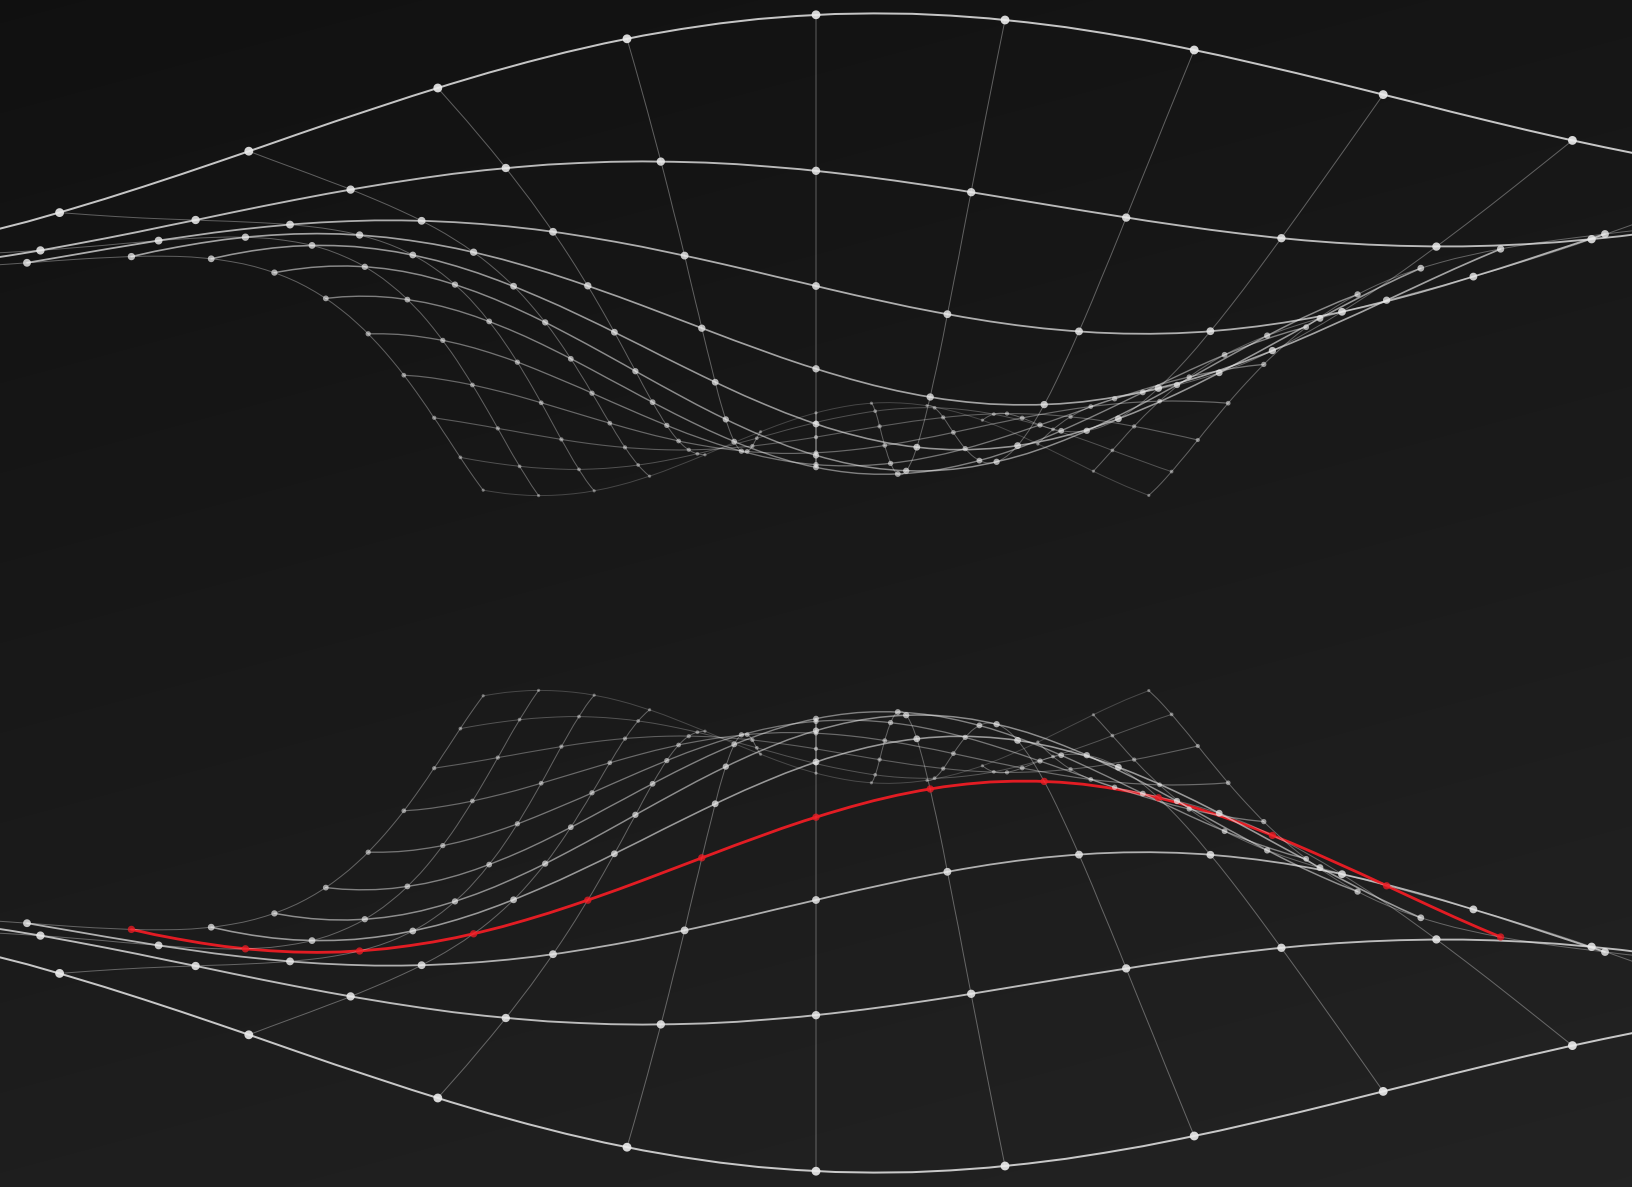


10 | Quantum Processor Benchmarking

It remains difficult to assess the performance of quantum computing today, especially for non-experts, which poses a challenge for making informed decisions about investment, commercial deployment, use case testing, and overall strategy. The technology is still nascent, several competing modalities exist, there is no independently verified universal performance metric, and the link between a hardware device and a given quantum algorithm is highly context-dependent.

Benchmarks

To give better insight into where quantum computing performance actually stands, we indexed and analyzed published data on nearly 300 Quantum Processing Units (QPUs) from 17 countries, spanning retired, prototype, currently available, and announced QPUs. As of August 2026, we observed 56 commercially available QPUs from 24 manufacturers.[1]

It is worth stating up front that most experts agree that quantum computers remain below the capabilities of classical computers for all useful tasks today. A helpful way to think about QPU benchmarking is to consider racecars. A racecar benefits from high horsepower and torque, light construction, independent suspension, and good aerodynamics, but lap time on a specific circuit is the best single measure of how those elements work together. QPUs have comparable construction, design and performance idiosyncrasies. Three main categories of benchmarks are useful to consider:

Physical benchmarks: e.g., qubit count and the fidelity of qubit gates.

These are the core metrics of a QPU, akin to a racecar's weight and torque. They are objective, but in isolation they reveal only part of the likely overall performance. Manufacturers are most willing to disclose physical benchmarks, and they are the focus of the analysis below.

Aggregated benchmarks: e.g., Quantum Volume, CLOPS, and Logical Qubits.

These combine physical benchmarks, much as power-to-weight ratios do for cars. They are more informative than any single measure but still fall short of capturing a QPU's full performance.

Application-level benchmarks: e.g., Q-Score and RACBEM.

These measure how QPUs perform on specific problems, in the way LINPACK is used to rank classical supercomputers. They allow QPUs to be compared with one another and with classical devices, and are analogous to a racecar's lap time at a given track in defined conditions. As with lap times, different benchmarks reward different algorithmic strengths: a Formula 1 car set up for Monaco would be ineffective on a NASCAR oval.

Manufacturers still rarely publish application-level benchmarks, as today's QPUs cannot run sizable applications. We expect to track them in future reports as devices become more capable.

Benchmarks

Our QPU dataset was assembled primarily from manufacturer announcements, publicly available datasets, online searches, and earlier direct queries to providers. It captures a range of benchmarks, of which the most important are:

Qubit Counts: Qubit count can mislead, because error rates tend to rise alongside it. This has at times prompted manufacturers to reduce qubit counts to improve overall performance. IBM followed its 1,121-qubit Condor in 2023 with the higher-performing 133-qubit Heron, and continued in that direction with the 120-qubit Nighthawk in 2025, a device built for circuit quality rather than scale. In short, qubit count alone is not a definitive measure of capability and must be read alongside other benchmarks.

Coherence: This refers to how long a qubit holds its quantum state. Through interaction with their environment, qubits inevitably lose their quantum information, a process known as decoherence, characterized by two timescales: T_1 (energy relaxation) and T_2 (dephasing). Both matter because they determine the execution time for a calculation. Trapped-ion and, to a slightly lesser degree, neutral-atom modalities show T_2 times several orders of magnitude longer than superconducting qubits, an inherent advantage for applications needing longer coherence.

Fidelity: Quantum computers still experience error rates orders of magnitude higher than classical systems. These errors originate in imperfect control pulses, inter-qubit couplings, and qubit state measurement, and reflect the engineering limits of current hardware. Several benchmarks characterize quantum error rates: single-qubit gate fidelity, two-qubit gate fidelity, readout fidelity, state preparation and measurement (SPAM) error, decoherence-related errors, and crosstalk error. Manufacturers describe these differently (mid-circuit, median, or average), which hinders direct comparison, and the problem has grown rather than receded as more vendors publish. Each of these benchmarks captures inaccuracies in a specific quantum operation essential to running a quantum circuit.

For instance, two-qubit gate fidelity remains among the most critical metrics, as it is usually the bottleneck in large circuits, where two-qubit gates are both more prevalent and more error-prone than single-qubit gates.

Quantum Volume (QV): Introduced by IBM in 2017, QV reflects several physical-level benchmarks at once, including gate fidelity, qubit count, and connectivity. Unlike the volume of a cube, it is not a simple multiplication but the result of a set of statistical tests. QV identifies the largest square-shaped random circuits, where qubit count equals circuit depth, that a device can implement with high fidelity. It has been criticized by scholars and by IBM itself as unhelpful for larger devices and for relying on square circuits that poorly represent real applications. Despite this, many hardware providers have adopted it, including for spec sheets and marketing.

Gate speed: This is the time taken to perform a single-qubit or two-qubit gate operation. It connects directly to decoherence, since together they determine how many operations a system can complete before the qubit becomes ineffective. Systems with fast gate speeds, such as superconducting qubits and electron spins, generally have shorter coherence times than trapped ions or neutral atoms, which have slower gates but far longer coherence times. Superconducting quantum computers operate at raw gate speeds in the 1–100 MHz range, and 1–10 kHz when fully burdened with error correction and overhead, significantly slower than classical CPUs at 2–5 GHz, though the comparison is only partly useful given the fundamentally different computational approaches.

Execution time

Gate speed is critical but under-reported, and many hardware vendors do not disclose it at all. It nonetheless sets a hard limit on circuit runtime. Consider molecular simulation using a Quantum Phase Estimation algorithm, which can require circuits exceeding 10^{13} logical gates. On a trapped-ion processor with gate speeds around 10 microseconds, running such a circuit once would take several days. Since a single molecule may need thousands of full executions for statistical confidence, and since error correction significantly increases circuit depth and gate count, total runtime could stretch into years, well beyond practical limits for most applications. Businesses evaluating quantum computing should estimate execution time from circuit size, hardware gate speed, and error-correction overhead. Gate speed imposes a fundamental limit, but total runtime can be reduced by optimizing algorithms for parallelism, reducing circuit depth, and improving qubit fidelity to lower the cost of error correction.

Error correction

Error correction is fundamental to quantum computing. Methods such as surface codes require growing numbers of physical qubits to form a single logical one: thousands, or tens of thousands for very large circuits. In 2024, Google announced a relevant breakthrough, demonstrating a system operating below the fault-tolerance threshold, meaning that adding qubits and correction cycles produces a net decrease in logical error rates.

Composing a detailed QPU list remains challenging, made harder by vendors highlighting only their most favorable benchmarks. Trapped-ion gates are generally orders of magnitude slower than the fastest superconducting gates. IBM is still among the most transparent firms on QPU benchmarking, publishing most relevant metrics including data for individual QPU instances and even individual qubits.

Modalities

Physical quantum computing systems are built and operated using several distinct technical architectures, referred to as modalities. Although each modality utilizes a different method for encoding, processing and reading out quantum information, they achieve comparable functionality (such as gate-based or measurement-based models). Each comes with distinct strengths and limitations, measurable across metrics like qubit count, speed and fidelity. Below, the table outlines the leading commercial or prototype devices for each modality, where a clear frontrunner has yet to be established.

	Trapped Ions	Superconducting	Neutral Atoms	Photonics	Electron Spin / NV Center
Qubit Count	98	1,121	1,225	216	424
2Q Gate Fidelity	99.92%	99.70%	99.35%	93.80%	99.20%
1Q Gate Fidelity	99.998%	99.98%	99.90%	99.60%	99.995%
2Q Gate Speed (LogHz)	N/A	7.32	6.18	N/A	N/A
1Q Gate Speed (LogHz)	4.00	7.70	5.39	N/A	N/A
Coherence – T1 (μs)	10,000,000	305	1,000,000	N/A	N/A

Best Worst N/A

*Vendors report different statistics under different conditions, these figures mark the frontier of each modality, not a ranking of individual devices.

A quantum circuit is a sequence of operations, or quantum gates, that a QPU follows to solve a problem. It is the foundation of quantum algorithms, which use these circuits to process information.

Superconducting QPUs are electronic circuits made with the lithography techniques used in classical computing fabrication. They are cooled to millikelvin temperatures, which suppresses thermal noise and permits coherent quantum behavior. They excel in gate speed and qubit count with reasonable fidelities, but require extensive cooling. They remain the most common modality, accounting for half of commercially available QPUs.

Trapped-ion QPUs implement gate-based quantum computing using individual ions held by radiofrequency traps, with gates performed by lasers or microwaves that manipulate the ions' internal quantum states. Trapped ions offer high fidelity, coherence, and qubit connectivity, but have slower gate speeds. They have not yet scaled to large qubit counts: the leading trapped-ion device reached 98 qubits in 2025 without sacrificing its fidelity lead.

Photonic QPUs use photons as qubits, propagating through photonic integrated circuits containing linear optical elements such as beam splitters and phase shifters, which can implement certain quantum gates.

Neutral Atom QPUs use atoms, typically alkali or alkaline-earth metals, that are laser-cooled and confined in vacuum chambers by optical and magnetic trapping. The atoms are ultracold while the hardware operates near room temperature. Relative to other modalities, they show promise in high qubit counts but have slower gate speeds and lower fidelity, though development is rapid.

Electron Spin QPUs use the quantum state of single electrons as qubits, offering relatively long coherence times and the potential for high-fidelity control, with the added benefit of compatibility with existing semiconductor fabrication techniques. Scalable control and readout of large arrays remain active research areas.

Nitrogen-Vacancy (NV) centers in diamond are a promising solid-state platform, where a nitrogen impurity adjacent to a lattice vacancy hosts a localized electron spin used as a qubit. NV-center QPUs can operate at room temperature, unlike most other quantum computer hardware.

Nuclear Magnetic Resonance (NMR) QPUs use the nuclear spins of atoms within molecules as qubits, manipulated by radiofrequency pulses inside a strong magnetic field. Because they operate at room temperature and are comparatively simple to build, NMR systems produced some of the earliest experimental demonstrations of quantum algorithms. They are difficult to scale, however, since measurements are taken across a large ensemble of molecules and the usable signal weakens sharply as qubits are added, which has kept devices to a handful of qubits. Today they are used mainly for teaching and research.

Cats, Rails and Flux

Superconducting quantum processors are among the most promising modalities, and several sub-architectures are being pursued. Amazon and Alice & Bob are working on "cat qubits," which are harder to fabricate but improve qubit fidelity because only one type of error occurs: the phase flip (no bit flip error). Quantum Circuits Inc. takes a different approach with Dual Rail QPUs, adding physical redundancy to transmon qubits with a second cavity / rail to protect from errors. Atlantic Quantum uses Fluxonium qubits rather than traditional transmons; Fluxonium offers longer coherence and error-correction possibilities, again at the cost of fabrication complexity.

Annealers

A distinct class of quantum computer is the adiabatic quantum computer, or annealer, named after the metallurgical process. Quantum annealing rests on the adiabatic theorem, which holds that a quantum system remains in its lowest-energy state if its parameters change slowly enough and noise stays low. An optimization problem can be mapped onto an energy landscape of possible solutions, with the lowest energy the best solution; by annealing (adjusting the system parameters), the system is guided toward the lowest-energy state, which, if reached, yields the optimum. D-Wave produced the first commercial annealer in 2011 with 128 qubits. In this report, annealers are treated separately, as their architecture is not directly comparable to gate-based quantum computers. They reach much larger qubit counts but do not implement universal gate-based control, which restricts them to a narrower class of problems compared to gate-based QPUs.

Majorana Qubits

Another type of superconducting qubit is the Majorana qubit, which drew significant attention in early 2025. Microsoft has invested in the approach for over a decade and remains the primary industry player pursuing it. The design uses superconducting nanowires hosting Majorana zero modes, exotic quasiparticles predicted to appear at the ends of the wire under specific conditions. Microsoft's Majorana 1 QPU represents a milestone toward a Majorana-based processor[2], though the February 2025 announcement met skepticism from parts of the scientific community, as conclusive evidence for the topological nature of the zero modes remains under debate.

Logical Qubits

Most quantum algorithms assume a qubit is "perfect," behaving flawlessly throughout the algorithm. In reality, qubits are error-prone and short-lived, so many physical qubits are combined using error correction techniques such as surface codes to form a more stable unit known as a logical qubit. The term can mislead, since practical utility depends heavily on the size and complexity of the circuit the logical qubit can reliably support. To be viable for applications such as simulating complex molecules, a logical qubit would need to support circuits millions to billions of gates long, several orders of magnitude beyond current capability. When presented with a number of logical qubits for a QPU, the key follow-up question is what circuit depth it can support, which is essential for meaningful interpretation.

Quantum emulators

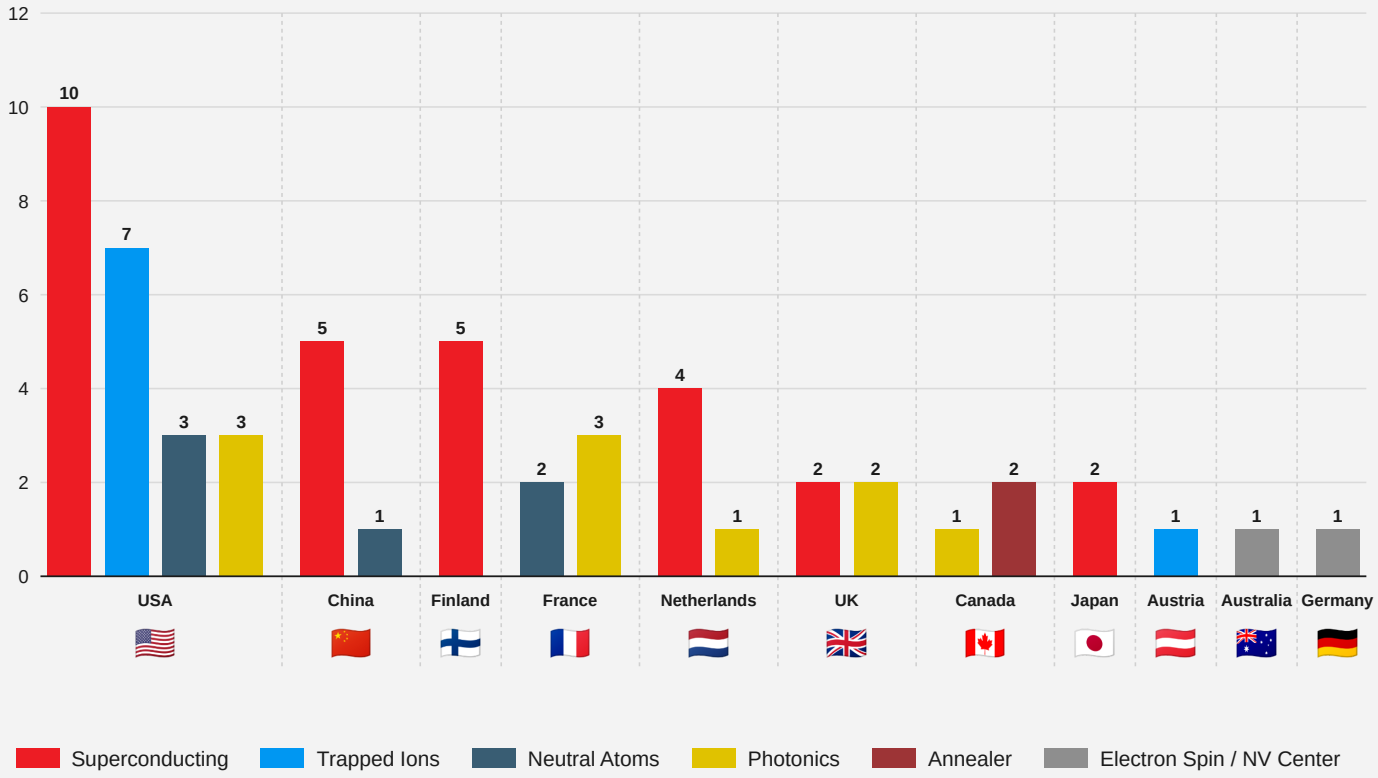
Quantum circuits can be simulated on classical computers, allowing quantum algorithms to be developed and tested without access to a quantum processing unit. Emulators do not physically use quantum effects such as entanglement or superposition. They are particularly useful for testing, debugging, and benchmarking quantum algorithms. Classical supercomputers can typically simulate general quantum circuits containing around 50 qubits, although the practical limit depends heavily on the circuit's structure and the simulation method used. For classical computers, each additional qubit is exponentially harder, whereas for QPUs it requires only incremental logical qubits. Today's best quantum computers remain orders of magnitude slower and more expensive to run than equivalent CPUs.

The number of commercially available QPUs globally now stands at 56, from 24 manufacturers. Countries are in a strategic race to achieve high-performance QPUs. The U.S. continues to lead this race in 2026.

10.1 | QPUs per country and modality

Fig. 10.a | Commercially available QPU models per country

MIT Quantum Index Report



Countries remain in a strategic race to field high-performance QPUs.

The number of commercially available QPUs globally is 56 from 24 manufacturers across eleven countries. The total number of commercially available QPUs is up from 40 in our 2025 report.

The United States retains a clear lead with 23 and is the only country present in four modalities, its superconducting fleet having grown from six devices to ten. China holds second place with 6 and is the only country to have added new manufacturers rather than new devices from existing ones.

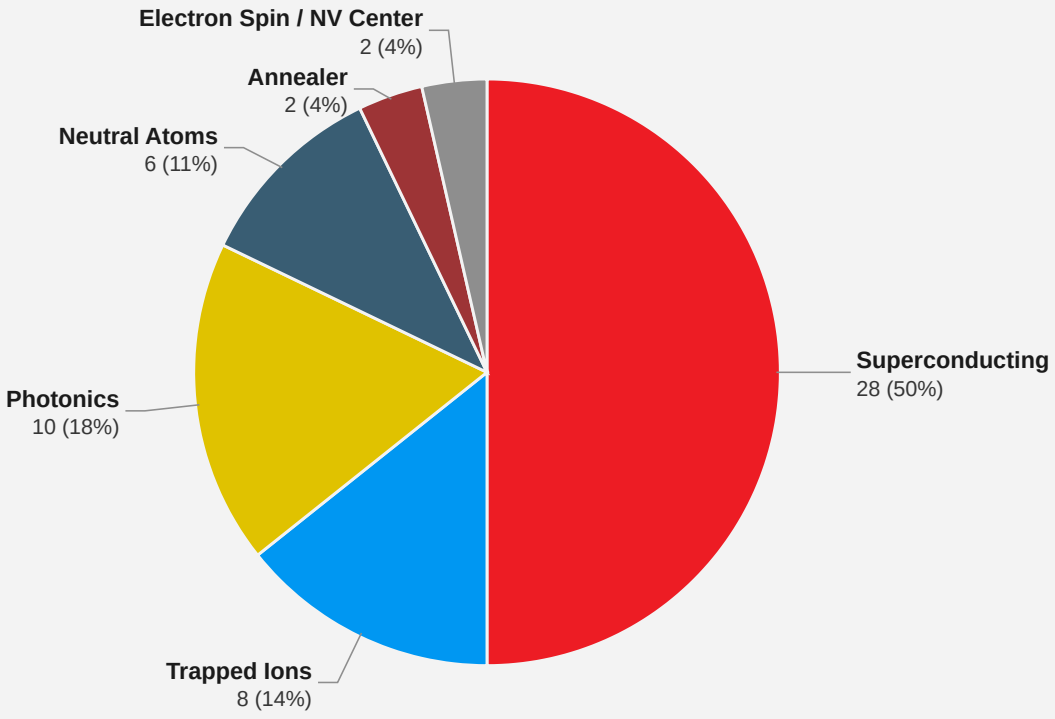
Below them, Finland, France and the Netherlands hold five each, the UK four, Canada three, Japan two, and Australia, Austria and Germany remain at one.

Data considerations: we list QPUs that are commercially available per country and per quantum computing modality. Country data is allocated based on the location of the manufacturer's headquarters as described in their official materials and website. QPUs are classified as commercially available if there is public access to the QPU either via on-premise or cloud. This also includes QPUs that may not be available on public clouds, but where access is provided to specific partner companies for commercial use. However, the device must be intended as a useful quantum computer for commercial use and not solely for experimental purposes. The number of QPUs is determined as uniquely differentiated products actively provided and marketed by the provider, e.g., IBM Eagle and IBM Heron are two distinct QPUs, but Eagle r3-Brussels and Eagle r3-Sherbrooke are considered as one QPU. On the same basis, a trapped-ion system newly listed on a public cloud in 2025 is recorded against the existing product rather than as a new one. The number of QPUs is not necessarily an indication of the progress of each country in quantum computing, as some manufacturers have made several very small QPUs available for basic academic research and teaching, while others have retired smaller but powerful QPUs from their offerings (e.g., IBM).

10.2 | QPUs per modality

Fig. 10.b | Commercially available QPUs per modality

MIT Quantum Index Report

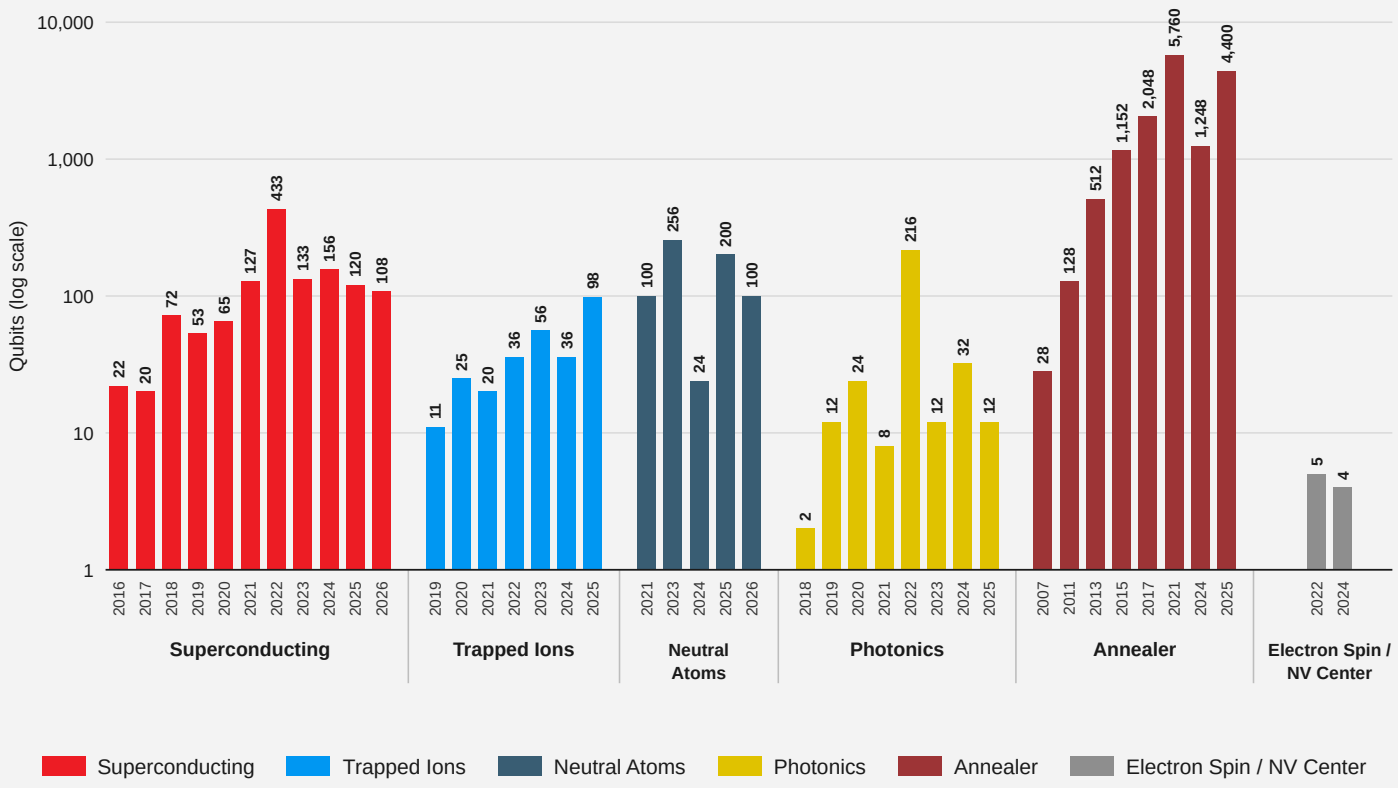


Superconducting remains dominant at 28 of 56 devices, exactly 50%, a share essentially unchanged (48% in 2025) despite the sample growing by 40%, which might be driven by the inherent manufacturing benefits and a historic head start in R&D. Photonics is next at 18%, followed by trapped ions at 14%. Neutral atoms hold 11%, annealers and electron spin / NV centers 3.6% each.

10.3.1 | Qubit count per modality

Fig. 10.c | Largest QPU released, prototyped, or planned per year

MIT Quantum Index Report



Superconducting counts peaked at 433 in 2022 and have generally declined since (133, 156, 120, 108), which might reflect a sustained shift toward error correction and fidelity rather than scale.

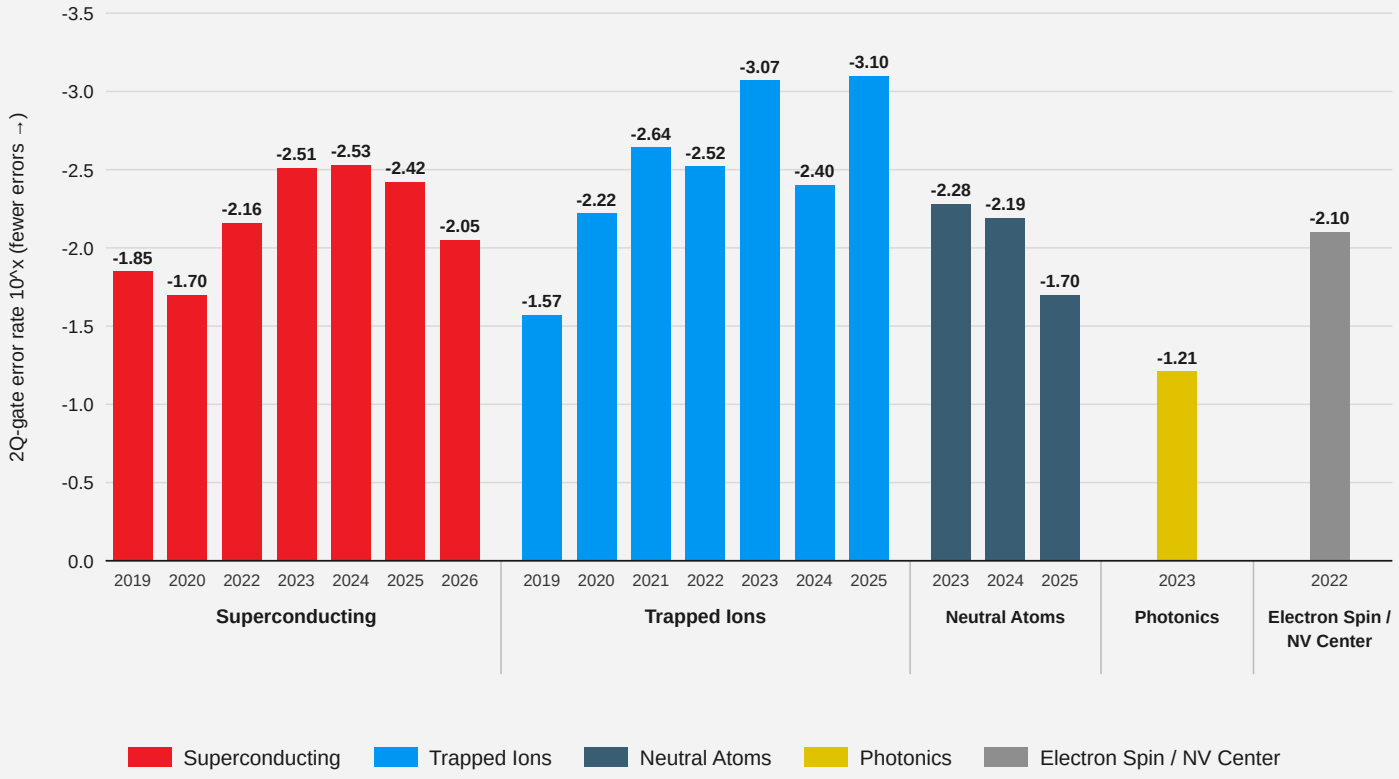
Trapped ions tell the opposite story, growing more consistently to 98 qubits in 2025, and remain on an upward trend. Neutral atoms reached 200 qubits in 2025 and 100 in 2026. Annealers sit at 4,400 for 2025, below the 5,760-qubit device of 2021. Photonics remains the least regular series, moving between 2 and 216 across eight years.

Data considerations: this graph shows the progression of the number of qubits in our dataset over time, considering only the largest QPU announced per modality in that year. For a QPU to be considered, it needed to be officially announced by a manufacturer and made commercially available in the given year (or expanded, e.g., Quantinuum H1 to H1-1). Retired QPUs are included, as they were available in the year shown. The data does not always show a steady increase, as some calendar years only contained new QPUs that were smaller than previously available.

10.3.2 | Fidelity per modality over time

Fig. 10.d | Best 2Q-gate error rate over time (commercially available QPUs)

MIT Quantum Index Report



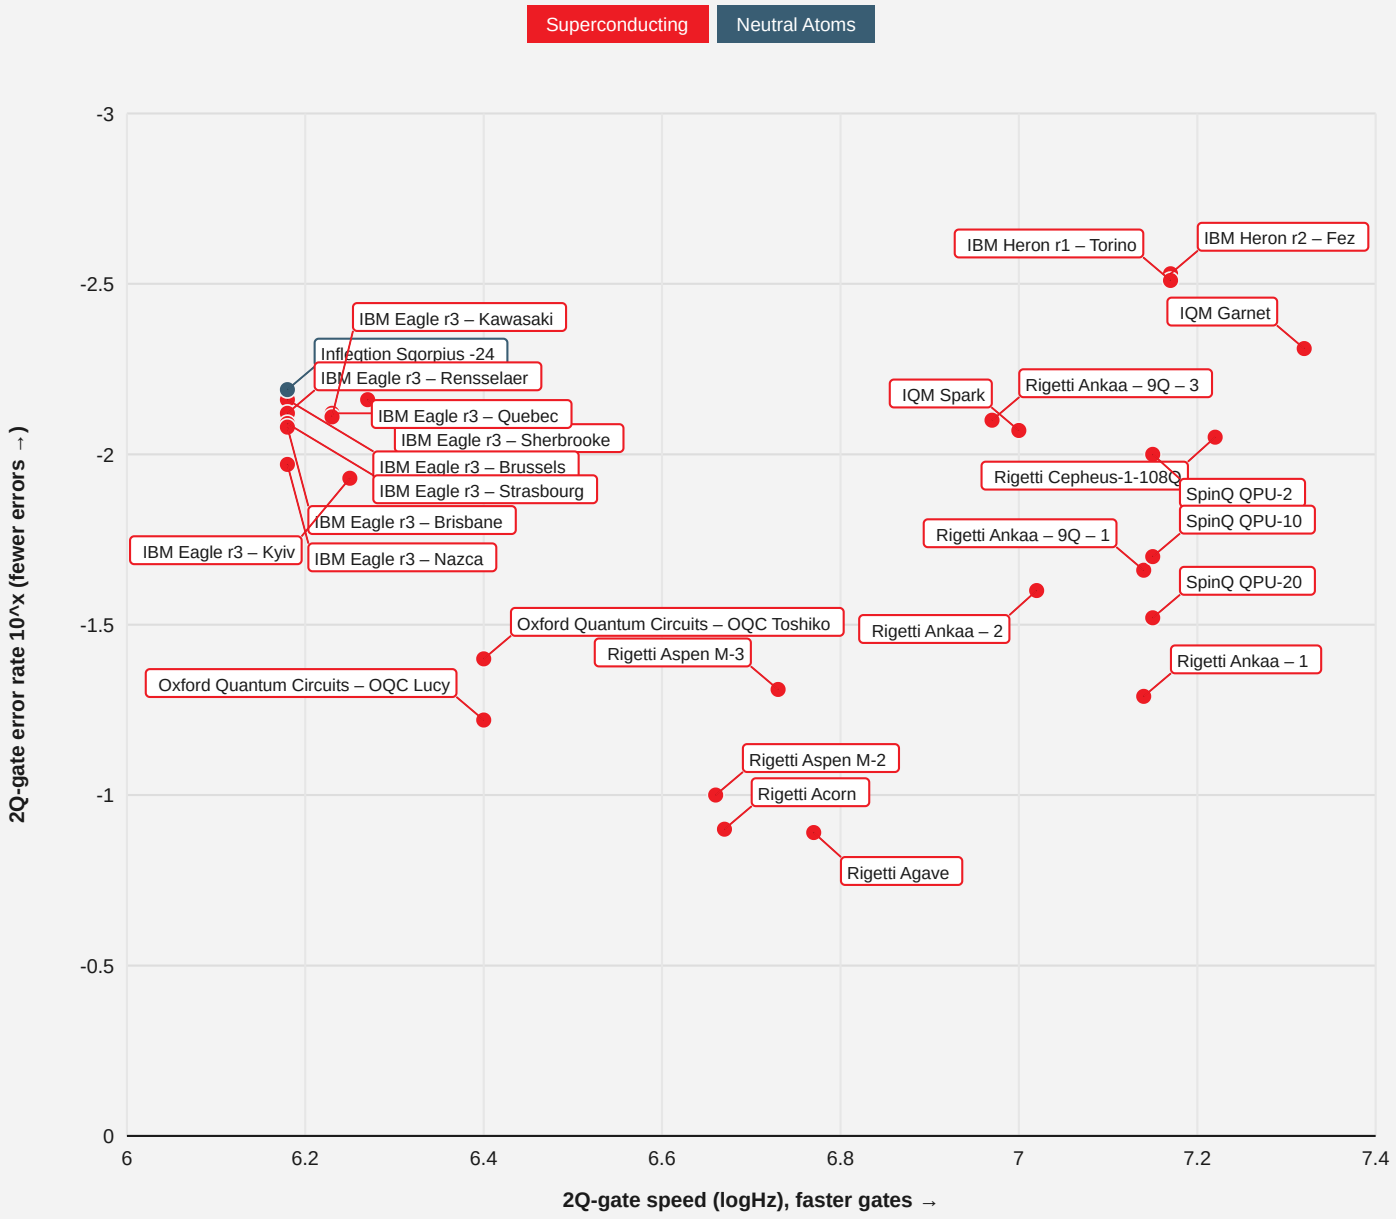
Fidelity for 2-Qubit is a key metric of performance improvement. Trapped ions have shown consistent growth and the highest overall fidelity.

Data considerations: this graph shows the progression of fidelity levels in new, commercially available QPUs in our dataset over time. We only considered the highest fidelity announced per modality per year, so a year's bar reflects that year's releases rather than the best result achieved to date. Error rates are given in a log10 scale, i.e. -3 translates to a 0.001 error rate which corresponds to a 99.9% or 0.999 fidelity. The noted error rates should be treated with caution as there are significant differences in the way they are measured for each QPU, e.g., mid-circuit vs first-gate-measurement, average vs median of several measurements across qubits, different gates (CZ, SWAP, etc.), and different iterations of the same QPU that give different values.

10.4 | Gate time versus gate fidelity

Fig. 10.e | Gate time versus gate fidelity (2-qubit)

MIT Quantum Index Report



To determine the maximum length of a circuit for a QPU, an important metric is the comparison of the speed of executing a single gate to how accurate this gate is (fidelity). For real-life scenarios such as Shor's algorithm for decrypting information utilizing RSA-2048, more than 10^{13} logical gates are required. Slow gate speeds at that size lead to calculation times of days or even months with some modalities. Speed alone is not the answer either: a fast gate that makes mistakes forces the machine to repeat work, so the useful QPUs are those that sit toward the right and the top of this chart at once.

Devices are shown for both one- and two-qubit gates; the two-qubit case matters more because such gates are more common in large circuits for most algorithms. The superconducting IQM Garnet is the fastest device to publish a figure, at 48 nanoseconds per two-qubit gate, with Rigetti's newly released Cepheus-1-108Q at 60 nanoseconds and IBM's Heron at 68. IBM Heron pairs its slightly slower gate with the best accuracy on the chart at 99.70%, so no single device now leads on both counts. Below them sits a group of processors running at comparable speeds but with fidelities between 94.9% and 99%, which illustrates the point directly: gate speed by itself says little about whether a machine can finish a long calculation correctly.

Notably, trapped ions are absent entirely and neutral atoms nearly so, as manufacturers tend not to disclose exact gate speeds, which are expected to be orders of magnitude slower than superconducting QPUs (as can be seen in the 1-Qubit graph).

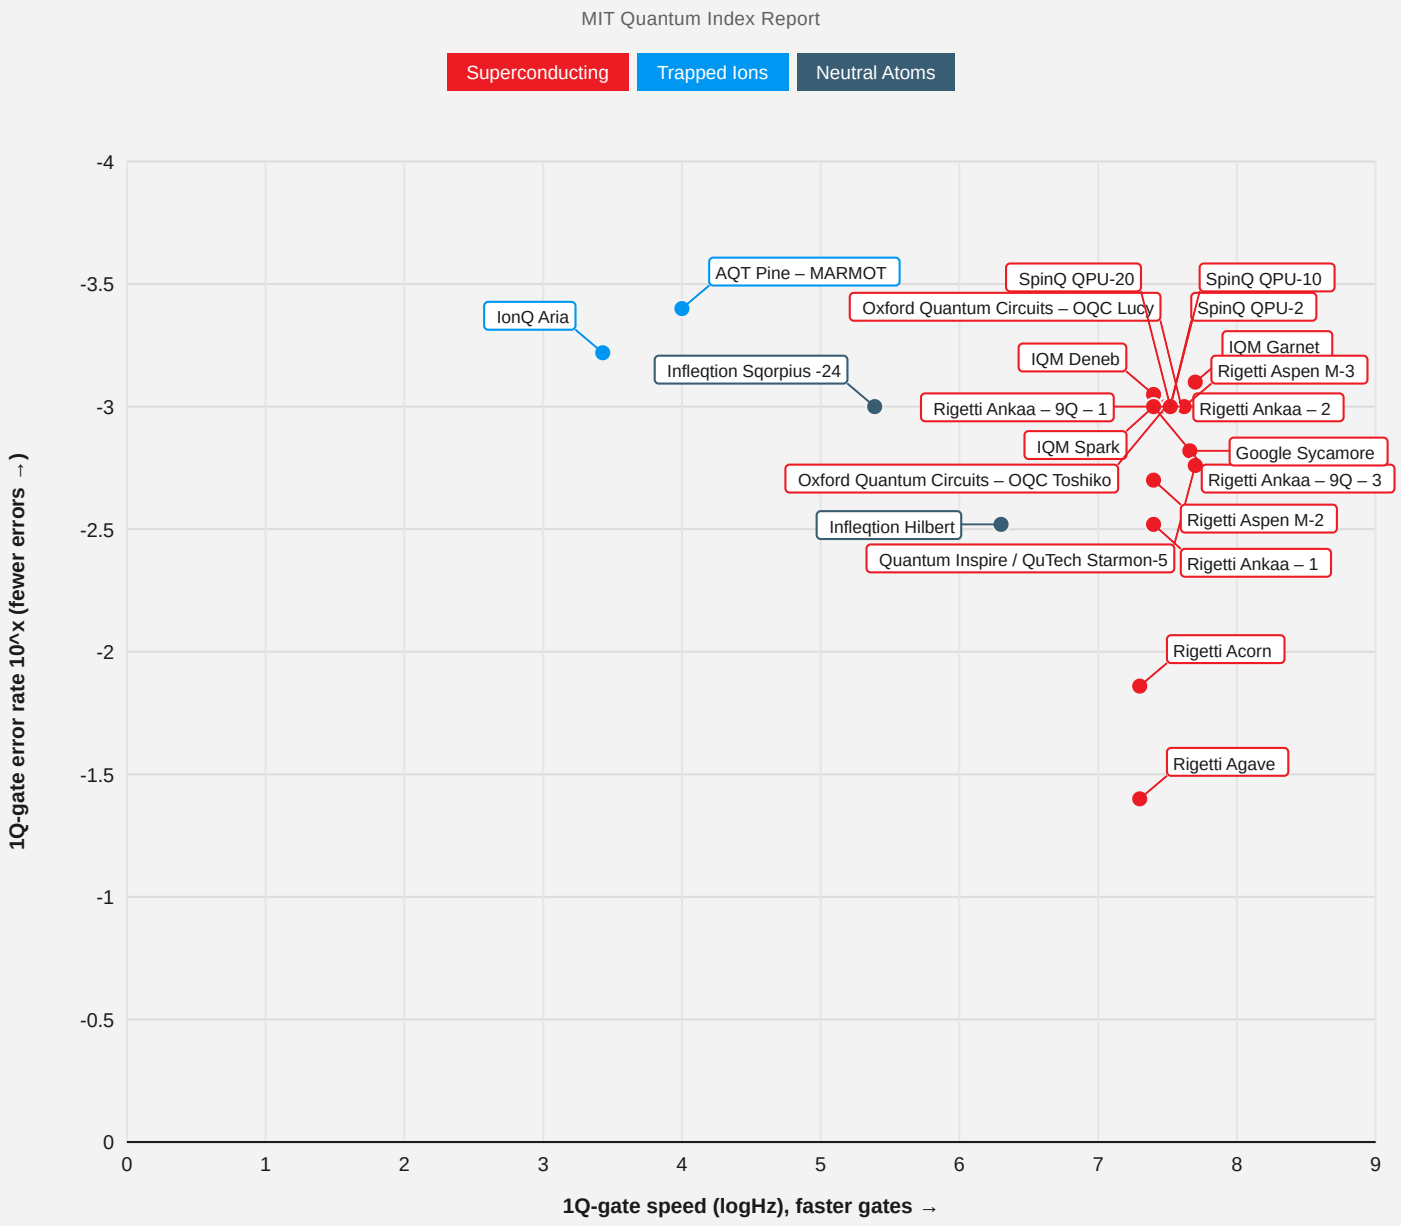
Of the 36 trapped-ion QPUs in our dataset, only three publish any gate speed at all, and none publish a two-qubit figure. The 1-Qubit graph shows what the gap looks like where the numbers do exist. Superconducting devices there run their fastest single-qubit gates in 20 to 32 nanoseconds, while the trapped-ion systems that publish a figure take around 100 microseconds – several thousand times longer, though those same systems hold their quantum states for seconds rather than fractions of a millisecond, which is the compensating advantage.

3 of 36

Of the 36 trapped-ion QPUs in our dataset, only three publish any gate speed at all, and none publish a two-qubit figure.

Data considerations: we chose 2-Qubit-Gates like CZ, CNOT, and SWAP as they are used for most quantum algorithms and make up the majority of gates for these circuits. The 2Q Gate time is the time required for the execution of a 2-Qubit gate and is given in Hertz (Hz) in a logarithmic scale, where higher values mean faster gate speeds, i.e., 7 logHz corresponds to a 100ns gate speed and 3 logHz to 1,000,000ns. The error rates are given in a log10 scale where higher is better, i.e., 10^{-3} translates to a 0.001 error rate, which is 99.9% fidelity. The data points with missing labels in the graphs are closely related QPUs (e.g., different instances of IBM Eagle). Roadmap devices are excluded, as a projected gate speed paired with a projected fidelity describes an intention rather than a measurement. To illustrate the performance comparison, we included the 1-Qubit Gate graph, which shows that these would likely land in the top-left quadrant, trading high fidelity for low gate speeds. Notably absent are ion traps and neutral atom QPUs, as manufacturers tend not to disclose exact gate speeds.

Fig. 10.f | Gate time versus gate fidelity (1-qubit)



The companion 1-Qubit graph now shows the missing modalities directly, with trapped ions at 3.4 to 4.0 logHz achieving the best fidelities on the chart, three to four orders of magnitude slower than superconducting.

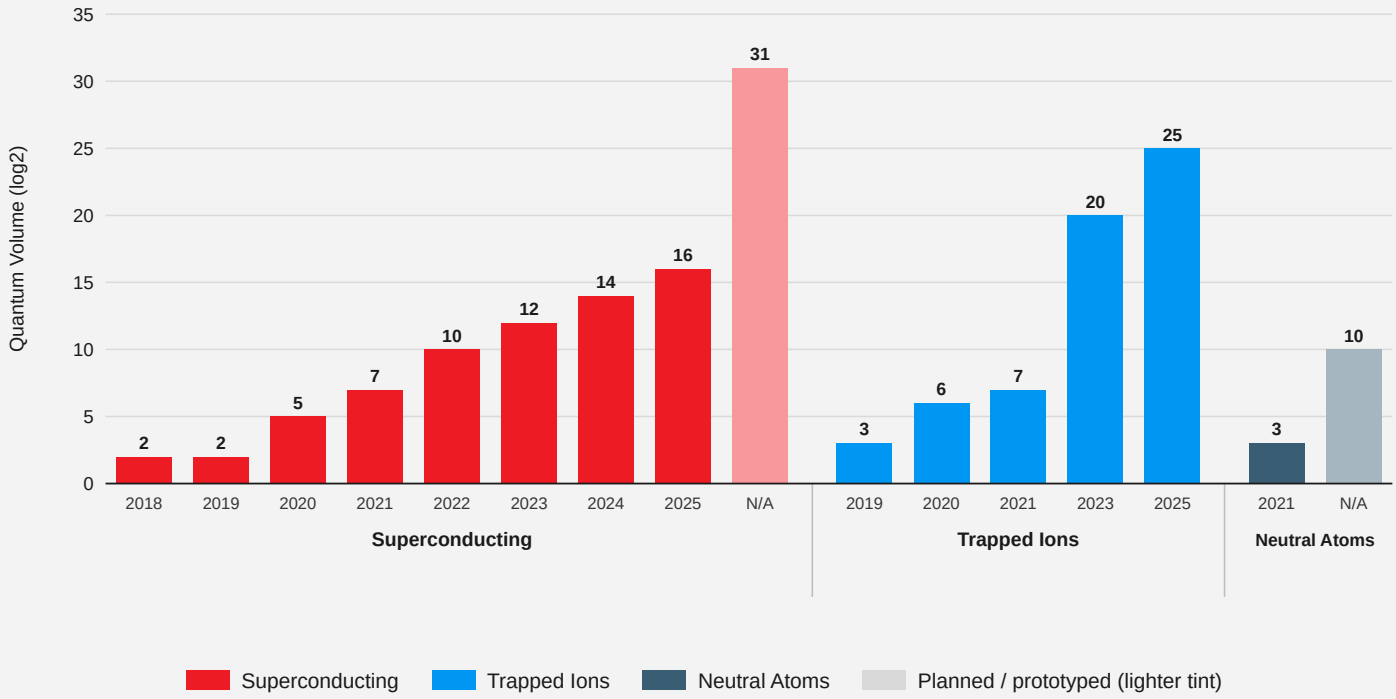
This is the direct expression of the relationship between gate speed and decoherence: fast gates and long coherence remain alternatives rather than a combination any modality has delivered.

Data considerations: we chose 2-Qubit-Gates like CZ, CNOT, and SWAP as they are used for most quantum algorithms and make up the majority of gates for these circuits. The 2Q Gate time is the time required for the execution of a 2-Qubit gate and is given in Hertz (Hz) in a logarithmic scale, where higher values mean faster gate speeds, i.e., 7 logHz corresponds to a 100ns gate speed and 3 logHz to 1,000,000ns. The error rates are given in a log10 scale where higher is better, i.e., 10^{-3} translates to a 0.001 error rate, which is 99.9% fidelity. The data points with missing labels in the graphs are closely related QPUs (e.g., different instances of IBM Eagle). To illustrate the performance comparison, we included the 1-Qubit Gate graph, which shows that these would likely land in the top-left quadrant, trading high fidelity for low gate speeds.

10.5 | Quantum Volume

Fig. 10.g | Max QPU Quantum Volume (log2) over time

MIT Quantum Index Report



As described in an earlier section, Quantum Volume (QV) is an aggregated metric designed to reflect a more holistic view of overall performance of a QPU.

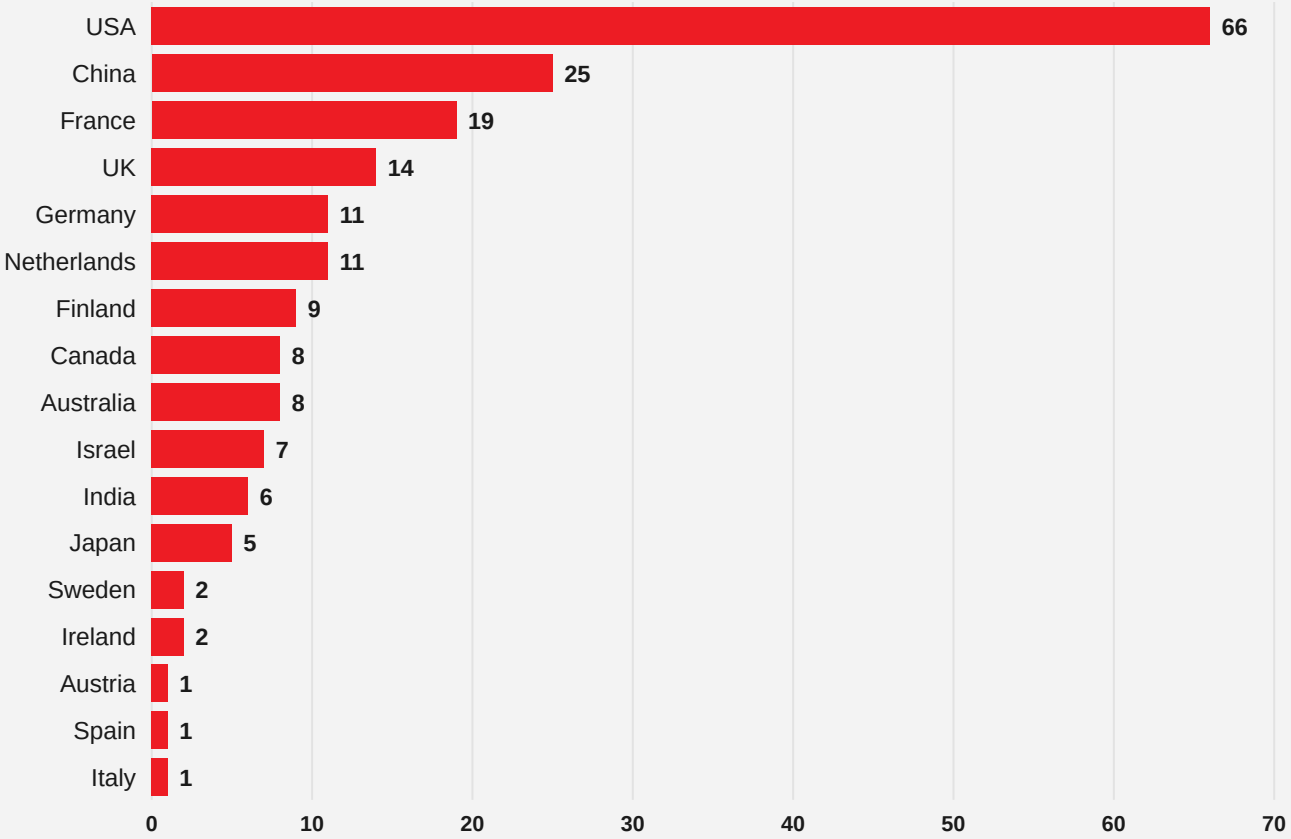
Although its usefulness is heavily debated, manufacturers have been reporting this metric. There is a substantial and consistent increase in QV performance of the top-performing QPU across modalities. However, the stated QV values are hard to validate, and in some cases the only data has been manufacturer claims.

Data considerations: The Quantum Volume (in log2 basis) is listed for each QPU at the given year. QV is itself a debated benchmark (even by IBM itself), which has been used less in the last few years. Despite all these caveats, we chose to include QV, as out of all aggregated metrics (like RACBEM, Algorithmic Qubit, CLOPS, etc.) it is the one that has published values for a sizable number of QPUs, and as such a progression on quantum computing capabilities can be roughly traced over time.

10.6.1 | A look into the future: QPUs per country and modality

Fig. 10.h | Planned, prototyped, and available QPUs per country

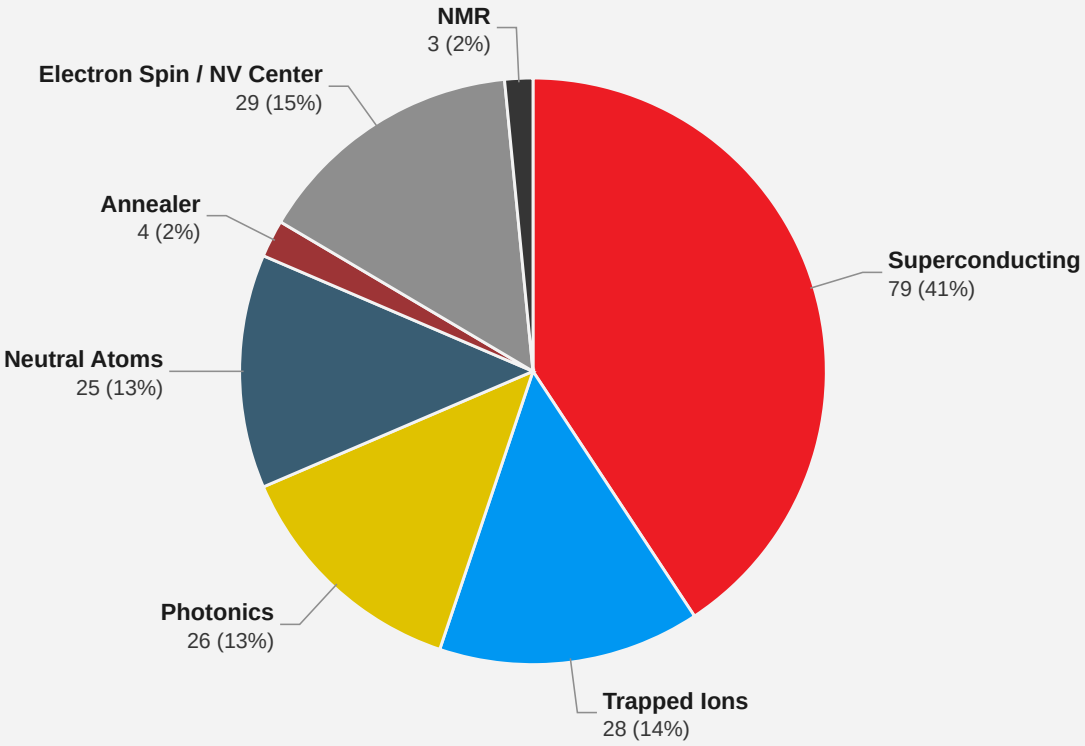
MIT Quantum Index Report



Quantum computing remains in the early stages of technical development. Numerous startups and established companies are actively developing prototypes and announcing roadmaps for future QPUs. A comprehensive overview of these efforts provides valuable insight into the evolving technological landscape and strategic directions within the field. The U.S. is the clear leader in the number and diversity of QPUs announced, with 66; China is a distant second with 25, followed by France with 19 and the UK with 14. Germany and the Netherlands have 11 each, Finland 9, and Canada and Australia 8 each. This data gives an approximate measure of country activity but is not an indication of QPU quality. It also reflects planned QPU announcements and does not fully reflect the probability of their successful release.

Fig. 10.i | Planned, prototyped, and available QPUs per modality

MIT Quantum Index Report

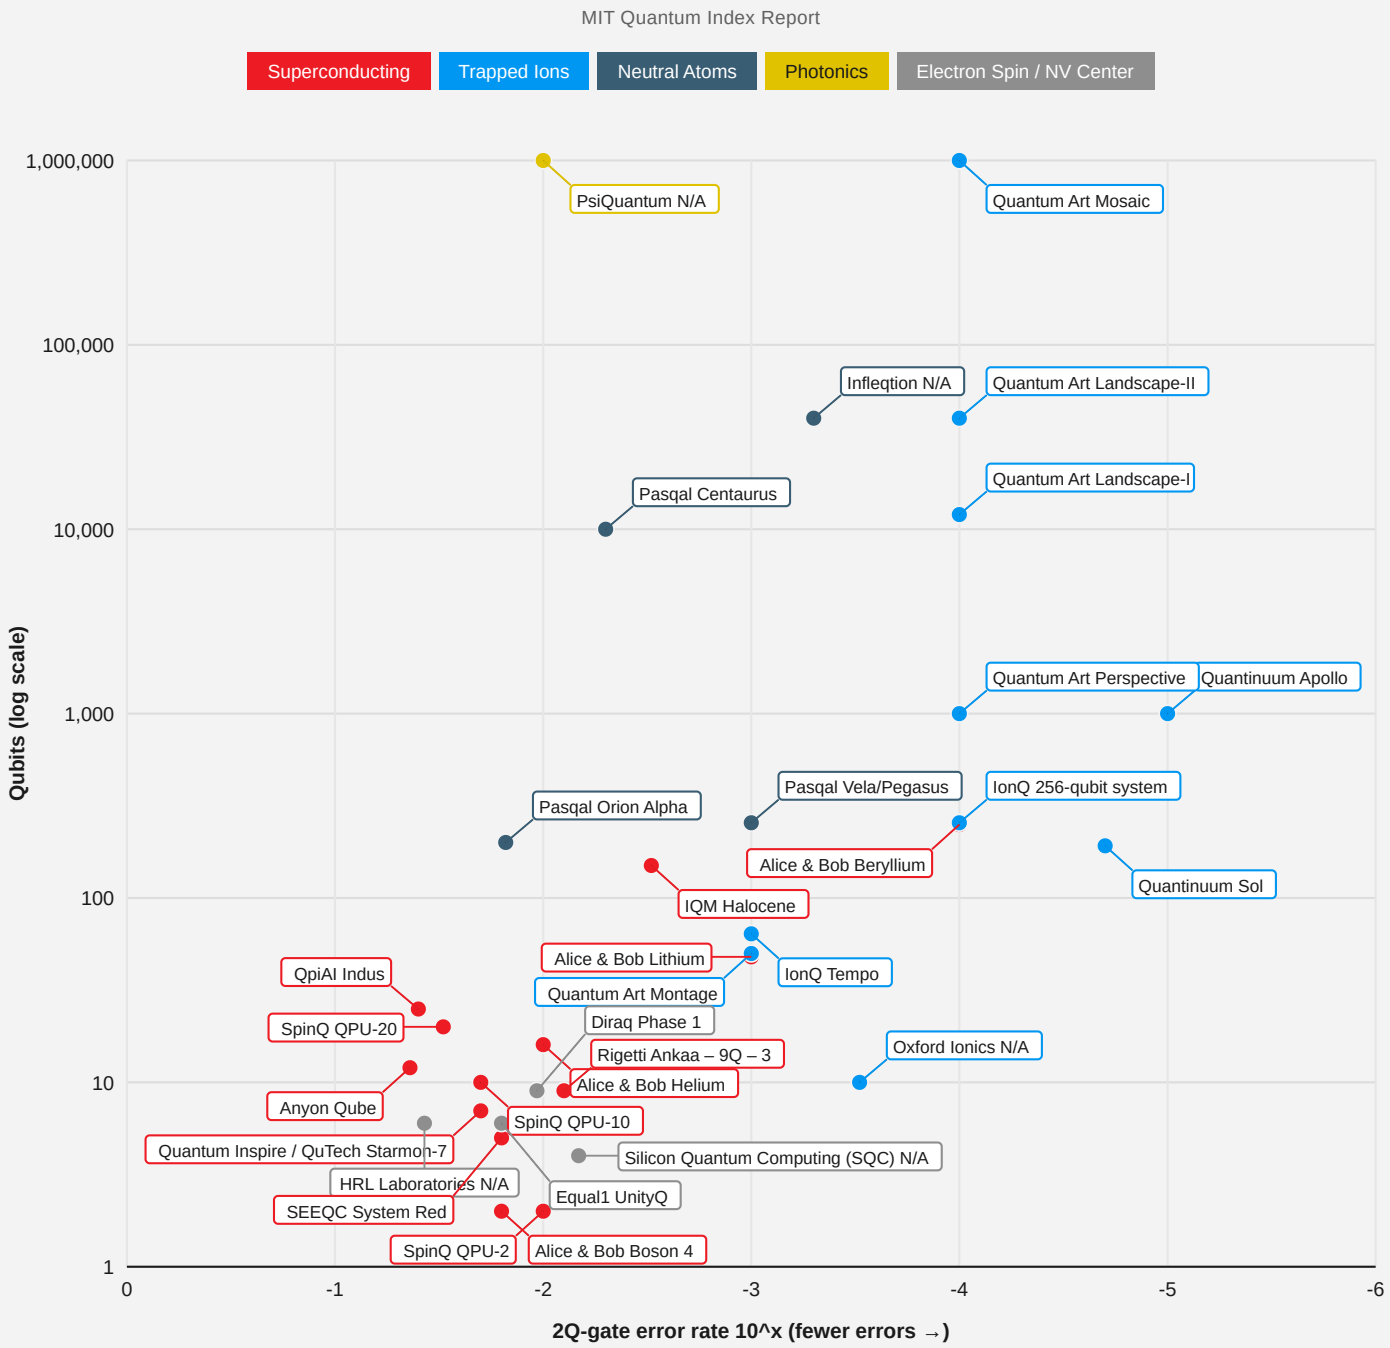


Looking at the future for different modalities, electron spin, NV Centers and neutral atoms are planned to become increasingly prevalent while NMRs and annealers are more stagnant and may be phased out. Photonics, superconducting, and trapped ion QPUs may have lower overall shares in the future due to the higher growth levels of other modalities. Superconducting retains the largest pipeline but a lower share of announced devices than of available ones, consistent with the expectation that other modalities will grow faster.

Data considerations: the data in these graphs includes prototype devices, which are not intended for commercial usage and are not available to the wider community of researchers but are used by a manufacturer for research to develop a new product. It also includes Future Planned QPUs, which are announced in a manufacturer roadmap or interview. Due to the dynamic evolution of startups in the space, recent announcements and changes in QPU roadmaps may not be fully captured in our dataset. The number of QPUs is not necessarily an indication of the progress of each country in quantum computing, as some manufacturers have made several very small QPUs available for basic academic research and teaching, while others have retired smaller but powerful QPUs from their offerings.

10.6.2 | A look into the future: qubit count and fidelity

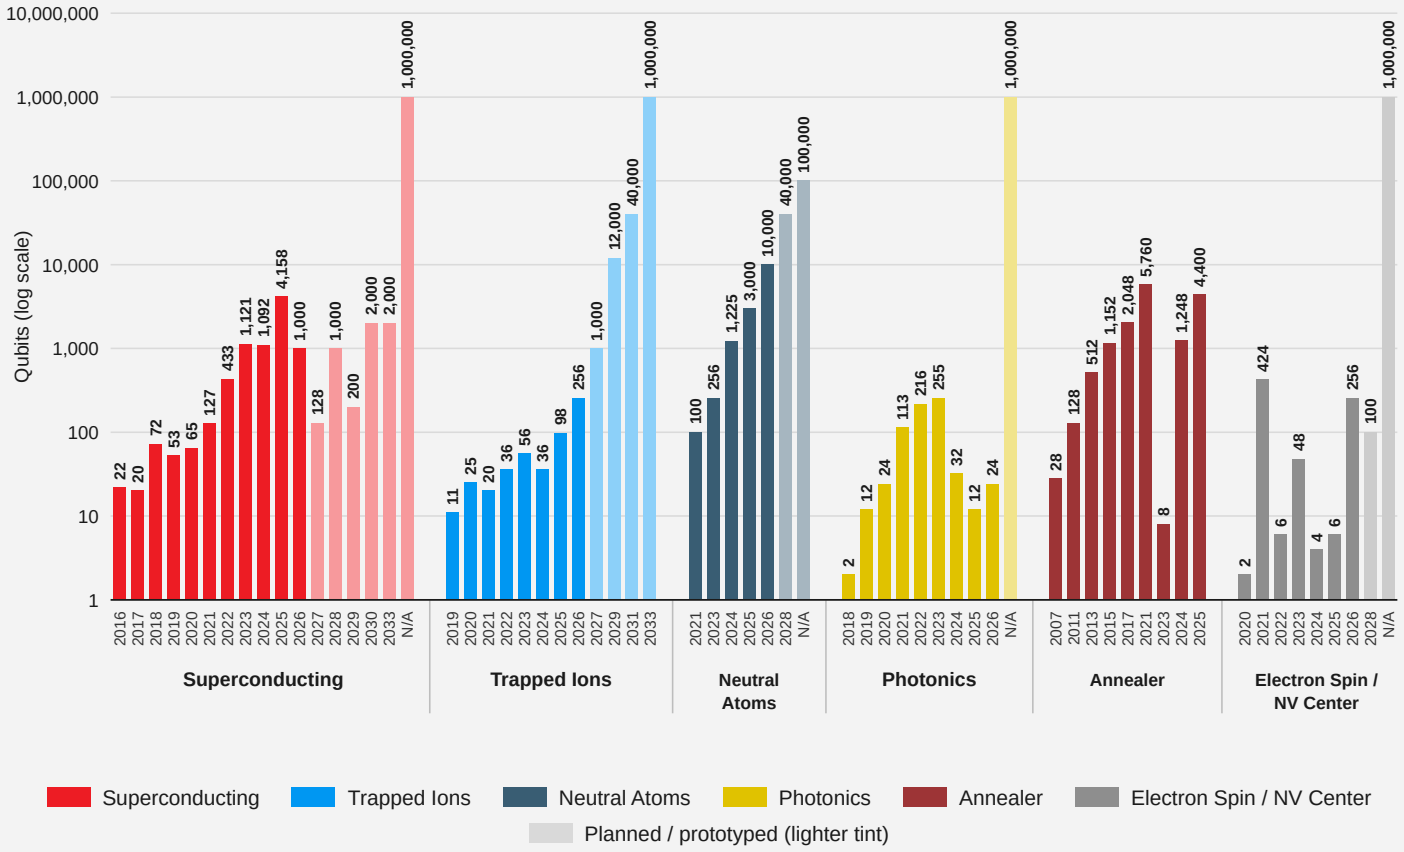
Fig. 10.j | A look into the future: qubit count and fidelity



QPU manufacturers are beginning to share more forward-looking roadmaps, offering insights into how their systems might compare across benchmarks. PsiQuantum leads in projected qubit count, Quantinuum excels in error rates, while Quantum Art and Infleqtion position themselves as balanced performers across both metrics.

Fig. 10.k | Largest QPU released, prototyped, or planned per year (incl. future plans)

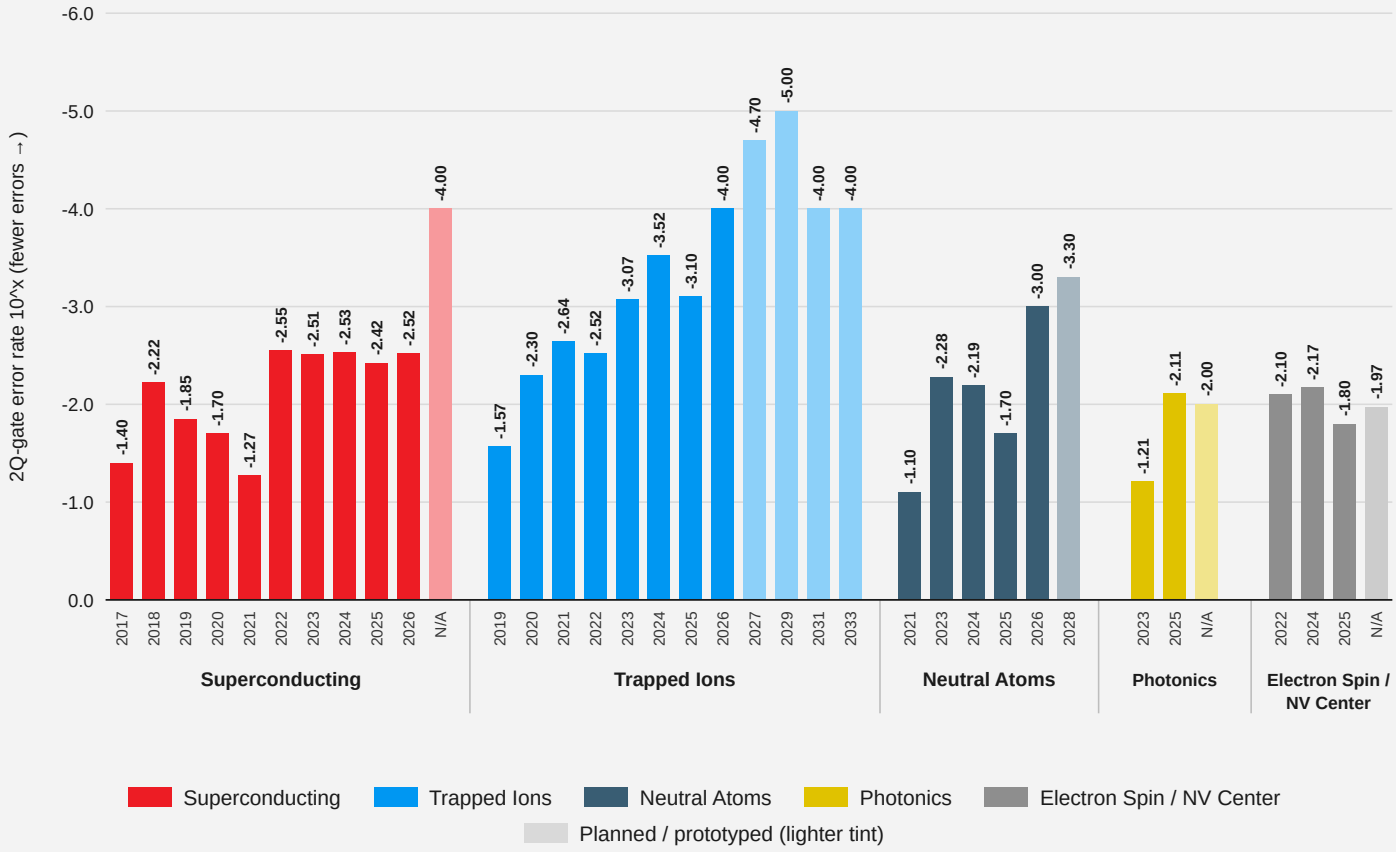
MIT Quantum Index Report



While qubit counts are expected to continue rising, the pace of growth is moderating. This reflects a shift in focus toward improving performance through better error correction and higher qubit fidelity rather than simply scaling up qubit numbers.

Fig. 10.l | Best 2Q-gate error rate over time (available, planned, and prototyped QPUs)

MIT Quantum Index Report



Trapped ion systems aim for exponential gains in fidelity and are on track to continue outperforming other modalities in that benchmark. Neutral atom platforms also show strong ambitions in this area, while other technologies appear more conservative in their likely fidelity trajectories.

Data considerations: these graphs show the progression of the fidelity and qubit counts for published QPUs over time, including any future plans, considering only the largest QPU announced per modality in that year. The data does not demonstrate a constantly increasing trend, as some calendar years saw smaller QPUs than previously available. Although approximately 60 manufacturers have announced approximately 95 future QPU models, only 19 QPUs have been provided with both target qubit count and fidelities, which is why there are fewer QPUs in the first graph, QPU vs 2-Qubit Fidelity. Devices announced without a target year appear in the “N/A” column at the end of each modality block; these hold the largest qubit counts in the dataset, so these would reflect the boldest projections that do not yet carry a delivery date.

10.7 | Future research

Benchmarking plays a vital role in advancing quantum technology, guiding strategic decisions, and establishing shared industry standards. This public overview is designed to increase transparency and spark meaningful dialogue across the community regarding how different systems perform. Because standardized data is still limited in this fast-moving field, we actively encourage industry stakeholders to share their metrics over time. Keeping this repository up to date will close existing data gaps while fostering greater collaboration and progress across the ecosystem.

Qubit count versus two-qubit gate fidelity graph is not included in this edition because vendors report this benchmark using different statistics, best pair, median, or mean, under different operating conditions, which leads to comparability issues. For future editions we look forward to working with the industry to present this data in a more detailed table that makes each vendor's reported statistic explicit.

You can reach us at contact@qir.mit.edu

Footnotes

[1] The gap between the 288 QPUs indexed here and the 56 counted as commercially available breaks down as follows. Of 288 records, 53 are separate instances of a QPU counted only once as a product (IBM Eagle r3-Brussels and Eagle r3-Sherbrooke are one QPU, while IBM Eagle and IBM Heron are two), leaving 235 uniquely differentiated products. Of those, 95 are planned but not yet released, 39 have been retired, 25 are prototypes not intended for commercial use, and 20 are available for research or teaching only. The remaining 56 are commercially available. These 56 come from 24 manufacturers headquartered in 11 countries, out of 86 manufacturers and 17 countries represented in the dataset as a whole.

[2] In June 2026, Microsoft followed with Majorana 2, a four-qubit array on a new InAs/Pb tetron material stack reporting parity lifetimes of around twenty seconds, roughly a thousandfold improvement on Majorana 1, and revised its fault-tolerance target from 2033 to 2029.

