QUANTINUUM

STRATEGIC BUYER'S GUIDE

A Strategic Guide to Selecting the Right Quantum Computing Solution

FOR ENTERPRISE AND PUBLIC SECTOR EXECUTIVES EVALUATING QUANTUM COMPUTING INVESTMENT

A Strategic Guide to Selecting the Right Quantum Computing Solution

For enterprise and public sector executives evaluating quantum computing investment

WHAT THIS GUIDE COVERS	Why quantum computing matters now, how to evaluate vendors objectively, how to design effective pilots, and how to take your first steps with confidence.
WHO IT IS FOR	CTOs, CIOs, CISOs, R&D leaders, procurement leaders, and program directors in enterprise and public sector organizations, at any stage of quantum familiarity.
HOW TO USE IT	Apply the evaluation framework to any vendor. A glossary of key terms is included at the back of this guide for readers new to the field.

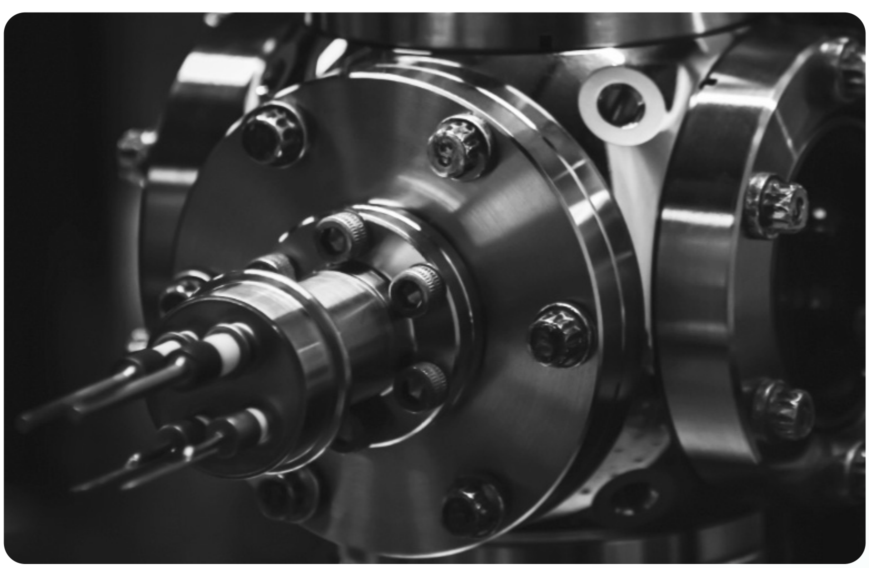

1. Why Quantum Computing – And Why Now?

Quantum computing has crossed a decisive threshold. After decades of theoretical promise, it is now delivering value across a range of industries. Organizations that engage with it today are building the institutional capabilities and proprietary insights that will define the next wave of technology-driven advantage.

As quantum computing begins to reshape industries, early movers will be best positioned to capture the lion's share of the value. For executive decision makers, the question is no longer whether quantum computing will matter. It is whether your organization will be positioned to benefit when it does — and that positioning begins now. However, the problem isn't so straightforward: quantum computing investments vary widely in value and potential. This guide aims to help readers maximize their probability of success in this rapidly evolving landscape.

A. Measurable business value today

Quantum computing is not a single technology that switches on at a defined moment of maturity. Value is being captured now, across multiple dimensions, by organizations that have engaged with intention. This value is taking the form of:

Operational learning and workforce readiness: Organizations that are engaging with quantum computing now are building internal expertise, developing quantum-aware talent pipelines, and establishing workflows that will scale as hardware capability improves. This institutional knowledge is genuinely difficult to acquire quickly. The organizations building it now will have a meaningful head start over those that wait.

Pilot programs with tangible outcomes: Leading enterprises are already running production pilots that are demonstrating quantifiable results that classical computers cannot replicate[1] .

B. Long-term strategic advantage for early movers

Beyond near-term ROI, organizations that engage with quantum computing early are building something more durable: strategic optionality. This first-mover advantage is taking the form of:

Proprietary IP and algorithmic advantage: Organizations developing quantum algorithms and workflows today are generating intellectual property that will be difficult for late movers to replicate. The quantum algorithms best suited to your specific business problems are not off-the-shelf. They are developed through experimentation, iteration, and deep domain knowledge built over time.

[1] Liu et. al., Certified Randomness Using a Trapped-Ion Quantum Processor, Nature 2025.

Talent and ecosystem positioning: The global community of quantum-capable talent is small and highly sought after. Early movers are attracting quantum specialists, building academic partnerships, and embedding quantum expertise into their technology organizations. This talent advantage compounds over time.

Vendor relationship depth: The quantum computing vendors who will define the next decade are accessible and collaborative today. Organizations working most closely with them — co-developing applications, sharing use case requirements, and participating in product roadmap conversations — will have disproportionate influence over how the technology evolves.

Access to cutting-edge quantum capabilities, talent, and expertise is likely to become an increasingly important source of competitive differentiation in the years ahead.

The executive summary on timing
Organizations that start their quantum computing journey in 2026 are making a practical investment in learning, risk reduction, and staying competitive. The cost of starting that journey is modest compared to the cost of being unprepared when quantum advantage becomes more widely accessible. The organizations best positioned to capture that advantage are those building real capabilities today.

2. Where Quantum Computing Is Delivering Value

The use cases generating the most executive attention — and the most credible commercial proof points — span both public and private sector organizations. The following is a practical guide to where value is being captured today and where the near-term opportunity is most clearly defined.

Public Sector

Technology independence and commercial leadership

Sovereign governments are investing in quantum computing because they don’t want to be dependent on other nations for a technology that could reshape entire industries. Beyond that, quantum promises to deliver real improvements to the financial health and well-being of their nations, and the countries that develop early expertise are well placed to shape the standards and platforms that define the field. Central to all of this is building a domestic talent pool, training the researchers, engineers, and entrepreneurs who will drive quantum innovation for decades to come.

Energy grid resilience and decarbonization

Energy is an increasingly strategic public sector quantum use case. Research into solutions like hydrogen fuel cells has stalled in the last decade due to classical compute’s inability to fully simulate the highly complex chemistry and materials science needed to understand them. Quantum computing can play a role in better elucidating the complex mechanisms underlying technologies like this, so that they can one day be deployed at scale.

Scientific research and national laboratory simulation

National laboratories are among the most advanced public sector quantum adopters globally. Quantum simulation is proving valuable across nuclear science, materials research, chemistry, and accelerator physics, tackling problems that are simply too complex for classical computers to model accurately.

Private Sector

Financial services

Financial services firms are among the most active explorers of quantum computing because the industry runs on exactly the kinds of complex calculations quantum does well. Banks are already piloting quantum algorithms for portfolio optimization, risk simulation, and fraud detection; areas where processing more variables faster translates directly into better decisions and lower losses. Quantum random number generation is also attracting interest, offering a level of unpredictability that current methods simply can't match. For an industry where speed and accuracy are everything, the potential upside is great.

Life sciences and drug discovery

Quantum simulation allows pharmaceutical and biotech organizations to model molecular interactions and chemical reaction pathways at a level of complexity and accuracy that is computationally intractable for classical systems. The commercial prize is significant: a meaningful reduction in drug discovery timelines translates directly to earlier revenue and reduced development cost.

Chemicals and advanced materials

The chemicals industry is one of the most compelling near-term beneficiaries of quantum simulation. Chemical manufacturers depend on the ability to model catalyst behavior, reaction pathways, and material properties to develop new products and optimize industrial processes at scale. Materials science applications in the automotive and energy sectors, particularly around fuel cells and catalysis, are advancing rapidly. Quantum computing allows users to model chemical reactions and properties at a level of detail and complexity that is not possible with other approaches, making it an invaluable emerging tool for the industry.

3. How to Evaluate a Quantum Computing Solution

The quantum computing vendor landscape has expanded rapidly, and the quality of commercial claims varies significantly. Selecting the right partner requires a structured evaluation framework applied consistently. The following nine criteria represent the most important dimensions for enterprise and public sector buyers to assess.

Criterion	Why it matters for your organization	What to look for
Compute power	Compute power comes from a combination of fidelity, connectivity, and qubit number. Higher fidelity (lower error rates) produces more reliable results. Low-fidelity systems are incapable of running the deep circuits needed for industrially important workflows. In combination with this, enough high-fidelity qubits are needed for those workflows. Finally, high connectivity improves compute power further[2].	Fidelity (or error rates) independently verified, combined with qubit count and all-to-all connectivity. A single pair of high-fidelity qubits is not useful, nor are thousands of qubits with low fidelities. The combination of fidelity and qubit count drives compute power, and both are enhanced when the qubits are movable and are capable of arbitrary connectivity geometries.
Error correction	Error correction is what separates today's capable NISQ systems from tomorrow's transformative fault-tolerant computers. Advances in reducing theoretical and engineering risk for realizing error correction are enabling roadmap continuity and reliable delivery of increasingly more powerful systems beyond NISQ.	A credible fault-tolerant roadmap anchored by demonstrated logical qubit performance, including fault-tolerant universal logical gate set and better than physical logical qubit error rates, real time control and error correction, high-efficiency codes, and the flexibility to realize state-of-the-art error correction protocols tailored to applications and user needs.
Connectivity	Connectivity determines which algorithms can be run natively. All-to-all connectivity (enabled by movable qubits) allows for whole classes of algorithms (and concomitant solutions) that fixed qubit architectures cannot support without significant overhead.	Physical qubit mobility, and the ability to connect any two qubits in arbitrary geometries. Closely related is the fidelity of two qubit gates; make sure that quoted fidelities are an average across all qubit pairs and not just for a "best pair" or from a non-commercial test stand.

[2] M. DeCross et. al., Computational Power of Random Quantum Circuits in Arbitrary Geometries, Physical Review X 2025

Criterion	Why it matters for your organization	What to look for
Delivery model	Public sector requirements often mandate on-premise. Enterprise buyers need cloud flexibility. Vendors offering only one mode limit your long-term options.	Genuine hybrid or quantum-only options — cloud access and on-premise deployment.
Developer ecosystem	A mature ecosystem dramatically reduces the internal investment required to build quantum-capable teams and integrate the technology into existing workflows.	SDK maturity, appropriate open-sourced elements, a high-quality functional user language, quality emulator, documentation quality, active research community, published use case libraries.
Full-stack capability	A full-stack provider reduces integration complexity, accelerates time-to-value, and ensures hardware and software are optimized to work together.	Hardware, middleware, user software, and domain-specific applications from a single integrated platform.
Supply chain integrity	Quantum hardware depends on highly specialized components (ion traps, lasers, cryogenics, control electronics) sourced from a complex, often global supply base. Gaps in provenance or security create risk for regulated and mission-critical deployments and can jeopardize continuity of supply as demand scales.	Documented control over sourcing and provenance of critical components, vetted and audited suppliers, resilience to single points of failure, and compliance with relevant security and export control requirements.
Manufacturing scale	Architectures that scale on a validated path to fault tolerance represent real technical progress, but that progress only translates into deployable systems if the vendor can manufacture at scale. Manufacturability and industrial discipline, turns lab results into deployable systems with predictable lead times, rather than one-off prototypes.	Evidence of repeatable, quality-controlled production processes, a track record of on-schedule system delivery, in-house manufacturing capability (versus reliance on unproven external fabrication), and a credible plan for scaling production alongside the technical roadmap.
Commercial proof points	Commercial deployment evidence is the most reliable signal of vendor maturity. A vendor with no named enterprise customers represents a higher implementation and partnership risk.	Named referenceable customers in sectors comparable to yours. Business outcomes, not just technical benchmarks. Support from governments and federal agencies.

Criterion 1: Compute Quality

The most commonly cited quantum metric is qubit count. It is also one of the least meaningful in isolation. A quantum computer's practical usefulness is determined primarily by the fidelity of its qubit operations in combination with the qubit count and the connectivity. How accurately operations are implemented and how reliably quantum states are maintained over the duration of a computation are crucial for realizing complex computations, as is the ability to connect any two qubits into arbitrary geometries (thus reducing limits on algorithmic approaches).

A system with 1,000 low-fidelity qubits will not be able to produce results of value, as the computer can only run very shallow circuits that have no industrial relevance. In contrast, a computer with 100 high-fidelity qubits may be able to tackle problems of interest. Similarly, a computer with 100 high-fidelity qubits with fixed connectivity will only be able to run a subset of the algorithms that a computer with 100 high-fidelity movable qubits can. When evaluating vendors, ask about connectivity, request independently verified two-qubit gate error rates, ask about system level benchmarks like quantum volume, and ask after benchmarks on circuit types relevant to your use cases.

QUESTIONS TO ASK:

What are your error rates (or fidelities) on your commercial systems (not just your testbeds)? Are those figures representative across all qubit pairs in your commercial offering? What is your quantum volume? Are your qubits movable or fixes in space?

Criterion 2: Error Correction

Today's quantum computers are NISQ devices — noisy, intermediate-scale systems that suffer errors which accumulate over longer and more complex computations. Quantum error correction addresses this by encoding protected 'logical' qubits from multiple physical qubits in ways that detect and fix errors before they corrupt results.

Vendors differ enormously in their error correction capability. Look for multiple demonstrations of error correcting codes and primitives, evidence of error suppression below physical error levels, high logical "encoding rates", a fault-tolerant universal gate set, demonstrations of fully fault tolerant algorithms, and a detailed roadmap toward large-scale fully fault-tolerant computing. The physical-to-logical qubit encoding rate is a particularly meaningful metric, alongside the demonstration of a universal, fully fault tolerant gate set.

QUESTIONS TO ASK:

Can you show us examples where your logical qubits demonstrate meaningfully better error rates than your physical qubits? Can you show us peer-reviewed research demonstrating fault tolerance primitives with error rates better than physical and a fault tolerant universal gate set? Is your machine capable of handling multiple different error correcting codes? Have you demonstrated real-time error correction? Have you run real computations with logical qubits, or have you only demonstrated error correction primitives, like gates?

Criterion 3: Connectivity

Different quantum computing architectures offer different levels of connectivity. For example, superconducting systems have immovable qubits, thus constraining each qubit to interact only with its immediate physical neighbors (unless a new chip, with longer-range interactions, is designed and manufactured. In this case, the connectivity is still fixed by the chip). Trapped-ion and neutral atom systems offer movable qubits, which gives the user full all-to-all connectivity, allowing any qubit to interact with any other. Systems with immovable qubits will always suffer from limited connectivity, while systems with movable qubits will always offer high connectivity.

This architectural difference has significant practical consequences. All-to-all connectivity enables entire classes of algorithms — including many of the most commercially valuable error-correcting codes — that fixed-connectivity architectures cannot run natively. Architectural differences go beyond just connectivity. Architectural choices affect fidelity, scalability, engineering complexity, and much more.

QUESTIONS TO ASK:

What is the connectivity model of your system? Can any qubit interact with any other or are interactions limited to physical neighbors or limited by your chip design? Are your qubits movable?

Criterion 4: Delivery Model

Public sector and defense organizations frequently require on-premise deployment due to security classification, data sovereignty constraints, and regulatory requirements. Many users are particularly interested in combining the power of quantum computing with High Performance Computing (HPC) and Supercomputing environments. Enterprise organizations generally prefer cloud access for flexibility, ease of integration, and cost efficiency at the experimentation stage. The most capable vendors offer credible options across all models without material trade-offs on performance.

QUESTIONS TO ASK:

Can you provide examples of on-premise offerings that have been tested in customer environments?
Can you provide examples of your computers being integrated with HPC environments or supercomputers?

Criterion 5: Developer Ecosystem

The practical usability of a quantum platform is heavily shaped by the quality of its developer ecosystem. A mature SDK, friendly user language, functional emulator, comprehensive documentation, an active research community, and a library of domain-specific quantum algorithms significantly reduce the internal investment required to build organizational quantum capability.

QUESTIONS TO ASK:

Can you walk us through your developer ecosystem, and how active is the research community building on your platform? What proof points do you have to support your claims, and can you provide examples of peer reviewed research or white papers from community partners?

Criterion 6: Full-Stack Capability

A vendor that offers hardware access alone places the entire burden of software integration, middleware selection, and application development on the buyer. A full-stack provider dramatically reduces this burden and accelerates time-to-value. Evaluate whether the vendor's applications are commercially deployed or still at the research stage.

QUESTIONS TO ASK:

Which components of your full-stack offering are built and maintained in-house, and which relies on third-party partners or licensed technology?

Criterion 7: Supply chain integrity

A quantum computer is only as reliable as the components and partners behind it. Systems depend on specialized, often globally sourced materials and subsystems, precision optics, lasers, vacuum and cryogenic components, and custom control electronics, any of which can introduce risk if provenance, security, or continuity of supply isn't tightly controlled. For enterprise and public sector buyers, this isn't just an operational concern; it's a trust and compliance issue, particularly where national security, data sensitivity, or long-term deployment commitments are involved.

Vendors differ in how much of their supply chain they control directly versus outsource, and in how rigorously they vet, audit, and secure their suppliers. Look for demonstrated ownership of critical manufacturing steps, documented supplier vetting processes, and evidence the vendor can maintain continuity of supply even under geopolitical or market disruption.

QUESTIONS TO ASK:

How much of your critical hardware supply chain do you control directly versus source externally? What processes do you have in place to vet, audit, and secure your suppliers? Can you demonstrate provenance and traceability for critical components? How do you mitigate single points of failure in your supply chain?

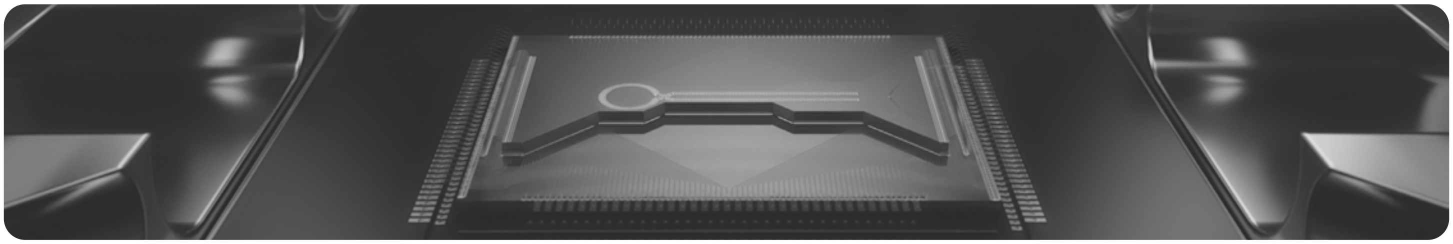

Criterion 8: Manufacturing scale

Architectures that scale on a validated path to fault tolerance represent real technical progress, but that progress only translates into deployable systems if the vendor can manufacture at scale. Manufacturability, industrial discipline, is what turns promising lab results into commercial-grade systems with predictable lead times, consistent quality, and a credible path to volume production.

Vendors vary widely in manufacturing maturity: some rely on hand-built, lab-scale processes that don’t generalize, while others have invested in repeatable, quality-controlled production lines. Look for a demonstrated track record of on-time system delivery, in-house manufacturing capability, and a scaling plan that matches the pace of the technical roadmap, not just aspirational claims about future qubit counts.

QUESTIONS TO ASK:

What is your manufacturing capacity today, and how does it need to scale to meet your roadmap? Do you manufacture critical components in-house, or rely on external fabrication partners? What is your track record of delivering systems on schedule? How do you plan to scale manufacturing capacity alongside your technical roadmap?

Criterion 9: Commercial Proof Points

Commercial deployment evidence is the most reliable signal of vendor maturity. Ask for named case studies, not anonymized ones. Require that outcomes be expressed in business terms. Look for customer testimonials. A vendor with genuine commercial traction will facilitate this readily.

QUESTIONS TO ASK:

Can you provide named case studies or reference customers in our industry?

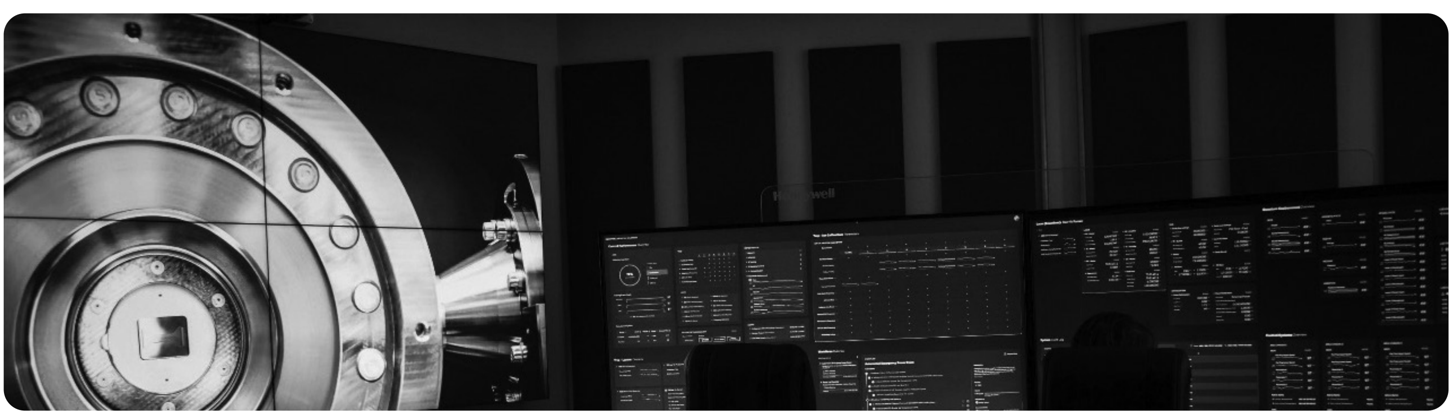

4. How Quantinuum Maps to the Evaluation Framework

The following section applies the nine evaluation criteria from Section 3 directly to Quantinuum, drawing on publicly available information, published benchmarks, and documented customer deployments. The intent is to give buyers a concrete, evidence-based view of how Quantinuum performs against each criterion — and to illustrate how the framework can be applied to any vendor you are assessing.

Buyers are encouraged to apply this same framework to any other vendors under consideration and to validate all assessments through direct engagement, reference calls, and independent technical review.

Criterion	How Quantinuum maps to this criterion
Compute quality	Quantinuum's Helios processor has demonstrated leading performance in benchmarks like physical and logical qubit fidelity and quantum volume among commercially available systems. Its trapped-ion architecture consistently delivers leading results in independently verified assessments[3], producing reliable results on circuit types that challenge lower-fidelity systems.
Error correction	Quantinuum is an industry leader in fault-tolerance[4], with results ranging from logical qubits outperforming physical qubits[5], to real-time error correction[6], to demonstrations of universal fault tolerant gate sets[7]. Quantinuum has demonstrated a wide breadth of QEC codes[8], has one of the highest encoding rates in the industry[9], and consistently has low logical error rates across fault tolerance primitives[10].

[3] J. A. Montanez-Barrera et. al., Evaluating the performance of quantum processing units at large width and depth, arXiv 2502.06471

[4] Ryan-Anderson et. al., High-fidelity and Fault-tolerant Teleportation of a Logical Qubit using Transversal Gates and Lattice Surgery on a Trapped-ion Quantum Computer, Science 2024, See also: IARPA ELQ Challenge

[5] Dasu et. al., Order-of-magnitude extension of qubit lifetimes with a decoherence-free subspace quantum error- correction code, arXiv 2503.22107

[6] Ryan-Anderson et. al., Realization of real-time fault-tolerant quantum error correction, Physical Review X 2021

[7] Dasu et. al., Breaking even with magic: demonstration of a high-fidelity logical non-Clifford gate, arXiv 2506.14688

[8] See this link

[9] Dasu et. al., Computing with many encoded logical qubits beyond break-even, arXiv 2602.22211

[10] Daguerre et. al., Experimental Demonstration of High-Fidelity Logical Magic States from Code Switching, arXiv 2506.14169, Ryan-Anderson et. al., High-fidelity and Fault-tolerant Teleportation of a Logical Qubit using Transversal Gates and Lattice Surgery on a Trapped-ion Quantum Computer, Science 2024

Criterion	How Quantinuum maps to this criterion
Connectivity	Quantinuum's QCCD trapped-ion architecture provides movable qubits, allowing any qubit to interact with any other. This enables algorithms and error-correcting codes that are not natively executable on immovable qubit (fixed-connectivity) architectures without a much higher overhead. This translates to a utility scale computer that requires fewer physical qubits and can save computational time.
Delivery model	Quantinuum's systems are accessible via cloud through Quantinuum Nexus as well as via partners such as ORNL and EPFL. In addition, Quantinuum systems have been deployed on-premise in national laboratory environments — including integration with the Fugaku supercomputer at RIKEN in Japan[11]. This hybrid delivery capability addresses both enterprise and public sector access requirements.
Developer ecosystem	The Guppy-based SDK is a widely adopted quantum development toolkit, with a substantial research community and active enterprise user base. Quantinuum's annual Q-Net Connect forum draws attendance from enterprise, academic, and government users, reflecting the breadth of its deployment community. Quantinuum also has a Startup Partner Program, which connects Quantinuum with pioneering startups to accelerate ecosystem innovation. Through access to our hardware, software, and services, startup partners can co-develop new solutions, explore differentiated use cases, and help grow the next generation of quantum applications. Finally, Quantinuum has use-case libraries such as InQuanto, our quantum chemistry and materials platform.
Full-stack capability	Quantinuum's commercial portfolio includes hardware (System Model H1, H2 and Helios), developer tooling (Guppy, HUGR, TKET, Selene), quantum chemistry applications (InQuanto), and new algorithms libraries coming soon.
Supply chain integrity	Quantinuum is vertically integrated: designing, manufacturing, and programming its quantum computers in-house rather than relying on third-party processors or outsourced fabrication for critical hardware. This gives the company direct control over the components that matter most to system performance and security, from ion traps and photonics to control electronics, reducing exposure to the provenance and continuity risks that come with a fragmented, multi-vendor supply chain. Hardware engineering is concentrated at dedicated, purpose-built facilities in Colorado and Minnesota, supporting close oversight from design through production. This commitment extends to Quantinuum's external suppliers. The company's Supplier Code of Conduct sets out its expectations for integrity and compliance across its global supply chain, requiring suppliers to provide, amongst other things, a safe working environment, treat workers with dignity and respect and comply with the law in every country in which they operate.

[11] Yamamoto et. al., Quantum-HPC hybrid computation of biomolecular excited-state energies, arXiv 2601.15677

Criterion	How Quantinuum maps to this criterion
Manufacturing scale	Quantinuum has demonstrated that its trapped-ion architecture can scale in practice, solving long-standing "wiring" and "qubit sorting" problems that constrain other architectures as qubit counts grow. The company has already delivered multiple hardware generations on a defined cadence, System Model H1, H2, and Helios, with its next systems, Sol and Apollo, already in development and targeting delivery in 2027 and 2029 respectively. This track record of moving successive systems from lab result to production device, including achieving industry-leading two-qubit gate fidelity in a production machine, reflects a manufacturing process built for repeatable, predictable delivery rather than one-off demonstrations.
Commercial proof points	Quantinuum's publicly referenceable enterprise customers include JPMorganChase (financial applications and certified randomness), BMW Group (materials science and catalysis research), BP (energy exploration), and RIKEN (national laboratory and classical supercomputer integration). These partnerships are supported by publicly available documentation and named customer advocacy and represent some of the most commercially advanced quantum deployments in their respective sectors.

A NOTE ON VENDOR ASSESSMENT:

Quantum computing is a fast-moving field. Vendor capabilities evolve quickly, and the landscape today is materially different from two years ago. Any vendor assessment — including the one above — should be validated with current benchmarks, reference calls, and direct technical engagement. The evaluation framework in Section 3 is designed to help you do exactly that, regardless of which vendors you are assessing.

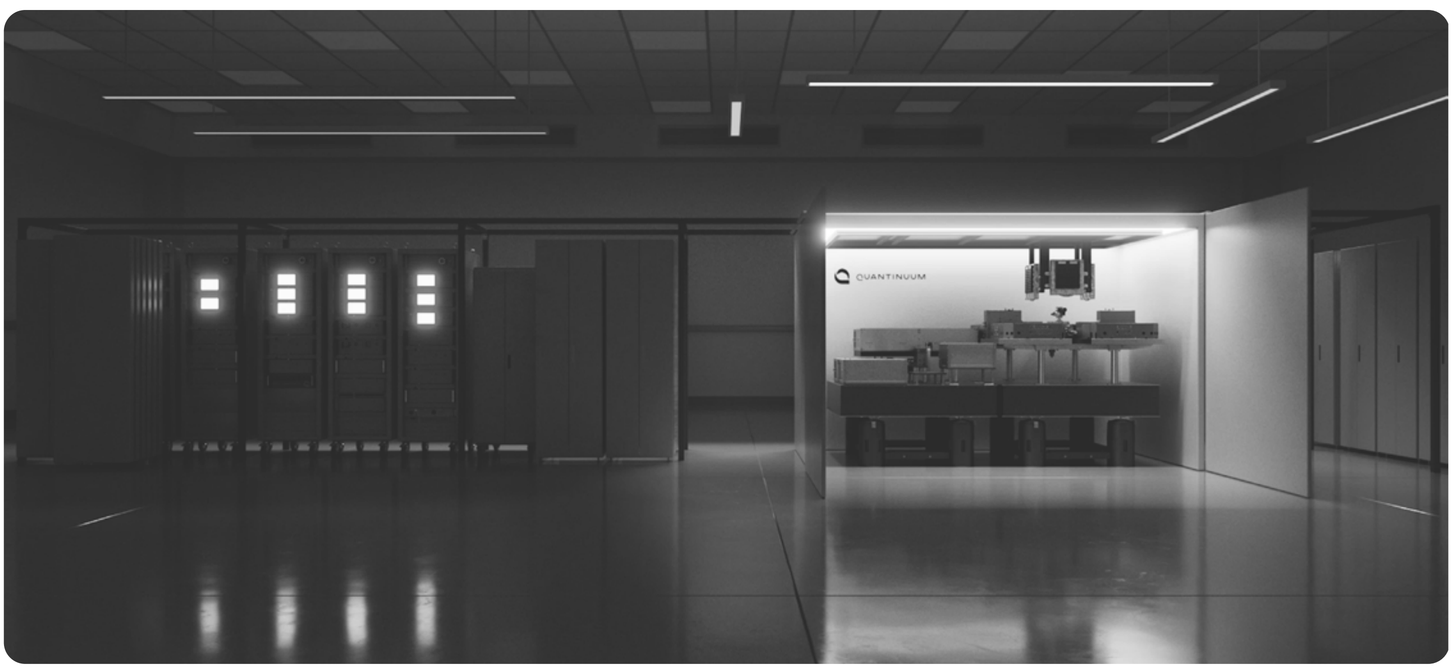

Quantum Vendor Evaluation Checklist

Criterion	Question to Ask	Quantinuum	Vendor B	Vendor C
1. Compute Quality	What are your error rates (or fidelities) on your commercial systems (not just your testbeds)? Are those figures representative across all qubit pairs in your commercial offering?			
2. Error Correction	Can you show us examples where your logical qubits demonstrate meaningfully better error rates than your physical qubits? Can you show us peer-reviewed research demonstrating fault tolerance primitives with error rates better than physical and a fault tolerant universal gate set?			
3. Connectivity	What is the connectivity model of your system? Can any qubit interact with any other or are interactions limited to physical neighbors?			
4. Delivery Model	Can you provide examples of on-premise offerings that have been tested in customer environments?			
5. Developer Ecosystem	Can you walk us through your developer ecosystem, and how active is the research community building on your platform? What proof points do you have to support your claims, and can you provide examples of peer reviewed research or white papers from community partners?			

Quantum Vendor Evaluation Checklist

Criterion	Question to Ask	Quantinuum	Vendor B	Vendor C
6. Full-Stack Capability	Which components of your full-stack offering are built and maintained in-house, and which relies on third-party partners or licensed technology?			
7. Supply chain integrity	How much of your critical hardware supply chain do you control directly versus source externally? What processes do you have in place to vet, audit, and secure your suppliers? Can you demonstrate provenance and traceability for critical components? How do you mitigate single points of failure in your supply chain?			
8. Manufacturing scale	What is your manufacturing capacity today, and how does it need to scale to meet your roadmap? Do you manufacture critical components in-house, or rely on external fabrication partners? What is your track record of delivering systems on schedule? How do you plan to scale manufacturing capacity alongside your technical roadmap?			
9. Commercial Proof Points	Can you provide named case studies or reference customers in our industry?			

Notes:

5. Quantinuum as a Co-Development Partner

Selecting a quantum computing vendor is not the same as selecting a software platform or a cloud provider. At this stage of the technology's development, the quality of the partnership matters as much as the quality of the hardware. The organizations achieving the most commercially meaningful results are those that have invested in deep, collaborative vendor relationships — not transactional ones.

Quantinuum is built around a co-development model. Rather than selling access to a platform and leaving buyers to figure out the rest, Quantinuum works alongside its customers to design use cases, develop quantum algorithms, build hybrid architectures, and translate technical results into business outcomes. This distinction is material, and worth understanding in detail.

What co-development looks like in practice

Joint use case design

Quantinuum's engagement model begins with a structured discovery process in which its domain specialists and quantum scientists work alongside the customer's technical and business teams to identify and prioritize the use cases most likely to yield near-term value. This is not a sales exercise — it is a substantive technical collaboration designed to ensure that the problems being addressed are genuinely suited to quantum treatment at the current stage of hardware maturity.

Algorithm development and customization

Off-the-shelf quantum algorithms rarely produce the best results for specific enterprise or public sector problems. Quantinuum's team works with customers to develop and optimize quantum algorithms tailored to their specific use cases, data characteristics, and computational constraints. This co-development of algorithmic intellectual property is one of the most enduring sources of value from a Quantinuum partnership.

Hybrid architecture design

Quantinuum's engineering teams support customers in designing hybrid quantum-classical architectures where necessary. This includes defining the interfaces between quantum and classical components, optimizing data pipelines for quantum workloads, and ensuring that the overall system architecture is production-ready rather than demo-grade.

Pilot support and knowledge transfer

During pilot programs, Quantinuum embeds technical specialists alongside customer teams — not to run the pilot on the customer's behalf, but to build internal capability while delivering results. The explicit goal is knowledge transfer: by the end of a well-run Quantinuum pilot, the customer's team should be meaningfully more quantum-capable than it was at the start. This approach directly addresses one of the most common failure modes of technology pilots — the expertise walking out the door when the vendor engagement ends.

Roadmap alignment and long-term planning

Quantinuum shares its hardware and product roadmap with strategic partners, and actively incorporates customer use case requirements into its development priorities. This means that organizations working closely with Quantinuum today have a meaningful voice in shaping the capabilities that will be available on tomorrow's hardware — a relationship dynamic that is simply not available with vendors where the partnership is transactional.

THE PARTNERSHIP DIFFERENTIATOR

The most productive quantum partnerships are those in which the vendor is as invested in the customer's success as the customer is. Quantinuum's co-development model is designed around this principle — the goal is not to sell platform access but to help organizations build durable quantum capability. For early-stage buyers, this distinction is particularly important: the vendor who will help you learn the most from your first pilot is the one most worth engaging with for the long term.

6. How to Get Started: Recommended First Steps

The most common barrier to quantum computing adoption is not budget or technology readiness. It is organizational uncertainty about where to begin. The following framework provides a practical, phased approach to building quantum capability with appropriate rigor and risk management.

Step 1: Build internal quantum literacy

Investing in internal education at three levels creates the organizational foundation necessary to evaluate vendors credibly, design effective pilots, and scale successful programs:

Executive level: Focused briefings on quantum computing's commercial relevance, near-term risk profile, and strategic implications for your sector.

Technical leadership: Deeper engagement with quantum algorithms, hardware trade-offs, and integration architectures.

Developer teams: Hands-on training with quantum development environments, SDK tooling, and algorithm development workflows.

Step 2: Identify and prioritize use cases

Conduct a structured internal assessment of where quantum computing could deliver meaningful business or operational value. Focus on problems with the following characteristics:

- High computational complexity that strains or exceeds classical systems
- Large variable spaces where optimal solutions have significant economic or operational value
- Simulation requirements at molecular, atomic, or sub-atomic scale
- Long-term data confidentiality requirements where post-quantum cryptographic risk applies

Prioritize use cases where early progress is demonstrable, stakeholder appetite exists, and where a quantum pilot could be structured with clear success criteria within six to twelve months.

Step 3: Launch a targeted pilot program

Apply the pilot design framework from Section 7.

- Define the business question clearly
- Establish classical benchmarks before you start
- Design the hybrid architecture deliberately
- Build in structured learning checkpoints
- Capture talent and workflow learnings explicitly throughout

Step 4: Establish a quantum security baseline

Regardless of where your organization is on its broader quantum journey, post-quantum cryptographic readiness should be treated as a parallel priority:

- Conduct a cryptographic inventory to identify data assets with long-term confidentiality requirements
- Assess current encryption standards against published NIST post-quantum cryptographic guidelines
- Develop a migration roadmap with prioritized timelines based on data sensitivity and longevity
- Evaluate quantum-hardened key generation solutions as a near-term, commercially available risk mitigation measure

Step 5: Select a strategic quantum partner

Apply the evaluation framework from Section 3 rigorously. Engage reference customers directly. Assess the vendor's co-development model using the questions in Section 6. Look for evidence that the vendor's roadmap aligns with the trajectory of your use cases — and that their partnership approach is designed to build your internal capability, not create dependency on their team.

THE SINGLE MOST IMPORTANT FIRST STEP

If your organization takes one action as a result of reading this guide, make it this: identify your highest-priority quantum use case, assign an internal owner, and initiate a conversation with a quantum computing vendor about what a 90-day scoping engagement would look like. The organizations that have moved furthest in quantum adoption started not with a large program investment but with a well-defined question and the curiosity to begin exploring it.

Need Guidance On Getting Started?

Quantinuum's Consulting Services can help to build your quantum advantage by meeting your organization where it is today and helping you move forward. Please contact us for more information.

7. Getting Maximum Value From Your Pilot Program

For many organizations, a structured pilot program is the most productive way to begin a quantum computing journey. A well-designed pilot does more than test a technology — it generates the organizational intelligence needed to make confident decisions about where, how, and how fast to scale. This section provides a detailed framework for designing pilots that produce tangible, actionable learnings.

What a good quantum pilot actually delivers

The most valuable pilots are not proof-of-concept demonstrations designed to impress a board presentation. They are structured learning exercises designed to answer specific strategic questions that your organization genuinely needs to answer before committing to a scaled program.

A well-designed quantum pilot should deliver clear findings across five dimensions:

Quantum workflow design intelligence: How does quantum computation integrate with your existing classical infrastructure? What data pipelines need to be built or modified? Where does the quantum processor fit within a broader hybrid architecture, and what does the handoff between quantum and classical components look like in practice? These architectural insights are among the most valuable outputs of a first pilot, regardless of whether the quantum results themselves are groundbreaking.

Talent and capability gap assessment: Who in your organization can engage productively with quantum tools today, and where are the gaps? A pilot surfaces the specific skills your team needs - whether that is quantum algorithm development, quantum-classical integration engineering, or domain-specific quantum application expertise. This intelligence is essential for making defensible resource allocation decisions and building a realistic hiring and training plan.

Use case validation or redirection: Does the problem you chose for the pilot actually benefit from quantum treatment at the current stage of hardware maturity? Some use cases that appear to be strong quantum candidates are better addressed by near-term quantum algorithms; others are better suited to classical approaches for now and quantum approaches in 18 to 24 months. A pilot answers this question with evidence rather than theory.

Vendor partnership quality: How does your vendor perform as a co-development partner, not just as a technology supplier? A pilot is the most reliable test of whether a vendor's support model, technical depth, and collaborative approach match what you need for a scaled program. The quality of co-development during a pilot is the strongest predictor of long-term partnership success.

A clear, evidence-based next steps framework: A pilot should conclude with a structured set of recommendations — not just a set of technical results - that directly inform go/no-go decisions, resource allocation, and program scoping for the next phase.

THE PILOT MINDSET SHIFT

The most productive frame for a quantum pilot is not 'did quantum win?' It is 'what did we learn that we could not have learned any other way?' The organizations extracting the most value from early quantum engagement are those that treat pilots as learning investments — with the same rigor they apply to any strategic capability-building initiative.

Designing a pilot that produces real learnings

⟶ Step 1: Define the business question, not just the technical problem

Every effective pilot starts with a business question that a decision maker genuinely needs answered. ‘Can quantum computers optimize our logistics network?’ is a research question. ‘Can quantum-assisted optimization reduce our inter-facility routing cost by more than 8% within our current planning cycle constraints?’ is a business question. The latter gives you a clear success criterion and a direct line to a resource allocation decision.

⟶ Step 2: Establish classical benchmarks before you start

To measure quantum advantage meaningfully, you need to know exactly what your best classical approach achieves on the same problem. Establishing rigorous classical baselines before the pilot begins is essential. Without them, you cannot distinguish genuine quantum advantage from the natural improvement that comes from reframing a problem more carefully.

⟶ Step 3: Design the hybrid architecture deliberately

Even on today’s hardware, the most effective quantum deployments are hybrid. Decide before the pilot begins which components will run on quantum processors, which will run classically, and how data will move between them. Leaving the architecture to emerge organically during execution almost always results in a pilot that is slower, messier, and harder to learn from.

⟶ Step 4: Build in structured learning checkpoints

Rather than running a pilot as a single continuous experiment, structure it as a series of sprints, each with defined learning objectives and a checkpoint where findings are documented, and the next sprint is adjusted accordingly. This iterative approach produces richer learnings and avoids the common failure mode of discovering at the end of a twelve-week pilot that a key assumption was wrong in week two.

⟶ Step 5: Capture talent and workflow learnings explicitly

Assign someone on your team the specific responsibility of documenting workflow design decisions, integration challenges, and capability gaps as they emerge during the pilot — not just at the end. The most frequently overlooked pilot output is the institutional knowledge embedded in the heads of the small team that ran it. Making this knowledge explicit and transferable is what allows a pilot to inform a scaled program rather than remaining a one-time exercise.

Translating pilot outputs into next steps and resource allocation

A quantum pilot should conclude with a structured debrief that directly addresses four strategic questions:

Should we scale this use case?

Based on the pilot results and classical benchmarks, does this problem warrant a scaled quantum program, a hybrid classical-quantum approach, or a return to the queue for a future hardware generation? The answer should be grounded in evidence, not enthusiasm.

What capabilities do we need to build or acquire?

Based on the talent and workflow gaps identified during the pilot, what is the most cost-effective path to the capabilities required for the next phase — internal hiring, vendor support, academic partnership, or a combination?

What is the right scale and timeline for the next phase?

A pilot should produce enough data to scope a follow-on program with reasonable precision. This includes hardware access requirements, team size, integration investment, and a realistic timeline to production deployment.

What is our vendor relationship strategy going forward?

Is the vendor that supported your pilot the right partner for a scaled program? Do their roadmap, support model, and co-development approach align with what you need? The pilot is the most reliable basis for answering this question.

THE MOST COMMON PILOT MISTAKE

The most common mistake organizations make with quantum pilots is treating them as technology demonstrations rather than strategic learning investments. A pilot whose primary output is a slide deck showing that quantum ran faster than classical on a toy problem has not generated the organizational intelligence needed to make a confident next-steps decision. Define the learning objectives before you start and hold the pilot accountable to deliver against them.

Conclusion

The quantum computing landscape has changed fundamentally. This is no longer a technology to monitor from a distance. It is one that forward-looking organizations are actively deploying, learning from, and building competitive advantage with today.

The evaluation framework in this guide is designed to cut through vendor noise and focus your assessment on the dimensions that most reliably predict commercial success: qubit quality, error correction maturity, architectural capability, delivery flexibility, ecosystem depth, full-stack integration, supply chain integrity, manufacturing scale and verifiable commercial proof points.

Equally important is the quality of the partnership. At this stage of quantum computing's development, co-development is not a nice-to-have — it is what separates organizations that build durable quantum capability from those that accumulate impressive pilot results without a clear path to scaled impact.

The organizations best placed to capture quantum's competitive advantage are those building their internal capabilities, forging deep vendor partnerships, and launching thoughtfully designed pilot programs now. The technology is ready for serious organizational engagement. The question is whether your organization is ready to engage with it.

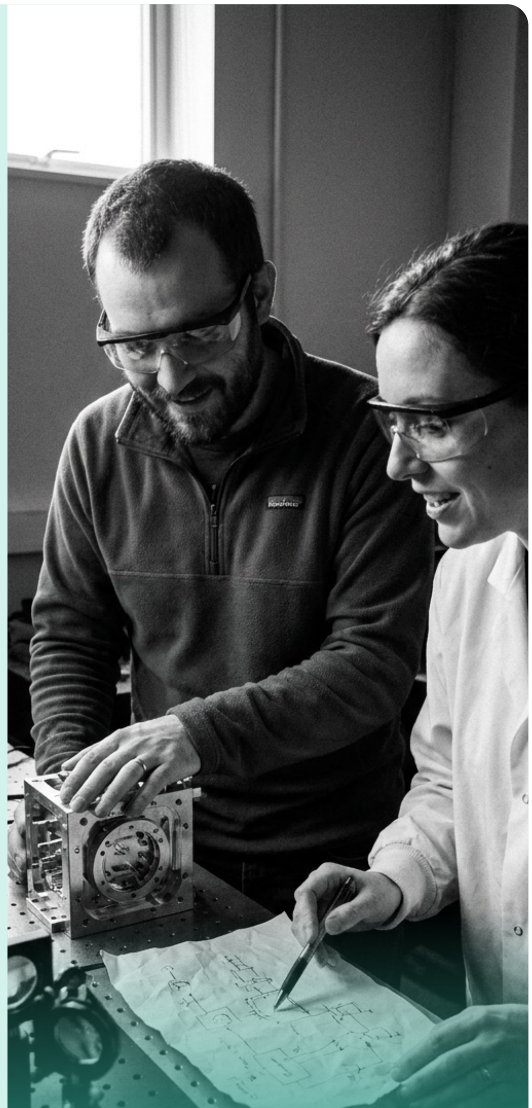

To explore how Quantinuum can support your quantum computing strategy, visit **quantinuum.com**

Disclaimer

This guide is provided by Quantinuum for informational and evaluation purposes only and does not constitute a warranty, guarantee of performance, or commitment to deliver any specific capability or result. All technical claims and benchmarks reflect publicly available information as of 2026 and are subject to change. Statements regarding product roadmaps and future capabilities are forward-looking and based on current expectations; actual results may differ materially. Descriptions of co-development engagements and intellectual property arrangements are illustrative only. The terms governing any customer engagement, including IP ownership and licensing, are determined solely by the applicable customer agreement. Comparisons between quantum computing architectures reflect Quantinuum's assessment of publicly available data at the time of writing and should not be relied upon as a comprehensive evaluation of any third-party technology. Readers are encouraged to conduct independent technical and commercial due diligence.

Nothing in this guide constitutes legal, regulatory, financial, or investment advice.

Glossary of Key Terms

Quantum computing has its own vocabulary. This glossary explains the most commonly used terms in plain language for readers new to the field.

Term	What it means in plain language
Qubit	The basic unit of quantum information. Unlike a classical computer bit, which is always either 0 or 1, a qubit can represent 0, 1, or both simultaneously (a property called superposition). This allows quantum computers to solve problems in a completely different way from classical computers.
Superposition	A qubit's ability to exist in multiple states at the same time until it is measured. Think of it like a coin spinning in the air before it lands — while it's spinning, it's neither heads nor tails, but has a probability of being either. This property underlies much of quantum computing's power.
Entanglement	When two or more qubits become linked so that the state of one instantly affects the others, regardless of the distance between them. This allows quantum computers to process information in ways that are fundamentally impossible for classical systems.
Gate fidelity	A measure of how accurately a quantum computer executes a basic operation on qubits. High fidelity means fewer errors per operation. Low fidelity means errors accumulate quickly, producing unreliable results.
Error correction	A set of techniques for detecting and fixing errors in quantum computations before they corrupt the result. Because qubits are fragile, error correction is essential for running long, complex calculations reliably.
Logical qubit	A reliable, error-corrected qubit built from multiple physical qubits working together. Logical qubits are more stable than raw physical qubits and are required for the most demanding quantum computations.

Term	What it means in plain language
Physical qubit	A single, raw qubit on quantum hardware before error correction is applied. Physical qubits can be noisy and error-prone. Physical qubits must be combined to produce logical qubits. The number of physical qubits you use for each logical qubit is referred to as the "encoding rate".
NISQ	Noisy Intermediate-Scale Quantum. The term used to describe today's quantum computers, which have a limited number of qubits and make errors that accumulate over longer computations. NISQ systems are already commercially useful for specific tasks but cannot yet handle the most complex problems reliably.
Fault-tolerant quantum computing	The next stage of quantum development, in which quantum computers can correct their own errors reliably enough to run extremely complex, long computations. This is the target state that will unlock quantum computing's full transformative potential.
Quantum advantage	The point at which a quantum computer can solve a specific problem faster or more accurately than the best available classical computer. This threshold is already being crossed for specific use cases and will expand as hardware matures.
Hybrid quantum-classical	A computing architecture that combines quantum processors with classical computers, using each where it performs best. A number of commercially productive quantum deployments today are hybrid, with quantum handling specific high-complexity tasks within a broader classical workflow.
Trapped-ion qubits	A type of qubit in which individual charged atoms (ions) are suspended using electromagnetic fields. Trapped-ion systems typically achieve very high qubit fidelity and full all-to-all connectivity. Trapped-ion qubits are typically cooled to ~1 millionth of a degree above absolute zero using laser cooling.
Superconducting qubits	A type of qubit in which special superconducting circuits mimic the properties of an atom, enabling superposition and entanglement between the qubits. Superconducting systems do not have movable qubits, so are typically limited to nearest-neighbor connectivity, limiting the algorithmic and error-correcting approaches available. Superconducting qubits are typically cooled to ~1 thousandth of a degree above absolute zero using dilution refrigerators, which limit the size of superconducting chips and need a rare isotope of Helium, He-3, to run.

Term	What it means in plain language
Neutral atom qubits	A type of qubit in which individual atoms are suspended using laser beams. Neutral atom systems achieve full all-to-all connectivity. Neutral atom systems are typically cooled to ~1 millionth of a degree above absolute zero using laser cooling.
Post-quantum cryptography	Encryption methods designed to remain secure even after large-scale quantum computers become capable of breaking today's standard encryption. Organizations with long-term data confidentiality requirements should already be planning their migration to post-quantum standards.
SDK	Software Development Kit. A set of tools, libraries, and documentation that developers use to build applications for a specific platform. In quantum computing, SDKs allow teams to write quantum algorithms and deploy them on hardware without building from scratch.
All-to-all connectivity	A hardware property in which any qubit can directly interact with any other qubit on the chip. This enables more efficient algorithm execution than fixed-connectivity architectures, where qubits can only interact with their immediate physical neighbors.
Quantum volume	A single-number benchmark that attempts to capture overall quantum computer performance, combining qubit count, fidelity, and connectivity into a single score. Higher quantum volume generally indicates a more capable system.

QUANTINUUM

QUANTINUUM.COM

AUTHORED BY

Quantinuum Advisory Services

Quantinuum Consulting Services helps organizations develop practical quantum strategies, build the capabilities needed for adoption, and identify high-value use cases. The team combines deep quantum expertise with business insight to help companies move from exploration to measurable impact.