still in the 1–10 Hz range, limited mid-circuit measurement capability)—the two cannot substitute for each other. The neutral-atom route is advancing on both scientific simulation and error-correction engineering at once.

Diamond Color Centers • Acoustic Protection of Qubits

Harvard researchers used microscopic sound waves to protect quantum information: by having mechanical vibrations continuously envelop a diamond-based qubit, coherence time (how long a quantum state stays intact) was extended to roughly three times its original value[20]. The same phonons may in future both transmit and protect quantum information, pointing toward compact on-chip acoustic quantum networks[20].

Technical significance: The threefold gain is an improvement relative to the device itself, not a cross-platform record—diamond color centers do not compete on two-qubit fidelity leaderboards to begin with; their value lies in room-temperature operation and optical interfaces. The real novelty is the protection mechanism: instead of dynamical decoupling pulse sequences, a continuous acoustic field is used, effectively building noise shielding into the device's physical environment.

Landscape impact: What this affects is quantum networking and sensing, not the mainline of universal computing; if phonons work as an on-chip bus, they could eliminate conversion losses between microwave and optical domains. The timescale is measured in years—this remains a single-qubit proof of principle, a full engineering layer away from multi-node links.

Superconducting • Crosstalk Suppression

Work on IBM's heavy-hex (a sparse layout in which each qubit connects to only 2–3 neighbors) processors has realized an anisotropic (3,5) surface code using a "fold-unfold" embedding combined with robust dynamical decoupling [11].

Technical significance: The inherent problem with heavy-hex layouts is low connectivity; the surface code natively requires square-grid connections, and forcing it in wastes a large number of qubits. Fold-unfold embedding plus crosstalk-resistant pulse sequences is an engineering workaround for squeezing code distance out of existing hardware. An anisotropic code distance like (3,5) means error-correction capability is asymmetric in the two directions—a direct reflection of the layout compromise.

Landscape impact: Measured against the publicly reported state of superconducting quantum error correction, this work's code distance and suppression capability are both at an earlier stage; it is an adaptation advance —"letting non-square-grid hardware run surface codes"—rather than a record. Notably, IBM's own new roadmap (Nighthawk) has shifted to square-grid topology, indicating that the error-correction cost of heavy-hex has been conceded internally at the vendor as well.

Photonics • A Quantum Machine from Classical Light

Physicists have reworked classical light fields into a quantum machine usable for information processing, an approach that sidesteps the obstacle that "multi-particle quantum systems are extremely hard to build" [2]. Noise, loss, and even the tiniest disturbances destroy the fragile behavior that quantum systems depend on to function [2].

Technical significance: The source gives only the conclusion and the background—classical light reworked into a quantum machine usable for information processing, and that quantum systems are inherently fragile while multi-particle systems are hard to build—without disclosing the implementation mechanism [2]. So exactly which class of information-processing tasks this "quantum machine" can perform, and whether genuine many-body entanglement is involved, requires going back to the original paper to judge.

Landscape impact: The hard barrier for the photonic route has always been the end-to-end loss budget, not fidelity figures themselves—the fusion fidelities published for this route are typically conditioned on "a photon having been detected." Any optical scheme that sidesteps the loss problem must first be clear about what it has given up.

Optoelectronic Components • The Overlooked Foundation

Hamamatsu's Klea Dhmitri argues that photonics is the invisible backbone behind the scaling of quantum hardware [23].

Landscape impact: Detectors, light sources, and optical paths are a shared dependency across all platforms—trapped ions and neutral atoms rely on efficient fluorescence collection for readout, and superconducting systems rely on optical interconnects for cross-module links. As a due-diligence threshold, detection efficiency and optical path loss are metrics that must be interrogated item by item. Companies in the supply chain often see shifts in order composition earlier than full-system vendors do.

Algorithms and Software

The "Clock Speed" of Quantum Computers Gets Standardized for the First Time

A sustained execution rate of over one million circuit layer operations per second now sets a new benchmark for evaluating quantum computer performance[8]. Prior benchmarking work focused solely on fidelity, overlooking the practical speed at which these machines run complex computations repeatedly[8]. The CLOPS_h metric introduced by this work ties system speed directly to hardware limits[8].

Technical significance: This is a long-missing piece of the quantum benchmarking framework. Quantum volume (QV) measures "how complex a circuit you can run," and fidelity measures "whether you run it correctly," but variational algorithms, error mitigation, and error-correction decoding all require thousands upon thousands of repeated samples—run slowly, and no amount of fidelity buys a usable result. As for what the "h" suffix denotes, the source offers no explanation, stating only that CLOPS_h ties system speed directly to hardware limits [8].

Landscape impact: Once this yardstick is established, the hidden weaknesses of each platform get put on the table. The well-known afflictions of trapped ions are precisely repetition rate and ion transport overhead, and typical neutral-atom cycle repetition rates are only 1–10 per second—the two routes that look best on fidelity leaderboards will have to be reranked on the speed table; superconducting systems come out ahead thanks to gate times in the 20–100 nanosecond range. For buyers, this means tender documents can add a column that previously could not be written; for vendors, the pressure propagates into engineering roadmaps within six months to a year. One caveat: quantum volume has already saturated at large scale and can serve only as a mid-scale lateral anchor—whether the speed benchmark repeats that fate depends on whether it can withstand targeted vendor optimization.

Distributed Quantum Computing: Gate Teleportation Beats Circuit Cutting

Researchers at memQ Inc. compared two routes to distributed quantum computing—

gate teleportation and circuit cutting—and concluded that the latter carries exponential overhead[15].

Technical significance: Circuit cutting splits one large circuit into several smaller ones that run separately on small machines, then stitches the results back together with classical post-processing—it requires no quantum link and is a pure software approach, but the number of samples needed for the reassembly grows exponentially with the number of cuts. Gate teleportation, by contrast, requires actually distributing entanglement, a high hardware bar; the source states only, as a conclusion, that the exponential overhead falls on the circuit-cutting side, and offers no quantitative conclusion about gate teleportation's overhead[15]. The comparison draws a directional dividing line between "make do in software" and "solve it properly in hardware."

Landscape impact: For every vendor betting on "networking many small machines to stand in for one big one," this is a cost comparison sheet. Circuit cutting remains worthwhile when the number of cuts is very small and can serve as a transitional measure; but to truly scale up, the hardware investment in entanglement distribution is unavoidable—which is precisely the reason companies like Qoro and memQ, working on quantum interconnect and orchestration, exist.

A New Fix for Slow Convergence in Quantum Key Distribution

A new method called FastRényiQKD uses Rényi entropy (a generalized family of information measures) to construct security proofs, aiming to overcome the slow convergence of security proofs in quantum key distribution [13].

Technical significance: This is not hardware getting faster but the mathematics getting cheaper. Quantum key distribution key rates are limited by finite-size effects—when the sample count is insufficient, the security proof must leave a conservative margin, so the key actually extractable falls far short of the theoretical value. Switching to a different family of entropies for the proof aims to shrink that "security discount."

Landscape impact: The potential beneficiaries are operators and financial institutions that have already deployed quantum key distribution links—if the method holds up, they could raise key rates without replacing equipment; but the source says only that the method "aims to" overcome slow convergence, and offers no verified key-rate improvement [13]. The timescale is short—

deploying a proof method requires only an update to the software post-processing layer, measured in months.

Simulation Fidelity Gains a Controllable Error Bound

A "tape-recorder coarse-graining" technique yields exact results under Lorentzian and flat-band bath conditions, where existing hierarchical equations of motion (HEOM) calculations fail to converge[7]. The method dynamically manages computational overhead during quantum nanostructure simulation without sacrificing fidelity, and yields the splitting of Coulomb blockade peaks under interacting-electrode conditions[7].

Technical significance: The key is not "faster" but "bounded"—only when simulation error is explicitly boxed in can the results serve as evidence. Hierarchical equations of motion are the workhorse tool for open quantum system simulation, and their failure to converge under certain bath spectra is a long-known pain point.

Landscape impact: What this affects is the device design stage. Qubit design on every platform relies on open-system simulation to predict decoherence, and non-convergent simulation means design falls back on trial and error. Tooling improvements like this never make headlines, but they quietly shorten the iteration cycle of every tapeout round.

Materials Models Get a New Exam Paper

Michele Simoncelli's team, working with colleagues at the University of Cambridge, has released a new benchmark for evaluating models that predict material properties [9].

Technical significance: The value of a benchmark lies in exposing overfitting—machine-learning potentials that look great within the training distribution and collapse in extrapolation regions are an old problem in this field.

Landscape impact: Materials simulation is one of the repeatedly promised killer applications for quantum computing. Only by first quantifying the capability boundary of classical methods can you know where exactly quantum methods need to take over—the more accurately this baseline is drawn, the harder it becomes to keep claims of quantum advantage vague.

Industry and Ecosystem

Pasqal Teams Up with LG CNS on a Three-Year Hybrid Computing Platform

Pasqal and LG CNS have signed a three-year memorandum of understanding to jointly build a scalable hybrid AI and quantum computing platform to support next-generation AI infrastructure [12].

Business significance: LG CNS is a major Korean systems integrator, holding the data center and operations contracts of enterprise customers. For Pasqal, this is a channel for getting neutral-atom machines into commercial data centers, not another research collaboration. Its MoU status means there is as yet no binding procurement commitment; neither value nor deployment scale has been disclosed.

Landscape impact: Korea has made few public moves in quantum hardware until now, driven mainly by the compute demands of groups like

Samsung and LG. The neutral-atom route placing a stone in the Asian data center market forms a contrast with IQM's positioning in Europe and IBM's in North America—a three-year term is the typical window for this kind of partnership, and the real checkpoint is the timeline for the first machine to land.

Qoro Quantum Plugs into the UK's Hartree Centre

Qoro Quantum has signed a memorandum of understanding with the STFC Hartree Centre to jointly build a UK quantum-supercomputing demonstration platform, integrating Qoro's software stack into the Hartree Centre's high-performance computing clusters [5]. The solution uses Qoro's Divi SDK and Maestro simulation engine, with seamless integration via the open-source QRMI interface [5]. The goal is to provide enterprise-grade hybrid quantum-HPC capability for UK industry and public-sector research [5].

Business significance: Qoro sells the orchestration and scheduling layer and does not build hardware—the value of this layer rises as the number of backend machines grows, a classic "the more fragmented, the more profitable" position. Going with an open interface (QRMI) rather than a proprietary protocol is meant to avoid lock-in to a single hardware vendor.

Landscape impact: Judging by recent moves across countries, putting quantum machines into national supercomputing centers has gone from isolated case to standard practice. The UK's step here fills in the software middle layer, indicating that national build-out sequences are converging: procure hardware first, then backfill the orchestration stack. The most directly affected are domestic enterprise users, whose threshold for gaining experimental access keeps coming down.

Quandela Places a Stone in Vietnam

Quandela and Quanova formally established a partnership at the Vietnam-France leaders' forum, with backing from Nam A Bank, to advance quantum computing in Vietnam[16].

Business significance: A bank appearing on the list of backers of a quantum hardware partnership is worth noting—this usually means the funder treats quantum capability as a long-term strategic asset rather than a procurement item. The value of the partnership and specific deliverables have not been disclosed.

Landscape impact: A French photonics-route vendor exporting to Southeast Asia is going through state-level diplomatic channels. Quantum build-outs in emerging markets typically start with talent development and demonstration machines, with commercial returns measured in years, but first movers gain a voice in standards and training systems.

Q.ANT, QuantrolOx, and Two Other Companies Reshuffle Leadership

Q.ANT has appointed Joerg Behrend as VP of Hardware, responsible for NPU production [6]. Qtonic Quantum has brought in Gina Fratarcangeli as a senior advisor for AI transformation and global alliances [6]. QuantrolOx has appointed Srishti Mahhajan as Head of Global Strategy and Corporate Development [6]. Leiden Cryogenics' new chief executive is Eric Kievit, with founder Giorgio Frossati moving to a senior scientist role [6].

Business significance: Two directional signals stand out in this batch of appointments. Q.ANT creating a "VP of hardware responsible for production" indicates it has moved from the prototype stage into capacity ramp-up; Leiden Cryogenics' founder handing over to a professional manager and returning to science himself is the classic handoff as a cryogenics-type supply chain company moves toward scale.

Landscape impact: Dilution refrigerators are a shared bottleneck supplier for the superconducting and spin routes, and while management changes at this link look unremarkable, shifts in delivery lead times propagate to every downstream system vendor.

Chicago State University Establishes the CQuEST Center

Chicago State University has launched the CQuEST Center as a research and workforce hub on Chicago's South Side, integrating the university's quantum and semiconductor degree programs [4]. The center offers a quantum information science and engineering (QISE) minor and certificate programs, giving students access to research facilities, internships, and employment pathways, alongside regional STEM outreach [4].

Landscape impact: The quantum industry's talent gap is concentrated at the technician level—people who can operate and maintain systems—not at the PhD level; certificate and minor programs are precisely a response to that gap. Chicago already has the infrastructure of the Illinois Quantum park, and local institutions filling in the talent pipeline is a standard component of a maturing industrial cluster. Returns are measured in three to five years.

MIT Quantum Initiative Launches Postdoctoral Scholars Program

The MIT Quantum Initiative (QMIT) has launched a new postdoctoral scholars program, funded by the Gordon and Betty Moore Foundation, starting December 2025[19]. QMIT says the program is intended to connect quantum science with other disciplines and cultivate the next generation of research leaders[19].

Landscape impact: This forms a contrast with the previous item—one end fills the technician gap, the other the cross-disciplinary research leadership gap. Private foundation funding of postdoctoral programs is an important buffer at a time when federal research funding is tight. This kind of investment has the longest payoff cycle, but it determines who will be defining the questions a decade from now.

05 Other

Academic Frontier

"Exponential Protection" in Majorana Qubits Breaks Down in Real Devices

New work shows that in proposed Majorana qubits (the core building block of the topological route, theoretically immune to noise by nature), the initial protection against environmental noise they should enjoy—manifested as signal decay falling exponentially with a parameter—vanishes as the energy splitting between quantum states grows [17]. This departure from ideal behavior originates in real-device limitations such as quasiparticle poisoning, factors that prior analytical models did not address [17]. The new work provides an explicit derivation [17].

Technical significance: The entire appeal of the topological route rests on a single promise—that protection comes from topological properties themselves, with errors suppressed exponentially in system parameters, so that the steep error-correction overhead borne by other routes is unnecessary. This work writes out "under what conditions exponential protection degrades" as an explicit expression, effectively attaching boundary conditions to that promise. Quasiparticle poisoning is the acknowledged number-one killer in Majorana devices, and theoretical models have long treated it as a negligible term.

Landscape impact: For investors in the topological route, this is a risk document worth reading carefully. It does not negate the route, but it converts "topological protection" from a qualitative selling point into a quantitative metric to be verified item by item—the due-diligence question is no longer "is there protection" but "how large is the energy splitting, what is the quasiparticle poisoning rate, how much of the exponential regime is left." Against the maturity of other routes: trapped ions and neutral atoms are already competing on two-qubit gate fidelity and logical qubit counts, while the topological route is still at the stage of demonstrating whether its basic building block holds up—the gap is more than one generation.

Surface Code Embedding on Heavy-Hex Topology

The aforementioned "fold-unfold" embedding and dynamical decoupling work on IBM heavy-hex processors belongs to a class of efforts adapting error-correcting codes onto hardware with non-ideal connectivity [11].

Technical significance: Dynamical decoupling here does more than extend coherence time—it suppresses crosstalk. Crosstalk patterns in sparse-connectivity layouts differ from those in square grids and require purpose-designed pulse sequences. Work of this kind determines whether the dozens of superconducting machines already installed can still be used for error-correction experiments.

Landscape impact: What this affects is the research user base on IBM's cloud platform. In the short term it raises utilization of the installed hardware base; in the long term, IBM's shift to square-grid topology has already given the answer—error-correction adaptation on heavy-hex is a transitional measure, not an endgame.

---

06 Impact

Today's Impact

1. **Buyers gained a parameter today that they can write into a tender.** CLOPS_h establishes the sustained execution rate of "one million circuit layer operations per second" as a new benchmark [8], where

previously only fidelity and quantum volume were comparable. The most directly affected are supercomputing centers and enterprise customers preparing to procure quantum machines—in the past they had no way to write "runs slowly" into a contract. What to watch next: whether each vendor publishes its own CLOPS_h figure, and who stays silent. The two routes strongest on fidelity (trapped ions, neutral atoms) are precisely the ones most disadvantaged on repetition rate; silence is itself information.

1 **The cost accounting for the two routes to distributed quantum computing has been settled.** The comparison of gate teleportation and circuit cutting concludes that the latter carries exponential overhead [15]. Affected are all vendors with "multi-machine interconnect" on their roadmaps, as well as orchestration-layer companies like Qoro [5]. What to watch next: the specific upper bound on the number of cuts for which circuit cutting remains viable—that number determines how much longer the transitional window for pure-software approaches lasts, and when investment in quantum interconnect hardware must begin.

1 **Quantum key distribution has a path to higher throughput without swapping**

hardware. FastRényiQKD rewrites the security proof using Rényi entropy, aiming to tackle the key-rate discount caused by slow convergence [13]. Affected are telecom operators and financial institutions that have already laid quantum key distribution links. What to watch next: when this proof enters standardization—software-layer updates are measured in months, but writing it into industry specifications requires independent review and takes considerably longer.

1 **The topological route's core selling point has been given boundary conditions.** Exponential protection in Majorana qubits breaks down under growing energy splitting and quasiparticle poisoning [17]. Affected are investors and teams betting on the topological route. What to watch next: whether follow-up experiments provide measured ranges for energy splitting and poisoning rates—once the theoretical boundary is drawn, the next question is inevitably which side of it real devices fall on.

1 **Two data-center-scale partnerships landed on the same day, giving a new answer to "where do quantum machines go."** Pasqal and LG CNS will jointly build a hybrid platform over three years [12], and

Qoro and the UK's Hartree Centre will jointly build a quantum-supercomputing demonstration platform [5]. What this affects is the access threshold for enterprise customers—once machines are inside commercial and national data centers, every user no longer has to negotiate a dedicated line with a vendor. What to watch next: when these two MoUs convert into binding procurement, and when the first machines arrive. MoUs themselves carry no dollar figures; the fulfillment rate is the real test for partnerships of this kind.

07 Other

Editor's Note

There is an understated throughline in the day's news: the industry is starting to set rules for itself. The speed benchmark turns "how fast does it run" from vendor rhetoric into a measurable number, the materials model benchmark quantifies the capability boundary of classical methods, the distributed computing overhead comparison turns the cost difference between two technical paths into an explicit conclusion, and the analytical work on Majorana qubits attaches failure conditions to one route's core promise. None of this work produces new records, but what it changes is how records should be counted. A field shifting from "whose number is bigger" to "what is the number's measurement convention" is usually in the transition from demonstration to procurement—once buyers appear, sellers have to accept third-party units of measure.

What bears watching is that the units of measure themselves get optimized against. Quantum volume saturated beyond large scale and degraded from a differentiating metric into a marketing number; that lesson is still fresh. How long the new speed benchmark holds up depends on whether it can resist targeted tuning—making a test circuit run fast is easy, making real workloads run fast is hard. For readers, the practical test hasn't changed: look at the conditions under which a number was measured, whether it is a single best-case sample or a median across all, and whether raw data was attached. Today's pseudogap observation and acoustically protected coherence time are both solid science, but they remain an entire engineering chain away from products; today's two data center MoUs both lack dollar figures, and whether they are fulfilled will take one to three years to find out. At this stage, work that states things clearly is worth more than work that states numbers large.

References

[1] Find Your Favorite Speakers (Literally)

[2] Physicists turn classical light into a quantum machine for information processing

[3] Alignment Is Whether AI Wants to Kill Us. Cybersecurity Is Whether We Let It.

[4] SEALSQ Reports H1 2026 Financial Results: Revenue Surges 131% to $11.2M, $486M Cash Balance, and Post-Quantum Hardware Progress

[5] Chicago State University Launches CQuEST Center to Anchor South Side Quantum Workforce Pipeline

[6] Qoro Quantum Signs MoU with STFC Hartree Centre to Build UK Quantum-HPC Demonstrator

[7] Who’s News: Strategic Appointments at Q.ANT, Qtonic Quantum, QuantrolOx, and Leiden Cryogenics

[8] QClairvoyance Quantum Labs Signs MoU with IITM–C-DOT Samgnya Foundation to Advance Indian Quantum Ecosystem

[9] Qubic Secures CA$1.5 Million ($1.08 Million USD) Canadian Federal Contract for Cryogenic Amplifiers

[10] Researchers Bound Quantum Infidelity during Simulation Periods

[11] Researchers Measure Quantum Processor Speed Using New Benchmark

[12] Columbia & Cambridge unlock quantum AI for new material models

[13] Security in the Post-Quantum Era, What Changes Once the Migration Is Done

[14] Chalmers speeds quantum operations 1,000x, nearing fault tolerance

[15] Dynamical Decoupling Shields Qubits From Heavy-Hex Crosstalk

[16] Pasqal and LG CNS build quantum-ready AI data centers

[17] A new cone, FastRényiQKD, boosts quantum key distribution speed

[18] Quantum simulator finds a ‘pseudogap’ in a model material

[19] memQ compares gate teleportation to circuit cutting for quantum computing

[20] Quandela and Quanova team up to grow quantum skills in Vietnam

[21] Qubic lands $1.5M Canadian deal for quantum amplifiers

[22] Fisica Team Quantifies Decay Rates in Majorana Qubits

[23] Quantum Zeitgeist Weekly Digest

[24] QMIT launches quantum fellowships to train next-gen leaders

[25] Tiny sound waves could help solve a major quantum computing problem

[26] How AI can help with early schizophrenia diagnosis

[27] Why doing nothing is good for your brain

[28] Klea Dhmitri (Hamamatsu): Photonics as the hidden backbone of quantum hardware scaling