Morgan Stanley | RESEARCH

NORTH AMERICA INSIGHT

October 15, 2026 07:58 AM GMT

Semiconductors

Quantum Pure Plays: Who Leads the Race

Quantum computing could unlock a $100bn+ annual market by 2040, but commercial adoption remains uncertain. Our proprietary framework grades companies on demonstrated technology, execution and manufacturing. Quantinuum ranks highest yet trades at the lowest multiple, driving our upgrade to OW. AlphaWise α

Key Takeaways

- We upgrade Quantinuum to OW with a $62 PT (~88% upside), retain EW on IonQ and initiate Infleqtion and Xanadu at EW.
- Quantinuum leads our covered pure plays on published technical results, followed by IonQ, with Infleqtion and Xanadu at earlier development stages.
- Quantinuum offers the most attractive valuation, trading at the lowest multiple of our 2030E computing revenue despite leading on technology and execution.
- 2027 is a key test year across the pure plays; Sol is Quantinuum's next catalyst, while Apollo's 2029 timing is the main risk.

We upgrade Quantinuum to Overweight (PT $62, ~88% upside) as our top quantum play, stay Equal-weight on IonQ and initiate coverage on Infleqtion and Xanadu at Equal-weight. Quantum computing could open a $100bn+ market by 2040, but different hardware approaches make companies hard to compare, and revenue is still too small to separate them. We built our proprietary framework to judge each company on what it has demonstrated across four main areas: technical progress, execution, manufacturability, and broader indicators of progress. We weight technology most heavily today, because that is where the evidence is strongest and the differences between companies are clearest. Quantinuum stands out on technology and execution: it is furthest along among the pure plays and has the strongest execution record, yet the market assigns it the lowest multiple of our 2030 revenue in the group.

Large potential market, uncertain timing. We see a quantum market of about $10bn in 2030 that could grow to more than $100bn a year by 2040 if fault-tolerant machines reach broad commercial adoption. No company has built one yet, and current revenue tells investors little about where these businesses end up. Company roadmaps converge on early fault-tolerant systems in 2029-30, so we judge the stocks on how credibly each is progressing toward that window.

- **Logical qubits are the metric that matters.** Physical qubit counts and gate fidelities are reported differently across architectures, making direct comparisons difficult. Our eight stage proprietary framework tracks

Joseph Moore
Equity Analyst
Joseph.Moore@morganstanley.com +1 212 761-7516

Ella Tulchinsky
Equity Analyst
Ella.Tulchinsky@morganstanley.com +1 212 761-2222

Erik W Woodring
Equity Analyst
Erik.Woodring@morganstanley.com +1 212 296-8083

Brian Nowak, CFA
Equity Analyst
Brian.Nowak@morganstanley.com +1 212 761-3365

Rauf Ural
Research Associate
Rauf.Ural@morganstanley.com +1 212 761-5958

Wen Zhang, CFA
Equity Strategist
Wen.Zhang@morganstanley.com +1 212 761-1732

Cate Folan
Research Associate
Cate.M.Folan@morganstanley.com +1 212 296-3520

Mason Wayne
Research Associate
Mason.Wayne@morganstanley.com +1 212 761-6012

Nicole Kozhukhov
Research Associate
Nicole.Kozhukhov@morganstanley.com +1 212 761-1636

Shane Brett
Equity Analyst
Shane.Brett@morganstanley.com +1 212 761-1022

SEMICONDUCTORS	
North America Industry View	**Attractive**

WHAT'S CHANGED		
IonQ Inc (IONQ.N)	From	To
Price Target	$49.00	**$43.00**
Quantinuum (QNT.O)	From	To
Rating	Equal-weight	**Overweight**
Price Target	$78.00	**$62.00**

investors should be aware that the firm may have a conflict of

Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

demonstrated progress toward reliable logical qubits capable of useful computation. On this measure, Quantinuum leads the pure plays, followed by IonQ, with Infleqtion and Xanadu at earlier stages. We also look at assessments of roadmap execution, where Quantinuum again leads, and manufacturability, where IonQ has a slight edge.

Exhibit 1: Ratings and key findings

				WHAT THEY HAVE DEMONSTRATED			2030E		
COMPANY	RATING	PT	UPSIDE / (DOWNSIDE)	TECH READINESS	ROADMAP EXECUTION	MANUFACTURABILITY	VISIBILITY FACTOR [1]	COMPUTING REVENUE	RISK-ADJ. EV / SALES [1]
QUANTINUUM (QNT)	Overweight ↑	$62	88%	Stage 7	Relatively high	Moderate	75%	$2.0bn	12.8x
IONQ (IONQ)	Equal-weight	$43	37%	Stage 4	Moderate-High	Moderate-High	70%	$1.8bn	11.9x
INFLEQTION (INFQ)	Equal-weight (new)	$12	28%	Stage 3	Moderate	Moderate	55%	$252mn	13.8x
XANADU (XNDU)	Equal-weight (new)	$3.75	18%	Stage 1	Earlier-stage	Earlier-stage	50%	$150mn	12.5x

Note: = upgrade; (new) = initiation. Technology readiness = highest stage demonstrated on our eight-stage framework. Visibility factors are our judgment, informed by the framework; technology readiness carries the most weight.
[1] Risk-adjusted EV / Sales = growth-stage multiple (17x QNT/IONQ, 25x INFQ/XNDU) × our visibility factor; applied to 2030E computing revenue and discounted at a 14.6% WACC

How we value the stocks. We apply a risk-adjusted EV/Sales multiple to 2030E computing revenue and discount back at a 14.6% WACC. The multiple combines a growth-stage multiple, which reflects what that revenue is worth if a platform succeeds, with a visibility factor from our proprietary framework, which reflects how likely that is. Today's share prices imply that the market is pricing in about 43% visibility in Quantinuum, the lowest in the group, against our 75%. We forecast the highest 2030E revenue in the group for Quantinuum, but even if we assume it matches IonQ's, the next largest pure play, the market's implied visibility rises only to about 47%, still below IonQ's 52%. Our higher Quantinuum forecast, which depends on Apollo arriving in 2029, drives the rest of the upside to our $62 PT.

Exhibit 2: Base-case vs. market-implied EV/Sales multiple on 2030E computing revenue

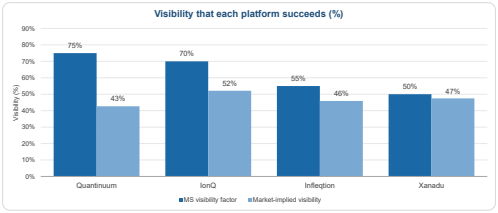
2030E EV / SALES BUILD	QNT	IONQ	INFQ	XNDU
Growth-stage multiple (if the platform succeeds)	17.0x	17.0x	25.0x	25.0x
× Our visibility factor	75%	70%	55%	50%
= Our base multiple (growth-stage × visibility)	**12.8x**	**11.9x**	**13.8x**	**12.5x**
Market-implied multiple (latest share price)	7.2x	8.9x	11.5x	11.9x
Market-implied visibility	43%	52%	46%	47%
Gap (percentage points)	**+32**	**+18**	**+9**	**+3**

Note: Market-implied multiple is the 2030E EV/Sales multiple at which our base-case price target equals the current share price. Market-implied visibility = market-implied multiple ÷ growth-stage multiple.

2027 is the proof year. Each company has a major system milestone due next year, and each is a test of manufacturing as well as technology. Sol is the next most important hardware milestone for our call, as Quantinuum's first test of scaling its trap and laser control. Timing of Sol and Apollo (2029) is the main risk to our Overweight thesis.

Key risks. Quantum computing remains a highly speculative investment area. No company has yet built a fault-tolerant quantum computer, and many of the results we track come from controlled experiments that have not been shown at commercial scale. Our valuations rest on revenue that is still several years away, so slower technical progress or later commercial adoption would push value further out. It is also too early to know which hardware approach will prove most scalable,

Exhibit 3 : Base-case vs. market-implied EV/Sales multiple on 2030E computing revenue



Note: Market-implied visibility is the market-implied EV/Sales multiple on 2030E computing revenue divided by our growth-stage multiple.

and the market could end up more fragmented than our forecasts imply. These are pre-profit businesses, and delays would increase their reliance on outside funding and the risk of dilution.

What else is inside:

- Executive Summary
- Our Proprietary Technology Evaluation
- Valuation Framework
- Company Profiles: Quantinuum, IonQ, Xanadu, Infleqtion, Big Tech and Private Companies
- Quantum Industry Primer

relation to a definitive agreement to acquire Elisity, Inc. as announced on October 12th, 2026. This transaction is subject to customary closing conditions, including receipt of required regulatory approval. IonQ has

to the notes at the end of the report.

Executive Summary

Our Call: Quantinuum is our Top Quantum Play

Exhibit 4: Ratings, price targets and valuation summary

COMPANY	RATING	PRICE	PRICE TARGET	UPSIDE / (DOWNSIDE)	BASE MULTIPLE	MARKET-IMPLIED MULTIPLE	VISIBILITY FACTOR	MARKET-IMPLIED VISIBILITY	BEAR	BULL
QUANTINUUM	Overweight	$33.03	$62.00	+88%	12.8x	7.2x	75%	43%	$10.00	$104.00
IONQ	Equal-weight	$31.42	$43.00	+37%	11.9x	8.9x	70%	52%	$17.00	$69.00
INFLEQTION	Equal-weight	$9.11	$12	+28%	13.8x	11.5x	55%	46%	$3.60	$23.90
XANADU	Equal-weight	$3.19	$3.75	+18%	12.5x	11.9x	50%	47%	$0.70	$9.10

Note: Market-implied multiple is the 2030E EV/sales multiple at which today's share price equals the present value of our base case; market-implied visibility divides it by the growth multiple

We upgrade Quantinuum to Overweight with a $62 PT (~88% upside). It has progressed furthest among the pure plays on our error-correction framework and has the strongest execution record in the group. Yet today's share price implies only about a 43% visibility factor for Quantinuum, the lowest in the group, against our 75% assumption. IonQ has a slight edge on manufacturability, which keeps its visibility factor close to Quantinuum's, but the gap between our view and the market's is far larger for Quantinuum. We stay Equal-weight on IonQ and initiate on Infleqtion and Xanadu at Equal-weight. The market's implied visibilty in all three is closer to our own.

How to Think About Quantum

Quantum computing is a moonshot, and timing is the key uncertainty. If fault-tolerant quantum computers reach broad commercial adoption, the market for hardware, software and services vendors should grow from about $10bn in 2030 to more than $100bn a year by 2040. No company has built one yet. Company roadmaps converge on initial fault-tolerant systems in 2029-30, and today's revenue is too small to show which businesses will succeed. The investment question is which companies are most likely to reach fault tolerance first. For a primer on quantum computing and our market sizing, see the Appendix: Quantum Primer .

How We Assess the Companies

Because financial results say little today, we compare the pure plays on what they have demonstrated, across four areas:

- **Technology readiness:** We focus on logical computing maturity: how far each company has progressed toward large numbers of low-error logical qubits that can run useful calculations. We track this with a proprietary eight-stage framework and give credit only for published hardware results, including preprints, rather than simulations or roadmap targets.
- **Roadmap Execution:** We assess whether each company has delivered on past commitments and analyze current roadmaps.
- **Manufacturability:** We assess how readily each company can build its systems at scale, based on the systems it has in use today, its next-generation system, what gets harder as its systems grow and outside evaluation of its scaling plan.

- **Broader indicators:** We look at government, commercial relationships, and other company specific factors as supporting evidence.

Together, these form our proprietary framework, which sets the visibility factor in each company's valuation. We weight technology readiness most heavily today, because that is where the evidence is strongest and the differences between companies are clearest. We also benchmark the pure plays technology against IBM, Google (IBM covered by Erik Woodring and Google by Brian Nowak).

What We Found

Exhibit 5: Where each company stands on technology, execution, funding and manufacturability

		TECHNOLOGY READINESS	EXECUTION	FUNDING	MANUFACTURABILITY	VALUATION
		Demonstrated / Partial / Not demonstrated	Credibility / History	*Liquidity position and potential additional funding needed* (No / Yes)	*Main scaling hurdle*	*2030E computing revenue × adj. growth multiple* (visibility factor)
PURE PLAYS IN OUR COVERAGE	QUANTINUUM (QNT) Trapped ions	Stage 7: Logical operations	Relatively high	$2.1bn no debt	Laser scaling, integrated optics	12.8x (17x × 75%)
	IONQ (IONQ) Trapped ions	Stage 4: Repeated QEC	Moderate-high	$3.0bn no debt	On-chip CMOS integration for Superion 10K	11.9x (17x × 70%)
	INFLEQTION (INFQ) Neutral atoms	Stage 3: Error correction	Moderate	$582mn no debt	Laser control and parallel operations at scale	13.8x (25x × 55%)
	XANADU (XNDU) Photonic	Stage 1: Encoding	Earlier Stage (short history)	$313mn $32.5mn debt	Optical loss	12.5x (25x × 50%)
TECH BENCHMARKS	GOOGLE Superconducting	Stage 6: Below-threshold scaling	n/a	Alphabet balance sheet n/a	n/a	n/a
	IBM Superconducting	Stage 4: Repeated QEC	n/a	> $10bn quantum investment (5 years) n/a	n/a	n/a

Note: Dots show whether each stage is demonstrated (green), partial (yellow) or not demonstrated (red), based on published hardware results only; stages are capability markers, not a strict sequence (see Exhibit 9; no company has demonstrated stage 8). Execution, manufacturability and visibility factors are our judgments, assessed for the four pure plays only; Google and IBM are technology benchmarks. Liquidity is cash and investments at 2Q26 (IonQ after the July 2026 SkyWater outflow); funding dot is red if we model new equity before 2030. Valuation is base-case EV/Sales on 2030E computing revenue (growth-stage multiple × visibility factor); see methodology section.

Technology readiness: Quantinuum and Google lead on technology, with different strengths. Quantinuum has gone furthest toward computing with protected qubits, with protected operations about 4.3x more reliable than unprotected ones and no runs discarded. Google has the clearest evidence that adding protection keeps lowering errors. IBM and IonQ form a second group, and Infleqtion and Xanadu are earlier. No company has yet demonstrated scaled fault-tolerant computation.

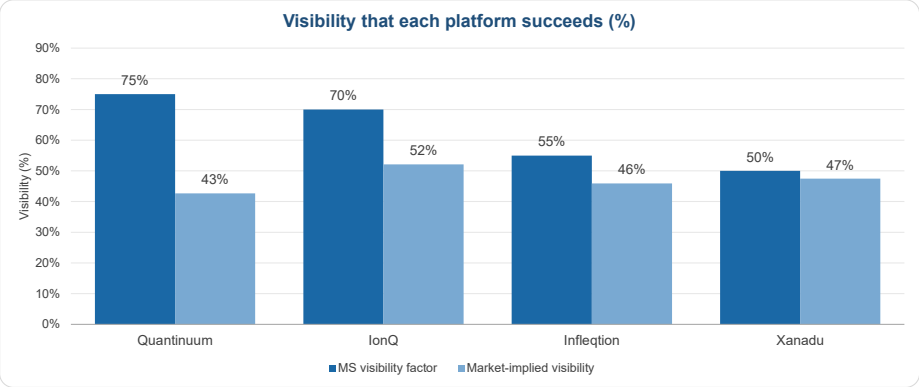
Roadmap execution: Quantinuum also has the strongest execution record. It has delivered three hardware generations on the same architecture and beat its five-year Quantum Volume commitment. IonQ has generally beaten its near-term targets but is moving to a new architecture, Superion, where its error-correction results have not yet been reproduced. Infleqtion and Xanadu have short public roadmap histories, which limits the credit we can give either. Funding is adequate for three of the four: on our model, Quantinuum, IonQ and Infleqtion can fund their roadmaps without new equity. Xanadu has the smallest balance sheet, and we assume will need outside funding to get to 2030 in our model.

Manufacturability: Manufacturing is a closer race, with a slight edge for IonQ. DARPA advanced it to Stage C of its Quantum Benchmarking Initiative, five customers and partners have committed to Superion 256 before it ships, and its design moves gate control from lasers into on-chip electronics. Quantinuum has the deepest record of

building and deploying systems on the architecture it plans to scale, but laser delivery is its main hurdle as systems grow and Lumos (2030+) will require integrated optics. Infleqtion has the largest system installed at a customer site among the four but is still building its manufacturing base. Xanadu's chip-based design has yet to reach a customer site. Manufacturability is harder to assess, so it adjusts our visibility factors rather than drives our analysis at this point.

Valuation: What's Priced In

Exhibit 6: Base-case vs. market-implied EV/Sales multiple on 2030E computing revenue



Note: Market-implied multiple is the 2030E EV/sales multiple at which today's share price equals the present value of our base case; market-implied visibility divides it by the growth multiple.

We value each company on 2030E computing revenue, with a multiple that combines growth and visibility. 2030 is the first year of meaningful commercial scale in our models, and we discount the result back at a 14.6% WACC. A growth-stage multiple values 2030 revenue at each company's stage of the ramp, assuming its platform succeeds. We use 17x for Quantinuum and IonQ, anchored to profitable high-growth semiconductor and hardware companies, and 25x for Infleqtion and Xanadu, which would still have the steepest part of their ramp ahead of them at the end of 2030. A visibility factor from our proprietary framework reflects how likely each company is to reach its 2030 position and keep growing from there. We put Quantinuum at 75%, IonQ at 70%, Infleqtion at 55% and Xanadu at 50%. Our bull case applies the full growth-stage multiple to revenue 30% above our forecast.

The market is least confident in Quantinuum. Dividing each market-implied multiple by our growth-stage multiple gives the visibility today's share prices imply: about 43% for Quantinuum, against about 52% for IonQ, 46% for Infleqtion and 47% for Xanadu (Exhibit 6). Quantinuum leads on our technology and execution work, and we put it at 75%. The gap holds even if we give Quantinuum IonQ's lower 2030 revenue forecast: its implied visibility would rise only to about 47%, still below IonQ's 52%.

Stock Calls

Quantinuum (OW, PT $62). Quantinuum is our top quantum play. It leads the pure plays on error correction and execution yet trades at the lowest multiple in the group. Our forecast has revenue rising from about $270mn in 2028 to $2.0bn in 2030 as Sol and Apollo arrive, so timing is the main risk. Laser delivery is its main manufacturing hurdle as

systems grow, which makes on-time delivery of Sol in 2H27 the key catalyst on both fronts.

IonQ (EW, PT $43). IonQ has the strongest outside validation of its scaling plan, including DARPA's Stage C selection and five Superion 256 customers and partners, but it sits behind Quantinuum on technology. Its best error-correction results come from today's laser-based systems and have not yet been shown on Superion. The market's implied visibility on their quantum business is closer to our view than for Quantinuum, but it does have downside protection for the non-quantum businesses. Superion 256 deliveries in 2027 and a first logical-qubit demonstration on Superion are the key catalysts.

Infleqtion (EW, PT $12). Infleqtion has run protected calculations on up to 30 logical qubits, but its results rely mainly on discarding flagged runs, and it has not yet shown repeated error correction. Its manufacturing base is still being built, and the market's implied visibility is about 9 points below ours. Its sensing business and net cash provide a floor of about $2.70 a share. Repeated correction and on-time delivery of its Illinois system in 2027 are the key catalysts.

Xanadu (EW, PT $3.75). Xanadu's chip-based photonic approach has potential scaling advantages, but it is the earliest of the four on our framework and we model that they may need further funding to reach commercialization. It has a wide range of outcomes, with a ~$9 bull case and a ~$1 bear case. Reducing optical loss, its main technical and manufacturing hurdle, is the next test.

What to Watch

Exhibit 7: Key milestones by company: 2026-2029

COMPANY	2026	2027: KEY TEST YEAR	2028	2029
PURE PLAYS IN OUR COVERAGE				
QUANTINUUM (QNT)	- Helios installation in Singapore (in progress) - Helios sale to Oracle (installation date not given)	- Sol launch (2H27): 192 physical and ~100 logical qubits at ~10^{-5} logical error	—	- Apollo: universal fault-tolerant system with hundreds of logical qubits
IONQ (IONQ)	- Superion 256 system demonstrated (in progress)	- First Superion 256 customer deliveries, including NVIDIA - 10,000-qubit Superion lab system; 800 logical qubits	- Commercial Superion 10K systems - 20,000 physical and 1,600 logical qubits	- 8,000 logical qubits - Photonic interconnect
INFLEQTION (INFQ)	- Full paper on the result (expected in coming weeks)	- Fault-tolerant system for the Illinois Quantum & Microelectronics Park, designed for >50 logical qubits	100 logical qubits	—
XANADU (XNDU)	- Loss ratio of 24.1x (management says on track)	- Qubit Factory technology developed (2026–27 target) - Loss ratio of 10.2x	- Fault tolerance (2028–29 window) - Loss ratio of 7.2x	- Up to 200 logical qubits; loss ratio of 1.8x - Quantum data center (2029–30 window)

2027 is the first real test for every pure play. Each has a major system milestone due by the end of next year, and Quantinuum's Sol is the next hardware proof point for our call. Systems delivered on time, with error correction running through calculations, would raise our confidence. Delays, or logical qubits that still depend on discarding runs, would lower it.

Key Risks: Quantum computing is still an early-stage technology, and outcomes for these companies remain highly uncertain. No company has built a fault-tolerant quantum computer, and many of the results we track come from controlled experiments not yet shown at commercial scale. Our valuations rest on 2030 revenue, so slower technical progress or later adoption would push value further out. A competing approach, could prove more scalable than we expect, and the market could end up more fragmented

than our forecasts assume. These are pre-profit businesses, so delays would increase reliance on outside funding and the risk of dilution, particularly for Xanadu. Valuations are also sensitive to market sentiment and interest rates. On the upside, faster progress toward fault tolerance or consolidation around fewer winning platforms could support higher revenue and visibility factors than we assume.

Analysis Summary

We assess the pure plays across three areas: technology readiness, roadmap execution and manufacturability. Together these form our proprietary framework, which informs the visibility factors in our valuation. We also look at broader indicators as supporting evidence.

Technology Readiness: Logical Computing Maturity

Why logical computing maturity: There are many ways to assess progress across quantum computing efforts. The field is still early, and no company has yet built a commercial fault-tolerant quantum computer. Hardware approaches also differ fundamentally, and companies report different metrics, definitions and performance benchmarks. In this report, we focus on logical computing maturity: how far each company has progressed toward producing large numbers of low-error logical qubits capable of useful computation.

We believe logical computing maturity is the most useful technology measure because it brings many underlying hardware challenges together into a common end goal. Physical qubit count, gate fidelity, connectivity, speed, measurement quality and control all matter, but only insofar as they let a system create low-error logical qubits that can be used for computation. A machine with more physical qubits or better component metrics is not necessarily closer to useful quantum computing if those gains do not translate into better logical performance. Conversely, evidence that a quantum computer can correct errors and become more reliable as it scales is a more direct sign that a platform is progressing toward fault tolerant computation.

We assessed each company's progress in quantum error correction (QEC), logical computing and fault tolerant operations, drawing on published research and company disclosures. Because these technologies remain experimental, many of the reported results were produced in controlled environments and come with important qualifications. To compare progress across companies, we developed an eight stage framework that organizes key milestones in quantum error correction and logical computing into a common measure of maturity (Exhibit 8). We summarize each company's position below, with detailed analysis in the individual company profile sections.

How we measure progress: We segment logical-computing progress into eight stages, which answer three questions in rough order of difficulty:

- *Can the system protect information at all? (Stages 1–3)* The first step is encoding: spreading one piece of quantum information across several physical qubits so that it can be protected (Stage 1). The system then needs to detect when an error has occurred (Stage 2). It also needs to use that information to fix the error, rather than simply throwing out the affected run (Stage 3).
- *Does the protection actually work? (Stages 4–6)* The system must keep correcting the same information over many rounds, as a real calculation would require (Stage 4). The protected information should then perform at least as well as unprotected qubits, a point known as breakeven (Stage 5). Finally, adding more protection should keep lowering errors (Stage 6). We place particular weight on Stage 6,

because it is the clearest sign that an approach will keep improving as systems grow.

- *Can the system compute while protected? (Stages 7–8)* A useful computer must run operations on protected qubits, not just store them, and those operations must stay more reliable than unprotected ones (Stage 7). The end goal is many protected qubits running long calculations together (Stage 8), which no company has yet demonstrated.

Exhibit 8: Our eight-stage framework for technology progress

STAGE	WHAT IT MEANS	WHAT WE NEED TO SEE
1. ENCODING	Spread one piece of quantum information across several physical qubits	Protected (logical) qubits created and read out on hardware
2. ERROR DETECTION	Spot when an error has occurred	Checks on protected qubits that flag real errors
3. ERROR CORRECTION	Use detected errors to fix the protected information	Errors worked out and fixed, live or after the run, with affected runs kept
4. REPEATED QEC	Keep protecting the same information over many rounds	At least three rounds of correction on the same information
5. LOGICAL IMPROVEMENT / BREAKEVEN	Protected information performs at least as well as unprotected	A clear advantage without throwing out flagged runs
6. BELOW-THRESHOLD SCALING	More protection produces fewer errors	Errors fall consistently across at least three levels of protection
7. FAULT-TOLERANT LOGICAL OPERATIONS	Compute on protected qubits while keeping an advantage	Protected operations, including one linking qubits, beat unprotected versions under active correction
8. SCALED FAULT-TOLERANT COMPUTATION	Many protected qubits run long, universal calculations	Universal computation across many protected qubits with continuous correction

Note: Progress against each stage is assessed on published hardware results, including results from outside teams on a company's hardware, not simulations or roadmap targets.

How we rate each stage: We rate each stage as demonstrated, partial or not demonstrated, based on published hardware results rather than simulations or roadmap targets. Results produced by outside teams on a company's hardware also count. Results that depend on throwing out runs where errors were flagged generally earn only partial credit, because that approach becomes impractical as calculations get longer. There are other qualifications that warrant a partial ranking as well. The stages are capability markers rather than a strict sequence, so a company can show a later capability while its evidence at an earlier stage is still partial.

What the framework does not capture: The stages measure what has been demonstrated and published, not everything that will determine long-term success. They don't reflect how fast each system runs, whether results came from the architecture a company plans to scale, or progress that hasn't been published yet. There is also differentiation within stages. We address these limitations in the individual company profiles.

For a full list of terms, please see Quantum Computing Glossary.

Summary of Technical Results

Across the companies we reviewed, no platform has yet demonstrated scaled fault-tolerant computation (Stage 8), but the group spans a wide range of maturity.

Exhibit 9: Technology progress by company

Stage	Stage name	What it means	Quantinuum	Google	IBM	IonQ	Infleqtion	Xanadu
1	Encoding	Create encoded logical information	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Demonstrated
2	Error detection	Detect when errors occur	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Partial
3	Error correction	Use detected errors to correct / track the logical state	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Partial
4	Repeated QEC	Protect the same logical information over multiple correction cycles	Demonstrated	Demonstrated	Demonstrated	Demonstrated	Not demonstrated	Not demonstrated
5	Logical improvement / breakeven	Protected information performs at least as well as relevant physical / unencoded baseline	Demonstrated	Demonstrated	Partial	Partial	Partial	Not demonstrated
6	Below-threshold scaling	More protection produces lower logical error	Partial	Demonstrated	Not demonstrated	Not demonstrated	Not demonstrated	Not demonstrated
7	Fault-tolerant logical operations	Perform protected logical operations while preserving a measured QEC advantage	Demonstrated	Partial	Partial	Partial	Partial	Not demonstrated
8	Scaled fault-tolerant computation	Many low-error logical qubits operate together in sustained universal computation	Not demonstrated	Not demonstrated	Not demonstrated	Not demonstrated	Not demonstrated	Not demonstrated

Exhibit 10: Demonstrated logical computing progress

COMPANY	QEC MATURITY	LARGEST ENCODED SCALE	STRONGEST QEC DEMONSTRATION	BEST RESULT
QUANTINUUM	Stage 7 demonstrated (Stage 6 partial)	94 error-detected logical qubits; 48 error-correcting*	2 logical qubits / 20 data qubits; 20 rounds of correction without discarding runs	Adding protection cut logical-memory error ~4.5x; protected operations ~4.3x more reliable than unprotected hardware
GOOGLE	Stage 6 demonstrated (Stage 7 partial)	2 logical qubits in a separate logical-operation experiment	1 logical qubit / 101 physical qubits; repeated QEC at three code sizes	Logical error fell each time protection increased; largest logical memory lasted ~2.4x longer than its best physical qubit
IBM	Stage 4 demonstrated (Stages 5 and 7 partial)	64-qubit error-detected encoded circuit (~0.04% of runs kept)*	1 logical qubit on Heron; up to 30 rounds of correction	Corrected errors repeatedly, but did not clearly outperform physical memory or show that a larger code lowers logical error
IONQ	Stage 4 demonstrated (Stages 5 and 7 partial)	Up to 4 logical qubits per code, across nine codes	2 logical qubits / 38 active ions; up to 6 rounds of correction	Best logical memory lasted ~3.95 sec vs ~3.3 sec for physical qubits, within overlapping error ranges; a reanalysis moved all five high-efficiency codes past breakeven by discarding a few percent more runs
INFLEQTION	Stage 3 demonstrated (Stages 5 and 7 partial)	30 error-detected logical qubits*	4 logical qubits / 16 data atoms; one correction step at the end of the circuit	Encoded state-preparation error fell from ~13% to ~1.5%, with correction applied after measurement rather than through repeated rounds
XANADU	Stage 1 demonstrated (Stages 2 and 3 partial)	12 physical optical modes (Aurora); no logical qubits yet*	Distance-2 error detection with real-time decoding and feed-forward	Demonstrated real-time error detection and control, but not yet that encoding reduces logical error on hardware

Quantinuum and Google lead in technology progress, with different strengths. Quantinuum has gone furthest toward computing with protected qubits. It has corrected errors over 20 rounds without discarding runs, and its protected operations were about 4.3x more reliable than the same operations on unprotected hardware, meeting our Stage 7 bar. Its results also come from the same architecture it plans to scale. Google has the clearest evidence that error correction improves as more protection is added. Its logical error fell each time it increased protection, across three levels, and its largest protected memory lasted about 2.4x longer than its best physical qubit. Each leader's main gap is the other's strength. Quantinuum has shown errors falling with added protection only once. Google has shown that improvement for stored information, but not yet for protected operations that beat unprotected ones.

IBM and IonQ form a second group, at a similar stage with different strengths. Both have demonstrated Stages 1 through 4, protecting information through repeated rounds of correction, with partial evidence at Stages 5 and 7, but neither has yet shown that adding protection lowers error (Stage 6). IBM has the deeper and larger demonstrations. Its hardware has corrected errors for up to 30 rounds, versus six for IonQ, and it has run thousands of protected operations across large encoded circuits with about 10x lower error than unprotected versions, though only by discarding almost all runs. IonQ has come closer to showing that protection beats unprotected qubits. Its best memory result roughly matched its physical qubits, within error bars and after discarding many runs, and a reanalysis moved codes past breakeven while discarding a small additional share of runs

each round. Both now face the same test: their results come from today's hardware, and neither has yet reproduced them on the new architecture it plans to scale, qLDPC-based modules for IBM and Superion for IonQ.

Infleqtion and Xanadu are earlier in the progression, though each has features that could matter at scale. Infleqtion has encoded information, detected errors and corrected them (Stages 1–3), and has run protected calculations on up to 30 logical qubits. That is rapid progress for a younger approach. Its correction so far has been a single step at the end of a calculation, though, and its best results rely on discarding runs, so it has not yet shown repeated correction (Stage 4). Its two-atom design is intended to make repeated correction easier. Xanadu is earlier still. It has generated its target type of protected qubit on a chip (Stage 1). It has also built a modular photonic system that responds to errors in real time, a capability few peers have shown on hardware. However, its protected qubits are not yet high enough quality, and it has not yet detected and corrected errors on them. Reducing optical loss is its critical next step.

Roadmap and Execution

There are two distinct questions in assessing execution: whether a company has delivered against the commitments it previously made, and whether it has the resources to deliver the roadmap ahead. We assess delivery by comparing dated public commitments with actual results, distinguishing true misses from roadmap changes made as technical strategy evolved, and considering how relevant that record is to the architecture each company now intends to scale. We assess resources through funding, manufacturing access, facilities and partnerships. Our overall view combines both. We focus here on the four pure plays.

Exhibit 11: Execution assessment by company

COMPANY	SYSTEM DELIVERY TRACK RECORD	ROADMAP ADHERENCE	FUNDING AND RESOURCES	OVERALL VIEW
QUANTINUUM	High	High	Large balance sheet, no debt	Relatively high
IONQ	Moderate-High	Moderate	Large balance sheet, owns foundry	Moderate-High
INFLEQTION	Moderate-High	Moderate	Mid-sized balance sheet, no debt	Moderate, short history
XANADU	Earlier Stage	Limited history	Smaller balance sheet - may need additional funding	Earlier stage, short history

Note: Overall view reflects our combined assessment of system delivery track record, roadmap adherence, and funding and resources.

In our view, Quantinuum has the strongest execution record among the pure plays. Quantinuum has delivered successive H1, H2 and Helios hardware generations and completed its five-year Quantum Volume improvement commitment ahead of the implied trajectory. Helios missed one original gate performance target, but the system arrived on schedule with the planned architecture and has subsequently supported substantially larger logical-computing demonstrations. Importantly, Quantinuum has maintained relatively high architectural continuity across generations, making its historical execution record particularly relevant to Sol and Apollo, their next generation computers. We therefore assign Quantinuum relatively high execution credibility, with Sol representing the next major proof point.

IonQ has a mixed long-term record, but strong early momentum on Superion. IonQ delivered several nearer term AQ performance milestones and continued to improve its legacy trapped ion systems, but a number of longer dated system and fault tolerance commitments were delayed, revised or ultimately replaced. The relevance of that history

is also reduced by IonQ's transition to the Oxford Ionics electronic control architecture underlying Superion. Early execution on Superion has been fast. IonQ has fabricated its first 256-qubit chips at SkyWater, trapped ions in prototype systems, and announced agreements to deploy Superion 256 systems with the University of Cambridge, Horizon Quantum, NVIDIA, SDT and Florida International University, with first deliveries expected in 2027. We therefore assign IonQ Moderate-High execution credibility. Those deliveries will test both whether IonQ can execute against its current roadmap and whether it can turn the acquired architecture into a scalable commercial system.

Xanadu and Infleqtion have demonstrated meaningful technical execution, but neither has a long public roadmap history. Xanadu progressed from earlier photonic systems to Borealis and Aurora, but its first detailed investor roadmap was published only in March 2026 and was revised within roughly six months. We view those revisions more as evidence that the roadmap is still being refined than as outright execution missteps. Infleqtion's detailed roadmap dates only to 2024, but it has delivered a 100-qubit system to the UK's National Quantum Computing Centre and the quantum processor for Japan's Shunkai system, and has beaten its logical-qubit count targets, though on an easier bar than its original roadmap set. We therefore assign Infleqtion Moderate execution credibility and Xanadu Earlier Stage, with both grades limited by short track records.

Funding positions vary across the group. Quantinuum and IonQ are on similar footing, each with large balance sheets. IonQ also recently announced a $2bn equity raise. Both are also investing in manufacturing: Quantinuum through its CHIPS-backed US ion-trap supply chain, and IonQ through its ownership of SkyWater's foundry. Infleqtion sits in between, with about ~$500mn of cash and investments and no debt, enough to fund its roadmap through 2030 on our model without new equity. It is expanding manufacturing through its Oxford facility, and its pending CHIPS award of up to $100mn would add further support, though it remains a letter of intent. Xanadu has a much smaller balance sheet and we model them may feeding additional funding by 2030. Canada's backing to date, including C$195mn for Project OPTIMISM, helps but does not remove that need. Overall, these resources reduce financing risk, but they do not remove the technical and organizational challenges of system integration, manufacturing scale and roadmap delivery.

Manufacturability

Logical computing maturity tells us whether a platform works, but not whether it can be built at the scale and cost commercial fault tolerance requires. Company roadmaps call for hundreds of thousands to millions of physical qubits within the next five years, orders of magnitude beyond today's systems. At that scale, the challenge shifts toward production: manufacturing chips repeatably, controlling more qubits without proportionally more hardware, and deploying and supporting systems at customer sites at on the cloud.

How we assess manufacturability: We look at four areas for each company:

- **Systems in use today:** What each company has built, offered to outside users on the cloud and installed at customer sites.
- **Next-generation system:** The system each plans to scale, its timing and any customer commitments.

- **What gets harder as systems grow:** The hardware that must scale with qubit count, such as lasers, control electronics or optical loss, and how each company plans to address it.
- **Outside evaluation of the scaling plan:** Independent review through programs such as DARPA's Quantum Benchmarking Initiative, CHIPS awards and government due diligence.

Exhibit 12: Manufacturability at a glance

	QUANTINUUM	IONQ	INFLEQTION	XANADU
SYSTEMS IN USE TODAY	Three generations on one architecture; cloud since 2020; two generations installed at RIKEN; Helios being installed in Singapore	Cloud on all major providers; installed at QuantumBasel, the NQCC and EPB, mostly on its earlier architecture	100-qubit system at the NQCC; processor for Japan's Shunkai	Systems built in-house; Borealis offered on the cloud (2022–23); no customer-site installations
NEXT-GENERATION SYSTEM	Sol (2H27); current commitments are for Helios, launched November 2025	Superion 256 (2027 deliveries); five customers and partners committed	Illinois system (2027)	Qubit Factory (2026–27)
WHAT GETS HARDER AS SYSTEMS GROW	Laser delivery; integrated optics planned for Lumos (2030+)	Control electronics on the chip; CMOS integrated on the chip needed for the 10K	Laser control and parallel operations across larger arrays	Optical loss across chips, fiber and packaging
OUTSIDE EVALUATION OF THE SCALING PLAN	DARPA QBI Stage B; CHIPS award finalized	DARPA QBI Stage C, testing whether Superion can be built as designed	CHIPS letter of intent, not final; not in QBI	DARPA QBI Stage B; Canadian federal due diligence
BIGGEST OPEN QUESTION	Whether laser delivery scales to Apollo's thousands of qubits before integrated optics arrive	Whether Superion performs as a full system at 256 qubits, then at 10,000	Whether its manufacturing base and laser control scale with the Illinois system	Whether loss falls to target; no measured value disclosed
OVERALL VIEW	Moderate	Moderate-High	Moderate	Earlier-stage

Manufacturability is a closer race and harder to assess than technical progress alone. Quantinuum, IonQ and Infleqtion have all deployed systems at customer sites, but it remains too early to judge their ability to manufacture much larger systems at scale. The challenges differ by architecture and are closely tied to each company's technology roadmap, making this an assessment that will evolve as systems scale. Based on available evidence, we see a slight edge for IonQ, driven by outside validation of its scaling plan and a design that relies less on lasers as systems grow. However, the race remains close, with 2027 milestones likely to provide a clearer picture of each company's ability to scale.

IonQ has the strongest outside validation of its scaling plan. DARPA advanced it to Stage C of its Quantum Benchmarking Initiative, where an independent team will test multiple Superion generations against IonQ's roadmap. Quantinuum and Xanadu remain in Stage B, and Infleqtion is not in the program. Five customers and partners, including NVIDIA and the University of Cambridge, have committed to Superion 256 ahead of 2027 deliveries, showing demand for its newest system before it ships. Superion's design also moves gate control from lasers into electronics on the chip, reducing the laser-delivery challenge that grows with system size. IonQ has deployed systems at QuantumBasel, at EPB in Chattanooga and, through Oxford Ionics, at the UK's NQCC, mostly on its earlier architecture. It has fabricated its first 256-qubit processors at SkyWater, but full-system performance on Superion 256 and on-chip CMOS integration at the 10,000-qubit scale remain unproven.

Quantinuum has the strongest delivery record on the architecture it plans to scale. It has built three generations using the same trapped ion approach, replaced an older system at a customer site, and demonstrated promising methods for scaling trap control and laser sharing. Its CHIPS award, finalized in September 2026, supports a US supply chain for traps, lasers and optics through partners including GlobalFoundries and Monarch Quantum. Laser control and delivery remain its biggest scaling challenges. Quantinuum plans to address these in stages: Sol introduces a larger grid trap in 2027, Apollo (2029) integrates more laser and microwave control electronics onto the chip, and Lumos

(2030+) adds integrated optics and multichip packaging. These advances remain unproven at scale, but Quantinuum's consistent architecture and delivery history support confidence in execution. A successful Sol launch or further outside validation could narrow IonQ's manufacturing edge.

Infleqtion has meaningful deployment experience, but its manufacturing capacity is still expanding. It delivered a 100 qubit system to the UK's NQCC and supplied the processor for Japan's Shunkai system, with a modular system under contract for Illinois in 2027. Unlike chip based peers, Infleqtion manufactures neutral atom systems built around lasers, optics and vacuum chambers rather than semiconductor qubit chips. Its sensing business adds experience producing field deployable hardware, although larger computers introduce greater optical complexity. Infleqtion is expanding manufacturing capacity through a new Oxford facility and has a pending CHIPS investment. The Illinois system will test its ability to manufacture larger modular computers.

Xanadu has a potentially scalable semiconductor manufacturing approach, but the least deployment evidence. It designs photonic chips and uses external foundries for wafer fabrication, while developing internal packaging, testing and system integration capabilities. Its October 2026 GlobalFoundries agreement targets production of photon detectors and silicon nitride circuits on 300mm wafers, alongside existing work with Tower Semiconductor. Xanadu is also developing its Qubit Factory technology for 2026 to 2027 and a larger Toronto manufacturing facility. However, it has not installed systems at customer sites. Its main scaling hurdle, optical loss, depends heavily on fabrication and packaging quality, making manufacturability and technical progress effectively the same test.

Broader Indicators of Progress and Other Considerations

Beyond technical results and execution, we look at external indicators of each company's progress: research and patent activity, government and public-sector relationships, commercial and industry partnerships, and hiring. We treat these as supporting evidence rather than proof of technology leadership. The underlying publication, patent, contract and job-posting data are likely imperfect, so we use them to indicate direction and scale rather than to rank companies.

Exhibit 13: External validation indicators

COMPANY	RESEARCH AND IP	GOVERNMENT AND PUBLIC SECTOR	COMMERCIAL AND INDUSTRY	JOBS TREND
QUANTINUUM	Deep research record; patent activity up sharply since 2024	Broad across the US, UK, Japan and Singapore, including DARPA's QBI Stage B and a CHIPS R&D award	Multi-year Helios sale to Oracle for its cloud	Rising sharply since mid-2024
IONQ	Broad research record; large patent portfolio, a meaningful share acquired	DARPA's QBI Stage C and HARQ program, AFRL networking work, KISTI Tempo agreement	Superion 256 deployment agreements with five customers and partners, including NVIDIA	Risen several-fold since early 2025
INFLEQTION	Steady research record; patents covering computing and sensing hardware	Systems in the UK (NQCC) and Japan (Shunkai), DARPA HARQ, DOE Genesis Mission projects, Illinois system planned for 2027	NVIDIA partnership	Rose sharply into early 2026, eased since
XANADU	Deep research record, especially algorithms and software; smaller patent footprint	Concentrated in Canada, including C$195mn for Project OPTIMISM and the Quantum Champions Program; DARPA's QBI Stage B	Mainly software (PennyLane) and industry partnerships with AMD, ASML and Mitsubishi Chemical	Picked up in 2026 ahead of the Qubit Factory development

Note: Data shown here is not comprehensive. Job trend measured for 1H2026

External indicators broadly support the picture from our technical and execution work. All four companies show active research programs, meaningful government backing and stepped-up hiring. The clearest differences are in the type of validation. Quantinuum and IonQ increasingly have outside institutions buying and hosting their systems, including large AI and cloud infrastructure players such as Oracle and NVIDIA. Infleqtion's strongest

signals are national programs in the UK and Japan running on its hardware. Xanadu's validation comes mainly from government funding for its manufacturing build-out and broad adoption of its PennyLane software.

Exhibit 14: Google and IBM lead research output, supported by larger R&D programs...

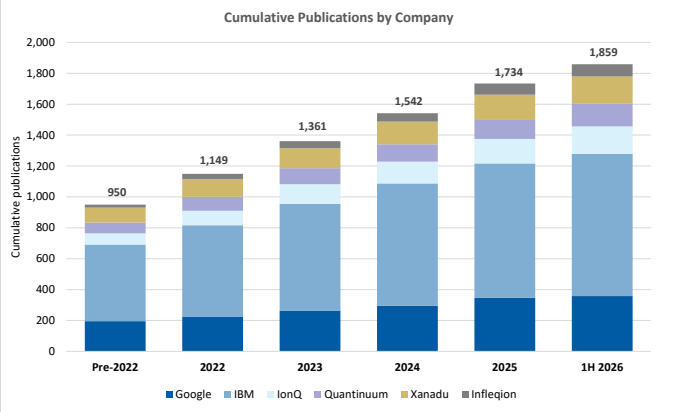


Exhibit 15: ...while patent activity highlights growing investment in quantum IP across the industry...

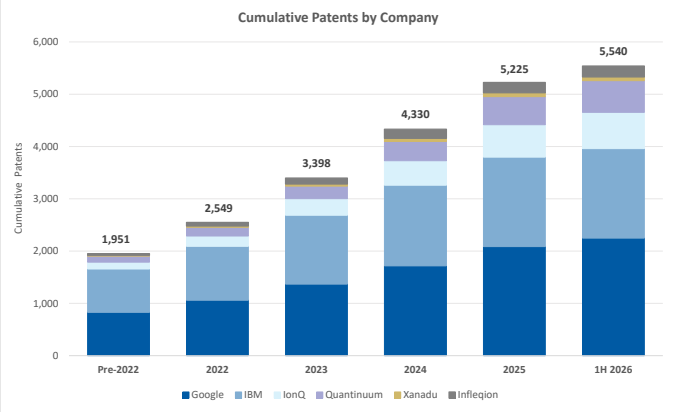
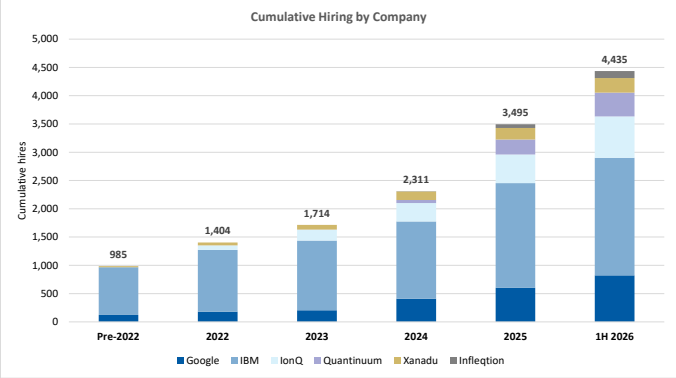


Exhibit 16: ...and accelerating job posting signals an expanding industry



Valuation

Methodology

We value the quantum pure plays using risk adjusted multiples that reflect each company's stage of commercialization in 2030 and our visibility in its long term trajectory. With commercial quantum computing still in its infancy, current revenue provides limited insight into the potential value of these businesses. We therefore anchor our valuations to 2030E computing revenue and build our multiples in three steps: where each company is expected to be in 2030, a growth-stage multiple reflecting its position on the commercialization curve at that point, and a visibility factor assessing the progress of its technology roadmap.

Exhibit 17: Our valuation in three steps

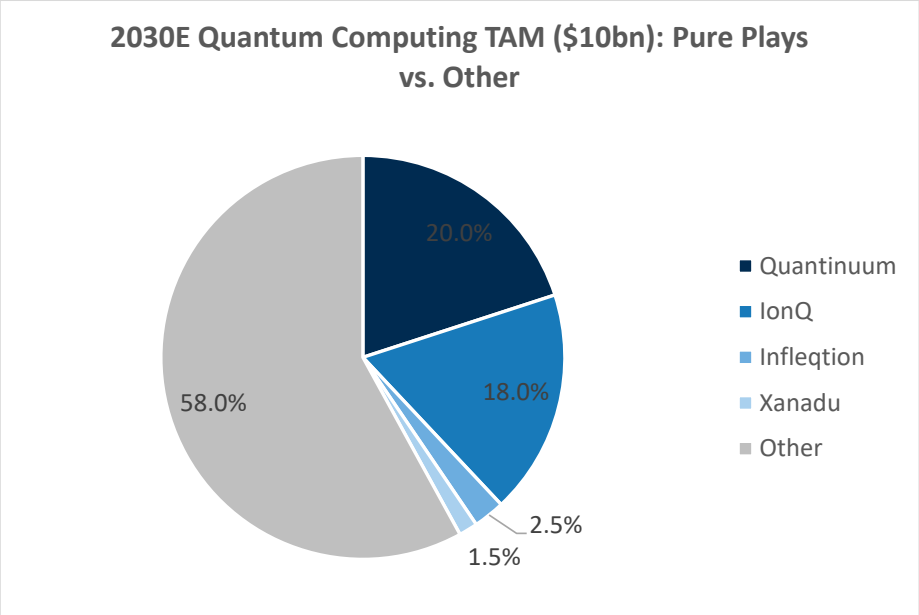
STEP	WHAT IT CAPTURES	HOW WE SET IT	VALUES (QNT / IONQ / INFQ / XNDU)
1. 2030E computing revenue	How far each company can get by 2030	Company models, held within a ~$10bn 2030 market	$2.0bn / $1.8bn / $252mn / $150mn
2. Growth multiple	Value of growth after 2030 if the platform succeeds at scale	Trailing EV/Sales of listed companies growing at similar rates today	17x / 17x / 25x / 25x
3. Visibility factor	Our visibility the 2030 position	Our proprietary framework: QEC stage, execution and manufacturability	75% / 70% / 55% / 50%
Base multiple (2 × 3)	**Risk-adjusted EV/Sales on 2030E computing revenue**	**Growth multiple × visibility factor**	**12.8x / 11.9x / 13.8x / 12.5x**

Note: Success multiples are trailing EV/Sales. Visibility factors derived from our proprietary framework.

Where companies sit in the 2030 TAM

Our 2030 forecasts reflect where each company can reach if its roadmap broadly delivers, given its stage of commercialization today. We hold them within a 2030 quantum computing market of about $10bn, in line with our previous estimate. That estimate is highly sensitive to timing. If roadmaps slip by a year or more, which we see as a real possibility, the 2030 market would be closer to $5bn or less. If the leading roadmaps deliver, we think technology rather than demand is the binding constraint on 2030 revenue, since customers are already committing ahead of capability. On that basis, our four pure plays account for about 40% of the 2030 market (Exhibit 18). The remaining 60% sits with IBM (covered by Erik Woodring), Google, PsiQuantum, Atom Computing, Microsoft, other system vendors, and software and services providers.

Exhibit 18: Implied market share of our covered pure plays



Quantinuum and IonQ lead. We expect both to be delivering early fault tolerant systems by 2030, Apollo for Quantinuum and Superion for IonQ, with computing revenue of about $2bn and $1.8bn, respectively. IonQ has the larger total revenue base today, but much comes from existing systems and its broader quantum platform, including networking, sensing and security, which we value separately. Quantinuum's forecast also includes software and services built on its existing stack. We model the steepest revenue ramp for Quantinuum,from about $270mn in 2028 to $2bn in 2030, making Apollo's timing a main risk to our Overweight thesis.

Infleqtion and Xanadu are earlier in commercialization. We expect Infleqtion's next-generation systems to reach customers by 2030 at a smaller scale, supporting about $250mn of computing revenue. Xanadu's commercial revenue begins in 2029 in our model and reaches about $150mn in 2030, near the start of its ramp. Their smaller 2030 forecasts reflect later commercialization rather than a smaller long-term opportunity.

These different revenue trajectories are important for valuation. While all four companies operate in an emerging market, their forecasts represent different points along the commercialization curve. By 2030, Infleqtion's approximately $250mn of computing revenue is comparable to Quantinuum's approximately $270mn in 2028, suggesting a roughly two year difference in commercialization maturity. This difference in timing determines the growth-stage multiples we apply.

Growth-stage multiple

Our growth-stage multiple reflects where each company is expected to sit on its commercialization curve in 2030. Companies further along in their revenue ramps receive a lower multiple, while those still entering the steepest phase of commercialization receive a higher multiple. We benchmark these assumptions against trailing EV/Sales multiples of publicly traded companies at comparable stages of revenue growth. The higher growth-stage multiples do not imply stronger technologies or greater

ultimate commercial potential. They reflect the later timing of commercialization.

Quantinuum and IonQ: 17x. Both are the furthest along in their commercial ramps by 2030 in our models. We compare them with profitable, high-growth semiconductor and hardware companies, which trade at roughly 12–23x trailing sales. We use 17x, around the midpoint, which we consider a high growth multiple.

Infleqtion and Xanadu: 25x. Both are about two years behind the leaders on our forecasts. If they succeed, we expect them to follow a similar path about two years later, so at the end of 2030 the steepest part of their ramp would still be ahead of them. We compare them with companies in a phase of very rapid revenue growth, many of them early in their ramp, which trade at a median of about 25x trailing sales. We apply 25x because these companies are growing at rates close to what we expect for Infleqtion and Xanadu in the early 2030s if they succeed.

Visibility factors from our proprietary framework

The visibility factor reflects how likely we think each company is to reach the position in our 2030 forecast and keep growing from there. We set it with our proprietary framework, which brings together our assessments of technology readiness on our eight-stage technology scale (Exhibit 8), roadmap execution (Exhibit 11) and manufacturability (Exhibit 12).

Exhibit 19: Our proprietary framework informs each company's visibility factor

COMPANY	TECHNOLOGY READINESS (HIGHEST QEC STAGE)	EXECUTION	MANUFACTURABILITY	OTHER CONSIDERATIONS	VISIBILITY FACTOR
QUANTINUUM	Stage 7: Logical operations	Relatively high	Moderate	–	75%
IONQ	Stage 4: Repeated QEC	Moderate-High	Moderate-High	Strongest outside validation of scaling plan	70%
INFLEQTION	Stage 3: Error correction	Moderate	Moderate	No new equity needed on our model	55%
XANADU	Stage 1: Encoding	Earlier Stage	Earlier-stage	PennyLane software; ~$385mn new equity needed	50%

Note: Visibility factors are our judgment, informed by the assessments shown. Technology readiness scale (Exhibit 8) carries the most weight; execution (Exhibit 11) and manufacturability (Exhibit 12) adjust within the range it supports. Stage ratings reflect published hardware results only.

Our visibility factors are a judgment, informed by the framework. Technology readiness carries the most weight, since whether a platform reaches fault tolerance remains the largest uncertainty. Execution and manufacturability then move each company up or down within the range its technology supports, and we take account of company-specific factors where they bear on the odds of success. We will revisit the factors as companies report results against the next stages of our framework.

Quantinuum has the highest visibility factor, at 75%. It has progressed furthest on our technology progress scale, including protected logical operations that outperform the underlying hardware, and it has the strongest execution record in the group, with three hardware generations delivered on one architecture. **IonQ follows at 70%.** It has shown that error correction holds over repeated rounds and has the strong
of its scaling plan, including DARPA's Stage C selection and five Sup
commitments. Its design also relies less on lasers as systems grow. T
edge keeps IonQ close to Quantinuum, but its best error-correction
today's laser-based systems and have not yet been reproduced on S

Infleqtion 55% and Xanadu 50% have yet to show repeated error c

人在草木间

both well below the leaders. Infleqtion is further along on our QEC scale and has the largest system installed at a customer site among the four. It can also fund its roadmap without new equity on our model. Xanadu has a shorter roadmap record and has yet to install a system at a customer site. We give it some credit for PennyLane, its open-source software platform, which the Unitary Foundation's 2025 survey ranked as the most-used quantum software platform after IBM's Qiskit.

Our Price Targets

Multiplying the growth-stage multiple by the visibility factor gives our base case multiples, which we apply to 2030E computing revenue (Exhibit 20). Infleqtion has the highest base multiple in the group, at 13.8x. That reflects the runway it would still have after 2030 if it succeeds. Our visibility factor in Infleqtion, at 55%, is well below our visibility in either leader. Quantinuum's value comes from the highest visibility factor in the group applied to the largest 2030 revenue.

Exhibit 20: Our price targets

COMPANY	RATING	PRICE	PRICE TARGET	UPSIDE / (DOWNSIDE)	BASE MULTIPLE	MARKET-IMPLIED MULTIPLE	VISIBILITY FACTOR	MARKET-IMPLIED VISIBILITY	BEAR	BULL
QUANTINUUM	Overweight	$33.03	$62.00	+88%	12.8x	7.2x	75%	43%	$10.00	$104.00
IONQ	Equal-weight	$31.42	$43.00	+37%	11.9x	8.9x	70%	52%	$17.00	$69.00
INFLEQTION	Equal-weight	$9.11	$12	+28%	13.8x	11.5x	55%	46%	$3.60	$23.90
XANADU	Equal-weight	$3.19	$3.75	+18%	12.5x	11.9x	50%	47%	$0.70	$9.10

Note: Market-implied multiple is the 2030E EV/sales multiple at which today's share price equals the present value of our base case; market-implied visibility divides it by the growth multiple. Prices as of market close, Oct. 14th 2026

Sum of the parts: a floor for IonQ and Infleqtion

IonQ and Infleqtion own businesses outside quantum computing that we value separately and hold at base-case values in every scenario. Together with net cash, these businesses give both companies a floor that does not depend on their quantum computers succeeding (Exhibit 21).

We value IonQ's SkyWater foundry at 3.9x sales, consistent with its acquisition multiple. We apply 7x to IonQ's acquired networking, sensing and space businesses and Infleqtion's sensing business, between defense electronics peers (~5x) and space imaging peers (~9.6x), reflecting their specialized technologies and growth potential. Note we did not include Elisity in our IonQ revenue projections.

For IonQ, the SkyWater foundry is worth about $2.75 a share and the acquired networking, sensing and space businesses about $2.00. With 2027E net cash of about $7.50 a share, IonQ's floor is about $12.30, or roughly 39% of the current share price. For Infleqtion, the sensing business is worth about $0.85 a share and 2027E net cash about $1.85, a floor of about $2.70, or roughly 30% of the current price.

Quantinuum and Xanadu have no separately valued businesses, so net cash is their only floor, at about 14% and 13% of their share prices. The difference explains why a delay hurts Quantinuum more: its computing business accounts for about 92% of our value, against about 71% for IonQ and 77% for Infleqtion. This floor is not a liquidation value. It assumes our separately valued businesses achieve their forecast revenue and valuation

multiples, and that net cash follows our modeled trajectory.

Exhibit 21: Sum of the parts: non-computing businesses and net cash provide a floor for IonQ and Infleqtion

$ PER SHARE	QUANTINUUM	IONQ	INFLEQTION	XANADU
Computing business (base case)	57.19	30.39	8.97	3.35
SkyWater foundry (3.9x 2030E sales)	–	2.76	–	–
Acquired businesses* (7.0x 2030E sales)	–	2.00	–	–
Sensing (7.0x 2030E sales)	–	–	0.85	–
Net cash, 2027E	4.65	7.52	1.83	0.42
Value per share	**61.84**	**42.67**	**11.66**	**3.76**
Floor: non-computing businesses + net cash	**4.65**	**12.28**	**2.69**	**0.42**
Floor as % of current price	14%	39%	29%	13%
Computing as % of value per share	92%	71%	77%	89%

Note: Note: 2030E EV discounted to Oct-27 at a 14.6% WACC; net cash and per-share values on 2027E diluted shares. Price as of 10/14/26.
*IonQ acquired businesses: networking, sensing and space.

What is priced in: Quantinuum stands out

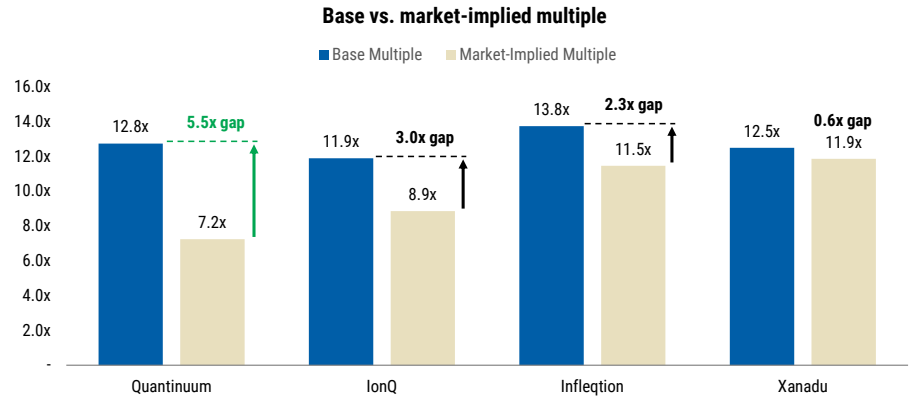
Our framework also allows us to reverse engineer what current share prices imply. We calculate the 2030E EV/Sales multiple required to reconcile each company's current share price with our valuation assumptions, then divide that multiple by our growth-stage multiple to derive a market implied visibility factor.

The results reveal a notable disconnect. Quantinuum has the highest visibility factor in our proprietary framework, yet the market assigns it the lowest implied multiple and visibility level among our covered pure plays. The gap does not rest on our higher revenue forecast for Quantinuum. If we give Quantinuum IonQ's 2030E computing revenue of about $1.8bn, its market-implied visibility factor rises only to about 47%, still below IonQ's.

Exhibit 22: The market assigns the greatest discount to the company furthest along our framework

	QUANTINUUM	IONQ	INFLEQTION	XANADU
Market-implied multiple (EV / 2030E sales)	7.2x	8.9x	11.5x	11.9x
Success multiple	17x	17x	25x	25x
Market-implied visibility	43%	52%	46%	47%
Our visibility factor	75%	70%	55%	50%
Gap (percentage points)	**+32**	**+18**	**+9**	**+3**

Exhibit 23: Our view vs. what the market is pricing



Note: Price as of Oct. 14th 2026 market close

Bull and bear cases

Our bull case assumes the platform succeeds and 2030E computing revenue comes in 30% above our forecast. We apply the full success multiple, without the visibility discount, to that higher revenue. Our bear case assumes a two-year delay: we apply each company's base-case multiple less four turns to its 2028E computing revenue. Non-computing businesses are held at base-case values in both.

Exhibit 24: Bull, base and bear cases

	QUANTINUUM	IONQ	INFLEQTION	XANADU
Bear: base multiple less 4 turns on 2028E revenue	8.8x on $272mn	7.9x on $429mn	9.8x on $35mn	8.5x on $15mn
Bear value ($)	**$10**	**$17**	**$4**	**$0.70**
Base: success multiple × visibility on 2030E revenue	12.8x on $2.0bn	11.9x on $1.8bn	13.8x on $252mn	12.5x on $150mn
Base price target ($)	**$62**	**$43**	**$12**	**$3.75**
Bull: success multiple on 2030E revenue +30%	17x on $2.6bn	17x on $2.3bn	25x on $327mn	25x on $195mn
Bull value ($)	**$104**	**$69**	**$24**	**$9.10**

Stock Calls

Quantinuum: Upgrade to Overweight, PT $62. Quantinuum has the highest visibility factor in our framework, at 75%. It has progressed furthest among the pure plays on our technology scale, including protected logical operations that outperform the underlying hardware, and it has the strongest execution record. Applying 12.8x (17x × 75%) to 2030E computing revenue of about $2.0bn gives a $62 price target, about 88% above the current price. The market implies about a 43% visibility factor, the lowest in the group. The main risk is timing, since our forecast assumes revenue rises sharply as Sol (2H27) and Apollo (2029) arrive. Laser delivery is also Quantinuum's main manufacturing hurdle as systems grow, which makes Sol a test of both its technology and its ability to scale.

IonQ: Equal-weight, PT $43. IonQ's 70% visibility factor reflects demonstrated repeated error correction and the strongest outside validation of its scaling plan. Applying 11.9x (17x × 70%) to 2030E computing revenue of about $1.8bn, and valuing SkyWater and the acquired businesses separately, gives a $43 price target, about 37% above the current

price. The market implies about a 52% visibility factor, an 18 point gap to our view. Our bear and bull cases are $17 and $69, and its non-computing businesses and net cash provide a floor of about $12.30 a share. Superion 256 deliveries in 2027 and a first logical-qubit demonstration on Superion are the catalysts that could make us more positive.

Infleqtion: Equal-weight, PT $12. Infleqtion has run protected calculations on up to 30 logical qubits, but its results rely mainly on discarding flagged runs, and it has yet to show repeated error correction. Its 55% visibility factor and 25x success multiple give a base multiple of 13.8x. Applying that to 2030E computing revenue of about $250mn, and valuing sensing separately, gives an $12 price target, about 28% above the current price. The market implies about a 46% visibility factor, 9 points below our view. Our bear case of $3.60 reflects a two-year delay, and our bull case of $23.90 reflects the value if its neutral-atom approach scales as planned. Its sensing business and net cash provide a floor of about $2.70 a share. Infleqtion has lower cash burn, and we do not anticipate external funding needed at this time to get to 2030. Repeated error correction and on-time delivery of the Illinois system in 2027 are the key catalysts.

Xanadu: Equal-weight, PT $3.75. Xanadu is the earliest of the four on our technology progress scale and may need additional funding before 2030 in our model. After credit for PennyLane, its visibility factor is 50%. Applying 12.5x (25x × 50%) to 2030E revenue of $150mn gives a $3.75 price target, about 19% above the current price. Its range of outcomes is wide, with a $9.10 bull case and a $0.70 bear case. Reducing optical loss, its main technical and manufacturing hurdle, is the next test.

Key valuation risks

Our valuation depends on revenue that is still several years away, and the main risk is timing. If fault-tolerant systems arrive later than our models assume, or commercial adoption builds more slowly once they do, 2030 revenue would be lower and value would shift further out. Our bear case reflects that risk with a two-year delay. A delay matters more for Quantinuum, whose value rests almost entirely on its computing business, than for IonQ and Infleqtion, whose non-computing businesses and net cash provide a floor.

Our growth multiples and visibility factors are judgments. Confidence factors could change quickly as companies report results against the next stages of our framework. A competing approach could prove more scalable than we expect, and the market could end up more fragmented than our forecasts imply. Valuations for long-duration, pre-profit technology companies are also sensitive to market sentiment and interest rates, independent of company progress.

On the other side, faster progress toward fault tolerance, an earlier commercial ramp or consolidation around fewer platforms could support higher revenue and higher visibility factors than we assume. That applies most to the companies that demonstrate the next stages of our framework first.

Quantinuum Company Profile

Exhibit 25: QNT Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Trapped ion
Physical qubits (flagship)	98
Logical qubits (best count)	48
Technology framework stages	Demonstrated: 1-5, 7 | Partial: 6 | Not demonstrated: 8
2Q gate fidelity, avg (%)	99.921%
FINANCIAL HIGHLIGHTS	
Year public	2026
IPO gross proceeds ($mn)	$1,680.00
FY25 revenue ($mn; FY = Dec)	$30.90
Market cap, latest ($bn)	$9.91
Liquidity: cash + marketable securities ($mn)	$2,106.69

Note: Market cap as of October 14, 2026.

Technology

Our view: Quantinuum has demonstrated the broadest set of fault tolerance capabilities among the pure plays. It has corrected errors repeatedly without discarding runs and performed protected logical operations that outperform the underlying physical hardware. Its results also come from the same trapped ion architecture it intends to scale, providing greater continuity between its experimental achievements and future systems. We have relatively high confidence in its technical direction, although combining these capabilities at much larger scale remains a significant challenge.

In our eight stage technology framework, Quantinuum has demonstrated Stages 1 through 5 and Stage 7, with partial evidence for Stage 6. Its progress spans multiple hardware generations, from real time error correction in 2021 to larger logical circuits and more reliable protected operations on Helios in 2026. The main gaps are showing that errors continue to fall as protection increases and combining all the operations needed for universal quantum computing under continuous correction.

Exhibit 26: QNT: Key Technology Results

QEC stage	Assessment	Key demonstration / source	Year
1. Encoding	Demonstrated	Encoded 1 logical qubit on H1 (Ryan-Anderson); protected circuits with 8 logical qubits (Self); later up to 94 error-detected and 48 error-correcting-code logical qubits (Dasu)	2021, 2022, 2026
2. Error detection	Demonstrated	Detected errors on encoded information (Ryan-Anderson); later extended detection to circuits with 8 logical qubits (Self)	2021, 2022
3. Error correction	Demonstrated	Applied error corrections in real time on H1 (Ryan-Anderson); later combined correction with discarded unsuccessful runs on H2 (Paetznick)	2021, 2024
4. Repeated QEC	Demonstrated	Repeated correction on H1 and H2 (Ryan-Anderson; Paetznick); then protected 2 logical qubits through 20 rounds without discarding runs (Berthusen)	2021, 2024, 2026
5. Logical improvement / breakeven	Demonstrated	Reduced logical errors by up to 800x with discarded runs (Paetznick); then achieved 4.3x lower error than unprotected operations without discarding runs (Berthusen)	2024, 2026
6. Below-threshold scaling	Partial	Threshold behavior under added noise (Sommers); stronger protection reduced errors approximately 4.5x, but not across 3 protection levels (Berthusen)	2025, 2026
7. Fault-tolerant logical operations	Demonstrated	Non-Clifford gate outperforming physical equivalent with some runs discarded (Dasu); then protected logical operations under active correction without discarded runs (Berthusen)	2025, 2026
8. Scaled fault-tolerant computation	Not demonstrated	None qualifying	—

Source: Ryan-Anderson et al., Physical Review X (2021); Self et al., research preprint (2022), Nature Physics (2024); Paetznick et al., research preprint (2024), Nature (2026); Sommers et al., Physical Review Research (2025); Dasu et al., research preprints (2025, 2026); Berthusen et al.,

Note: The 94 and 48 logical qubit demonstrations relied on discarding unsuccessful runs and should be distinguished from the smaller Helix experiment, which demonstrated repeated active correction without discarding runs. Assessments reflect our eight stage technology framework.

What's different about Quantinuum's approach? Quantinuum uses trapped ions, individual charged atoms held in place by electric fields. Its quantum charge coupled device (QCCD) architecture physically moves ions between storage and operating regions on a chip. This gives the system flexible connectivity, allowing any pair of qubits to interact without requiring a fixed wiring pattern. Helios, launched in November 2025, uses 98 barium ions and achieves average two qubit gate fidelity of 99.921%.

Quantinuum has maintained the same underlying QCCD approach across H1, H2 and Helios, although future generations introduce more complex chip designs. The main tradeoff is speed: moving ions and performing error correction takes longer than on superconducting hardware. Logical error rates therefore need to be considered alongside the time required to execute calculations.

Logical Computing Progress: Quantinuum demonstrated repeated real time error correction as early as 2021 on H1, encoding one logical qubit using seven physical qubits. The system checked for errors and applied corrections while preserving the stored quantum information, establishing Stages 1 through 4.

In February 2026, Quantinuum demonstrated calculations involving up to 94 logical qubits using error detection and 48 logical qubits encoded in a stronger code capable of correcting errors. Both outperformed comparable unprotected calculations, but relied on discarding runs where errors were detected. These experiments demonstrated the ability to encode and compute with many logical qubits, rather than continuous correction across all of them.

The September 2026 Helix experiment provided stronger evidence of working error correction. Using Helios, Quantinuum protected two logical qubits through 20 rounds of active correction without discarding runs, achieving an error rate of approximately 4.6×10^{-5} per logical qubit per round. Its protected logical operations were also approximately 4.3 times more reliable than comparable physical operations, meeting our requirements

for Stages 5 and 7. Although this is among the strongest logical error correction results reported, correction cycles on trapped ion hardware take longer than on faster architectures, which is a consideration.

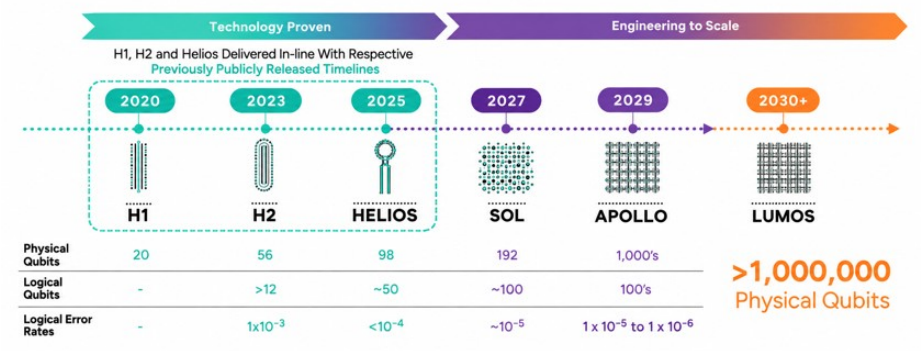
Evidence for below threshold scaling (Stage 6) remains partial. Helix showed that increasing protection reduced logical error by approximately 4.5 times, while separate 2025 research demonstrated threshold behavior under deliberately added noise. However, Quantinuum has not yet shown that errors consistently decline across three or more protection levels under normal operating conditions.

Quantinuum has also demonstrated a protected non Clifford operation, a harder class of operation needed for universal quantum computing. In 2025, this operation outperformed its physical equivalent, although the experiment relied on error detection and discarded approximately 15% of runs. It has therefore demonstrated both reliable logical operations under active correction and a high fidelity non Clifford operation, but not yet these capabilities together in one continuously corrected computation.

Next Steps: The key technical milestones we are watching for are combining non Clifford operations with active error correction without discarding runs, demonstrating continued improvement as protection increases, and extending reliable computation to larger numbers of logical qubits. Achieving these would provide stronger evidence that Quantinuum can move from individual fault tolerance demonstrations toward sustained universal computation.

Roadmap and Execution

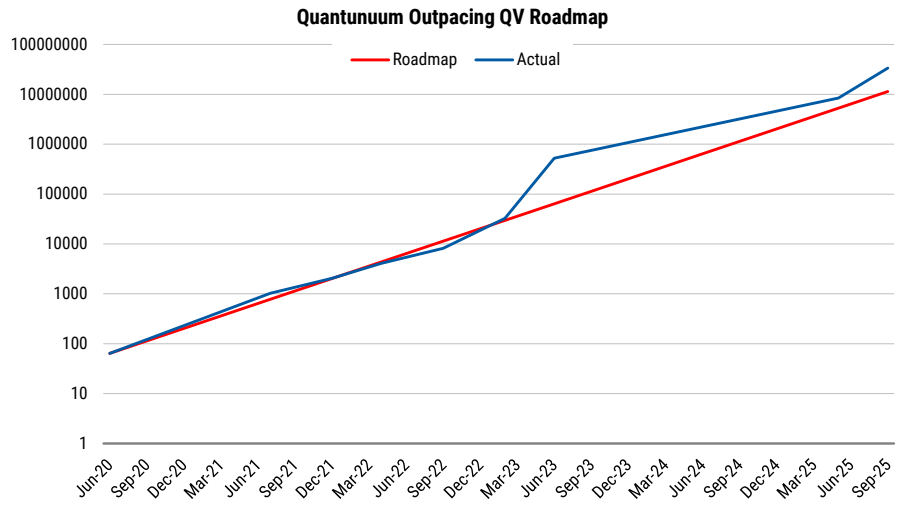
Our view: Quantinuum has one of the strongest execution records among the pure plays. It exceeded its clearest multiyear performance commitment and delivered three successive hardware generations using a consistent trapped ion architecture. Helios launched broadly in line with its original roadmap, although it missed one gate performance target. We assign relatively high execution credibility, with Sol in 2027 representing the next major test as Quantinuum moves to a larger two dimensional chip architecture.



Execution Track Record: In March 2020, Honeywell's quantum business committed to launching a system with a Quantum Volume of at least 64 and increasing that benchmark by approximately ten times annually over five years. Quantum Volume combines the size

and reliability of quantum circuits into one performance measure. Honeywell achieved QV64 in June 2020, and Quantinuum subsequently reached QV8.4 million in May 2025, fulfilling its five year commitment. It then reached QV33.6 million in September 2025, a further fourfold improvement. Although Quantum Volume is not a direct measure of fault tolerance, this provides evidence of sustained improvements in system performance.

Exhibit 27: Quantinuum exceeded their 5-year quantum volume roadmap



Note: In 2020, Honeywell committed to increasing Quantum Volume approximately tenfold annually over five years, starting from QV64. Quantinuum identified QV8.4 million as the corresponding five year milestone, which it achieved in May 2025. It subsequently reached QV33.6 million in September 2025. The red line illustrates the implied tenfold annual growth trajectory extended through September 2025, rather than a separately announced target for that date. Quantum Volume measures overall system performance based on circuit size and reliability, but is not a direct measure of fault tolerance.

Quantinuum has also delivered successive generations of hardware, with H1 launching in 2020, H2 in 2023 and Helios in November 2025. Each has retained the underlying QCCD architecture, allowing Quantinuum to build on previous engineering work rather than adopting an entirely new approach.

Helios was broadly delivered as planned. It launched with 98 physical qubits against an earlier target of 96, introduced barium qubits and added junction based ion routing. However, its average two qubit gate error of approximately 7.9×10^{-4} remained above the original target of less than 5×10^{-4}, while the planned improvement in operating speed was not publicly confirmed. We view these as minor performance shortfalls within an otherwise successful platform launch, particularly given the more advanced logical computing experiments subsequently performed on Helios.

Looking ahead: Sol, targeted for the second half of 2027, introduces Quantinuum's first commercial two dimensional grid trap. The company targets 192 physical qubits, approximately 100 logical qubits and a logical error rate around 10^{-5}. Quantinuum reported in August 2026 that its first production candidate Sol trap chips had returned from fabrication and were progressing through validation.

Apollo, planned for 2029, is the larger milestone. Quantinuum targets thousands of physical qubits, hundreds of logical qubits and universal fault tolerant computation capable of running millions of operations. The company is developing the required subsystems in parallel with Sol, although their performance together at that scale remains

unproven.

What we are watching: Sol is the next major execution test. We will look for an on time launch, measured gate performance at full system scale and evidence that Quantinuum can extend its demonstrated error correction capabilities to approximately 100 logical qubits without relying heavily on discarding runs. Beyond Sol, progress toward Apollo's larger chip architecture and integrated control systems will be critical. Successful delivery would strengthen an already favorable execution record, while a delay or material performance shortfall would raise questions about the pace of scaling.

Exhibit 28: Quantinuum: milestones we are watching

TIMING	TYPE	MILESTONE (AS STATED BY QUANTINUUM)	STATUS	STATUS DETAIL
NEAR TERM (2026 TO 2027)				
2026	Customer deployment	Helios installed at Singapore's National Quantum Computing Hub, the first Helios outside the US	In progress	Singapore R&D and Operations Centre opened in March 2026; on- site deployment underway, with Singapore researchers already using Helios via the cloud
NOT DATED	Customer deployment	Helios installed in a US Oracle Cloud Infrastructure AI data center and offered as an OCI quantum service	Pending	Multi-year system sale announced August 2026; no installation date given. OCI service preview expected within months; most revenue comes on delivery
2H 2027	System launch	Sol: 192 physical qubits and ~100 logical qubits at a logical error rate of ~10^{-5}, on Quantinuum's first two-dimensional (grid) trap chip	On track (per Quantinuum)	First production-candidate Sol traps, made by Honeywell Aerospace, back from fabrication and in validation with no showstoppers so far (Q2 2026); a smaller Sol prototype was built and validated earlier
2027	Performance target	Sol two-qubit error below 2×10^{-4}, roughly 2× Helios's speed, and circuits of up to 10,000 operations	Pending	Set in the September 2024 roadmap and not restated since. Helios averages ~7.9×10^{-4}, so the error target implies a ~4× improvement
2026 ONWARD	Manufacturing	US ion-trap and laser supply chain backed by a $100M CHIPS R&D award: 300 mm trap fabrication with GlobalFoundries, lasers and optics with Monarch Quantum	In progress	Award finalized in September 2026 after a May letter of intent; no dated manufacturing milestones disclosed
2029				
2029	System launch	Apollo: universal, fully fault-tolerant system with thousands of physical and hundreds of logical qubits, running circuits with millions of operations	Pending	Company says Apollo remains on schedule, with key subsystems in prototyping. Plans call for a CMOS-built 2D trap chip and new laser delivery instead of integrated optics
2029	Roadmap target	Apollo logical error rate of 1×10^{-5} to 1×10^{-10} (current roadmap graphic)	Pending	The September 2024 roadmap set a 10^{-6} target, with 10^{-10} possible; the current graphic's 10^{-5} upper bound is looser (see Notes)
2030 AND BEYOND				
2030+	Roadmap target	Lumos: utility-scale system with more than 1,000,000 physical qubits	Pending	Concept design developed for DARPA's Quantum Benchmarking Initiative (Stage A); now in the year-long Stage B review, which assesses whether utility scale is achievable by 2033

Manufacturability

Our view: Quantinuum has built three generations of quantum computers using a consistent architecture and demonstrated its ability to deploy hardware at customer sites. Its ion trap chips already use semiconductor fabrication, but scaling lasers and optical controls remains a significant challenge. Sol will test its first larger grid architecture, while Apollo and Lumos require further integration of control electronics and optics. These advances remain unproven at commercial scale.

Building and deploying systems: Quantinuum has offered cloud access since 2020 and installed hardware at RIKEN in Japan, where its H1 system became operational in 2025 and an H2 upgrade was delivered in April 2026. Helios installations are also planned for Singapore and Oracle. These deployments establish experience manufacturing and

supporting quantum systems, although building substantially larger machines consistently and at volume remains unproven.

Scaling the chip: Quantinuum's ion traps are manufactured on silicon wafers. Sol introduces a larger two dimensional grid, using an approach demonstrated in a 2024 prototype that can control more ion trapping zones without proportionally increasing the surrounding electronics. Its first production candidate chips, manufactured by Honeywell Aerospace, are undergoing validation. Apollo will introduce further changes to chip fabrication, controls and system integration, with much larger numbers of qubits to manage.

Lasers and optics: The larger challenge is delivering precise laser control across increasing numbers of qubits. Helios uses eight operation zones that require coordinated laser beams and electrical controls. Quantinuum has improved its design through laser sharing and the transition to barium qubits, which use more readily available laser technology. However, Apollo still requires significant optical hardware, with integrated optics planned for the later Lumos generation.

Supply chain and funding: Quantinuum relies on manufacturing partners rather than owning its own semiconductor foundry. Its $100mn CHIPS R&D agreement, finalized in September 2026, supports development of a broader US manufacturing supply chain. GlobalFoundries will work on next generation ion traps and control electronics using 300mm wafers, while Monarch Quantum will develop lasers and optical components.

What we are watching: Sol is the next major manufacturing test, particularly its larger grid design and ability to operate more qubits without a proportional increase in control complexity. Beyond Sol, we will watch progress on optical integration, Apollo's manufacturing readiness and evidence that production can scale beyond individual systems. Consistent system installation and performance will be important measures of manufacturing execution.

Broader Indicators of Progress

Our view: Quantinuum's research record, patent portfolio and relationships with government and commercial customers support its technology and execution assessments. The strongest external validation comes from institutions committing to use or host Quantinuum systems, alongside continued government investment in the platform.

Research and intellectual property: Quantinuum has more than 100 quantum related publications in our dataset, reflecting research across trapped ion hardware, quantum error correction, algorithms and software. Its patent activity has accelerated since 2024, with approximately 600 published patent documents. Together, these support a broad technical foundation spanning hardware and software.

Government and research relationships: Quantinuum works with leading research institutions and government programs across the US, UK, Japan and Singapore, including NASA Ames, Sandia, Oak Ridge and RIKEN. Its participation in DARPA's Quantum Benchmarking Initiative provides another opportunity for independent evaluation of its

scaling plans. Quantinuum remained in Stage B following DARPA's October 2026 announcement of four new Stage C participants, with further advancement still possible.

Commercial partnerships: Quantinuum's August 2026 agreement with Oracle represents an important development, bringing Helios to a hyperscale cloud environment and providing a potential channel to enterprise users. The company also works with customers and partners across financial services, pharmaceuticals, materials science and technology. Its software ecosystem is expanding as well, with approximately 180 organizations using its Nexus developer platform as of 2Q26. These relationships provide growing evidence of demand, although commercial quantum adoption remains early.

Hiring and investment activity: Job postings have risen sharply since mid 2024 as Quantinuum invests in engineering, manufacturing and commercialization. While postings do not directly measure headcount, the trend is consistent with continued investment in its technology roadmap.

Taken together, Quantinuum's research, customer relationships and institutional partnerships provide meaningful external validation of its platform. The next challenge is translating these advantages into larger commercially useful systems and sustained revenue growth.

Exhibit 29: Quantinuum has built a deep research base

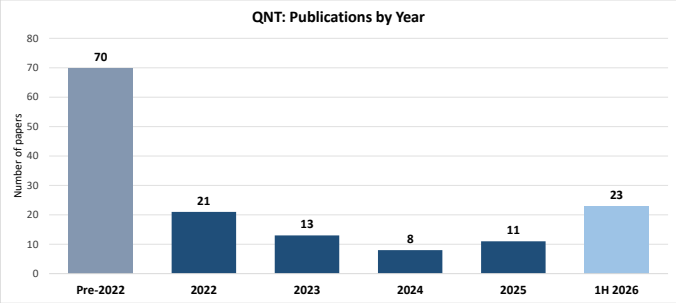


Exhibit 30: Patent activity has accelerated since 2024

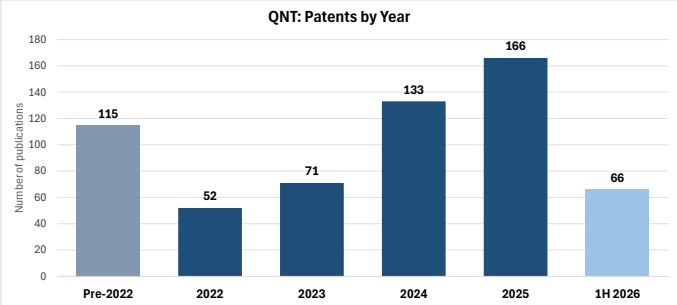
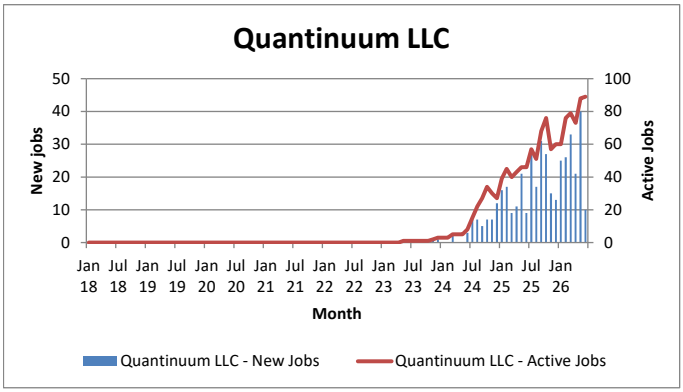


Exhibit 31: Hiring has ramped since mid-2024, supporting its transition from Helios to Sol.



Revenue Outlook and Funding

Our view: Quantinuum has one of the strongest technical and execution records among the pure plays, but its commercial revenue remains small relative to its long term opportunity. We forecast revenue accelerating as Sol expands cloud capacity and Apollo enables more commercially valuable workloads, reaching approximately $2bn in 2030. The company's large cash balance should allow it to fund its roadmap without additional equity in our base case. The main risk is Apollo's timing, given how much of our forecast depends on its successful commercialization.

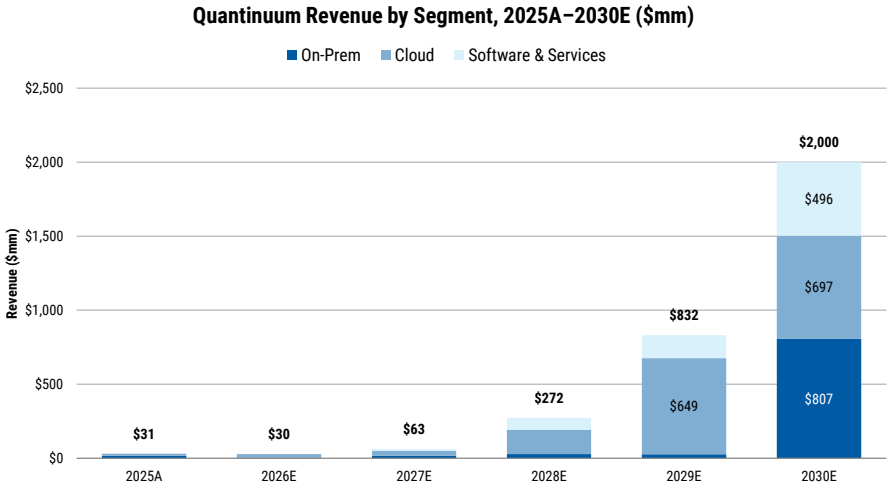
Where revenue stands today: Quantinuum generated $30.9mn of revenue in 2025, and we forecast a similar level in 2026, consistent with management's guidance of $28mn to $32mn. Most of our near term estimate is supported by signed backlog and committed programs, with government and sovereign customers accounting for approximately three quarters of our 2026 forecast. Quantinuum also describes a sales pipeline exceeding $5bn, although these opportunities span different stages of maturity and should not be treated as contracted revenue.

How the ramp builds: Quantinuum sells access to its computers through cloud subscriptions and capacity reservations, software and services, and physical systems deployed at customer sites. Its pricing is designed around the value delivered to customers rather than a fixed hardware price, allowing revenue per customer to expand as its systems support more useful calculations.

Near term revenue is supported by existing Helios engagements, including Singapore and Oracle. We expect Sol to enter commercial service in 2027, expanding cloud capacity and supporting revenue growth to approximately $270mn in 2028. Apollo, targeted for 2029, represents the more significant inflection. Its planned fault tolerant capabilities could enable customers to move beyond research and early applications toward larger production workloads, supporting increased spending on cloud access, software and strategic system deployments.

We forecast revenue reaching approximately $830mn in 2029 and $2bn in 2030. This would represent about 20% of the $10bn quantum computing market in our updated base case, making Quantinuum one of the leading providers as fault tolerant systems begin entering commercial use. That does not assume Quantinuum will be the only winner, and we expect longer term market shares to become more competitive as other architectures mature.

Exhibit 32: Quantinuum: revenue by segment



Funding and runway: Quantinuum ended 2Q26 with approximately $2.1bn of cash and no debt, following its June IPO, which raised $1.68bn in gross proceeds. We forecast free cash outflow of approximately $0.5bn in 2026, $0.45bn in 2027 and $0.32bn in 2028 as the company invests in Sol, Apollo and manufacturing capacity.

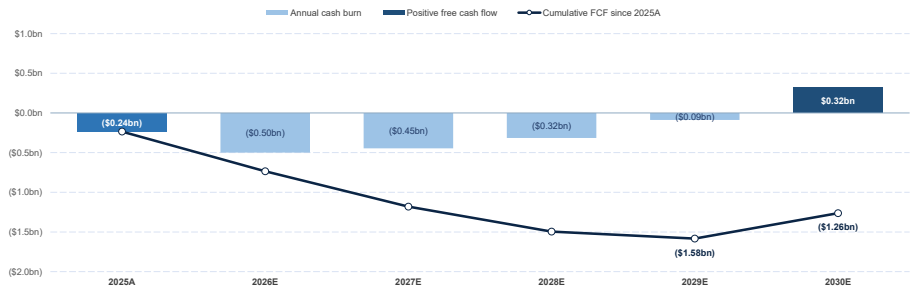
We expect burn to moderate as revenue scales, with cash reaching a low point of approximately $0.9bn at the end of 2029 and free cash flow turning positive in 2030. On our estimates, Quantinuum can fund its existing roadmap without new equity through the end of the decade.

The main sensitivity is commercialization timing. A one year delay to Apollo would defer a significant portion of our expected revenue and profitability, although the balance sheet provides a cushion against near term funding pressure. Manufacturing costs and capital spending for larger systems represent an additional variable as the company scales.

Exhibit 33: Quantinuum: funding snapshot

Metric	Value
Cash (2Q26)	~$2.1bn; no debt
Free cash flow (MSe)	2026E ($0.50bn); 2027E ($0.45bn); 2028E ($0.32bn)
Low point in cash (MSe)	~$0.9bn, end-2029E
Free cash flow breakeven (MSe)	2030E

Exhibit 34: Quantinuum: cash burn



IonQ Company Profile

Exhibit 35: IonQ: Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Trapped ion
Physical qubits (Tempo, demonstrated)	40
Logical qubits (Tempo, with discarded runs)	4
Technology framework stages	Demonstrated: 1-4 | Partial: 5, 7 | Not demonstrated: 6, 8
2Q gate fidelity (demonstrated)	99.990%
FINANCIAL HIGHLIGHTS	
Year public	2021
Gross proceeds at listing ($mn)	$636.00
FY25 revenue ($mn; FY = Dec)	$130.02
Market cap, latest ($bn)	$15.21
Liquidity: cash + marketable securities ($mn)	$2,959.33

Note: pre-SkyWater. Closing used about $0.9bn, leaving about $2b. Market cap data as of 10/14/26; $2bn funding announced 10/14/26.

Technology: Logical Computation

Our view: IonQ made meaningful error correction progress in 2026, demonstrating repeated rounds of correction and early progress toward computing with protected qubits. We have moderate confidence in its technical direction. However, the results remain relatively small scale and were achieved on existing laser controlled systems, while IonQ's roadmap depends on Superion, a new electronically controlled architecture.

In our eight stage technology framework, IonQ has demonstrated Stages 1 through 4, with partial evidence for Stages 5 and 7. Its evidence comes mainly from a June 2026 memory experiment, a September reanalysis of the same data, and an August logical-operation experiment on Forte. The partial ratings reflect two limitations: its logical memory result is at breakeven within overlapping error bars, with the clearer advantage depending on discarding runs (Stage 5), and its logical operation has not yet been shown to beat a physical baseline (Stage 7). IonQ has also not yet demonstrated Stage 6, where adding protection consistently lowers logical error.

Exhibit 36: IonQ Key Technology Results

QEC stage	Assessment	Key demonstration / source	Year
1. Encoding	Demonstrated	Tested 3-qubit repetition code circuits on Aria to study hardware noise (Ma et al.); later encoded 4 logical qubits using 18 physical qubits (Tham et al.)	2024, 2026
2. Error detection	Demonstrated	Detected errors on encoded qubits through repeated checks (Tham et al.)	2026
3. Error correction	Demonstrated	Determined corrections from measured errors after the experiment, rather than applying them live (Tham et al.)	2026
4. Repeated QEC	Demonstrated	Protected the same logical information through up to 6 correction rounds, with corrections determined afterward (Tham et al.)	2026
5. Logical improvement / breakeven	Partial	Logical memory reached approximately breakeven; stronger improvements required discarding additional runs (Tham et al.; Aydin et al.)	2026
6. Below-threshold scaling	Not demonstrated	None qualifying	—
7. Fault-tolerant logical operations	Partial	Adjustable logical rotations on Forte demonstrated, but not an advantage over unprotected operations (Huang et al.)	2026
8. Scaled fault-tolerant computation	Not demonstrated	None qualifying	—

Source: Ma et al., research preprint (2024), Physical Review Research (2025); Tham et al., research preprint (2026); Huang et al., research

Note: Assessments reflect our eight stage QEC framework. The 2024 Aria experiment studied noise using repetition code circuits rather than demonstrating improved logical reliability. IonQ's stronger QEC results emerged in 2026. Stage assessments reflect our eight stage QEC framework.

What's different about IonQ's Approach? Like Quantinuum, IonQ uses trapped ions, individual charged atoms that can store quantum information for long periods and support high quality operations. Its earlier systems use lasers to control ions held in a chain. Superion represents a significant change. Using technology acquired with Oxford Ionics, it replaces much of the laser based gate control with electronics integrated onto semiconductor manufactured chips. This could make larger systems easier to control and manufacture, but the company's published logical computing results have not yet been reproduced on Superion.

Logical Computing Progress: IonQ's main QEC demonstration came in June 2026, using a 40 ion Tempo engineering system. The researchers tested a range of different error correcting codes on the same hardware, demonstrating the flexibility of its trapped ion approach. One code encoded four logical qubits using only 18 physical qubits. More importantly, IonQ detected and corrected errors over as many as six rounds, demonstrating Stages 1 through 4. Corrections were determined in software after the experiment rather than applied live, however, and six rounds remain relatively few compared with the longer experiments reported by some peers.

The next question is whether error correction makes protected qubits more reliable than unprotected ones, known as breakeven (Stage 5). IonQ's best logical memory lasted approximately 3.95 seconds versus 3.3 seconds for its physical qubits, a marginal advantage within overlapping uncertainty ranges. The experiment also discarded runs where ions left their intended states, with nearly 75% of attempts rejected for the largest code after six rounds. A September reanalysis more than doubled logical lifetimes and moved five codes beyond breakeven, but required discarding an additional 2.6% to 5.6% of runs per correction round. We therefore rate Stage 5 as partial.

IonQ has not yet demonstrated below threshold scaling (Stage 6). In the June experiment, larger versions of the same code performed worse than smaller ones, likely because the additional measurements introduced more errors than the extra protection removed. This is an important remaining test of whether its approach can become more reliable as

systems grow.

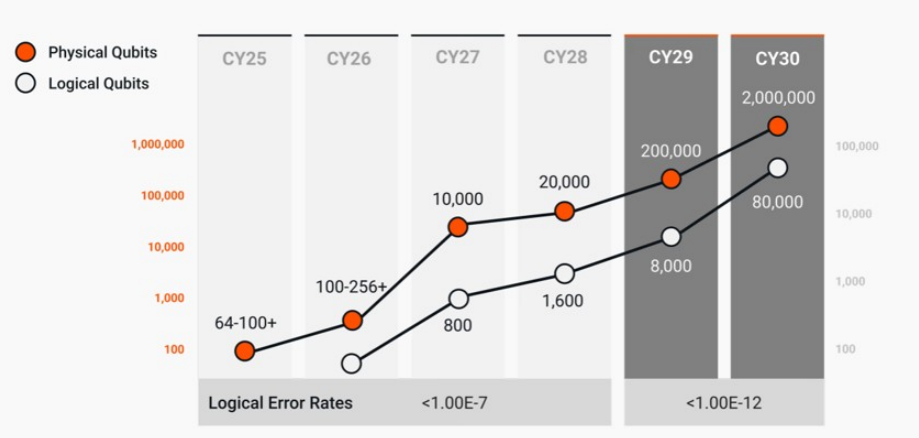
IonQ has also taken an early step toward computing with protected qubits. In August 2026, researchers used Forte to demonstrate adjustable logical rotations, a class of operations needed for universal quantum computing. The experiment was a useful proof of principle but did not show that the protected operations were more reliable than unprotected ones, supporting partial credit for Stage 7. Separately, IonQ demonstrated software capable of decoding simulated workloads involving up to 408 logical qubits on a standard CPU. This is encouraging evidence that its classical control system may be able to keep pace with future fault tolerant computers, although the decoder has not yet been tested at that scale on actual quantum hardware.

Next Steps: We are looking for IonQ to show that logical errors consistently decline as more protection is added, reduce its reliance on discarded runs and offline decoding, and extend logical operations under active error correction. Most importantly, it needs to reproduce its recent QEC progress on Superion. Demonstrating these capabilities on the architecture it intends to scale would provide stronger evidence that IonQ can progress toward sustained fault tolerant computation.

Roadmap and Execution

Our view: IonQ has a mixed longer term execution record but strong early momentum on Superion. It has consistently met or beaten its nearer term algorithmic qubit targets, while several earlier commitments were delayed, revised or achieved through different technical approaches. Its current roadmap also requires a substantial transition from laser controlled systems to Superion. We assign Moderate to High execution credibility, supported by early manufacturing progress, customer commitments and a strong balance sheet, with Superion 256 deliveries in 2027 the next major test.

Exhibit 37: IonQ Roadmap



Source: Years reflect IonQ's targets. IonQ plans to demonstrate the 256-qubit system in 2026 with customer deliveries in 2027, and a 10,000-qubit lab system in 2027 with commercial systems in 2028.

Execution Track Record: IonQ has performed well against its shorter term algorithmic qubit (AQ) milestones, its own benchmark intended to measure how large a quantum

calculation its systems can reliably perform. It reached AQ29 in mid-2023, seven months ahead of schedule; AQ35 in January 2024, a year early; and AQ64 in September 2025, three months ahead of its target.

Its record against longer dated commitments has been less consistent. IonQ's December 2020 roadmap anticipated modular quantum computers small enough to be networked together in a data center by 2023, and broad quantum advantage by 2025. The company subsequently delivered a rack mounted system to QuantumBasel in 2024, but connecting quantum computers into larger systems remains a future goal, although IonQ has recently demonstrated some of the building blocks in a recent October 9 result. IonQ also achieved AQ64 through improvements in physical hardware, control and software rather than through the error correction architecture originally contemplated. Broad quantum advantage has not yet been demosntrated. We view these developments as credible technical adaptation, although meeting a headline performance target has not always meant delivering the originally proposed capability.

IonQ's current roadmap sets a more demanding test. An older 2026 target calls for 12 logical qubits with an error rate below one in ten million, whereas its published experiments have demonstrated up to four logical qubits at a time and much higher error rates. That leaves a substantial gap to the stated target, thought IonQ has realigned their roadmap.

Superion progress: Since acquiring Oxford Ionics in 2025, IonQ has moved quickly toward its new architecture. It reported 99.99% two qubit gate fidelity on electronically controlled test hardware in October 2025 and pre sold its first Superion 256 system in early 2026. The processor went through six tapeouts during the first half of 2026, and in September IonQ announced that its first integrated 256 qubit chips had been fabricated at SkyWater, with ions successfully trapped in prototype systems.

First customer deliveries are expected in 2027. In parallel, IonQ targets a 10,000 physical qubit laboratory system in 2027, followed by commercial systems in 2028. Its roadmap also specifies 800 logical qubits with error rates below one in ten million in 2027, a substantial increase from the logical computing results demonstrated to date. Achieving this will require capabilities that have not yet been established on Superion.

What we are watching: Superion 256 is the next execution milestone. We will look for on time customer deliveries, disclosed gate performance across a functioning 256 qubit system and error correction results on the new architecture. Progress toward the 10,000 qubit laboratory system, including measured logical performance against the roadmap targets, will be equally important. Successful execution would strengthen confidence that IonQ can translate its previous engineering achievements into scalable fault tolerant systems.

Exhibit 38: IonQ Milestones we are watching

TIMING	TYPE	MILESTONE (AS STATED BY IONQ)	STATUS	STATUS DETAIL
NEAR TERM (2026 TO 1H 2027)				
2026	Roadmap target	12 logical qubits at a logical error rate below 10^{-7}, with 99.99% physical two-qubit fidelity, electronic (microwave) gates and a 2D qubit array	Not yet shown	Best published logical error rates are ~1-3% per logical qubit per cycle, with up to four logical qubits per code (June 2026)
2026	Demonstration	Superion 256 system demonstrated	In progress	Integrated 256-qubit QPUs back from SkyWater and under test; ions trapped in prototype systems at multiple facilities; full- system performance not yet disclosed
2026	Customer deployment	100-qubit Tempo systems for KISTI (Korea) and QuantumBasel (Switzerland)	In progress	Subsystems shipped to KISTI and being assembled on site; final on-site assembly at QuantumBasel (Q2 2026)
1H 2027	Commissioning	Superion 256 system commissioning and manufacturing-line ramp	On track (per IonQ)	Management reiterated the first-half 2027 timing on the Q2 2026 call
2027 TO 2028				
2027	Customer delivery	First Superion 256 customer deliveries, including installation at NVIDIA's Accelerated Quantum Research Center	Pending	Available to order; first system pre-sold in Q1 2026
2027	Lab demonstration	Fully fault tolerant 10,000 physical qubit system (Superion 10K) in a lab setting, as described by IonQ	Pending	SkyWater beginning tape-out of the 10,000-qubit chip (per Q2 2026 call)
2027	Roadmap target	800 logical qubits at a logical error rate below 10^{-7}	Pending	Roadmap does not specify lab or commercial; we read it as the target for the 2027 lab system
2028	Commercial availability	Manufacturable, commercial Superion 10K systems	Pending	IonQ positions 2028 as the commercial follow-on to the 2027 lab system
2028	Roadmap target	20,000 physical and 1,600 logical qubits at a logical error rate below 10^{-7}; June 2025 plan described two interconnected chips and a commercially available interconnected system	Pending	IonQ says its ~20,000-qubit resource estimate for breaking 256-bit elliptic-curve signatures is consistent with systems in this timeframe
2029 TO 2030				
2029	Roadmap target	8,000 logical qubits at a logical error rate below 10^{-12}; photonic interconnect; ~200,000 physical qubits	Pending	Roadmap page lists the photonic interconnect in 2029, whereas the June 2025 plan had interconnected systems in 2028
2030	Roadmap target	2 million physical and 80,000 logical qubits at a logical error rate below 10^{-12}	Pending	The June 2025 announcement gave a range of 40,000-80,000 logical qubits; the roadmap page now states 80,000

Manufacturability

Our view: IonQ has experience building and delivering quantum systems, while its acquisition of SkyWater provides direct access to semiconductor manufacturing for its next generation architecture. The shift from laser based gates to electronic control could simplify scaling, and early Superion manufacturing progress has been encouraging. However, complete system performance remains unproven at 256 qubits, with substantial integration challenges ahead at 10,000 qubits. First Superion deliveries in 2027 are the next major test.

Building and deploying systems: IonQ has provided cloud access to its computers since 2021 and has experience installing hardware at customer sites. It delivered Forte Enterprise to QuantumBasel in 2024 and launched a system at EPB's Quantum Center in Chattanooga in September 2026. Through Oxford Ionics, it also has experience supplying systems to the UK's National Quantum Computing Centre. Tempo systems are being deployed for QuantumBasel and South Korea's KISTI. These demonstrate the ability to manufacture and commission existing systems, although Superion represents a new architecture.

Owning chip manufacturing: IonQ's July 2026 acquisition of SkyWater brought semiconductor fabrication capabilities in house. The companies have already produced integrated 256 qubit processors, with six chip tapeouts during the first half of 2026. Management says working with SkyWater has reduced development cycles from approximately nine months to two, allowing faster design iterations.

Superion uses electronic gate controls integrated onto the chip, reducing the need for the

laser systems used in earlier IonQ computers. Lasers remain necessary for qubit preparation and measurement, but the new design could simplify manufacturing larger systems. The main remaining test is whether the chip and its surrounding hardware can deliver the targeted performance as a complete system.

Scaling to 10,000 qubits: Superion 10K would represent an approximately 40 times increase in physical qubit count from Superion 256 and is expected to incorporate additional cryogenic control electronics. IonQ is developing the two generations in parallel, targeting a 10,000 qubit laboratory demonstration in 2027 and commercial availability in 2028. These advances remain unproven at scale, while longer term growth will also require reliable photonic connections between multiple chips.

Customer validation: Five customers and partners have announced commitments involving Superion 256: the University of Cambridge, Horizon Quantum, NVIDIA, South Korea's SDT and Florida International University. These provide early demand validation before the new system ships, although they do not yet establish production performance or repeatable manufacturing.

What we are watching: Superion 256 commissioning and customer deliveries in 2027 are the next major manufacturing tests. We will be watching full system gate performance, manufacturing yields and installation timelines. Beyond that, successful integration of control electronics for Superion 10K and independent hardware validation through DARPA's Quantum Benchmarking Initiative Stage C will provide important evidence of whether IonQ's approach can scale.

Broader Indicators of Progress

Our view: IonQ has built a broad research, intellectual property and institutional footprint, supported by a growing network of government, academic and industry relationships. Acquisitions have materially expanded its technology portfolio and commercial reach. We view these indicators as supporting evidence of IonQ's ability to execute its roadmap.

Research and intellectual property: IonQ has more than 150 quantum related publications in our dataset, including foundational trapped ion research from its founders' academic laboratories. More recent publications focus on error correction, hardware engineering, applications and quantum networking. Its patent portfolio is substantially larger, at approximately 700 publications, with acquired intellectual property accounting for a meaningful share.

Government and research partnerships: IonQ participates in DARPA's HARQ program on quantum interconnects and works with the Air Force Research Laboratory on quantum networking. Its relationships also include the University of Maryland and South Korea's KISTI, reflecting growing public sector demand for computing and networking technology.

A particularly important development came in October 2026, when DARPA advanced IonQ to Stage C of its Quantum Benchmarking Initiative. The program will independently evaluate whether future generations of IonQ hardware perform as its technical roadmap anticipates. IonQ's advancement provides external validation of its development plan,

although successful performance in the upcoming evaluations remains to be demonstrated.

Broader quantum platform: IonQ's acquisitions have expanded its capabilities beyond computing into networking, security, sensing, photonics and satellite systems. These businesses broaden its customer relationships and intellectual property, while also increasing the complexity of integration. IonQ also recently announced they plan to acquire Elisity, a Zero Trust security company.

Hiring and investment activity: Active job postings have increased severalfold since early 2025, with significant activity around its manufacturing facilities and College Park headquarters. While postings are not a direct measure of headcount, the trend is consistent with increased investment ahead of Superion commercialization.

Taken together, IonQ's research capabilities, government relationships and industry partnerships support its position in the quantum ecosystem. The next step is converting this institutional support and expanding manufacturing capacity into reliably performing commercial systems.

Exhibit 39: IonQ has a broad research base focused on hardware engineering and networking

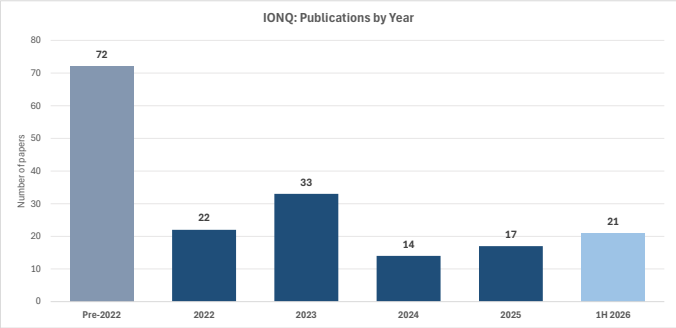


Exhibit 40: Acquisitions have expanded its IP portfolio to roughly 700 patent publications

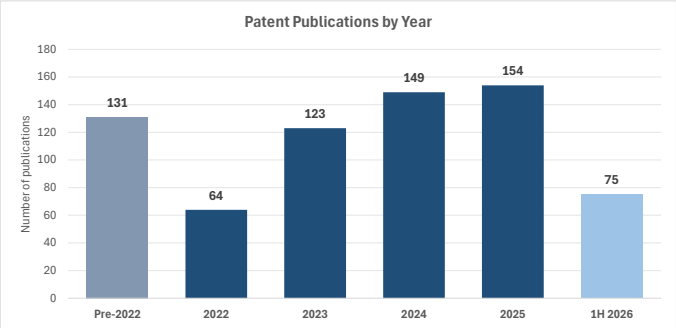
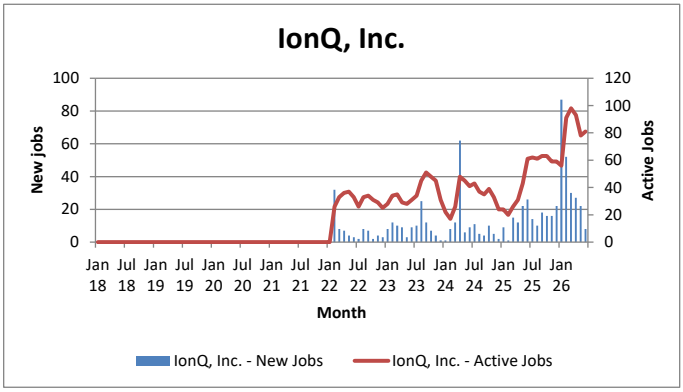


Exhibit 41: Hiring has accelerated since early 2025, supporting its manufacturing scale-up and Superion commercialization



Revenue Outlook and Funding

Our view: IonQ has the largest revenue base among the quantum pure plays,

although much of its reported revenue comes from acquired businesses rather than quantum computing. We forecast computing revenue increasing from approximately $136mn in 2026 to $1.8bn in 2030 as Superion systems reach customers and cloud adoption expands. IonQ also has a strong balance sheet, which we expect will allow it to fund its roadmap without additional equity through 2030.

Where revenue stands today: IonQ generated $130mn of revenue in 2025, and we forecast approximately $455mn in 2026. Our estimate comprises roughly $136mn of computing revenue, $150mn from acquired businesses and $170mn of net SkyWater foundry contribution following its July acquisition. The total is consistent with IonQ's September 2026 consolidated revenue guidance of $450mn to $460mn. We have not included the contribution from Elisity. Computing revenue today comes from hardware systems, cloud access and government programs, while acquisitions have broadened the company's revenue base across several other quantum and adjacent markets.

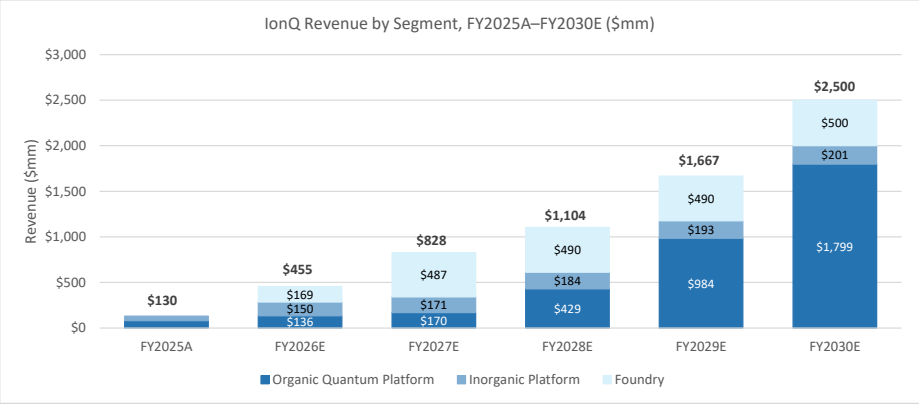
How the computing ramp builds: We expect computing revenue to increase to approximately $170mn in 2027, supported by customer deployments and continued government demand. The more meaningful acceleration begins in 2028 as IonQ targets commercial availability of its larger Superion systems. We forecast computing revenue of approximately $430mn in 2028, $980mn in 2029 and $1.8bn in 2030, driven by system sales, cloud access and broader customer adoption.

The main risk is timing. Our forecast assumes that IonQ can translate its early Superion manufacturing progress into commercially operating systems and eventually demonstrate the error correction performance required for larger fault tolerant computers. Delays would push revenue into later years.

Foundry and acquired businesses: SkyWater adds a US semiconductor foundry serving government, aerospace and commercial customers, while also manufacturing processors for Superion. We model its annual revenue at approximately $0.5bn from 2027, before accounting for sales between SkyWater and IonQ that must be eliminated on consolidation.

IonQ's other acquisitions have expanded its presence in quantum networking, security, sensing, photonics and satellite imagery. We forecast approximately $150mn of revenue from these businesses in 2026, rising to $200mn in 2030. We value these activities separately from computing and exclude their revenue from our quantum computing market share calculation.

Exhibit 42: IONQ: revenue by segment



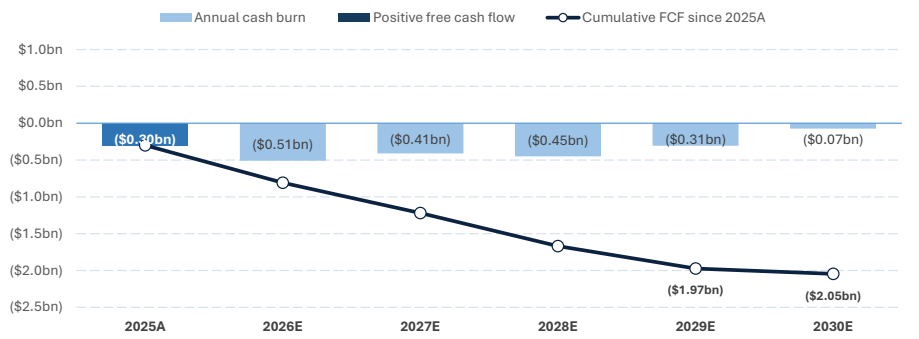
Funding and runway: IonQ ended 2Q26 with approximately $3.0bn of cash and investments and no debt, supported by substantial equity issuance in 2025. Completing the SkyWater acquisition in July required approximately $0.9bn of cash, including the acquisition payment and repayment of SkyWater's revolver, leaving pro forma cash and investments of approximately $2.0bn. On 10/14/2026, IonQ priced an additional ~$2bn equity offering, which we include in our 4Q cash.

We expect annual free cash outflows of approximately $0.4–0.5bn from 2026 through 2028 as IonQ invests in Superion development and manufacturing. Burn should moderate as computing revenue scales, with adjusted EBITDA turning positive in 2030 in our model. On our estimates, IonQ can fund its existing roadmap without additional equity through the end of the decade, although that conclusion depends on achieving the expected commercialization ramp.

Exhibit 43: IonQ: funding snapshot

Metric	Value
Cash & investments (2Q26)	~$3.0bn; no debt
Free cash flow (MSe)	2026E ($0.51bn); 2027E ($0.41bn); 2028E ($0.45bn)
Low point in cash & investments, post-offering (MSe)	~$2.6bn, 2Q30E
Free cash flow breakeven (MSe)	~2030E (first positive quarter: 3Q30E)

Exhibit 44: IonQ: cash burn



Xanadu Company Profile

Exhibit 45: XNDU Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Photonic
Physical qubits (flagship)	12
Logical qubits (Aurora 2025 modular system)	0
Technology framework stages	Demonstrated: 1 | Partial: 2-3 | Not demonstrated: 4-8
Fault-tolerance squeezing threshold (dB)	0.62
FINANCIAL HIGHLIGHTS	
Year public	2026
Gross proceeds at listing ($mn)	$301.65
FY25 revenue ($mn; FY = Dec)	$4.62
Market cap, latest ($bn)	$1.14
Liquidity: cash & equivalents ($mn)	$312.78

Note: market cap as of October 14, 2026.

Technology: Logical Computation Analysis

Our view: Xanadu's photonic approach offers potential advantages in manufacturing and scalability, but its path to fault tolerance remains at an early experimental stage. Aurora demonstrated the ability to connect photonic modules and respond to error measurements in real time, while a separate experiment produced GKP encoded qubit states on a chip. However, these capabilities have not yet been combined in a system that corrects errors on protected logical qubits. We therefore see more uncertainty in Xanadu's technical roadmap than for more advanced peers.

In our eight stage technology framework, Xanadu has demonstrated Stage 1 encoding, with partial evidence for Stages 2 and 3. Its strongest evidence comes from two experiments published in 2025. One generated GKP qubit states on an integrated photonic chip; the other, Aurora, demonstrated the hardware and real time control needed to detect and respond to errors. Neither demonstrated reliable error correction on protected logical qubits, and Stages 4 through 8 remain unproven.

Exhibit 46: Xanadu: key technology results

QEC stage	Assessment	Key demonstration / source	Year
1. Encoding	Demonstrated	Generated GKP encoded qubit states on a photonic chip (Larsen et al.)	2025
2. Error detection	Partial	Aurora performed error checks with real-time decoding, but mainly on light not encoded into protected qubits (Aghaee Rad et al.)	2025
3. Error correction	Partial	Aurora used error-check results to adjust subsequent measurements, but did not demonstrate correction of protected qubits (Aghaee Rad et al.)	2025
4. Repeated QEC	Not demonstrated	None qualifying	—
5. Logical improvement / breakeven	Not demonstrated	None qualifying	—
6. Below-threshold scaling	Not demonstrated	None qualifying	—
7. Fault-tolerant logical operations	Not demonstrated	None qualifying	—
8. Scaled fault-tolerant computation	Not demonstrated	None qualifying	—

Note: Assessments reflect our eight stage QEC framework. Xanadu has demonstrated GKP qubit generation and real time error detection and control, but has not yet demonstrated error correction on protected logical qubits.

What's different about Xanadu's approach: Xanadu uses light to carry quantum information, rather than atoms or superconducting circuits. Its photonic chips can be manufactured using semiconductor processes and connected through optical fibers, while most components operate at room temperature. These characteristics offer a potential route to large, modular systems. The main challenge is photon loss: photons can disappear as they pass through chips, fibers and other optical components, creating errors that become harder to manage as systems grow.

Xanadu plans to protect quantum information using two layers of error correction. The first uses GKP qubits, which encode information in specially structured states of light to make certain errors easier to identify and correct. The second would use a quantum low density parity check (qLDPC) code to provide additional protection across groups of GKP qubits.

Logical Computing Progress: In June 2025, Xanadu demonstrated GKP qubit states generated on a photonic chip fabricated using 300mm wafer technology. This was an important step toward scalability, as previous attempts to make GKP qubits using light relied on complex laboratory optical equipment rather than semiconductor chips that could eventually be manufactured in volume. However, the qubits were still too noisy for reliable error correction. They achieved 0.62 dB of effective squeezing, a measure of how well quantum noise is suppressed, compared with the roughly 9.75 dB that Xanadu has estimated is needed for fault tolerance in one version of its design. The experiment therefore demonstrated a potentially scalable way to produce GKP qubits, but significant improvements in optical loss and qubit quality are still needed.

Aurora, published in January 2025, demonstrated a different capability. It connected 35 photonic chips across four server racks and produced 12 physical optical modes per clock cycle. Importantly, its decoder processed error measurements in roughly 64 nanoseconds, with the full feedback process completed in 976 nanoseconds, within the system's one microsecond operating cycle. This demonstrated the ability to respond to errors fast enough to influence the next measurement, an important capability for photonic computing where photons move quickly through the system.

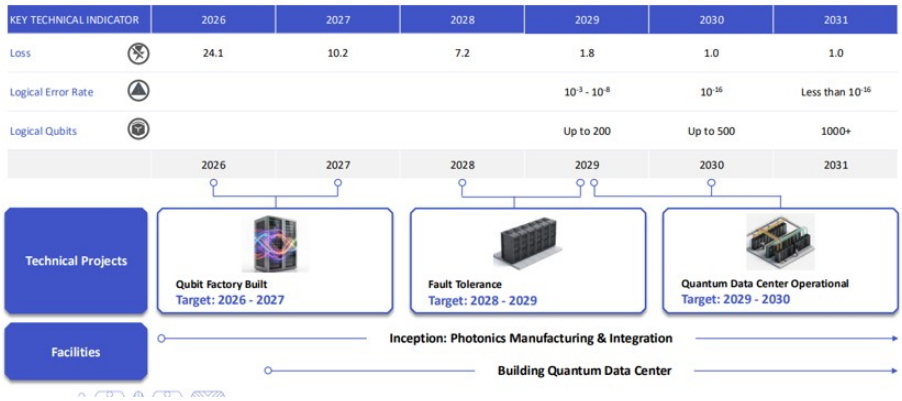
Two limitations keep Aurora from receiving full credit for Stages 2 and 3. First, only about 3% of the optical modes contained rudimentary GKP like states, with most of the experiment using simpler squeezed light. Second, its error detection code could identify certain errors but not correct them. Aurora therefore demonstrated much of the machinery needed for error correction, but not successful correction of protected logical information.

Next steps: The key technical milestone is bringing these capabilities together: producing higher quality GKP states, detecting errors on those states and correcting them during a calculation. Xanadu would then need to show that correction improves reliability, continues over repeated rounds and becomes more effective as the system scales. These are the capabilities needed to move from promising engineering demonstrations toward fault tolerant computation.

Roadmap and Execution

Our view: Xanadu has demonstrated meaningful engineering progress, from Borealis in 2022 to the modular Aurora system in 2025. However, its record against dated public commitments remains limited. The company introduced a detailed investor roadmap in March 2026 and revised several longer term targets just six months later, generally toward later dates or lower qubit counts. We therefore assign Earlier Stage execution credibility, with the next two years providing the first meaningful test of its roadmap.

Exhibit 47: XNDU Roadmap



Source: Company data (August 2026 roadmap). Note: Loss is the ratio of hardware loss to the level Xanadu estimates is required for fault tolerance, aggregated across 17 key optical components along the most demanding optical path. A ratio of 1.0x equals the required level.

Execution Track Record: Xanadu's earlier achievements demonstrate technical capability but provide limited evidence of performance against public deadlines. Borealis, introduced in 2022, made a claim of quantum computational advantage on a specialized sampling task, while Aurora demonstrated modular networking and real time control in 2025. These were significant engineering milestones, although neither established fault tolerant computation.

In March 2026, shortly before its public listing, Xanadu published a roadmap targeting completion of its Qubit Factory during 2026 to 2027, fault tolerance in 2028 and up to

500 logical qubits alongside the start of a quantum data center in 2029. In August, the company updated that roadmap, moving its fault tolerance milestone to 2028 to 2029, reducing its 2029 logical qubit target to up to 200 and moving the 500 logical qubit target to 2030. It also extended the data center window to 2029 to 2030 and added a goal of more than 1,000 logical qubits in 2031.

We would not consider these missed commitments, since none of the affected deadlines had passed. However, revisions so soon after the initial roadmap highlight uncertainty around the pace of development.

Looking ahead: Optical loss reduction is the clearest measure of execution against Xanadu's current plan. Its August roadmap targets a loss indicator of 24.1 times the required level in 2026, improving to the required level by 2030. Achieving these targets will depend on improvements in manufacturing processes, individual optical components and system integration. The company has said it is on track for its 2026 milestone, but has not publicly disclosed a measured result on the same basis as its roadmap.

What we are watching: The main near term tests are progress Qubit Factory technology within the targeted 2026 to 2027 window and reporting measurable progress against the optical loss roadmap. Beyond that, the key execution milestone is the company's targeted transition to fault tolerance during 2028 to 2029. Delivery against these commitments would strengthen confidence in its longer term logical qubit and data center targets. Delays, further timeline revisions or a lack of comparable performance data would increase execution uncertainty.

Exhibit 48: Xanadu: milestones we are watching

TIMING	TYPE	MILESTONE (AS STATED BY XANADU)	STATUS	STATUS DETAIL
NEAR TERM (2026 TO 2027)				
2026	Roadmap target	Key loss technical indicator of 24.1x (hardware loss vs. the level required for fault tolerance, across ~17 optical components; 1.0x = required level)	On track (per Xanadu)	Management says it is on track for year-end, though progress may be non-linear; no measured value disclosed. The ~5-10x remaining gap cited on the Q2 2026 call does not obviously reconcile (see Notes)
2026-2027	Technical project	Qubit Factory built (described by management as volume production of photonic GKP qubits)	In progress	Timing unchanged since March 2026. Q2 2026 fab runs rose ~50-75% to speed Qubit Factory chip iterations; no completion criteria defined. Best published on-chip GKP states still need lower loss to reach fault-tolerant quality
2026 ONWARD	Facility	Inception: 158,000 sq ft Toronto site for testing, heterogeneous integration, chip packaging and module assembly	In progress	Supported by CAD 195M of federal Project OPTIMISM funding (August 2026); no completion or commissioning date given
2027	Roadmap target	Loss indicator of 10.2x	Pending	Chart only; the press release gives just the 2026 and 2030 endpoints. Implies a ~2.4x reduction during 2027
2028 TO 2029				
2028	Roadmap target	Loss indicator of 7.2x; logical error rate of roughly 10^{-3} to 10^{-8}	Pending	Press release puts the error-rate range in 2028; the roadmap chart shows it under 2029. The range spans five orders of magnitude
2028-2029	Technical project	Fault tolerance milestone	Pending	Fault-tolerant operations by 2028 in April-May 2026; widened to 2028-2029 in August (CEO cited 2029 or 2030 on the Q2 call). What will be demonstrated is not defined; loss reaches 1.0x only in 2030
2029	Roadmap target	Up to 200 logical qubits; loss indicator of 1.8x	Pending	Down from up to 500 in 2029 at the March Analyst Day (500 in 2029-2030 in April-May). The 7.2x to 1.8x step (~4x) is the largest single-year loss reduction on the roadmap
2029 TO 2031				
2029-2030	Quantum data center	Quantum data center operational	Pending	March deck targeted 2029 (with ~100k physical and up to 500 logical qubits listed for 2029+); 2029-2030 since the Q1 2026 call. Chart shows construction starting before the operational window
2029-2030	Commercialization	Meaningful end-customer commercialization	Pending	2029+ in March; 2029-2030 in August. Paths: cloud access, enterprise PennyLane, system sales and photonic IP licensing. Q2 2026 revenue was $1.5M, mainly DARPA QBI Stage B
2030	Roadmap target	Up to 500 logical qubits; logical error rate of 10^{-16}; loss indicator of 1.0x (required level)	Pending	First year the loss indicator reaches the required level, after the 2028-2029 fault tolerance window opens
2031	Roadmap target	1,000+ logical qubits; logical error rate below 10^{-16}	Pending	Added in August. Loss held at 1.0x; management says that below threshold, more hardware and error correction lower logical error further

Manufacturability

Our view: Xanadu's photonic architecture offers a potentially scalable manufacturing path through semiconductor fabrication, modular systems and optical fiber connections. The company has demonstrated photonic chip fabrication on 300mm wafers and is assembling a broad network of manufacturing partners. However, it has not yet demonstrated volume production or announced quantum computers installed at customer sites. We assign Earlier Stage manufacturability, with reducing optical loss and establishing repeatable production processes the key challenges.

Manufacturing Approach: Unlike quantum computing approaches that require specialized qubit fabrication or extensive cryogenic infrastructure, Xanadu designs photonic chips that can leverage established semiconductor manufacturing processes. Most of its optical processing operates at room temperature, with cryogenic cooling required for certain photon detectors. The architecture is designed to connect multiple modules over optical fibers, potentially allowing capacity to scale by adding more chips and systems.

Xanadu has built several generations of hardware in its own laboratories, while Borealis was accessible to outside users through cloud services. However, the company has not announced customer site installations.

The Qubit Factory: Xanadu's near term manufacturing priority is its Qubit Factory, targeted for completion during 2026 to 2027. It is intended to support repeatable production of the GKP qubits required by its architecture. The company's 2025 chip experiment demonstrated an initial manufacturing path using 300mm silicon nitride wafers, while increased fabrication activity in 2026 is supporting faster design iterations, although Xanadu has not disclosed specific completion or production qualification criteria for the Qubit Factory.

Optical loss and manufacturing partners: Photons are lost as they travel through chips, fibers and optical connections, making fabrication accuracy, packaging and component quality central to Xanadu's ability to scale. Its partnerships with established semiconductor companies are intended to address these challenges.

Tower Semiconductor is supporting photonic chip development, while an October 2026 agreement with GlobalFoundries targets future production of low loss silicon nitride circuits and advanced photon detectors on 300mm manufacturing lines. Xanadu is also working with ASML on lithography processes to reduce optical loss. These partnerships provide access to established manufacturing expertise.

Expanding manufacturing capacity: Xanadu is developing Inception, a 158,000 square foot facility in Toronto for photonic chip packaging, testing, integration and quantum module assembly. The facility forms part of its approximately C$893mn Project OPTIMISM, supported by a signed agreement for up to C$195mn of conditionally repayable federal funding. The investment should expand Xanadu's manufacturing capabilities, although production yields, qualification milestones and the economics of manufacturing at scale remain uncertain.

What we are watching: The key manufacturing tests are completion of the Qubit Factory technology, measured improvements in optical loss, successful production of components through its foundry partnerships and commissioning of the Inception facility. The first

customer site system installation would provide additional evidence that Xanadu can move beyond laboratory systems toward commercial deployment.

Broader Indicators of Progress and Other Considerations

Our view: Xanadu has established a strong research base, particularly in quantum algorithms and software, while PennyLane provides an existing developer ecosystem ahead of meaningful hardware commercialization. Government programs and partnerships with established technology companies offer further validation of its research and development efforts. We view these as encouraging indicators of the company's long term potential.

Research and intellectual property: Xanadu has produced more than 170 quantum related publications in our research dataset, with a particular strength in quantum machine learning, algorithms and photonic hardware. Its work includes peer reviewed demonstrations of Borealis, Aurora and on chip GKP qubit generation. Its patent footprint is smaller than its publication record, with fewer than 100 patent publications in our dataset, reflecting a more focused research intensive platform still progressing toward commercialization.

Government and public sector relationships: Xanadu's government backing is concentrated in Canada, but also includes meaningful engagement in the US. The company was selected for Canada's Quantum Champions Program, offering up to C$23mn of support for quantum computing development. It is also participating in Stage B of DARPA's Quantum Benchmarking Initiative, which evaluates the technical plans and risks of companies pursuing utility scale quantum computing. Xanadu was not among the companies selected for Stage C in DARPA's October 2026 announcement, although they are still part of the program.

Canada's commitment to Project OPTIMISM, discussed in Manufacturability, provides a further source of support for Xanadu's planned expansion. Together, these programs demonstrate government willingness to fund the company's development, but do not establish that its technical roadmap will be achieved.

PennyLane and commercial partnerships: One of Xanadu's strongest assets today is PennyLane, its open source quantum software platform. It allows developers to create quantum applications across different hardware architectures, giving Xanadu a presence in the developer ecosystem before its own computers reach commercial scale. Approximately 31% of respondents to the Unitary Foundation's 2025 quantum software survey reported using PennyLane over the preceding year, second only to IBM's Qiskit. Xanadu also reports approximately 1.1mn monthly package installs across its software family.

Xanadu has also established partnerships spanning quantum applications and technology development, including work with Rolls Royce, Oak Ridge National Laboratory and Lockheed Martin. We view these relationships as potentially useful channels for future software and hardware commercialization, although the revenue opportunity remains early.

Hiring and investment activity: Active job postings rose sharply in the first half of 2026, with hiring activity concentrated in photonics engineering, fabrication and process integration. Job postings are an imperfect measure of workforce growth, but the trend is consistent with increased investment in the manufacturing capabilities needed to support Xanadu's roadmap.

Taken together, Xanadu's research record, software adoption and external partnerships provide an established foundation for continued development. The next step is translating these strengths into commercially deployable hardware and paying customers.

Exhibit 49: Xanadu has a strong research base, particularly in quantum algorithms and software...

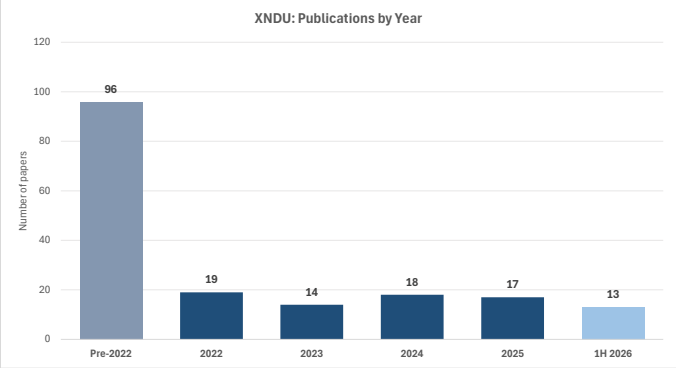


Exhibit 50: ...But a relatively small patent footprint

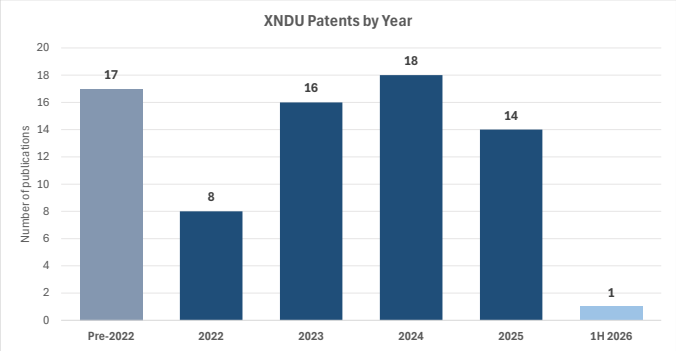
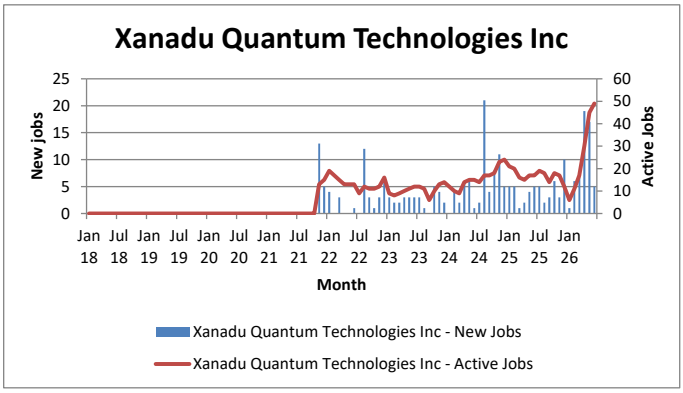


Exhibit 51: Hiring has accelerated in 2026, supporting its planned photonic manufacturing build-out



Revenue Outlook and Funding

Our view: Xanadu's revenue remains small and primarily reflects government programs and development services rather than commercial quantum hardware. We expect the start of commercialization to begin in 2029, with revenue reaching approximately $150mn in 2030. The company has the most constrained funding outlook among our covered pure plays, and we expect that they may need additional funding to get to 2030.

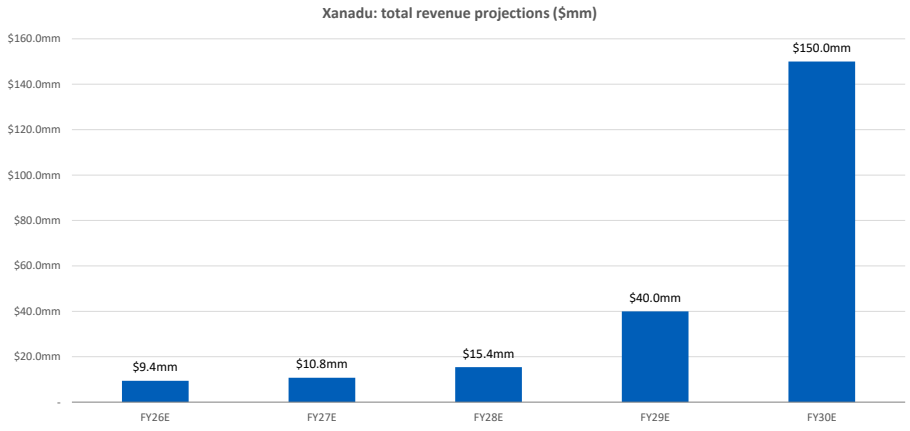
Where revenue stands today: Xanadu generated approximately $4.6mn of revenue in 2025, and we forecast approximately $9mn in 2026. Its customer base remains concentrated, with three customers accounting for 78% of 2025 revenue. Revenue largely consists of professional services and funded development programs rather than system sales or recurring cloud consumption.

DARPA's Quantum Benchmarking Initiative has been an important contributor to recent growth, with Stage B work contributing to revenue in 1H 26. We expect revenue to remain dominated by government and development programs while Xanadu advances its hardware toward commercialization, reaching approximately $15mn in 2028.

How the computing ramp builds: Our model assumes that Xanadu begins generating meaningful commercial revenue in 2029 through a combination of enterprise software, hardware sales to government and research customers, photonic IP licensing and eventually cloud access. We forecast revenue of approximately $40mn in 2029 and $150mn in 2030. These estimates are primarily based on an assumed share of the future quantum computing market rather than contracted system deliveries or an established customer revenue base. The forecast therefore depends heavily on Xanadu progressing toward fault tolerance and establishing commercial manufacturing capabilities within its stated roadmap window.

Our $150mn forecast is smaller than the other pure plays. This primarily reflects the timing of commercialization, with 2030 representing only the beginning of commercial revenue in our model. A delay of even one year would move a substantial portion of the expected revenue opportunity beyond our forecast horizon.

Exhibit 52: XNDU: revenue projection



Funding and runway: Xanadu ended the second quarter of 2026 with about $313mn of cash and $33mn of debt. This was helped by $67mn raised during the quarter through its $300mn equity facility with Yorkville. We forecast about $87mn of free cash outflow in 2026, rising to about $150mn–$215mn a year from 2027 through 2030 as the company invests in manufacturing and engineering. It does not reach positive free cash flow within our forecast period.

On our estimates, existing cash alone gives about two years of runway. We therefore

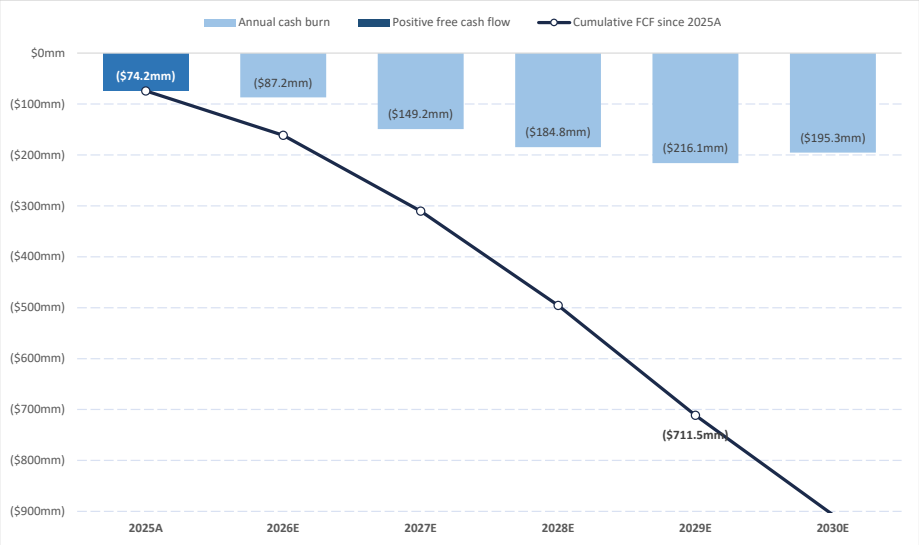
assume Xanadu will need some form of additional funding before 2030. We also assume about $140mn of repayable federal support for Project OPTIMISM is drawn over 2027-30, which takes total debt to about $180mn by 2030, though actual grant contribution may differ. Our model excludes the separate Ontario support of up to C$195mn, which is still pending. That is a potential further source of liquidity for eligible manufacturing investments.

The largest sensitivity remains commercialization timing. If Xanadu's technical or manufacturing milestones slip, its revenue ramp would be delayed while the company keeps incurring heavy development costs, increasing its dependence on external financing.

Exhibit 53: Xanadu: funding snapshot

Metric	Value
Cash (2Q26)	~$312.8mm; $32.5mm debt
Free cash flow (MSe)	2026E –$87.2mm; 2027E –$149.2mm; 2028E –$184.8mm
Low point in cash (MSe)	~$33.8mm, end-2030E
Free cash flow breakeven (MSe)	Not reached by 2030E

Exhibit 54: Xanadu: cash burn



Infleqtion Company Profile

Exhibit 55: INFQ Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Neutral atom
Physical qubits (NQCC system)	~100
Logical qubits (error detection / post-selection)	30
Technology framework stages	Demonstrated: 1-3 | Partial: 5, 7 | Not demonstrated: 4, 6, 8
2Q gate fidelity, CZ (%)	99.730%
FINANCIAL HIGHLIGHTS	
Year public	2026
Gross proceeds at listing ($mn)	>$550
FY25 revenue, as adjusted ($mn)	$31.11
Market cap, latest ($bn)	$2.47
Liquidity: cash + AFS securities ($mn)	$581.75

Note: The 100 physical qubit figure refers to Infleqtion's system delivered to the UK's NQCC, rather than the largest atom array demonstrated. Gate fidelities are from Radnaev et al., PRX Quantum (2025). Market cap as of October 14, 2026.

Technology: Logical Computation Analysis

Our view: We have confidence in Infleqtion's technical direction, but there is more uncertainty because neutral atoms are a younger approach and Infleqtion is earlier on our QEC framework. It has made meaningful progress in a short period: since 2024, it has expanded from two logical qubits to a company-reported 30 in 2026, and it has shown it can correct certain errors after a calculation. However, it has not yet shown that it can correct errors continuously during a calculation, or that increasing protection consistently improves reliability.

In our eight stage technology framework, Infleqtion has demonstrated Stages 1 through 3, with partial evidence for Stages 5 and 7. Its strongest results demonstrate logical encoding, error detection and limited correction, but improvements have often depended on discarding experimental runs where errors were detected. It has not yet demonstrated repeated error correction (Stage 4), improving reliability with stronger protection (Stage 6) or scaled fault tolerant computation (Stage 8).

Exhibit 56: Infleqtion: key technology results

QEC stage	Assessment	Key demonstration / source	Year
1. Encoding	Demonstrated	2 logical qubits (Bedalov/Chung); 12 logical qubits (Rines); 30 logical qubits (Infleqtion)	2024, 2025, 2026
2. Error detection	Demonstrated	Detected errors on 2 logical qubits and discarded affected runs (Bedalov/Chung)	2024
3. Error correction	Demonstrated	Corrected certain errors on 4 logical qubits after experiment while still discarding many runs (Rines)	2025
4. Repeated QEC	Not demonstrated	None qualifying	—
5. Logical improvement / breakeven	Partial	Selected logical calculations outperformed physical qubits, but relied on discarding runs (Bedalov/Chung; Rines)	2024, 2025
6. Below-threshold scaling	Not demonstrated	None qualifying	—
7. Fault-tolerant logical operations	Partial	Logical operations on up to 12 qubits (Rines); 30-qubit circuit including 4 non-Clifford gates (Infleqtion)	2025, 2026
8. Scaled fault-tolerant computation	Not demonstrated	None qualifying	—

Source: Bedalov et al., research preprint (2024), subsequently published as Chung et al., npj Quantum Information (2025); Rines et al., research

What's different about Infleqtion's approach? Infleqtion stores quantum information in neutral atoms, which are individual uncharged atoms held in place by tightly focused laser beams. Like trapped ions, neutral atoms can be physically moved to bring distant qubits together, and they can be packed into large two dimensional arrays. This flexibility could make error correction more efficient and reduce the number of physical qubits needed for each logical qubit as systems scale.

The main challenges are atom loss and measurement. Atoms can escape their traps during a calculation, while checking for errors can disturb nearby atoms and damage the information being protected. Infleqtion is developing a design using two types of atoms, rubidium and cesium, where one stores information and the other checks for errors. A 2026 study demonstrated some of the measurement capabilities needed for this approach, but Infleqtion has not yet shown repeated error correction using the two atom design.

Logical Computing Progress: Infleqtion's first logical results were reported in 2024 and published in a peer reviewed journal in 2025. The company protected two logical qubits using a code that could detect errors but not correct them. By checking for errors and discarding runs where something went wrong, the researchers improved results relative to unprotected qubits. This demonstrated Stages 1 and 2, but not error correction.

Its strongest evidence for Stage 3 came in 2025, when Infleqtion used 16 atoms to encode four logical qubits. After the calculation finished, software analyzed the results and worked out corrections for certain errors, reducing the error rate in preparing the logical state from roughly 13% to 1.5%. However, the experiment still discarded most attempts, retaining only about 9% of runs. This demonstrated that certain errors could be corrected after a calculation, but not continuous error correction while a calculation is running.

Infleqtion has also made progress in computing with protected qubits. In 2025, it ran circuits across 8 and 12 logical qubits with two to four times lower error than comparable unprotected circuits. However, these improvements came from detecting errors and discarding unsuccessful runs rather than fixing errors during the calculation. In the 12 logical qubit experiment, fewer than 1% of runs were retained. The September 2026 demonstration extended this work to 30 logical qubits and included four non Clifford gates, an important class of operations needed for universal quantum computing. That

experiment also relied on discarding runs where errors were flagged, although some errors from lost atoms could be corrected afterward. It therefore did not demonstrate continuous error correction. We therefore rate Stage 7 as partial.

Next steps: The main technical milestone is repeated error correction, where Infleqtion fixes errors throughout a calculation rather than only at the end (Stage 4). We are also looking for evidence that adding more protection consistently lowers error rates (Stage 6) and that protected calculations outperform unprotected ones without relying on discarding runs (Stages 5 and 7). Demonstrating these capabilities would provide stronger evidence that Infleqtion can translate its growing logical qubit count into reliable, scalable quantum computing.

Roadmap and Execution

Our view: Infleqtion's execution record is short but encouraging. It has met or exceeded recent logical qubit count targets and successfully delivered a 100 physical qubit system to the UK's National Quantum Computing Centre (NQCC). However, its demonstrated logical qubits have not yet achieved the error correction capabilities envisioned in its earlier roadmap, and its next generation architecture remains unproven. We therefore assign Moderate execution credibility, with the Illinois system planned for 2027 representing the next major test.

Exhibit 57: Infleqtion roadmap



Source: Company Data

Execution Track Record: Infleqtion has worked with neutral atoms since it was founded as ColdQuanta in 2007, but its record against dated computing commitments is relatively short. The company initially planned to launch Hilbert, a cloud based 100 qubit computer, in late 2021, but released it in beta in May 2022. In February 2024, Infleqtion demonstrated a 1,600 atom array and introduced a five year roadmap targeting a fully error corrected computer with 100 logical qubits capable of running calculations more than one million steps deep.

Since then, Infleqtion has increased its demonstrated logical qubit count from two in 2024

to twelve in 2025 and thirty in September 2026, meeting its revised numerical target. However, its original 2026 roadmap called for 10 logical qubits with error correction. While the company has exceeded that count, its larger logical qubit demonstrations have relied primarily on detecting errors and discarding unsuccessful runs, rather than correcting errors continuously during calculations. We therefore view its progress as encouraging, although it has not yet demonstrated the full error correction capabilities originally targeted.

Infleqtion has also demonstrated its ability to deliver systems to customers. It was the first of seven hardware providers in the NQCC testbed program to install a system in July 2024 and delivered the full 100 physical qubit machine in December 2025, meeting the program's year end goal. This provides a more tangible execution record than logical qubit announcements alone.

Looking ahead: The next major commitment is a system under contract for delivery to the Illinois Quantum & Microelectronics Park in 2027. Announced in July 2026, the new modular architecture is designed to scale beyond 1,000 physical qubits and demonstrate more than 50 logical qubits on the path toward 100 in 2028 and more than 1,000 in 2030. Infleqtion describes the Illinois system as fault tolerant, although the exact capabilities that will be demonstrated upon delivery have not been defined. The 2028 target is particularly ambitious given its requirement for fully error corrected calculations exceeding one million steps.

What we are watching: The Illinois delivery is the next major test of execution. We will be looking for on time delivery, progress toward the stated logical qubit capabilities and evidence that the system can correct errors during calculations. Nearer term, publication of the full technical results from the September 2026 demonstration would strengthen confidence in the company's recent progress. Achieving these milestones would support the credibility of its longer term roadmap, while delays or continued reliance on error detection would increase execution risk.

Exhibit 58: Infleqtion: Milestones we are watching

TIMING	TYPE	MILESTONE (AS STATED BY INFLEQTION)	STATUS	STATUS DETAIL
NEAR TERM (2026 TO 2027)				
2026	Roadmap target	30 logical qubits (set in September 2025, replacing an earlier target of 10 logical qubits with error correction in 2026)	Achieved	Announced 24 Sep 2026 (experiments completed in August): 30 logical qubits on 80 physical qubits in an IQP circuit that included four non-Clifford (CCZ) gates. Relies on error detection and discarding flagged runs, with loss correction; not yet peer reviewed
2026	Publication	Full paper on the 30-logical-qubit result, with experimental methods, resource counts and statistical analysis	Pending	Expected in the coming weeks, per the 24 Sep 2026 technical blog. Should show how many runs were kept and how encoded circuits compare with unencoded ones
2026	Facility	UK Quantum Innovation Centre at Oxford Technology Park for research, manufacturing and systems integration	In progress	Announced May 2026 and scheduled to open later in 2026; company says it triples its UK research, production and systems integration capacity
2026 ONWARD	Funding	Up to $100M in proposed U.S. Department of Commerce CHIPS R&D funding (letter of intent)	Pending	Contingent on development milestones and subject to definitive agreements; proposed terms include Commerce receiving Infleqtion common stock. Milestones not disclosed
2027	Customer delivery	Fault-tolerant, NVIDIA NVQLink-integrated Sqale system at the Illinois Quantum & Microelectronics Park, designed to demonstrate more than 50 logical qubits	Pending	Delivery planned for 2027 (announced 22 Jul 2026) and described as under contract; new modular architecture designed to scale beyond 1,000 physical qubits on a path to 100 logical qubits. What will be demonstrated at delivery is not defined
2028 TO 2030				
2028	Roadmap target	100 logical qubits	Pending	First set in February 2024 as a fully error-corrected computer running circuits more than 1 million steps deep; company says it remains on track (Sep 2026). Reaching it depends on repeated error correction, which is not yet demonstrated
2030	Roadmap target	More than 1,000 logical qubits in a full-stack fault-tolerant system	Pending	Added in the September 2025 roadmap update, which also raised the 2026 target to 30; reaffirmed September 2026

Manufacturability

Our view: Infleqtion has demonstrated an ability to build and deliver quantum systems, with hardware operating in the UK and Japan. Its sensing business adds relevant manufacturing experience, although larger quantum computers introduce substantially more complexity. Production capacity is expanding, and scaling laser controls remains a key challenge. The Oxford facility expansion and delivery of the Sqale system to IQMP in Illinois (planned 2027) are the main manufacturing milestones we are watching.

Manufacturing in house: Unlike semiconductor based quantum computing approaches, Infleqtion's systems use individual atoms controlled by lasers inside vacuum chambers rather than fabricated semiconductor qubit chips. The company manufactures its core quantum hardware in house, with facilities in Colorado and Oxford handling component production, system assembly, integration and testing. It also produces optical clocks, RF receivers and inertial sensors using related technology. Its Tiqker optical clock has operated aboard a UK Ministry of Defence submarine, demonstrating experience producing field deployable quantum hardware, although these products are less complex than large quantum computers.

Scaling lasers and optics: The main manufacturing challenge is scaling the optical systems used to trap and control atoms. Infleqtion demonstrated a 1,600 atom array in 2024, showing it can trap atoms at scale, although not perform full quantum calculations across them. As systems grow, they require more laser power, greater control and more complex optical integration. Infleqtion sources some optical components externally and is developing photonic chips to miniaturize laser controls, including a prototype fabricated with Honeywell Aerospace. It also leads a £2.2mn program with the NQCC targeting 10 to 100 times faster quantum gates and greater parallel processing. These efforts could make larger systems easier to manufacture, but their benefits at commercial scale remain unproven. Infleqtion also signed a letter of intent with the U.S. Department of Commerce in May 2026 for up to $100mn in proposed CHIPS funding. The funding would support development of the optical systems and system integration capabilities needed to scale neutral atom quantum computers, although it remains subject to definitive agreements and development milestones.

Expanding production capacity: Infleqtion is building a Quantum Innovation Centre at Oxford Technology Park, expected to triple its UK research, production and systems integration capacity. Construction began in August 2026, with opening planned for later in the year. The expansion should provide additional capacity for larger systems, although the company has not yet demonstrated repeatable production of its next generation modular architecture.

What we are watching: The main manufacturing test is whether Infleqtion can build and deliver its larger modular systems reliably. We will also be watching the Oxford facility ramp, progress in scaling laser controls and evidence that system assembly and testing can become more repeatable as complexity increases. Its existing deployments provide an encouraging starting point, but do not yet establish that the company can manufacture substantially larger quantum computers at scale.

Broader Indicators of Progress

Our view: Infleqtion's nearly two decades of neutral atom research have supported a meaningful patent portfolio and relationships across government, research and industry. Its computing deployments provide tangible external validation, while its sensing business has helped establish commercial and government partnerships. Research activity, public sector programs and hiring suggest continued investment in the platform.

Research and intellectual property: Infleqtion has a steady research record, with approximately 80 quantum related publications in our dataset spanning neutral atom hardware, error correction and quantum software. Its software work, including Superstaq, accounts for a meaningful share of recent output. Its patent portfolio is larger, with more than 200 patent publications covering technologies used across both its computing and sensing businesses.

Government and public sector relationships: Infleqtion has established relationships across the US, UK, Japan and Australia, extending beyond research into hardware deployments and funded development programs. Beyond its computing systems discussed earlier, the company has a DARPA HARQ contract for quantum software and participates in multiple DOE Genesis Mission projects with national laboratories.

Its sensing programs provide additional evidence of government demand. NASA's Quantum Gravity Gradiometer Pathfinder received an additional $20mn award in August 2026, bringing cumulative awards to approximately $40mn. Infleqtion also received $11mn through the Department of Defense's APFIT program for its optical clock technology. These contracts provide funding and external validation while broadening its customer relationships.

Industry partnerships: Infleqtion's most visible computing partnership is with NVIDIA, spanning logical qubit research and GPU accelerated error decoding. In sensing, its partnership with Safran offers access to established defense and commercial timing customers, while an agreement with Voyager Technologies targets deployment of its Tiqker clock aboard the International Space Station.

Hiring and external validation: Active job postings rose sharply from early 2025 to a peak in early 2026 before easing somewhat, consistent with investment in its computing and sensing businesses. Job postings are not a direct measure of headcount, but provide another indicator of investment activity. Infleqtion has not entered DARPA's Quantum Benchmarking Initiative yet, which could offer additional independent evaluation of its computing hardware.

Taken together, Infleqtion's research, government relationships and industry partnerships point to an established position in the quantum ecosystem. These support its ability to attract funding and technical partners, although they do not independently establish that its computing technology can scale commercially.

Exhibit 59: Infleqtion has a steady research record spanning neutral-atom computing, error correction and software...

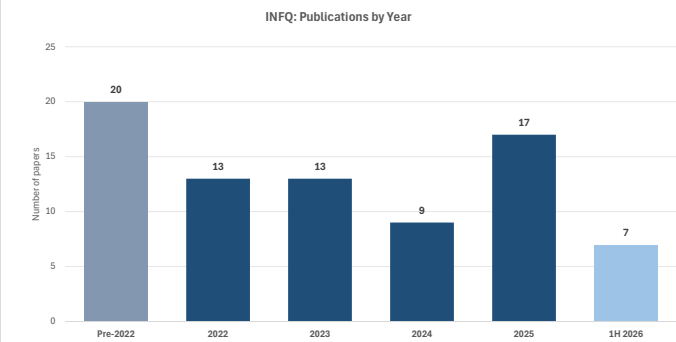


Exhibit 60: ... and patents supporting the hardware behind its computing and sensing businesses

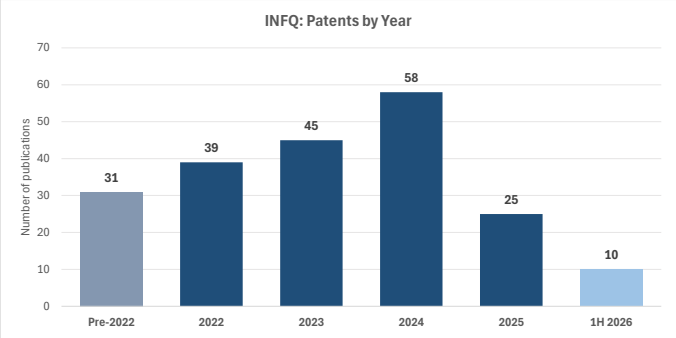
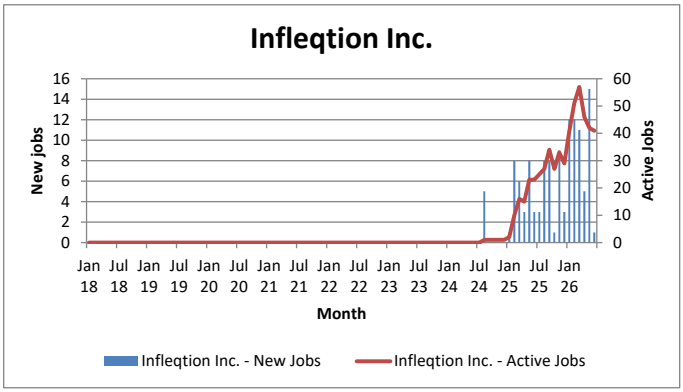


Exhibit 61: Hiring has ramped since 2025, supporting expansion ahead of its planned Illinois system



Revenue Outlook and Funding

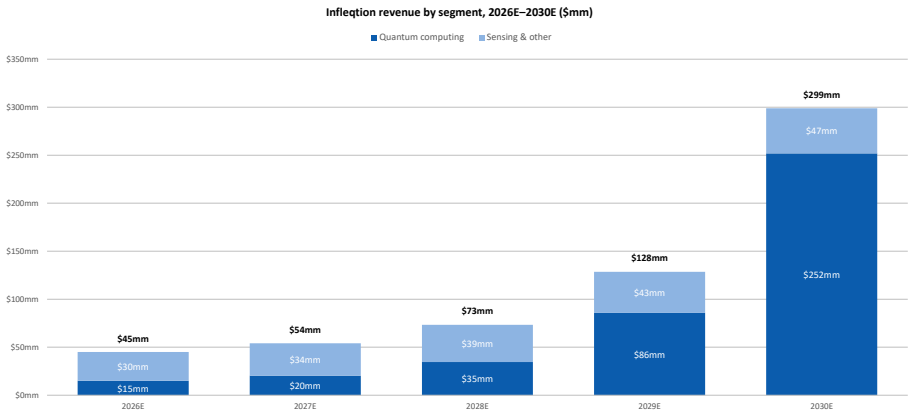
Our view: Infleqtion's quantum computing revenue today comes primarily from government programs, with sensing accounting for approximately two thirds of the total. We expect computing to become the main growth driver from 2028 and beyond as larger systems enter service and generate hardware, cloud and software revenue. Computing revenue rises from approximately $15mn in 2026 to $250mn in 2030 in our model. Its funding position is solid for its stage, with no new equity required through 2030 in our base case, although we assume their cash burn stays below peers.

Where revenue stands today: Infleqtion generated $31mn of revenue in 2025, and we forecast approximately $45mn in 2026, which we estimate was roughly $15mn of computing and $30mn of sensing. Government customers accounted for 89% of first half 2026 revenue. Computing revenue today largely reflects funded development programs rather than recurring system or cloud sales, including work for the UK's NQCC, ARPA-E and DARPA software contracts, and Japan's Shunkai program. Approximately half of our 2027 revenue estimate is supported by existing backlog and named programs, with the balance dependent on additional awards.

Sensing: a near-term revenue base: Infleqtion's sensing business provides commercial revenue while its computing business scales. Its portfolio includes commercially available Tiqker optical atomic clocks, alongside earlier stage quantum RF receivers and inertial and gravitational sensors. Government programs, particularly NASA's gravity sensing project, underpin near term demand, while commercial partnerships could broaden adoption over time. Revenue remains sensitive to the timing of government contracts. We forecast sensing revenue increasing from approximately $30mn in 2026 to $46mn in 2030, an 11% CAGR and roughly 1% of the $4bn to $5bn quantum sensing market estimated by McKinsey.

How the computing ramp builds: We expect computing revenue to accelerate as Infleqtion moves from funded research programs toward larger system sales, cloud access and software. Our forecast rises to $35mn in 2028, $86mn in 2029 and $252mn in 2030, supported by sales of a limited number of high value systems to government and research customers alongside increasing cloud and software revenue. The main uncertainty is the timing of system commercialization and customer adoption, given the technical milestones still required. We expect 2030 to represent the early stages of Infleqtion's commercial computing ramp.

Exhibit 62: INFQ: revenue by segment



Funding and runway: Infleqtion ended 2Q26 with approximately $582mn of cash and investments and no debt, supported by roughly $0.5bn of proceeds from its February 2026 listing. We expect annual free cash outflows of approximately $63mn to $94mn from 2026 through 2029 as the company invests in systems and manufacturing capacity. Burn moderates to approximately $50mn in 2030 as computing revenue grows and the adjusted EBITDA loss narrows to approximately $34mn.

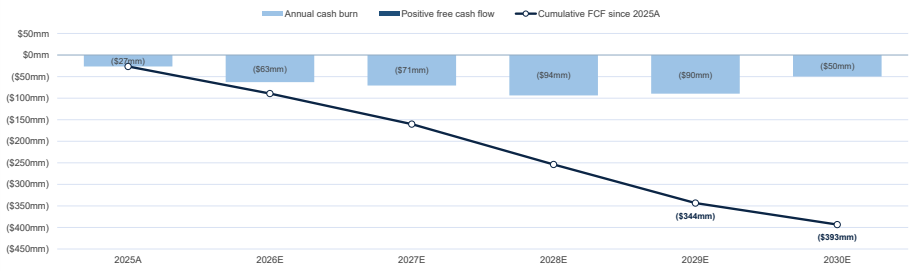
On our estimates, cash and investments bottom at approximately $220mn at the end of 2030, suggesting Infleqtion can fund its existing roadmap without additional equity through the end of the decade. However, the outlook is sensitive to the computing ramp.

Our model excludes the proposed CHIPS funding of up to $100mn, which remains subject to definitive agreements and could involve equity issued to the US government. Additional grants or strategic partnerships could provide funding upside, but are not included in our forecasts.

Exhibit 63: INFQ: funding snapshot

Metric	Value
Cash & investments (2Q26)	~$582mm; no debt
Free cash flow (MSe)	2026E –$63mm; 2027E –$71mm; 2028E –$94mm
Low point in cash & investments (MSe)	~$223mm, end-2030E
Free cash flow breakeven (MSe)	Beyond 2030E

Exhibit 64: INFQ: cash burn



Other Key Players: Big Tech and Private Companies

IBM

IBM is covered by Erik Woodring

Exhibit 65: IBM Quantum

Metric	Value
TECHNOLOGY SPECS	
Modality	Superconducting
Physical qubits (Heron, demonstrated)	156
Logical qubits (Heron, demonstrated, error-corrected)	1
Technology framework stages	Demonstrated: 1-4 | Partial: 5, 7 | Not demonstrated: 6, 8
Error per layered gate (EPLG), 100-qubit chain, Heron	2.15×10^{-3} (0.215%)
QUANTUM-SPECIFIC METRICS	
Quantum investment commitment ($mn, 5 years, June 2026)	>$10,000
Quantum signings since 2017 ($mn)	$1,100
Quantum systems deployed (#)	>90
Client and partner organizations running workloads (#)	>340

Our view: IBM is widely regarded as one of the industry's leaders, with a history in quantum computing going back more than four decades and, by its own count, the industry's largest fleet of quantum computers and its most widely used quantum software. It also has one of the most detailed roadmaps toward fault tolerance, built around qLDPC codes that could sharply cut the number of physical qubits needed for each logical qubit. On our framework, it has demonstrated repeated error correction and large encoded circuits, but not yet that adding protection lowers logical error. Its strongest results also come from today's processors rather than the qLDPC architecture it plans to scale, which makes Kookaburra, the first module designed to combine qLDPC memory with logical processing, the key test.

Where IBM stands on our framework: IBM has demonstrated Stages 1 through 4, with partial evidence at Stages 5 and 7 (Exhibit 66) In 2026 it corrected errors on a single logical qubit for up to 30 rounds, though without clearly outperforming the underlying physical memory or showing that a larger code lowers error. Separately, a 64-qubit encoded circuit substantially improved its accepted results, but only by discarding nearly all runs, with about 0.04% kept. We therefore treat it as a large-scale error-detection result rather than 64 corrected logical qubits.

Exhibit 66: Where IBM stands on our QEC framework

Company	1. Encoding	2. Detection	3. Correction	4. Repeated QEC	5. Breakeven	6. Below threshold	7. Logical operations	8. Scaled FT	Highest stage demonstrated
Quantinuum	Yes	Yes	Yes	Yes	Yes	Partial	Yes	No	Stage 7
Google	Yes	Yes	Yes	Yes	Yes	Yes	Partial	No	Stage 6
Atom Computing	Yes	Yes	Yes	Yes	Partial	Partial	Partial	No	Stage 4
IBM	Yes	Yes	Yes	Yes	Partial	No	Partial	No	Stage 4
IonQ	Yes	Yes	Yes	Yes	Partial	No	Partial	No	Stage 4
Infleqtion	Yes	Yes	Yes	No	Partial	No	Partial	No	Stage 3
Xanadu	Yes	Partial	Partial	No	No	No	No	No	Stage 1

What we're watching and what's different: IBM's longevity sets it apart. It co-hosted the 1981 Physics of Computation conference with MIT, widely seen as the starting point for the field. Its researchers co-invented the first quantum cryptography protocol in 1984 and ran the first demonstration of Shor's factoring algorithm in 2001. In 2016 it became one of the first companies to put a quantum computer on the cloud. Today, IBM says its Qiskit software is used by nearly 70% of quantum developers and that more than 340 client and partner organizations run workloads on its systems. In June 2026 it committed more than $10bn to quantum over five years. Technically, its distinguishing bet is qLDPC codes, which need the long-range connections its Loon chip introduced in 2025. Kookaburra is on the roadmap for 2026 but had not been announced as of late September. We are looking for a qLDPC memory that corrects errors repeatedly without relying on discarded runs and improves as protection increases, followed by protected logical operations ahead of Starling in 2029. We are expecting to hear updates at their upcoming Quantum day on December 10th.

Google

Google is covered by Brian Nowak

Exhibit 67: Google Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Superconducting
Physical qubits per logical qubit (Willow, demonstrated)	101
Logical qubits (Willow, demonstrated, error-corrected)	1
Technology framework stages	Demonstrated: 1-6 | Partial: 7 | Not demonstrated: 8
2Q gate error (Avg., simultaneous, CZ, Willow QEC chip) (%)	0.33%
COMMERCIAL MODEL	
Commercial & deployment model	Selective research access via the Quantum Computing Service (Google Cloud, approved list) and the Willow Early Access Program
Software & developer stack	Cirq: open-source Python framework for near-term (noisy) quantum circuits
Manufacturing infrastructure	Quantum processor fabrication facility in Santa Barbara, California
Long-term fault-tolerant goal	Architectural roadmap targets scaling to 1,000,000 physical qubits

Our view: Google's Willow result is one of the most important milestones in quantum computing to date, and one of the strongest results in our coverage group. It showed that adding more error-correction protection drove errors down every time, the clearest proof yet that error correction can scale, and a goal the field had pursued for nearly 30 years. The next step is extending that performance from storing information to running calculations.

Where Google stands on our framework: Google has demonstrated Stages 1 through 6, with partial evidence at Stage 7. On Willow, it ran three progressively larger protected memories, and logical error fell each time protection increased. The largest used 101 physical qubits for one logical qubit and lasted about 2.4x longer than its best physical qubit. Its Stage 7 evidence is partial. It has run protected single-qubit operations and a protocol linking two logical qubits, and prepared a high-quality magic state. However, it has not shown these operations improving with larger codes, and the best magic state was kept in only about 8% of attempts.

Exhibit 68: Where Google stands on our QEC framework

Company	1. Encoding	2. Detection	3. Correction	4. Repeated QEC	5. Breakeven	6. Below threshold	7. Logical operations	8. Scaled FT	Highest stage demonstrated
Quantinuum	Yes	Yes	Yes	Yes	Yes	Partial	Yes	No	Stage 7
Google	Yes	Yes	Yes	Yes	Yes	Yes	Partial	No	Stage 6
Atom Computing	Yes	Yes	Yes	Yes	Partial	Partial	Partial	No	Stage 4
IBM	Yes	Yes	Yes	Yes	Partial	No	Partial	No	Stage 4
IonQ	Yes	Yes	Yes	Yes	Partial	No	Partial	No	Stage 4
Infleqtion	Yes	Yes	Yes	No	Partial	No	Partial	No	Stage 3
Xanadu	Yes	Partial	Partial	No	No	No	No	No	Stage 1

What we're watching and what's different: Google's error-correction strategy is an evolution of its original superconducting design rather than a new platform, so its results come from the architecture it plans to scale. The trade-off is overhead: one logical qubit currently uses nearly an entire chip. Google does not publish dated intermediate milestones, so the next proof point is its Milestone 3, a long-lived logical qubit using roughly 1,000 physical qubits at a logical error rate near one in a million. After that, we are looking for a protected two-qubit operation that improves with larger codes, run with real-time decoding and feed-forward. In March 2026, Google also launched a neutral-atom hardware program alongside its superconducting work.

Atom Computing, Inc.

Exhibit 69: Atom Computing Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Neutral atom (optically trapped, nuclear-spin qubits)
Physical qubits (flagship AC1000)	1,200+
Logical qubits demonstrated (24 entangled, with Microsoft; 48 atoms; error-detected + loss-corrected)	24
Technology framework stages	Demonstrated: 1–4 | Partial: 5, 6, 7 | Not demonstrated: 8
2Q gate fidelity, AC1000 spec (%)	>99.6%
FINANCIAL HIGHLIGHTS (USD)	
Total disclosed equity raised ($mn)	$190.0
Latest disclosed raise, Series C 2026 ($mn)	$100.0
US Department of Commerce CHIPS LOI ($mn)	$100.0

Note: The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information

Atom Computing is one of the more advanced neutral-atom companies on error correction, but it still needs to show that its logical-qubit progress can scale into a commercially useful system. It has demonstrated 24 entangled logical qubits with

Microsoft (using error-detecting codes) and, in June 2026, reported up to 90 rounds of error correction on a toric code, with logical error lower for the larger code at short cycle counts. It has also moved beyond research-only systems, with one disclosed sale to Denmark's QuNorth. The next step is scale: Atom says economically valuable workloads may require roughly 1,000 logical qubits, vs. 50 targeted for Magne.

Where Atom Computing stands on our framework: Stages 1 through 4, with partial evidence at Stages 5, 6 and 7. Its June 2026 toric-code experiment corrected errors over up to 90 rounds while replacing lost atoms, meeting our Stage 4 test. Its Stage 5 evidence is partial. Atom demonstrated a breakeven-type result using a repetition code that protects against only one type of error, benchmarked against atom survival rather than qubit coherence. Stage 6 is also partial. Across two code sizes, logical error fell by approximately 1.3x on average at four cycles, but the improvement did not persist at longer cycle counts. For Stage 7, Atom demonstrated logical-gate operations with Microsoft in 2024 using error-detecting codes, with post-selection..

What we're watching and what's different: Atom's key attribute is the flexibility of its neutral-atom architecture. Because atoms can be rearranged, it can create connections between qubits that are harder to achieve on fixed-layout chips, which could make some forms of error correction more efficient. The next question is whether that flexibility can translate from protected memory into useful computation at scale. Atom advanced to Stage C of DARPA's Quantum Benchmarking Initiative in October 2026, where government evaluators will test its hardware. Next, we are watching for logical gates under sustained error correction, higher logical-qubit counts and first workloads on the 50-logical-qubit Magne system, targeted to begin running tasks around the start of 2027. We are also watching for additional system sales beyond QuNorth, Atom's only disclosed commercial customer.

Capital Raised (per FactSet): Atom Computing has raised roughly $190mn of disclosed equity, with recent funding focused on larger systems, error-correction software, deployments and the hardware infrastructure needed to scale beyond its current 1,200+ physical-qubit platform.

- **Series B** (2022): raised $60mn.
- **Extension** (2024): raised roughly $10mn at a flat share price.
- **Series C** (2026): announced a $100mn raise
- **Government support:** separately, Atom secured a $100mn Department of Commerce CHIPS letter of intent to support higher qubit counts, error-correction software and control systems, and global deployments

PsiQuantum

Exhibit 70: PsiQuantum Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Photonic
Physical qubits (design target)	~1,000,000
Logical qubits (demonstrated)	0
Technology framework stages	Not demonstrated: 1-8
Component fidelities (Conditional on photon detection; not accounting for loss) (%)	99.98% SPAM / 99.5% visibility / 99.72% chip-to-chip
FINANCIAL HIGHLIGHTS	
Total venture equity raised ($mn)	$1,660
Valuation, post-money Series E Sept 2025 ($bn)	$7.0
Latest disclosed raise, Series E Sept 2025 ($mn)	$1,000

Note: The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information

PsiQuantum stands out for its focus on manufacturing at scale. Rather than starting with smaller systems, the company is working toward a million-qubit photonic quantum computer using chips made in commercial semiconductor fabs. This could make it easier to scale production, but the company has yet to publish results showing a working logical qubit. On our framework, PsiQuantum remains at the earliest stage of QEC and still needs to show that its approach can deliver fault-tolerant performance.

Where PsiQuantum stands on our framework: PsiQuantum has not yet reached any stage of our QEC framework, as it has not published results showing encoded or logical qubits on hardware (Exhibit 9). Its progress so far has been at the component level. In 2025, its Omega chipset demonstrated 99.98% state preparation and measurement fidelity, 99.5% two-photon interference visibility, and 99.72% chip-to-chip fidelity. However, these results only measure photons that are successfully detected and do not account for photon loss, which remains a major challenge for photonic quantum computing.

What we're watching and what's different: PsiQuantum's approach stands out in three ways that could matter for scale. **(1)** Its fusion-based architecture connects smaller photonic components to build larger systems, potentially making it easier to scale without manufacturing increasingly complex chips. **(2)** Its chips are made at GlobalFoundries using 300mm commercial semiconductor processes, with added specialty materials, which could make it easier to scale production. **(3)** Its systems are designed to use industrial cooling plants rather than dilution refrigerators, which could make large-scale deployment easier. The key question is whether these attributes hold once the company addresses photon loss and demonstrates error correction.

Capital Raised (per FactSet): PsiQuantum has raised at least $1.66bn of venture equity, with its Series E (Sep. 2025) valuing the company at $7.0bn post-money. The capital is supporting large-scale prototype development, continued chip and architecture work, and the build-out of its Chicago and Australia sites.

- **Series C** (2020): raised $150mn at a $450mn post-money valuation.
- **Series D** (2021): raised $450mn at a $3.15bn post-money valuation.

- **Series E** (Sep. 2025): raised $1.0bn at a $7.0bn post-money valuation, led by BlackRock, Temasek and Baillie Gifford.
- **Government support:** separate from venture funding, PsiQuantum has also secured an A$940mn (~US$620mn) Australian equity, grant and loan package (per the company), a $125mn DARPA agreement, and a $100mn US Department of Commerce award (finalized Sep. 2026). It has also received a conditional $150mn loan commitment from the US Office of Strategic Capital (Oct. 2026).

Silicon Quantum Computing

Exhibit 71: SQC Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Silicon spin (phosphorus donor atoms)
Physical qubits (demonstrated)	11
Logical qubits (demonstrated)	0
Technology framework stages	Not demonstrated: 1-8
Nuclear 2Q gate fidelity, CZ (%, up to)	99.90%
FINANCIAL HIGHLIGHTS (AUD)	
Total disclosed funding raised since 2017 (A$mn)	~$193.0
Series A raise, July 2023 (A$mn)	$50.4
NRFC investment, 2026 (A$mn; A$20mn Mar. + A$40mn Jun.)	$60.0

Note: The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information

SQC's key attribute is the quality of its physical qubits. Its phosphorus-in-silicon approach has achieved single-qubit fidelities of up to 99.99% and two-qubit fidelity of 99.9% on an 11-qubit processor. SQC also reports that performance held steady or improved as it increased the number of qubits. The key question is whether it can maintain this quality in much larger systems and eventually demonstrate logical qubits.

Where SQC stands on our framework: SQC has not yet demonstrated any stage of our framework (Exhibit 9), since it has not published encoded or logical-qubit results on its own hardware. So far, the evidence is at the physical-qubit level. In 2025, it ran Grover's search algorithm on a four-qubit processor at about 98.9% of the theoretical maximum, followed by an 11-qubit processor published in Nature in December. A separate academic team demonstrated error-detecting logical qubits on a similar phosphorus-in-silicon device in March 2026, but SQC has not yet shown this itself.

What we're watching and what's different: SQC is differentiated in three ways that could matter for performance and commercialization. **(1)** It places individual atoms in silicon with reported accuracy of 0.13 nanometers, giving it precise control over qubit placement, which could help reduce errors. **(2)** It stores quantum information in the nuclear spins of atoms, which are less sensitive to noise and could allow qubits to hold information for longer. **(3)** Unlike many early-stage quantum companies, SQC already sells specialized products, including its Watermelon quantum machine learning processor, to customers

such as Telstra and Australian Defence, providing an early path to revenue. Looking ahead, we are watching for its first logical-qubit demonstration, progress beyond 11 qubits, and advancement from Stage B to Stage C of DARPA's Quantum Benchmarking Initiative.

Capital Raised (per FactSet): SQC has raised roughly A$193mn of disclosed equity since 2017, with recent funding increasingly focused on expanding its atomically precise manufacturing capacity as it works to scale beyond its current 11-qubit gate-based processor.

- **Founding capital** (2017): raised A$83mn from UNSW, the Australian Commonwealth, CBA, Telstra and the NSW Government.
- **Series A** (2023): raised A$50.4mn at a A$133mn post-money valuation
- **NRFC Investment** (2026): Australia's National Reconstruction Fund committed A $60mn across two tranches
- **Non-dilutive support:** separately, SQC has an Australian Defence contract of about A$3.2mn (per AusTender) to deliver a Watermelon quantum machine learning system, and was awarded a A$3.6mn government grant (Critical Technologies Challenge Program) for its energy-forecasting project with Schneider Electric.

Quantum Motion

Exhibit 72: Quantum Motion Factsheet

Metric	Value
TECHNOLOGY SPECS	
Modality	Silicon spin (CMOS / MOS quantum dots)
Integrated Si quantum dot devices on one chip (demonstrated)	1,024
Logical qubits demonstrated	0
Technology framework stages	Not demonstrated: 1-8
Spin-parity readout fidelity, dispersive (Best; 340 µs) (%)	99.92%
FINANCIAL HIGHLIGHTS	
Total equity raised to date ($mn)	$221
Series B raise, Feb 2023 (£mn)	£42.0
Latest disclosed raise, 2026 Series C	$160.0

Note: The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information

Quantum Motion's key differentiator is its manufacturing approach. It uses commercial semiconductor processes to build silicon spin qubits, which could reduce the cost and complexity of manufacturing larger quantum computers. So far, the company has demonstrated 99.92% spin-readout fidelity and characterized an array of 1,024 quantum dot devices. However, it has not yet published a measured two-qubit gate fidelity or demonstrated error correction.

Where Quantum Motion stands on our framework: Quantum Motion has not yet reached any stage of our QEC framework, as it has not published results showing encoded or logical qubits on hardware. Its progress so far has focused on qubit measurement, device testing and integration. It installed a complete quantum computing system at the UK's National Quantum Computing Centre in September 2025 and advanced to Stage B of DARPA's Quantum Benchmarking Initiative, but has yet to demonstrate error correction.

What we're watching and what's different: The key test is whether Quantum Motion can turn strong component-level results into a working multi-qubit system without losing performance as it scales. We are watching for a published two-qubit gate fidelity on its integrated chip, followed by repeated error correction and larger arrays. Further, progress from DARPA Stage B to Stage C would provide another external validation point. Its main differentiation is its stated approach of integrating quantum-dot devices, readout and control electronics on the same foundry-made chip, which could reduce wiring and packaging complexity at scale.

Capital Raised (per FactSet): Quantum Motion has raised about $221mn of equity to date, with its 2026 financing focused on custom silicon development, U.S. expansion, DARPA QBI work, and the infrastructure needed to scale beyond its current small-system demonstrations.

- **Series A** (2020): raised £8mn (~$10mn)
- **Series B** (2023): raised £42mn (~$51mn)
- **Series C** (2026): raised $160.2mn. Co-led by DCVC and Kembara, with £40mn from the British Business Bank
- **Non-dilutive support:** separately, Quantum Motion has received UK government (Innovate UK) awards totaling roughly £10mn since 2020, participates in the NQCC testbed programme, and DARPA's QBI Stage B (amounts undisclosed)