

U.S. IT Hardware and EU Semi**The Quantum Leap: Winners & Losers in the Quantum future**

Mark C. Newman
+1 212 845 7822
mark.newman@bernsteinsg.com



David Dai, CFA
+852 2918 5704
david.dai@bernsteinsg.com



Qingyuan Lin, Ph.D.
+852 2123 2654
qingyuan.lin@bernsteinsg.com



April Li
+1 917 344 8339
april.li@bernsteinsg.com



Phoebe Sun
+1 917 344 8481
phoebe.sun@bernsteinsg.com



Carmine Milano, CFA
+44 20 7762 1857
carmine.milano@bernsteinsg.com

This note is a deep dive into Quantum technology and in particular Quantum Computing and assesses different technical approaches and the playing field.

Quantum computing is set to become the next important step in computing.

We believe the future of advanced computing will be shaped by a tri-processor architecture composed of CPUs, GPUs, and QPUs. In this model, Quantum computing and QPUs will not replace classical hardware (CPUs and GPUs in traditional and AI servers) but operate as a specialized accelerator, tightly integrated into a broader hybrid quantum-classical system. This is similar to how GPUs have not replaced CPUs.

How does quantum computing work? In quantum computing, once physical qubits are created, quantum algorithms will manipulate these qubits through carefully sequenced operations that create superposition, entangle qubits, and use interference to amplify correct outcomes while canceling incorrect ones. After this controlled evolution, the system is measured, collapsing quantum states into classical bits that encode the solution.

Quantum computing has attracted several big tech companies as well as specialized pure plays. Among the big tech companies, IBM (rated Market-Perform) is arguably the leader, with Google, Microsoft and Intel with meaningful quantum computing initiatives. Geographically we see significant activity from companies and national governments in North America, Europe and China.

We analyze 6 quantum modalities and see three approaches as leading contenders for the Quantum future: 1) **Superconducting**, arguably today's most commercially advanced and execution-ready technology, driven by rapid development from IBM and Google. 2) **Trapped-ion**, the modality adopted by the biggest market cap quantum pureplay IONQ, offers the highest precision and stability. Infineon is the leading foundry in trapped ions, supplying QPUs to the largest players in this subsector; and 3) **Neutral atom**, which provides a highly flexible architecture capable of scaling. The other approaches we are watching closely include: **Photonic**, **Topological** and **Silicon spin**.

Real economic value in quantum computing is likely to be unlocked only after quantum advantage is achieved, which is anticipated to be around 2030. To estimate the implied market share of current quantum pure plays, we discount an estimated \$130bn TAM in 2040, apply SOX P/E multiple of 28x, and use a 30% net margin consistent with mature tech peers. This framework implies **acombined market share of the 6 publicly listed pure-plays at 24%**, with IONQ (the largest market cap of the six) accounting for ~11% share. These results appear reasonable and support our assumptions.

However, is it too early to call winners and losers? Given that this technology is so early stage, and with some companies still in semi-stealth mode, we believe it is too early to know for sure. For example Photonics, Topological and Silicon spin show promise but are far too early-stage to assess. Furthermore, **our analysis shows that quantum computing may not be a winner takes all industry.** Different modalities bring distinct strengths and tradeoffs, which should make certain architectures better suited for specific use cases, workloads, and time horizons.

BERNSTEIN TICKER TABLE

Ticker	Rating	5 Jun 2026		Price Target	TTM Rel. Perf.	Cur	Adjusted EPS			Adjusted P/E (x)		
		Cur	Closing Price				2025A	2026E	2027E	2025A	2026E	2027E
IBM (IBM)	M	USD	284.84	280.00	(17.6)%	USD	11.58	12.35	13.93	24.6	23.1	20.5
IFX.GR (Infineon)	O	EUR	74.42	74.00	97.7%	EUR	1.38	1.72	2.63	53.8	43.3	28.3
SPX			7,383.74									
EDME			1,544.13									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We rate IBM Market-Perform, PT = USD 280.

We rate Infineon Outperform, PT = EUR 74.

Table Of Contents

Part I. Future of Computing - Introducing the Tri-Processor Architecture (CPU, GPU, QPU).....	3
Part II: What is Quantum Computing? The Tech Basics.....	5
Part III: What are the Leading Quantum Modalities?.....	7
Part IV: Introducing the Major Players in Quantum Computing.....	13
Part V: Quantum Advantage and Real Economic Value.....	15
Part VI. Quantum TAM and Valuation.....	16
Part VII. Is it too early to call Winners and Losers?.....	18

DETAILS

PART I. FUTURE OF COMPUTING - INTRODUCING THE TRI-PROCESSOR ARCHITECTURE (CPU, GPU, QPU)

Quantum computing is set to become the next important step in computing. We believe the future of advanced computing will be shaped by a tri-processor architecture composed of CPUs, GPUs, and QPUs (see [Exhibit 1](#)). In this model, Quantum Computing and the QPU does not replace classical hardware (CPUs and GPUs in traditional and AI servers). Instead, it operates as a **specialized accelerator**, tightly integrated into a broader hybrid quantum-classical system. Tasks are dynamically distributed across the CPU, GPU, and QPU based on the nature and complexity of the workload ([Exhibit 2](#)). Similar to how GPUs (and AI servers) didn't replace CPUs (and traditional servers) - rather they are complementary and additive; we expect Quantum computing and QPUs will be complementary and additive to classical computing methods (CPUs and GPUs in traditional and AI servers).


- **The CPU remains the command center of the system, responsible for executing sequential logic and handling complex branching operations.** It oversees system-level coordination, including memory allocation and input-output management, ensuring that all components operate cohesively. In addition, the CPU decomposes large, complex problems into manageable sub-tasks and intelligently determines how to distribute these tasks across the GPU and QPU, assigning each portion of the workload to the processor best suited for efficient execution.
- **The GPU serves as the high-throughput workhorse for classical parallelism, excelling at large-scale tensor operations and intensive numerical computation.** It plays a central role in powering machine learning workflows, including training, inference, and optimization pipelines. In a hybrid architecture, the GPU also prepares and transforms data before it is sent to the QPU, ensuring it is in the correct form for quantum execution. When direct access to quantum hardware is limited, the GPU can simulate quantum circuits to approximate behavior and test algorithms. Additionally, it contributes to the reliability of hybrid computations by performing error mitigation, calibration, and verification of QPU outputs, acting as a crucial bridge between classical and quantum processing.
- **The QPU will tackle problems that are fundamentally intractable for classical systems due to their exponential complexity.** It excels at solving optimization challenges with vast combinatorial search spaces and simulating quantum systems such as molecules and advanced materials at a level of detail beyond classical capabilities. By executing specialized quantum algorithms - including variational circuits, the Quantum Approximate Optimization Algorithm (QAOA), and Shor-type routines - the QPU can explore complex solution landscapes that classical processors cannot efficiently navigate, offering a powerful tool for addressing some of the hardest computational problems.

IBM is the primary driver of this hybrid paradigm. They explicitly reject the idea of quantum as a standalone system. They have formalized this under a concept they call **Quantum-Centric Supercomputing**, with reference design provided in March 2026, which maps out a unified environment dividing the compute tiers across high-speed connections (like Ultra Ethernet and NVQLink). On the software front, IBM relies on **Qiskit**, their open-source software development kit. When a developer writes code in Python using Qiskit, the Qiskit Compiler automatically targets the different execution constraints. It parses the developer's high-level intents, maps the traditional programming loops to the local CPU/GPU cluster, and compiles the complex

abstract math directly into targeted hardware instructions for the physical qubits.

NVIDIA's **NVQLink** establishes an open, hardware-agnostic platform architecture that directly connects GPU-accelerated servers with third-party Quantum Processing Units (QPUs), enabling the QPU to operate as a native co-processor within the same rack. On the software front, NVIDIA provides the architectural allocation layer through **CUDA-Q**. This single-source programming framework allows developers to write unified C++ or Python code, using its built-in compiler to automatically split workloads—offloading massive AI tensor operations to the GPU, system logic to the CPU, and complex combinatorial algorithms straight to the QPU.

EXHIBIT 1: The Tri-Processor Architecture



Source: Bernstein analysis


EXHIBIT 2: CPU, GPU, QPU architecture, use case comparison

Dimension	CPU	GPU	QPU
Core concept	A few powerful cores optimized for sequential, general-purpose logic	Thousands of simpler cores optimized for massively parallel arithmetic	Qubits exploiting superposition & entanglement to explore many states at once
Architecture	Low core count (4–64), large caches, complex control logic, high clock speeds	Very high core count (thousands), SIMD/SIMT execution, high memory bandwidth	Qubits (gate-based or annealing); requires cryogenic cooling / error correction
Best at	Branching logic, decision-making, low-latency single-threaded tasks	Doing the same operation across huge data sets in parallel	Specific problem classes with exponential/combinatorial complexity
Compute model	Deterministic, serial → moderate parallelism	Deterministic, throughput-oriented parallelism	Probabilistic; results obtained from repeated measurement (sampling)
Parallelism	Low to moderate	Extremely high (data parallelism)	Quantum parallelism (fundamentally different — not just "more cores")
Tasks best addressed	OS/application logic, databases, business software, control flow, I/O orchestration	Deep learning training/inference, graphics rendering, simulations, matrix/vector math, crypto mining	Quantum chemistry/materials simulation, certain optimization, factoring, quantum sampling/ML research
Best-fit use cases	Running operating systems, web/app servers, transactional systems, spreadsheets, general everyday computing	AI/ML, scientific computing (CFD, weather, molecular dynamics), video/3D rendering, large-scale data analytics	Molecular & drug discovery, advanced materials, portfolio/logistics optimization, cryptography research, R&D (largely experimental today)

Source: Bernstein analysis

EXHIBIT 3: Quantum-classical workflow

Quantum-classical workflow



Source: IBM

PART II: WHAT IS QUANTUM COMPUTING? THE TECH BASICS

In quantum computing, once physical qubits are created, quantum algorithms will manipulate these qubits through carefully sequenced operations that create superposition, entangle qubits, and use interference to amplify correct outcomes while canceling incorrect ones. After this controlled evolution, the system is measured, which collapses the quantum states into classical bits that encode the solution. If quantum computing can become powerful and stable enough, their unique computational power will solve problems that are beyond the capabilities of today's supercomputers.

Superposition: Quantum computing differs fundamentally from classical computing because its basic unit, the qubit, can exist in superposition, representing multiple states simultaneously, rather than a single binary 0 or 1 like in classical computing. This allows quantum computing to have a much larger computational space and explore many possible solutions at the same time, whereas classical computing evaluates them sequentially.

Entanglement links qubits, so their states are defined jointly rather than independently, allowing quantum algorithms to encode relationships and constraints between variables that guide interference toward valid solutions, which classical computers must track explicitly and often less efficiently for certain problems.


Interference allows quantum algorithms to amplify correct solutions and suppress incorrect ones by controlling how probability amplitudes combine. Through carefully designed sequences of operations, amplitudes associated with valid outcomes reinforce each other, while those linked to invalid outcomes cancel out. This reshaping of the probability distribution ensures that, upon measurement, the correct or optimal solution is observed with higher likelihood than under classical computing for certain problems.

Measurement is the final step in quantum computation, where the quantum state is observed and collapses into a classical output that can be read by conventional computers. Because quantum information is inherently probabilistic, a single measurement yields the most likely outcome drawn from the probability distribution. By repeating the computation multiple times, the optimal solution emerges statistically.

Key potential use cases for quantum computing include:

- **Material science and chemistry:** simulating molecules and reactions for drug discovery, catalysts, and new materials
- **Optimization problems:** logistics, scheduling, portfolio optimization, and supply chain routing
- **Cryptography and security:** breaking or strengthening encryption schemes
- **Machine learning and AI:** accelerating specific linear algebra and sampling tasks
- **Financial modeling:** risk analysis and Monte-Carlo simulations where the set of possibilities explode exponentially.

EXHIBIT 4: IBM System Two



Source: Company website

PART III: WHAT ARE THE LEADING QUANTUM MODALITIES?

Quantum computing is currently being pursued through several hardware approaches, each based on a different physical system for creating and controlling qubits, and each with its own strength, limitation, and maturity curve.

The leading modalities today include:

- **Superconducting**, which has fast gate speed, scalability, and dominates near-term commercial roadmaps;
- **Trapped-ion**, known for exceptional stability but slower speed;
- **Neutral-atom** arrays, demonstrating high qubit counts but difficult to achieve high fidelity gates;
- **Photonic**, an early-stage technology using optical components;
- **Topological** qubits, a high-risk, high-reward path targeting inherently stable quantum information; and
- **Silicon spin qubits** have the potential for dense integration of qubits and control electronics, but are slower and harder to control since they rely on manipulating the spin of tiny electrons.

Together, these six approaches capture the core technological landscape shaping the quantum computing industry. We believe Superconducting, Trapped-ion, and Neutral-atom are the leading contenders for the Quantum future - more details below.

EXHIBIT 5: Quantum Computing Modalities

Modality	Technology	Pros	Cons	Major Players
Superconducting	Metal circuits on a chip are cooled to absolute zero, and microwave pulses put the circuit into a quantum state and read it out	Fast gate speeds and high qubit counts; leverages existing microelectronic lithography.	Short coherence and high sensitivity to noise; requires extreme cryogenic cooling (~10-20 mK); wiring and crosstalk make scaling harder	IBM, Google, Rigetti, D-Wave (annealing), IQM
Trapped-ion	Single charged atoms float in a vacuum, held by electric fields. Lasers or microwaves set each ion's state, and the ions' shared motion is used to entangle them	Identical natural ions ensure high qubit quality; long coherence times and extremely high gate fidelities.	Slower gate speeds than superconducting; complex laser/optical control systems hinder scaling.	IonQ
Photonic	The qubit is made from light. The quantum state is stored in a photon's path, polarization, or arrival time, and optical parts steer and measure it	Very fast operations; photons resist decoherence well; good for networking and modular scale; mostly can run near room temperature	Photon loss is a major problem; high-fidelity photon generation/detection are hard; some detectors still need cryogenics	PsiQuantum, Xanadu, Quantum Computing Inc., ORCA Computing
Neutral-atom	Neutral atoms are held in place by laser tweezers in a vacuum. Lasers prepare their states, move them around, and briefly excite nearby atoms so they become entangled	Strong scaling potential, long coherence times	Gates are slower than superconducting; fidelity is improving fast but trails trapped-ions, large optical control setups are complex	QuEra, Pasqal, inflection, Atom Computing
Topological	The goal is to store quantum information in a protected pattern of special quasiparticles called Majorana zero modes in nanowires, instead of one fragile local state	If it works, it could give built-in error protection (Majorana quasiparticles), better coherence/fidelity, and compact scaling.	Highest technical difficulty; materials science is largely unproven at scale and remains experimental.	Microsoft
Silicon Spin	A qubit is the spin of one electron trapped in silicon, usually inside a tiny quantum dot or around a single atom. Voltages plus microwave/electric pulses set and read the spin state	Extremely small footprint (high density); potential for billion-qubit scale on a single chip using mature semiconductor infra; long coherence	Strict fabrication tolerances and decoherence from material impurities; currently in earlier development.	Intel, Diraq, Silicon Quantum Computing

Source: Company websites, Bernstein analysis

Superconducting qubits is today's most commercially advanced and execution-ready technology, driven by rapid development from IBM and Google. These systems use tiny electrical circuits cooled near absolute zero, allowing electrons to behave like qubits. Then, the qubits are manipulated using precisely shaped microwave pulses. The appeal is straightforward: superconducting qubits are **fast, can be fabricated with semiconductor-style processes, and already support multi-qubit processors in the hundreds**. Because superconducting chips can be produced using modified semi processes, it

enables faster generation of larger, more complex processors. However, superconducting is **relatively unstable** because it is a quantum state that only survives when the environment temperature and noise level are very low. While unstable, industry leaders like IBM and Google compensate through faster gates, deeper engineering talent, and the most mature hardware-software stacks in the market.

Superconducting hardware underpins a second, more specialized quantum model known as **quantum annealing**, which is optimized for solving large-scale optimization problems rather than running general quantum algorithms. Quantum annealers, most notably D-Wave systems, use superconducting flux qubits arranged in an energy landscape that the machine gradually “anneals” to find low-energy solutions. Even though Quantum annealing is model-agnostic in theory, it is only implemented on superconducting in practice.

Trapped-ion qubits offer the highest precision and stability in the industry, though the architecture faces slower operational speed and scaling hurdles. The approach traps individual charged atoms in electromagnetic fields and manipulates them with lasers. Because atoms are naturally identical and exceptionally stable, this platform achieves long coherence times and industry-leading gate fidelity, which are critical metrics for error correction. The trade-off is slower operation and more complex optical systems which makes scaling more difficult, but trapped-ion players deliver the most reliable qubits available today. Trapped-ion is the modality adopted by the biggest market cap quantum pureplay, IONQ. Infineon is the leading foundry in trapped ions, supplying QPUs to the largest players in this subsector.

Neutral-atom platforms provide a highly flexible architecture capable of scaling, though fidelity at scale remains a challenge. These systems use lasers to hold and reposition individual atoms in optical “tweezers”, allowing rapid rearrangement. Neutral-atom hardware has demonstrated impressive qubit counts and potential for quantum simulation workloads. The main challenge is achieving consistently high-fidelity gates at scale, as complex optical setups introduce noise and calibration demands. While promising, neutral-atom systems still trail the industry maturity.


Photonic quantum computing offers long-term scalability potential through room-temperature operation, but the core building blocks are still early-stage. Using single photons routed through optical circuits (e.g. beam splitters, mirrors, phase shifters, etc.), photonic systems avoid cooling needs and align naturally with silicon photonic manufacturing. This positions the technology for eventual mass production if key components, such as deterministic single-photon sources and high-fidelity gates, reach maturity. However, since photons naturally do not interact with each other, error rates remain high due to probabilistic photon generation, photon loss, and imperfect interference. Today, error rates and component constraints keep photonic earlier stage relative to superconducting systems and trapped-ions.

Topological qubits represent a high-upside but highly experimental path, targeting qubits that would be naturally resistant to noise and far easier to correct. This approach seeks to leverage exotic quasiparticles, such as Majorana modes, to encode information in a way that is fundamentally protected from environmental errors. Each Majorana is effectively “half” an electron, and two halves can pair up to store quantum information in a way that naturally protects it from noise. If realized, topological qubits could dramatically reduce the overhead needed for fault-tolerant computing. However, the underlying physics remains unproven at scale, making this the most speculative of the major technologies. For now, topological efforts remain long-duration R&D, while superconducting and other modalities continue to advance more quickly.


Silicon spin qubits store quantum information in the spin of single electrons trapped in tiny silicon structures, using electrical signals to control them. Because they are built on silicon, this approach could eventually scale using familiar chip-manufacturing techniques and allow dense integration of qubits and control electronics. However, silicon spin qubits are slower and harder to control because they rely on manipulating the spin of individual electrons, an extremely small and delicate signal, making operations more sensitive to noise and fabrication imperfections. This means silicon spin qubits is a longer-term bet rather than a near-term commercial solution.

We view superconducting, trapped-ion, and neutral-atom architectures as the most advanced quantum modalities today, given their stronger technical progress, ecosystem support, and early commercial traction. **Still, it is too early to call a clear winner**, and we would not rule out photonics, topological, or silicon-spin approaches, especially given rapid innovation in quantum computing and limited milestone disclosure from private companies.


EXHIBIT 6: **Superconducting quantum information processor**



Source: Gambetta, J.M., Chow, J.M. & Steffen, M. Building logical qubits in a superconducting quantum computing system. npj Quantum Inf 3, 2 (2017). <https://doi.org/10.1038/s41534-016-0004-0>

EXHIBIT 7: **Schematic of IonQ trapped ion quantum processor**

Source: IonQ company website, <https://ionq.com/news/february-10-2022-duke-ionq-new-qc-gate>


EXHIBIT 8: **Silicon quantum photonic processor**

Source: R. Santagati, J. Wang, et. al.

Witnessing eigenstates for quantum simulation of Hamiltonian spectra. Sci. Adv. 4, eaap9646


(2018). <https://www.science.org/doi/pdf/10.1126/sciadv.aap9646>

EXHIBIT 9: **Neutral atoms quantum computing demonstrator**




Source: Munich Quantum Valley, <https://www.munich-quantum-valley.de/research/research-areas/neutral-atom-qubits>

EXHIBIT 10: **Conceptual Layout of a Topological Quantum Computing Device**



Source: Morteza Aghaee et al, InAs-Al hybrid devices passing the topological gap protocol, Physical Review B (2023). DOI: 10.1103/PhysRevB.107.245423

EXHIBIT 11: **Simplified silicon-spin qubit transistor cross section**

Source: Aspecore Network, EE Times, 05.30.2024, <https://www.eetimes.com/how-intel-quantum-chips-could-retransform-silicon-based-computing/>


PART IV: INTRODUCING THE MAJOR PLAYERS IN QUANTUM COMPUTING**Large companies:**

- IBM (rated Market-Perform)** has spent the past decade advancing its vision of large-scale fault-tolerant quantum computing with one of the most comprehensive [roadmap](#) in the industry. On the [hardware](#) front, IBM develops **superconducting** qubit quantum processors, including the 127-qubit Eagle, 133-qubit Heron r1, and 156-qubit Heron r2 and r3. These quantum processing units (QPUs) are then integrated into quantum systems and data centers. On the [software](#) side, IBM offers Qiskit, the world's leading open-source platform that efficiently translates high-level algorithms into instructions optimized for quantum devices. By combining cutting-edge hardware with powerful software, IBM Quantum delivers a full-stack quantum computing solution aimed at making practical quantum applications a reality.
- Google (covered by Mark Shmulik)** follows a quantum strategy similar to IBM's, employing a full-stack approach that integrates both hardware and software to provide quantum solutions. On the hardware side, Google focuses on **superconducting** qubits, with its latest quantum chip, [Willow](#), designed to exponentially reduce errors as the number of qubits scales up. On the software side, Google offers Cirq and Qualtran, to equip developers with the necessary frameworks to design quantum circuits and algorithms. Currently, Google is at milestone 2 of its six-step quantum computing [roadmap](#).
- Microsoft (covered by Mark Moerdler)** takes a different approach from IBM and Google by focusing on **topological** qubits. In February 2025, Microsoft introduced [Majorana 1](#), the first quantum processing unit (QPU) powered by a topological core. In June 2026, the company unveiled [Majorana 2](#), an upgraded chip that reportedly improves qubit reliability by ~1000x, and moved its scalable quantum-computer target to 2029, though the architecture still requires broader validation. On the software side, Microsoft provides Azure Quantum, which includes the Quantum Development Kit (QDK) - a specialized toolkit for quantum programming. The QDK supports Microsoft's own Q# language as well as other popular quantum programming frameworks such as Qiskit, Cirq, and OpenQASM.
- Infineon (rated Outperform)** is the leading foundry in trapped ions, supplying QPU to the largest players in this subsector. In quantum computing, Infineon's role is not to build full-stack quantum computers, but to provide the trapped-ion hardware platform that sits at the heart of them. Its roadmap is to build the enabling hardware — ion traps, integrated photonics, and control electronics, and partner with the pure-play companies most likely to scale trapped-ion systems commercially. Infineon started investing in trapped-ion QPUs since 2017, using its semiconductor manufacturing base to develop ion traps, integrated photonics, and control electronics, and to test and characterize those modules in its own quantum lab. In 2021, Infineon launched the OptoQuant project with the University of Innsbruck to work on ion-based processors with integrated optics and lay the basis for 100+ entangled qubits. In 2022, Infineon opened a quantum test lab

in Villach to speed up test cycles for industrially manufactured ion-trap modules. In the same year, it partnered with [Oxford Ionics](#) to develop high-performance, fully integrated trapped-ion QPUs with a path to hundreds of qubits. IonQ acquired Oxford Ionics in 2025. Infineon claims its ion traps, now in their third generation (Gen 3), are designed to be “predictable, repeatable, and reliable,” and that scaling to larger systems will require further integration of optics and cryogenic control electronics (Exhibit 12).

- **Intel (covered by Stacy Rasgon)** builds quantum computers based on **silicon spin qubits**. Intel approaches quantum computing primarily as a hardware and fabrication play, leveraging its core strength in advanced semiconductor manufacturing.

EXHIBIT 12: Infineon is now into the third generation of its ion traps; it adds electrodes in the third dimension to increase ion confinement by a factor of 10 compared with standard surface traps



Source: Infineon

Quantum pure plays: IonQ, Rigetti, D-Wave and Quantum Computing went public during 2021-2022. Infleqtion and Xanadu were listed in early 2026. Except for QUBT which was originally listed on OTC then up-listed to Nasdaq, the rest five quantum computing pure-plays went public via mergers with Special Purpose Acquisition Companies (SPACs) - reflecting the early-stage, research-intensive nature of quantum computing industry and long commercialization timelines.

- **IonQ (Not Covered)** uses **trapped ions** quantum computing. IonQ choose to focus on trapped ions because of its ultra-high fidelity, all-to-all connectivity and long coherence. Since launching Aria in 2022, IonQ has released Forte, Forte enterprise and aimed to launch its next-gen “Tempo” system in 2026, which is built to have faster gate speeds, mid-circuit measurement, and 99.9% fidelity. IonQ’s partners include Nvidia, Amazon Braket, Azure Quantum, Google Cloud, etc. IonQ 2026 revenue is estimated to be \$260mn (Bloomberg consensus) and has a current market cap of \$21bn.
- **D-Wave Quantum (Not Covered)** focuses on **quantum annealing**, a specialized approach designed to solve optimization problems rather than run general purpose quantum algos. Its latest Advantage2 System expands qubit count and connectivity for large-scale optimization tasks, and the company continues to emphasize near-term enterprise use cases where annealing offers practical performance benefits. D-Wave’s customers span across automotive, manufacturing, logistics, and government users. D-Wave 2026 revenue is estimated to be \$43mn (Bloomberg consensus) and a current market cap of \$8.7bn.

- **Rigetti Computing (Not Covered)** builds quantum computers using **superconducting**, similar to IBM and Google. The company incorporates a full-stack approach and offers access to its hardware via its Quantum Cloud Service (QCS), which can be integrated into any public, private or hybrid cloud. Rigetti has iterated on its multi-chip modular processor design and introduced improved error-mitigation techniques aimed at enabling larger, more commercially oriented quantum workloads. The company maintains partnerships with Amazon Braket and Azure Quantum. Rigetti 2026 revenue is estimated to be \$24mn (Bloomberg consensus) and a current market cap of \$6.8bn.
- **Xanadu (Not Covered)** is a prominent Canadian quantum computing company founded in 2016. The stock is dual-listed on the Nasdaq and Toronto Stock Exchange in March 2026, making it the market's first **pure-play photonic quantum computing company**. Xanadu develops on-chip, error-resistant photonic hardware alongside a comprehensive software stack. They are globally recognized for **PennyLane**, their highly popular, open-source quantum software library used widely by developers for quantum machine learning and application building. Xanadu 2026 revenue is estimated to be \$8mn (Bloomberg consensus) and a current market cap of \$3.8bn.
- **Infleqtion (Not Covered)** originally founded in 2007 as ColdQuanta, is a global leader in **neutral-atom quantum technology**. The company went public in February 2026 via a SPAC merger. Unlike companies focused exclusively on computing, Infleqtion offers a broad, commercial portfolio including **quantum computers** (Sqale), **quantum atomic clocks** (Tiqker), **RF receivers** (Sqywire), and **inertial sensors** for GPS-denied navigation. Everything is tied together by their proprietary Superstaq software platform. Infleqtion 2026 revenue is estimated to be \$41mn (Bloomberg consensus) and a current market cap of \$3.1bn.
- **Quantum Computing (Not Covered)** emphasizes **photonic** computing, with a focus on software, algos and room-temperature hardware concepts rather than large scale quantum processors. The company targets enterprise and government users seeking near-term quantum-enabled performance gains. QUBT 2026 revenue is estimated to be \$22mn (Bloomberg consensus) and a current market cap of \$2.2bn.

Except for Quantum Computing which is primarily hardware-led in quantum offering, the rest five quantum pureplays are full-stack, building both quantum processor and software layer to program.

There are also a few players in China working on the tech, please see more details in our previous report [Quantum Computing - Hype or Future?](#).

EXHIBIT 13: Quantum Companies financial info

\$mn except for EPS and # of employees	Market Cap	Revenue (2026 consensus est.)	# of employees	Asset	Cash	Debt	EBITDA	Adj. EPS	P/Sales	P/Employee
IONQ	21,194	260	1,132	6,570	1,031	30	(546)	(\$0.60)	81.5x	18.7x
QBTS	8,747	43	388	916	635	43	(97)	(\$0.26)	204.4x	22.5x
RGTI	6,874	24	162	667	45	7	(74)	(\$0.16)	291.3x	42.4x
XNDU	3,869	8	240	227	0	0	(3)	/	477.7x	16.1x
INFQ	3,177	41	250	425	1	0	(8)	/	77.5x	12.7x
QUBT	2,245	22	72	1,619	738	3	(47)	(\$0.11)	100.7x	31.2x

Source: Company reports, Bloomberg (updated as of 6/5/2026), Pitchbook, Bernstein analysis

PART V: QUANTUM ADVANTAGE AND REAL ECONOMIC VALUE

Published TAM estimates for Quantum are still all over the map, which reflects how differently industry experts are defining the market and underwriting the timing of commercialization. We believe BCG provides the best base case because it ties provider revenue to the underlying technology roadmap, with limited monetization before quantum advantage, broader adoption through 2030-2040 as advantage expands, and full scaling only once fault-tolerant systems emerge beyond 2040.

Quantum advantage is the point at which a quantum system can perform a commercially meaningful task that classical computers cannot feasibly match. That threshold matters because, until it is reached, demand is likely to remain concentrated in research, government, and pilot activity rather than broad enterprise deployment. The most promising use cases are already visible, particularly in simulation-heavy workloads such as drug discovery, advanced materials, catalyst design, financial risk, and cryptography, but meaningful commercial value depends on crossing from technical promise to repeatable

economic superiority. **In our view, that makes it important to frame near-term TAM conservatively.**

We believe BCG offers the most credible published view of provider TAM because it links commercialization directly to technology readiness, with limited near-term value before quantum advantage and broader monetization only as fault-tolerant systems emerge. That technology-based framing makes BCG more grounded than approaches with ambitious near-term TAM. Just as important, BCG separates total economic value from provider revenue, estimating \$450 billion-\$850 billion of economic value versus a \$90 billion-\$170 billion annual provider market by 2040. This distinction avoids overstating the revenue opportunity for providers by confusing ecosystem value creation with supplier capture.

How can we achieve quantum advantage? Rigetti Computing's CEO has pointed out four technical criteria that must be met simultaneously. First, systems must operate at sufficient scale, generally defined as at least ~1,000 logical qubits. Second, system-level fidelity must reach approximately 99.9% or better so that errors do not overwhelm computation. Third, gate speeds must be fast, on the order of tens of nanoseconds, so that full quantum algorithms can execute within the coherence window. Fourth, systems must support real-time error correction. All four criteria must be satisfied together; exceeding one dimension does not compensate for weakness in another. Recent results such as Google's Willow chip are viewed as encouraging primarily because they demonstrate that error correction performance can improve with scale, reinforcing confidence that these thresholds are achievable.

When will quantum advantage be achieved, and what are the bottlenecks? While qubit counts can be scaled, fidelity tends to collapse as systems grow due to interference effects and increased noise across large architectures. Maintaining long enough coherence times remains challenging, as qubits often decohere before algorithms of meaningful depth can be executed. Error rates also accumulate faster than they can currently be corrected at scale, preventing sustained, fault-tolerant operation. In short, the industry's constraint is the ability to preserve qubit quality - fidelity, coherence, and error-correction - as systems scale to the size required for quantum advantage. The timing of quantum advantage ultimately depends on resolving the remaining engineering bottlenecks, which Rigetti's CEO estimates will take **approximately 3-4 years.**

Real economic value will only be unlocked after quantum advantage is achieved. BCG estimates suggest \$1-2bn annual value for providers before 2030, driven primarily by government labs, universities, and early-stage enterprise experimentation. Once broad quantum advantage is achieved in 2030-2040, providers may see roughly \$15-30bn annual value, driven by initial commercial use cases such as drug discovery, materials modeling, and financial optimization. Last but not least, BCG estimates that after 2040, full-scale fault tolerance will be achieved, at which point we will see greater monetization from broad enterprise adoption, bringing \$90-170bn in annual provider revenue.

PART VI. QUANTUM TAM AND VALUATION

Based on the analysis above, we believe that meaningful commercialization of quantum computing will require time and multiple technical and economic breakthroughs before achieving broad adoption. As a result, near-term demand is likely to remain limited, with the most material value creation occurring further out in the adoption curve. Accordingly, our TAM framework is explicitly long-term in nature.

We anchor our analysis on the estimated quantum computing TAM in 2040, which we size at approximately \$90-170bn ([Exhibit 14](#)), reflecting expectations for widespread commercial viability across select enterprise and specialized use cases. We then discount this long-term TAM back to 2026 to account for the extended commercialization timeline, execution risk, and the time value of money. On a present value basis, this yields an implied quantum TAM of approximately \$26.6bn today ([Exhibit 15](#)).

We then extend this framework to assess whether current valuations of publicly listed quantum computing companies can be reasonably justified. We acknowledge upfront that valuing quantum companies at this stage is inherently difficult, if not speculative, given their nascent technologies, limited revenues, negative earnings profiles, and highly uncertain commercialization timelines. Any valuation exercise at this point necessarily relies on simplifying assumptions and long-dated projections. That said, we attempt to approach this analysis in a structured and internally consistent manner to frame expectations rather than to derive precise valuations.

Given the early-stage nature of these businesses, we assume that meaningful profitability only emerges over the long term. Specifically, we assume that quantum companies have achieved net income positive by around 2040. At that stage, we assume these companies trade at market-average valuation multiples, reflecting normalized growth, profitability profiles once the technology matures.



Using current market capitalizations and applying an assumed steady-state market multiple, we back out the level of net income that today's valuations appear to be implying. To translate implied profitability into revenue, we reference the margin profiles of established hardware and computing infrastructure companies, using this average ~30% net-margin as a proxy for long-term quantum industry economics ([Exhibit 16](#)). This allows us to estimate revenue from the implied net income embedded in current valuations.

Finally, we compare this implied revenue to the present value of the overall quantum computing TAM derived earlier. By dividing implied company-level revenue by the discounted TAM, we arrive at the implied long-term market share ([Exhibit 17](#)) that investors are effectively assigning to each company according to today's trading market caps. While this exercise is highly sensitive to assumptions and should be interpreted with caution, it provides a useful lens to assess whether current valuations embed conservative, reasonable, or aggressive expectations relative to the long-term opportunity set.

EXHIBIT 14: Projected TAM from various market research sources

Source	2024-26 TAM (bn)	2030 TAM (bn)	2035 TAM (bn)	2040 TAM (bn)	Implied CAGR
Juniper Research	\$2.7	\$9.4	—	—	~23%
Markets and Markets	\$3.5	\$20.2	—	—	42%
Precedence Research	\$1.4	—	\$19.4	—	30%
McKinsey QT Monitor 2025	\$4.0	—	\$28-72	—	19-30%
BCG	—	—	—	\$90-170	—

Source: Juniper research, Markets and Markets, Precedence research, Mckinsey Quantum Technology monitor 2025, BCG

EXHIBIT 15: Quantum TAM discounted to 2026

Future value at 2040			
	Conservative	Aggressive	Average
BCG	90	170	130

Present value at 2026, with below discount rate assumptions			
Discount rate	Conservative	Aggressive	Average
8%	\$30.6	\$57.9	\$44.3
12%	\$18.4	\$34.8	\$26.6
15%	\$12.7	\$24.0	\$18.4
20%	\$7.0	\$13.2	\$10.1

Source: BCG, Bernstein estimates and analysis

EXHIBIT 16: Quantum Comparables in AI - Net Income Margins

Company	3Y Avg	5Y Avg
NVIDIA	54.7%	47.4%
IBM	15.3%	14.5%
AMD	19.9%	20.8%
Micron	0.0%	11.2%
Broadcom	50.0%	49.1%
Qualcomm	28.7%	29.5%
Arm Holdings	40.9%	34.3%
TSMC	41.5%	41.4%
Net Margin Peer Average	31.4%	31.0%
Net Margin Peer Median	34.8%	31.9%

Adjusted/Reported Net Income Margin (%) (TSMC shown on reported TIFRS)
Source: Bloomberg, Bernstein analysis

EXHIBIT 17: Implied market share from current day market cap

Company	Market Cap (\$, bn)	Assumed P/E	Implied Net Income (\$, bn)	Implied Revenue (\$, bn)	Implied Market Share
IONQ	25.0	28x	0.9	3.0	11%
QBTS	10.3	28x	0.4	1.2	5%
RGTI	8.1	28x	0.3	1.0	4%
XNDU	4.3	28x	0.2	0.5	2%
INFQ	3.8	28x	0.1	0.4	2%
QUBT	2.5	28x	0.1	0.3	1%
Total	54.0		1.9	6.4	24%

Source: Bloomberg, Bernstein estimates and analysis

PART VII. IS IT TOO EARLY TO CALL WINNERS AND LOSERS?

We believe that superconducting, trapped-ion, and neutral-atom are the most advanced quantum modalities today, reflecting a strong combination of technical progress, ecosystem development, and commercial traction among the approaches currently being pursued. In particular, superconducting systems appear to benefit from fast gate speeds and a relatively mature development ecosystem, trapped-ion systems are differentiated by long coherence times and high gate fidelity, and neutral-atom approaches stand out for strong scaling potential ([Exhibit 5](#)).

That said, we think it remains far too early in the industry's development to declare an ultimate winner, and we would not rule out photonics, topological, or silicon-spin architectures at this stage, particularly given the pace of innovation across the space and the limited degree of public disclosure from a number of private companies. In other words, while superconducting, trapped-ion, and neutral-atom appear to be leading on current evidence, we view that lead as meaningful but not definitive in a market that has not yet reached Quantum Advantage.

Also, we believe the market is unlikely to be winner-take-all. Different modalities bring distinct strengths and tradeoffs, which should make certain architectures better suited for specific use cases, workloads, and time horizons. Some may be better positioned for near-term commercial relevance, while others may offer the stronger long-term path to scaled, fault-tolerant quantum computing. In our view, that points to a multi-modal market structure for the foreseeable future, with different architectures likely to coexist as the technology and end-market demand evolve.

Against that backdrop, our implied market share framework suggests that certain companies have more attractive risk reward than others. We believe certain companies have current market caps that imply only a modest slice of the long-term quantum opportunity despite exposure to one of the leading modalities. In that context, RGTI (Not Covered) and INFQ (Not Covered) stand out to us: while RGTI is tied to superconducting and INFQ to neutral atoms, our framework indicates that current valuations imply only about 4% and 2% long-term market share, respectively. Put differently, the market does not appear to be pricing in significant share for either company today, so if either ultimately captures larger market opportunity, the upside could be meaningful. We therefore see the setup as more favorable on a risk-reward basis, as investors can gain exposure to leading quantum modalities at relatively attractive valuations. In our analysis, the major pure plays collectively account for just 24% of the long-term market opportunity, implying that the remaining 76% is currently associated with larger diversified companies such as IBM and Google, as well as private players including Alice & Bob and Pasqal. That said, we do not view this split as fixed: it may reflect that current pure-play market caps do not fully capture their ultimate market share potential, leaving room for select companies to gain additional share over time if they execute well and their modalities continue to advance.