

Quantum Technology: A Technical Review

Computing, sensing, communications, and security: hardware approaches, evidence, and engineering roadmaps.

Dikran S Meliksetian

dmeliksetian@newhaven.edu

Revised: 14 October 2026 **Horizon:** selected milestones from 2024 to 2035

Scope: technical capabilities, deployment maturity, and the limits of current evidence.

Quantum technologies share components, fabrication methods, control systems, and supply chains, but their engineering maturity differs substantially. Atomic clocks and selected magnetometers already support operational systems. Quantum key distribution (QKD) is deployed in selected networks, while general entanglement networks remain experimental. Quantum computing has demonstrated increasingly capable physical devices and encoded computations; turning those demonstrations into useful, scalable fault-tolerant computation remains an engineering objective.

How to read this review

- **Count useful computation.** State which qubits are counted, which operations are protected, and what error rate and runtime the complete task requires.
- **Match the evidence to the claim.** A research demonstration, a customer pilot, and a commercially advantageous production workflow establish different things.
- **Treat roadmaps as targets.** Vendor dates describe intended engineering progress. They do not establish that a capability has been delivered.

This review distinguishes parity lifetime from qubit coherence, error detection with postselection from sustained error correction, and performance benchmarks from demonstrations of quantum advantage. Numerical milestones carry citations to specific papers, announcements, or official documents. Historical announcements retain their dates; subsequent revisions are identified explicitly.

Evidence cutoff: 14 October 2026. This is a selective technical review, not an exhaustive vendor inventory or an independent replication of reported experiments.

Contents

Technology maturity and commercial evidence	3
From physical qubits to useful computation	4
Who's building what	5
Hardware approaches: tradeoffs by modality	6
Quantum sensing and the SQUID–transmon connection	8
Modality vs. use case: computing, sensing, networking	9
Programming languages & SDKs	10
How quantum computing is delivered	10
Applications: candidate workflows and reported examples	11
How performance and quantum advantage are measured	12
Selected milestones: 2024–2026	13
Interpreting the IBM sampling result	14
Objectives and conditional roadmaps	15
Quantum communications is not one technology	16
Supply chain and standards	16
Security and policy	17
What could go wrong	18
Evidence discipline and reproducibility	19
Sources and technical documentation	20

Technology maturity and commercial evidence

The word “quantum” does not imply a common deployment timetable. The following categories describe readiness, not market size or a ranking of technical importance.

Technology	Maturity	Interpretation
Atomic clocks; selected magnetometers	Mature / fielded	Established quantum effects already support timing, measurement, and instrumentation.
Gravimetry; inertial sensing	Product-dependent: research to deployment	Field robustness, size, power, drift, and cost determine readiness.
QKD	Selected commercial deployments	Link economics, authentication, and trusted-node assumptions matter.
Quantum annealing	Commercial / specialized	Accessible systems exist; performance depends on the problem, mapping, and comparison.
Noisy gate-model computing	Research / pilots	Useful for experimentation; a successful pilot is not broad production advantage.
Fault-tolerant universal computing	Research / prototypes	Encoded operations and logical memories must be assessed by protection, depth, and error rate.
Repeaters; entanglement networks	Research / prototypes	Memories, swapping, interfaces, and routing remain engineering challenges.
Post-quantum cryptography (PQC)	Deployment / migration	Runs on classical computers; deployment centers on inventory, integration, and crypto-agility.

Sources: [14, 34, 35, 1, 36].

A PRACTICAL EVIDENCE LADDER

Research demonstration

Establishes feasibility under stated experimental conditions.

Proof of concept or pilot

Implements and tests a customer-relevant formulation, possibly with real data.

Production

Integrates a workflow into an operational process with reliability and support requirements.

Practical quantum advantage

Beats the best relevant classical alternative on a useful task at matched accuracy.

Commercial quantum advantage

Retains value after total cost, latency, staffing, reliability, and integration are included.

These are working definitions used in this review. A partnership announcement establishes interest or collaboration; it does not, by itself, establish production deployment or quantum advantage.

From physical qubits to useful computation

Raw qubit count is one of the easiest industry numbers to quote and one of the easiest to misinterpret. A physical qubit, a logical qubit, and a useful fault-tolerant computational resource are different units.

Physical qubit → encoded logical qubit → fault-tolerant logical operations → useful algorithm

A **physical qubit** is a controllable two-level quantum degree of freedom. A **logical qubit** is an encoded unit of quantum information. Its actual protection depends on the code and implemented protocol; merely encoding a state does not establish fault tolerance. Sustained quantum error correction (QEC) typically requires repeated syndrome extraction and classical decoding. There is no universal physical-to-logical conversion ratio: overhead depends on physical error rates, connectivity, leakage and loss, the chosen code, decoder performance, target circuit depth, and the acceptable probability that the complete algorithm fails. Consequently, physical and logical counts on different roadmaps are not directly interchangeable.[2]

Vendor terminology varies: IBM's programmable/physical split

IBM distinguishes **programmable qubits** from a broader count of **physical quantum elements**, including supporting couplers and reset elements. Nighthawk r2 has 120 programmable qubits and 458 total physical elements. This is IBM's reporting convention, not a universal definition of a computational qubit. Cross-platform comparisons should identify data, ancilla, coupler, reset, and spare resources separately.[3, 4]

What the error-correction race is actually measuring

- **Threshold and below-threshold scaling** — whether larger codes suppress logical errors under a specified noise model and operating protocol.
- **Logical error rate** — the failure probability of an encoded operation or memory cycle, not merely physical gate fidelity.
- **Code family** — surface codes remain the best-known baseline; qLDPC codes promise lower asymptotic overhead if demanding connectivity and decoding requirements can be engineered; bosonic/cat encodings and erasure-biased architectures attempt to exploit hardware-specific noise structure.
- **Logical gates** — protecting memory is not enough. A useful machine needs a universal set of fault-tolerant logical operations, often including expensive non-Clifford resources such as magic-state distillation.
- **Decoder and control loop** — syndrome extraction, real-time classical decoding, feed-forward, calibration, leakage management, and logical compilation are integral parts of the computer.[2]

Near-term bridge: quantum error mitigation (QEM)

Quantum error mitigation (QEM), including zero-noise extrapolation and probabilistic error cancellation, aims to improve estimates from noisy executions, often at substantial sampling cost. Circuit cutting is a related decomposition technique with its own reconstruction overhead. Neither automatically provides arbitrarily deep reliable computation.[17]

Postselection must be reported

An error-detection protocol may retain only runs whose measured syndromes satisfy an acceptance rule. Report the acceptance probability, total shots, retained-shot accuracy, and total runtime. Such a result establishes a different capability from a memory or processor that repeatedly corrects errors while retaining the computation.[6]

Who's building what

Representative organizations are grouped by their computing approach or sensing and networking activity. These categories describe technical focus; they are not mutually exclusive corporate classifications.

Computing approach	Representative organizations and focus
Superconducting circuits	IBM, Google, Rigetti, OQC: gate-model processors; different connectivity, control, and modularity strategies.
Trapped ions	Quantinuum, IonQ: atomic-ion qubits; transport, local gates, and interconnect approaches differ.
Neutral atoms	QuEra, Pasqal, Atom Computing, Infleqtion: optical arrays spanning analog simulation and digital operations.
Photonics	PsiQuantum, Xanadu, ORCA: distinct encodings and architectures using optical components, sources, detection, and measurement.
Semiconductor spins	Intel, Diraq, Quantum Motion, Silicon Quantum Computing, Equal1: electron or nuclear-spin approaches with differing devices and control schemes.
Topological approach	Microsoft: semiconductor–superconductor devices intended to support protected qubits and measurement-based processing.
Quantum annealing	D-Wave: commercial superconducting annealers, alongside a separate developing gate-model program. ^[32]

Specific performance and maturity claims are supported in the following hardware sections. Membership in a category does not imply that every company uses the same architecture or has demonstrated the same capabilities.

SENSING

- Vector Atomic: timing and inertial-sensing systems.
- AOSense and Infleqtion: atom-based sensing and timing.
- Q-CTRL: control software and quantum navigation programs.
- SBQuantum: diamond-based magnetometry.
- QuSpin: atomic-vapor magnetometers, distinct from NV-diamond devices.^[16]

COMMUNICATIONS AND NETWORKING

- ID Quantique and Toshiba: commercial QKD systems.
- Qubitekk and other networking programs: photon sources, entanglement distribution, and node technology.
- Arqit: classical quantum-safe cryptographic services; these should be distinguished from QKD.
- Repeaters, memories, transducers, and general network nodes remain a separate development track.

Hardware approaches: tradeoffs by modality

No single metric ranks all platforms. Compare gate type and fidelity, connectivity, measurement and reset, coherent storage, loss, calibration, and system overhead under explicit conditions.

Superconducting circuits

IBM, Google, Rigetti, OQC

Strengths

- Fast electrical control and integrated circuits support rapid gate and measurement cycles.
- Lithographic fabrication and extensive control infrastructure support iterative device development.

Engineering constraints

- Dilution refrigeration, wiring, leakage, crosstalk, and calibration constrain scaling.
- Coherence varies by design and device; it must be quoted as a defined quantity such as T_1 or T_2 .

Sources: [1, 4].

Trapped ions

Quantinuum, IonQ

Strengths

- Atomic internal states support long storage and precise control.
- Within an architecture, ion transport or collective modes can provide flexible interactions.

Engineering constraints

- Gate, transport, cooling, and measurement times all contribute to runtime.
- Vacuum, optical delivery, motional control, and module connections must scale together.
- Laser cooling of ions is distinct from cryogenic cooling of the apparatus; some systems operate in room-temperature environments.

Sources: [11].

Neutral atoms

QuEra, Pasqal, Atom Computing, Infleqtion

Reconfigurable optical arrays offer flexible geometry and repeated use of the same atomic species. Analog simulation and digital gates are distinct operating modes: the number of trapped atoms does not by itself identify the number participating in high-fidelity universal operations. Atom loss, reloading, optical addressing, and control uniformity affect usable depth and throughput. Reconfigurable-array experiments have demonstrated encoded processing, but the error-detection or correction protocol must accompany any logical-qubit count.[12]

Read numerical comparisons with their conditions

Compare single-qubit gates with single-qubit gates, and entangling gates with entangling gates. Separate record results on selected devices from array-wide performance. A ratio of coherence time to gate time is incomplete without measurement, reset, transport, leakage, and correlated-error information. Total system costs should include fabrication, lasers or cryogenics, controls, facilities, maintenance, and classical computation.

PHOTONIC, SEMICONDUCTOR-SPIN, AND TOPOLOGICAL APPROACHES

Photonics

PsiQuantum, Xanadu, ORCA

Strengths

- Optical transmission naturally supports modular systems and links between remote nodes.
- Integrated photonics can draw on established manufacturing techniques.

Engineering constraints

- Photon loss, source quality, switching, and detector efficiency set resource requirements.
- Architectures may use probabilistic entangling operations, multiplexing, and large resource states.
- Room-temperature photons do not imply a room-temperature machine: PsiQuantum cools photonic packages with liquid helium, including superconducting detectors and control electronics.

Sources: [10].

Semiconductor spin qubits

Intel, Diraq, Quantum Motion, Silicon Quantum Computing, Equal1

Strengths

- Small physical devices offer a route to dense arrays.
- Semiconductor process knowledge is relevant to fabrication and integration.

Engineering constraints

- Uniformity, interconnects, cryogenic control, tuning, and readout remain scaling tasks.
- High fidelity on a small device is not a guarantee of array-wide performance. A 2022 silicon experiment reported 99.5% two-qubit fidelity under its particular benchmarking protocol.

Sources: [13].

Topological approach

Microsoft aims to reduce physical error rates through protected semiconductor–superconductor devices. Hardware protection would complement, rather than replace, QEC and fault-tolerant logical operations. The 2026 InAs–Pb tetron preprint reports a characteristic wire-parity switching time of about 20 seconds. It does not, by that measurement alone, establish a 20-second coherent lifetime for arbitrary qubit states or universal protected computation. These are separate experimental questions.[7, 8]

Different computing models: annealing and analog simulation

Quantum annealing encodes an objective in a Hamiltonian and seeks low-energy outcomes; deployed annealers are specialized systems, not universal gate-model processors. Analog quantum simulation emulates a target physical model without compiling a universal gate sequence. Each needs its own accuracy criteria and classical comparison. D-Wave's separate gate-model development should not be conflated with the capabilities of its commercial annealers.[32]

Quantum sensing and the SQUID–transmon connection

Quantum sensing spans several distinct families. Some are long-established instruments; others are prototypes whose laboratory performance still needs to survive field conditions.

Family	Physical platform	Representative use
Atomic clocks	Atomic transition frequencies	Timing and synchronization.
Atom interferometers	Matter-wave phase	Acceleration, rotation, and gravimetry.
Atomic magnetometers	Optically pumped atomic vapor	Biomagnetism and field measurement; QuSpin is an example.[16]
NV/color-center sensors	Optically addressed solid-state spins	Local magnetic, electric, and temperature sensing.
SQUID magnetometers	Josephson interference	Sensitive magnetic measurements and established biomagnetic instruments.[14]
Quantum RF sensors	Rydberg atomic states	Electric-field and radio-frequency metrology.
Quantum imaging	Nonclassical light or correlations	Selected imaging and metrology tasks.

NV-diamond / color centers

Room-temperature operation and close sample proximity support compact sensors and high spatial resolution. Sensitivity depends on sensing volume, spin preparation, optical collection, noise, and interrogation time. Spin–photon networking can require different temperatures.

Avoid a universal sensitivity ranking against SQUIDs. Specify noise-equivalent sensitivity, often in $\text{T}/\sqrt{\text{Hz}}$, plus bandwidth, geometry, sensing volume, and stand-off distance.

SQUID magnetometers

Low-temperature SQUIDs commonly operate near liquid-helium temperature (4.2 K); high-temperature devices can operate near liquid-nitrogen temperature (77 K). Those requirements differ from millikelvin qubit refrigeration. Flux-locked-loop feedback linearizes the periodic response and expands usable dynamic range. Sensitivity depends on noise and magnetic coupling, not solely on the flux quantum.[14]

Shared physics, different devices

A SQUID is a superconducting loop interrupted by Josephson junctions. Phase consistency around the loop produces a response periodic in applied flux, with period $\Phi_0 = h/(2e)$. In a current-biased dc-SQUID, the voltage can resolve small flux changes; the applied flux need not change in integer steps.[14]

A transmon combines Josephson nonlinearity with a large shunt capacitance. Its anharmonic spectrum allows the lowest two levels to be addressed as a qubit, while the capacitance reduces charge-noise sensitivity. A split-junction transmon uses a SQUID loop for frequency tuning, which also introduces sensitivity to flux noise. Flux qubits instead use states associated with circulating currents. A sensor's sensitivity and a computing qubit's coherence therefore call for different operating choices.[15]

Modality vs. use case: computing, sensing, networking

Platforms can contribute to more than one domain, but shared materials do not make a sensor, a memory, and a universal processor equivalent products. Entries below indicate representative roles, not commercial readiness.

Platform	Computing	Sensing	Networking
Superconducting circuits	Gate-model processing; separate annealing architectures	Related SQUID devices	Cryogenic links and optical transduction research
Trapped ions	Gate-model processing	Ion-based precision frequency standards	Ion-photon node interfaces
Neutral atoms	Analog simulation and digital gates	Clocks and atom interferometry	Atom-photon memories and nodes
Photonics	Optical encodings and measurement-based approaches	Quantum imaging and optical metrology	Transmission, entanglement distribution, and QKD
Atomic vapor	Specialized research uses	Established magnetometry and timing	Optical-memory research
NV/color centers	Small-scale spin processing research	Local and ensemble sensing	Spin-photon interface research
SQUID sensors	Distinct from transmon processors	Magnetic measurement	Not a general network-node platform
Semiconductor spins	Processor and control research	Related spin-sensing research	Interfaces remain architecture-dependent
Topological devices	Proposed route to protected processing	Outside this review's scope	Outside this review's scope

Sources: [11, 10, 16, 14, 7].

Two distinctions matter

Atomic clocks include multiple architectures; trapped-ion optical clocks should not be presented as the universal timing backbone of fielded navigation. Similarly, a photonic communications link does not establish that the attached computing hardware can perform distributed fault-tolerant operations. Interfaces, fidelity, timing, and operational protocols must be assessed together.

Programming languages & SDKs

Python is a common host language, but frameworks also differ in their circuit model, compilation, differentiation, simulation, device interfaces, and support for hybrid execution. The descriptions below are capabilities, not rankings.

SDK / language	Organization; interface	Typical role
Qiskit	IBM; Python	Circuit construction, transpilation, quantum information, and execution integrations.[18]
Cirq	Google; Python	Hardware-aware circuit manipulation, simulation, and experiments on noisy processors.[19]
PennyLane	Xanadu; Python	Differentiable quantum programs, quantum chemistry, and hybrid learning; maintained PyTorch and JAX interfaces.[20]
Q# / QDK	Microsoft; Q#, Python tooling	Quantum programming, simulation, and resource-estimation workflows through the QDK and VS Code.[21]
tket / pytket	Quantinuum; C++ / Python	Circuit compilation and optimization with backend extensions.[22]
CUDA-Q	NVIDIA; C++ / Python	Hybrid programs and accelerated simulation, with hardware integrations.[23]
Braket SDK	AWS; Python	Circuit construction and access to supported simulators and hardware; other SDKs require supported integration packages.[24]

Software support is version-sensitive. PennyLane states that TensorFlow support is no longer maintained from v0.44. Pin SDK and provider versions in reproducible work, rather than describing a moving release as permanently “latest.”[20]

How quantum computing is delivered

Application → orchestration → CPU/GPU/HPC ↔ QPU

Cloud services expose remote hardware through compilers, scheduling, execution APIs, and domain software. Hybrid workloads include classical optimization, data preparation, decoding or mitigation, and post-processing. Queueing, communication, device availability, and calibration can dominate nominal gate execution time.[23, 24]

Evaluate the complete workflow at a specified accuracy and success probability. Include state preparation, all attempted shots, discarded runs, parameter sweeps, compilation where relevant, and classical reconstruction. A fast kernel or a large QPU does not establish a fast application.

Applications: candidate workflows and reported examples

Application relevance is strongest when a report identifies a task, algorithm, dataset, accuracy target, hardware configuration, and classical comparator. The candidate areas below should not be read as established production advantage.

CHEMISTRY AND MATERIALS

Electronic-structure and quantum-dynamics calculations motivate quantum simulation. Near-term methods can combine quantum samples with classical subspace calculations; deeper fault-tolerant approaches require explicit logical-resource estimates. State preparation, measurement cost, and the strength of classical chemistry methods are central to evaluation.

FINANCE

Candidate tasks include sampling, risk estimation, portfolio optimization, and predictive models. Any theoretical query benefit must survive data loading, oracle implementation, finite precision, and operational latency. Prediction-quality improvements and computational speedups are different claims.

OPTIMIZATION

Gate-model variational algorithms, annealing, and hybrid heuristics can be tested on scheduling, logistics, and design problems. Compare against strong specialized solvers at matched solution quality. Quantum-inspired algorithms run classically and should be compared with other classical methods.

QUANTUM MACHINE LEARNING

Parameterized circuits and quantum feature representations are research tools. Evaluate data-encoding cost, trainability, shot requirements, generalization, and the classical baseline. Faster training or lower energy consumption must be measured for the complete workflow.

Reported example: HSBC and IBM

In September 2025, HSBC reported a hybrid quantum–classical experiment predicting whether bond orders would be filled. Its announcement reported up to a 34% improvement relative to the classical techniques used in that study. This is a task-specific, company-reported result; it does not by itself establish superiority to all classical methods, a runtime advantage, or positive production economics.^[25]

Reported example: quantum-assisted chemistry

IBM, Cleveland Clinic, and RIKEN describe sample-based quantum diagonalization workflows combining quantum hardware and classical computing for molecular calculations. Quantum samples guide a classical subspace calculation. The example illustrates a concrete hybrid method; assess accuracy, classical post-processing, and total resources before assigning an advantage claim.^[26]

How performance and quantum advantage are measured

Hardware metrics characterize a device or execution stack. A quantum-advantage claim additionally requires a computational task, an accuracy criterion, and a justified classical comparison.

Task or metric	What it establishes and what to disclose
Random circuit sampling (RCS)	A sampling task. A hardness argument and experimentally justified fidelity are needed for an advantage interpretation.
Cross-entropy benchmarking (XEB)	A statistic used to compare samples with ideal probabilities under stated assumptions; it is not the sampling task itself or a universal measure of computational power.
Quantum Volume (QV)	Performance on specified random square circuits, combining compilation, connectivity, and operations; it is not a count of usable logical qubits.[27]
Algorithmic Qubits (AQ)	IonQ's aggregate benchmark; specify its version, circuit suite, and success criteria.[29]
CLOPS	Circuit-layer throughput under a defined protocol; distinguish it from arbitrary circuits per second.[3]
rQOPS	Microsoft's proposed measure: logical width times logical clock frequency, specified with a logical error rate.[30]
Application-oriented suites	Structured workloads such as SupermarQ reveal performance on selected circuit families, without proving advantage for every application.[28]

The moving classical baseline

GPUs, tensor networks, specialized solvers, improved Monte Carlo, and quantum-inspired methods continue to improve. Match approximation error, success probability, and output requirements; disclose preprocessing, wall-clock time, energy or cost where relevant, and the strongest classical implementation considered. A complexity-theoretic claim must state its assumptions and its connection to the finite experiment.

Three claims to keep separate

Computational advantage: a specified task exceeds the stated feasible classical baseline; the task may be artificial.

Practical advantage: the task is useful and the comparison remains favorable.

Commercial advantage: the improvement survives total operational cost and constraints.

Selected milestones: 2024–2026

Dates refer to the original publication or announcement. Experimental results, preprints, and vendor claims are labeled separately.

Dec 2024	Willow: below-threshold surface-code memory Peer-reviewed experiment. Published online December 9, 2024. The logical error decreased by a factor 2.14 ± 0.02 when code distance increased by two. This is a memory-scaling result, not a complete universal fault-tolerant computer.[1]
Feb 2025	Microsoft introduces Majorana 1 Vendor announcement. Microsoft presented its topological-device strategy. Hardware protection, coherent qubit control, and scalable logical operations require distinct experimental evidence.[9]
Sep 2025	Quantinuum reports a QV milestone Vendor benchmark report. System Model H2 reached Quantum Volume 33,554,432. This is a QV score under its protocol, not a physical- or logical-qubit count.[27]
Dec 2025	Stanford reports a room-temperature optical device Institutional report of published research. Stanford's December 2 announcement concerns a nanoscale light-matter device relevant to quantum signaling. Later May 2026 coverage should not replace the original date or be read as a demonstration of a complete photonic computer.[31]
Jan 2026	D-Wave reports on-chip gate-model control Vendor engineering announcement. D-Wave described adapting cryogenic control techniques to gate-model qubits. This is a component-level scaling milestone, separate from a demonstration of universal fault-tolerant computation.[32]
Jun 2026	Majorana 2: reported parity-lifetime improvement Technical preprint and vendor announcement. An InAs–Pb device exhibited a characteristic wire-parity lifetime of about 20 seconds. The reported observable must not be relabeled as arbitrary-state qubit coherence.[7]
Jul 2026	IBM announces quantum-advantage experiments Preprints and vendor interpretation. The IBM–University of Chicago sampling experiment combines space-time codes, syndrome postselection, and a device-dependent fidelity certificate. Its July version describes 70 computational qubits and 468 T gates.[5]
Aug 2026	IBM announces Nighthawk r2 Vendor hardware report, August 31. IBM reports 120 programmable qubits, 458 physical elements, and over 100,000 circuit executions per second for suitable workloads. This is not a universal application throughput or a CLOPS figure.[4]

Interpreting the IBM sampling result

The July announcement and the revised technical preprint describe different reported circuit instances. Preserve the source version when quoting a numerical milestone.

Reported quantity	July 28: v1	September 2: v3
Computational circuit width	70 qubits	64 qubits
Clifford circuit depth	70	73
T gates	468	314
Total physical qubits used	97	76
Certified fidelity lower bound	0.284	0.349
Confidence level	95%	95%
Error handling	Syndrome postselection	Syndrome postselection

Sources: [5, 6].

What the experiment demonstrates

The authors construct structured sampling circuits with complexity-theoretic hardness guarantees under stated assumptions and use their code structure to certify experimental fidelity. The certificate is device-dependent. Reported effective gate-error suppression is obtained after accepting runs according to their syndromes. The experiment therefore needs to be described with its acceptance rule and total sampling cost.[6]

How to compare it fairly

An encoded sampling circuit, a below-threshold memory, and a processor running a universal sequence of corrected logical gates establish different capabilities. Compare code parameters, supported operations, logical error rates, accepted fraction, decoder behavior, and executable depth. A single league table of “logical-qubit counts” discards these distinctions.

IBM’s other July announcements use different validation approaches. Their shared announcement date does not make their evidence interchangeable. A review should assess each task and its assumptions separately rather than treating the announcements as a single proof of general-purpose or commercial advantage.

Objectives and conditional roadmaps

The dates below are stated engineering targets. They are not independent forecasts, and a passed target date does not establish that the milestone has been delivered.

Target	Program	Stated objective / interpretation
2026	IBM Kookaburra	qLDPC memory with an attached logical processing unit.
2027	IBM Cockatoo	Entanglement between modules through a universal adapter.
2028	IBM Starling prototype	Multi-module demonstration including magic-state injection.
2029	IBM Starling	Vendor target: 200 logical qubits and circuits of 100 million gates.
2029	Microsoft	Vendor target for a scalable machine; the demonstrated parity lifetime is not evidence that this target has been met.
2030	IonQ	Current vendor roadmap states 2 million physical and 80,000 logical qubits. These ambitious counts require scrutiny of architecture, error correction, interconnects, and executable operations.

Sources: [2, 8, 33].

Physical-to-logical ratios are architecture-specific projections. Do not infer cryptographic capability from a physical-qubit count or calendar year alone: the algorithm, code, error budget, non-Clifford resources, and runtime must also be specified.

SENSING: QUALIFY THE DEPLOYMENT

Clocks and some magnetometers already operate outside laboratories. Other sensors need improved robustness against vibration, acceleration, temperature variation, and environmental magnetic noise. DARPA's Robust Quantum Sensors (RoQS) program, whose first phase began in 2025, explicitly targets this laboratory-to-platform gap.[34]

Assess sensitivity together with drift, calibration, size, weight, power, uptime, and maintenance. Navigation is an important application, but it does not define the maturity of all sensing technologies.

NETWORKS: DISTINGUISH THE SERVICE

EuroQCI combines terrestrial and space-based secure-communication infrastructure. Deployment of QKD links is a different milestone from repeaters that support entanglement distribution between quantum processors.[35]

For network targets, specify distance, loss, secret-key or entanglement rate, fidelity, memory lifetime, trust assumptions, and availability. No single global date defines completion of a general-purpose quantum internet.

Quantum communications is not one technology

- **QKD** distributes key material using quantum signals and an authenticated classical channel. The key is then used by classical cryptographic systems. Security depends on the protocol, device assumptions, authentication, and implementation.
- **Trusted-node QKD networks** and **entanglement networks** have different trust models. Extending a key-distribution service does not automatically create end-to-end quantum-state transmission.
- **Entanglement networks** distribute entangled states for tasks such as teleportation and distributed sensing; more advanced distributed computation has additional requirements.
- A general **quantum internet** needs suitable memories, interfaces, entanglement swapping, repeaters, synchronization, routing, and error management. Microwave–optical transduction is relevant to some hardware platforms, not a universal requirement for every node.

Sources: [35, 38].

Supply chain and standards

Supply chain dependencies

Depending on the platform, critical inputs include dilution refrigeration, cryogenic electronics, microwave components, lasers, vacuum systems, single-photon detectors, photonic foundries, engineered materials, packaging, and calibration systems.

Facilities, power, maintenance, uptime, operators, networking, and classical compute for error handling are part of the operating stack. Component price alone is not the cost of a usable quantum computer.[10, 11]

Standards and interoperability

Useful standards define terminology, benchmark procedures, metrology, interfaces, and security requirements. NIST's PQC standards are concrete deployment specifications. Network and hardware integration also require compatible timing, formats, and operating protocols.[36, 35]

Procurement should identify the workload and acceptance criteria, including reproducible accuracy, throughput, availability, and total cost. An SDK integration is not proof that hardware backends are interchangeable.

Security and policy

PQC migration is a classical-computing deployment task. Export-control rules and proposed legislation are separate policy instruments; their scope and legal status must be stated precisely.

Post-quantum cryptography

NIST finalized ML-KEM (FIPS 203, key establishment), ML-DSA (FIPS 204, digital signatures), and SLH-DSA (FIPS 205, stateless hash-based signatures) in August 2024. Migration should address cryptographic inventory, protocol and certificate integration, implementation validation, and algorithm agility. Replacing algorithms is more than changing keys.[36]
The harvest-now-decrypt-later concern applies to information encrypted today whose confidentiality must survive future cryptanalytic quantum computers. This motivates migration without requiring a reliable prediction of the first such computer’s arrival.[37]

Date / framework	Scope and status
January 1, 2027	CNSA 2.0: new National Security System acquisitions must comply unless otherwise noted.
December 31, 2030	CNSA 2.0: equipment and services unable to support the suite must be phased out unless otherwise noted.
December 31, 2031	CNSA 2.0: algorithm use is mandated unless otherwise noted.
After 2035	NIST’s initial public draft IR 8547 proposes disallowing specified quantum-vulnerable public-key algorithms. It is a standards-transition proposal, not a universal statutory ban.

Sources: [38, 37].

Export controls and remote access

In September 2024, the US Commerce Department introduced controls on specified quantum computers and related equipment, materials, software, and technology. Applicability depends on classification, destination, parties, authorizations, and exceptions. It should not be summarized as a blanket restriction on all quantum knowledge or collaboration.[39]
The House passed H.R. 2683, the Remote Access Security Act, in January 2026; the January 13 record shows referral to the Senate. The proposal would extend export-control authority to specified remote access. That legislative step alone neither establishes enactment nor proves that particular quantum-cloud services are subject to an implemented prohibition.[40]

The CNSA dates above follow the December 2024 FAQ, version 2.1. Its exceptions and system scope remain important. For an operational compliance decision, consult the current controlling policy and implementing rules.

What could go wrong

Press-release physics outpaces peer review

A vendor announcement establishes what a company reports. A preprint supplies methods but may change before publication. Peer review does not establish universal applicability or independent replication. Assess the measured observable, protocol, assumptions, and available data before adopting a headline claim.

Scaling can expose a new bottleneck

Excellent fidelity or coherence on a small device does not guarantee that calibration, crosstalk, wiring, yield, optical routing, atom loss, cryogenic control or decoder throughput will remain manageable as the system grows.

Error-correction overhead may dominate practical feasibility

Useful fault-tolerant algorithms may require far more physical qubits, syndrome cycles, classical decoding and non-Clifford resources than headline logical-qubit targets suggest. Roadmaps should therefore be judged by logical error rates and executable logical operations, not physical scale alone.

Classical methods can catch up

Better GPU implementations, tensor networks, specialized solvers or new algorithms can erase an apparent quantum speedup after publication. Advantage claims must be re-tested against a moving classical frontier.

Supply-chain concentration and correlated roadmaps

Specialized refrigerators, lasers, detectors, foundries, materials and packaging can create single-source dependencies. At the same time, many technical roadmaps assume several milestones arrive on schedule; delays in error correction, networking or manufacturing can therefore invalidate multiple downstream roadmaps at once.

What progress should look like

For computation, track logical error rates, supported logical operations, acceptance probability where relevant, and executable algorithm depth. For sensors, track field performance and stability. For networks, track delivered service, rates, fidelity, and trust assumptions. These quantities make roadmaps testable.

Evidence discipline and reproducibility

Sources are matched to the claims they can support

1. **Technical papers and preprints** support experimental methods, reported observables, algorithm definitions, and stated assumptions. Record the version and distinguish a demonstration from an extrapolation.
2. **Standards and official policy documents** support their stated requirements, dates, scope, and legal or draft status.
3. **Official documentation** supports an API or product's described capabilities as of the cited version or access date.
4. **Vendor announcements and roadmaps** support attribution of a claim or target. They do not independently validate performance, comparative advantage, or delivery dates.

Minimum information for a quantitative comparison

- Hardware, date, device configuration, and the units being counted.
- Gate or memory protocol, measurement method, uncertainty, and the difference between representative and record performance.
- Error correction, mitigation, postselection, and all associated resource costs.
- Task definition, input and output requirements, approximation error, and success probability.
- Classical baselines, implementation effort, and total runtime or cost under comparable conditions.

Numbered references identify the evidence for the main numerical milestones and technical claims. Broad taxonomies and the evidence ladder are editorial synthesis. A claim of a complete fault-tolerant computing stack requires assessable methods and results for its protected operations, error handling, and executable depth.

Historical figures are not silently combined with later paper versions. The IBM sampling comparison records both versions; the Majorana result retains its measured parity observable. Specific results should not be generalized to an entire modality without additional evidence.

Sources and technical documentation

- [1] Google Quantum AI and Collaborators. *Quantum error correction below the surface code threshold*. Nature 638, 920–926 (2025); published online 9 December 2024.
- [2] R. Mandelbaum et al., IBM. *How IBM will build the world's first large-scale, fault-tolerant quantum computer*. 10 June 2025. Vendor roadmap.
- [3] IBM Quantum. *How to measure the performance of a quantum computer*. 2026. Vendor metric definitions.
- [4] H. Haas, D. McKay, and R. Davis, IBM. *IBM Quantum Nighthawk r2—more circuits, faster*. 31 August 2026. Vendor hardware report.
- [5] S. Martiel et al. *Sampling hard circuits with verifiably high fidelity*. arXiv:2607.25941v1, 28 July 2026. Preprint.
- [6] S. Martiel et al. *Sampling hard circuits with verifiably high fidelity*. arXiv:2607.25941v3, 2 September 2026. Revised preprint.
- [7] M. Aghaee et al. *20 Second Parity Lifetime in an InAs–Pb Tetron Device*. arXiv:2606.03884v2, 3 June 2026. Preprint.
- [8] Microsoft Quantum. *Majorana 2—Microsoft's Scalable Quantum Processor With Reliable, Long-Lasting Qubits*. 2026. Vendor interpretation and roadmap.
- [9] Microsoft. *Microsoft unveils Majorana 1*. 19 February 2025. Vendor announcement.
- [10] PsiQuantum. *Technology*. Architecture and cryogenic infrastructure; accessed 14 October 2026.
- [11] IonQ. *2025 Annual Report, Form 10-K*. Filed 2026; trapped-ion technology and operating environment.
- [12] D. Bluvstein et al. *Logical quantum processor based on reconfigurable atom arrays*. Nature 626, 58–65 (2024).
- [13] A. Noiri et al. *Fast universal quantum gate above the fault-tolerance threshold in silicon*. Nature 601, 338–342 (2022).
- [14] R. Kleiner, D. Koelle, F. Ludwig, and J. Clarke. *Superconducting Quantum Interference Devices: State of the Art and Applications*. Proceedings of the IEEE 92, 1534–1548 (2004); DOI: 10.1109/JPROC.2004.833655.
- [15] J. Koch et al. *Charge-insensitive qubit design derived from the Cooper pair box*. Physical Review A 76, 042319 (2007).
- [16] QuSpin. *Technology—Measuring Magnetic Fields with Atoms*. Product technology documentation; accessed 14 October 2026.
- [17] Z. Cai et al. *Quantum Error Mitigation*. Reviews of Modern Physics 95, 045005 (2023); arXiv:2210.00921.
- [18] IBM Quantum. *Qiskit SDK API reference*. Official documentation; accessed 14 October 2026.
- [19] Google Quantum AI. *Cirq*. Official documentation; accessed 14 October 2026.
- [20] PennyLane. *Gradients and training; TensorFlow interface*. Documentation v0.45.1; maintenance change from v0.44.

- [21] Microsoft Learn. [Set up the Quantum Development Kit extension](#). Official QDK, Python, and VS Code documentation.
- [22] Quantinuum. [TKET documentation](#). Official documentation; accessed 14 October 2026.
- [23] NVIDIA. [CUDA-Q documentation](#). Official documentation; accessed 14 October 2026.
- [24] Amazon Web Services. [Constructing circuits in the SDK—Amazon Braket](#). Official documentation; accessed 14 October 2026.
- [25] HSBC. [HSBC demonstrates world’s first-known quantum-enabled algorithmic trading with IBM](#). 25 September 2025. Company report.
- [26] IBM. [Modeling realistic chemistry with quantum computing](#). Case study with Cleveland Clinic and RIKEN; accessed 14 October 2026.
- [27] Quantinuum. [Quantum Volume](#). Vendor account of System Model H2’s September 2025 result.
- [28] T. Tomesh et al. [SupermarQ: A Scalable Quantum Benchmark Suite](#). arXiv:2202.11045 (2022).
- [29] IonQ. [How To Read an Algorithmic Qubits \(AQ\) Chart](#). Updated 8 January 2025. Vendor benchmark definition.
- [30] Microsoft Quantum. [Reliable Quantum Operations Per Second](#). Proposed system-performance metric; accessed 14 October 2026.
- [31] A. Myers, Stanford University. [Scientists achieve breakthrough on quantum signaling](#). 2 December 2025. Institutional report of optical-device research.
- [32] D-Wave. [D-Wave Demonstrates First Scalable On-Chip Cryogenic Control of Gate-Model Qubits](#). 6 January 2026. Vendor engineering announcement.
- [33] IonQ. [Roadmap](#). Vendor targets; accessed 14 October 2026.
- [34] DARPA. [From fragile to field-ready: RoQS program launches first phase](#). 27 August 2025. Program announcement.
- [35] European Commission. [European Quantum Communication Infrastructure \(EuroQCI\)](#). Official program description; accessed 14 October 2026.
- [36] NIST. [NIST releases first 3 finalized post-quantum encryption standards](#). 13 August 2024; FIPS 203, 204, and 205.
- [37] NIST. [Transition to Post-Quantum Cryptography Standards](#). NIST IR 8547, initial public draft, November 2024.
- [38] National Security Agency. [The Commercial National Security Algorithm Suite 2.0 and Quantum Computing FAQ](#). Version 2.1, December 2024; U/OO/194427-22, PP-24-4014.
- [39] National Quantum Initiative. [Department of Commerce releases export controls on quantum technologies](#). September 2024; summary linking to the BIS rule.
- [40] US Congress, Government Publishing Office. [H.R. 2683—Remote Access Security Act](#). 119th Congress, referred in Senate, 13 January 2026. Legislative text and status for that action.