Quantum Computing Architectures: A Plain-English Investor Guide

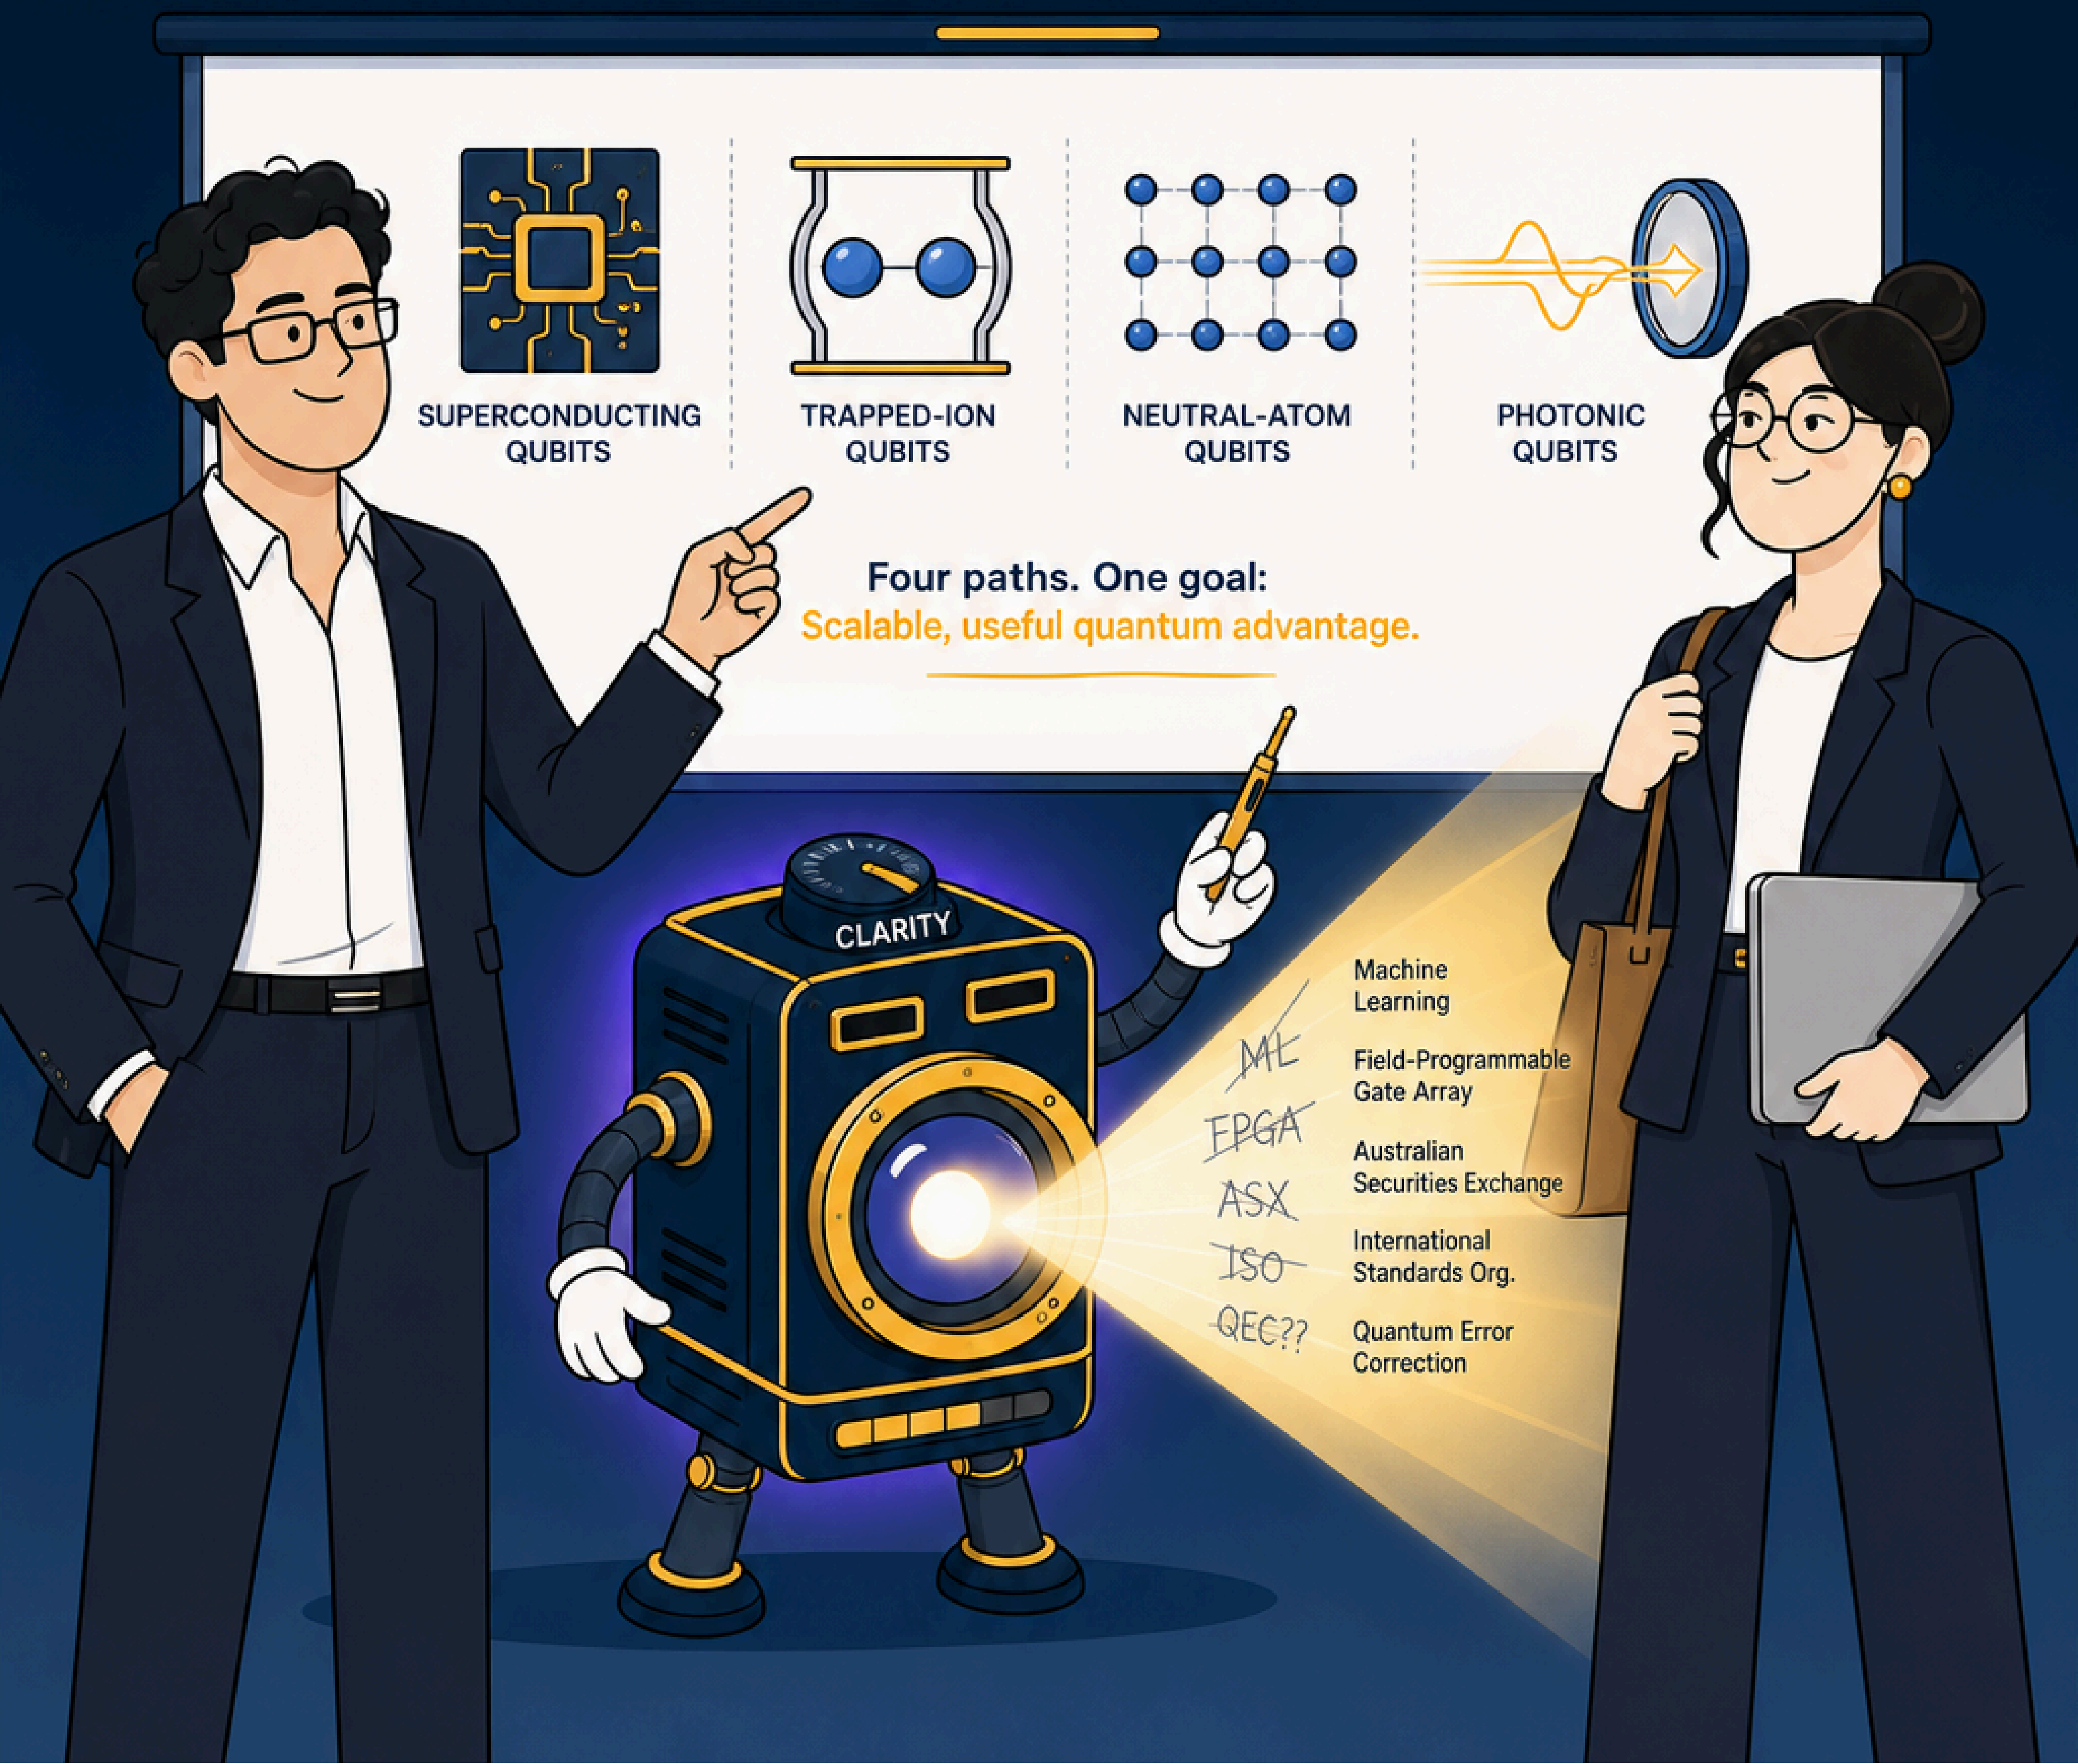

What a qubit actually is (and why the hardware choice matters)

A classical computer bit is either 0 or 1. A qubit can represent both states simultaneously — a property called superposition. Qubits can also become entangled with each other, meaning the state of one qubit correlates with another regardless of distance. These two properties allow quantum computers to explore many solutions in parallel for certain classes of problems.

The catch: quantum states are fragile. Any interference from the environment — heat, vibration, electromagnetic noise — collapses them. This is called decoherence. Every architecture below represents a different engineering bet on how to hold quantum states stable long enough to compute with them.

Superposition

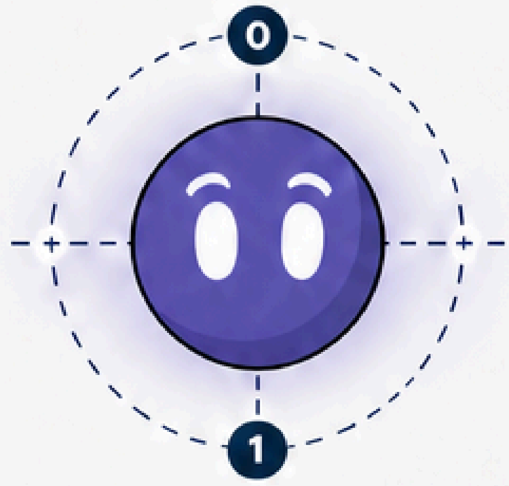

Entanglement

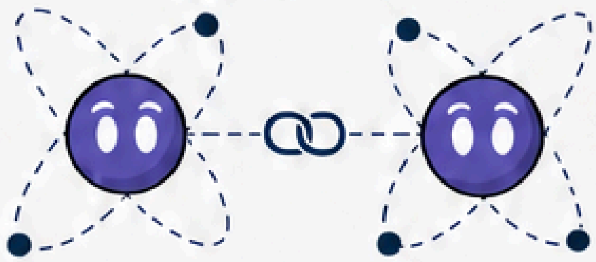

The four metrics every investor should understand before reading further:

Coherence time: how long a qubit holds its quantum state before decoherence destroys it. Longer is better.

Gate fidelity: the accuracy of each computational step. Expressed as a percentage. 99.9% sounds close to perfect; the 0.1% error rate compounds badly across hundreds of operations.

Qubit count and connectivity: how many qubits, and how freely they can interact with each other.

Fault tolerance: the ability to correct errors in real time using redundant "logical qubits" built from multiple "physical qubits." This is the milestone the entire industry is racing toward.

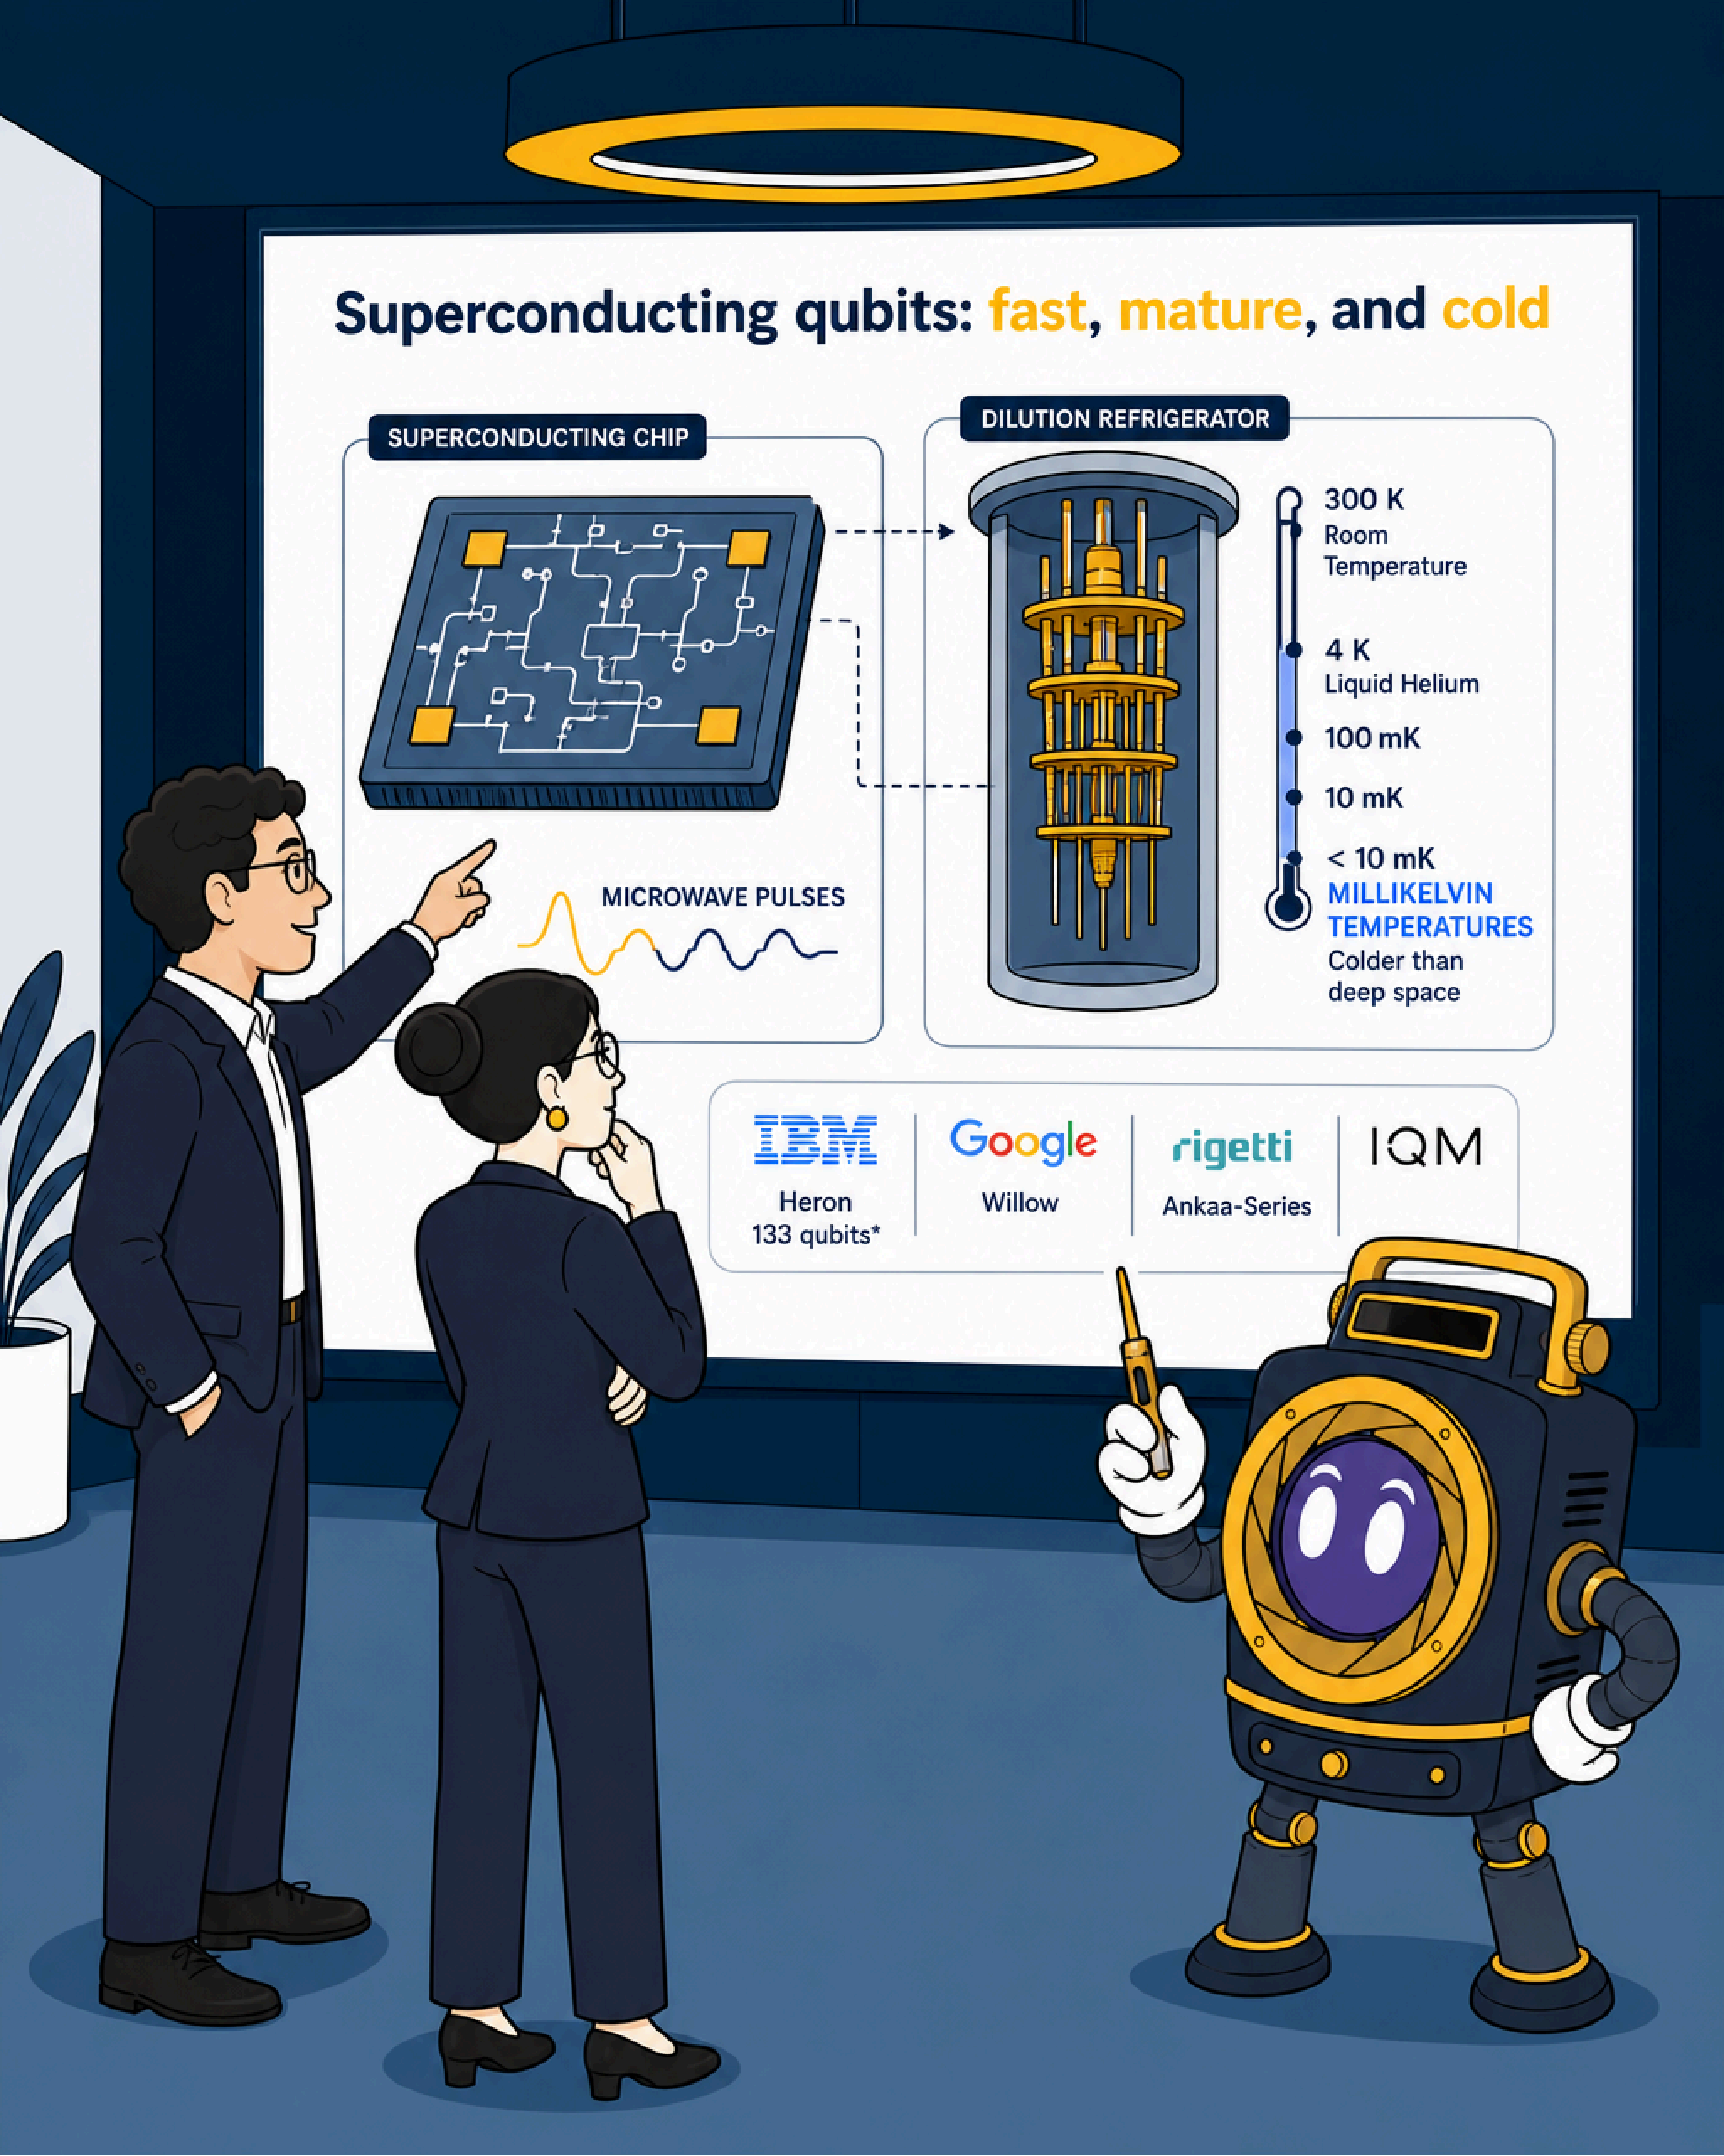

Superconducting qubits: fast, mature, and cold
SUPERCONDUCTING CHIP
DILUTION REFRIGERATOR
300 K
Room Temperature
4 K
Liquid Helium
100 mK
10 mK
< 10 mK
MILLIKELVIN TEMPERATURES
Colder than deep space
MICROWAVE PULSES
IBM
Heron
133 qubits*
Google
Willow
rigetti
Ankaa-Series
IQM

Trapped-ion qubits: high fidelity, slow gates

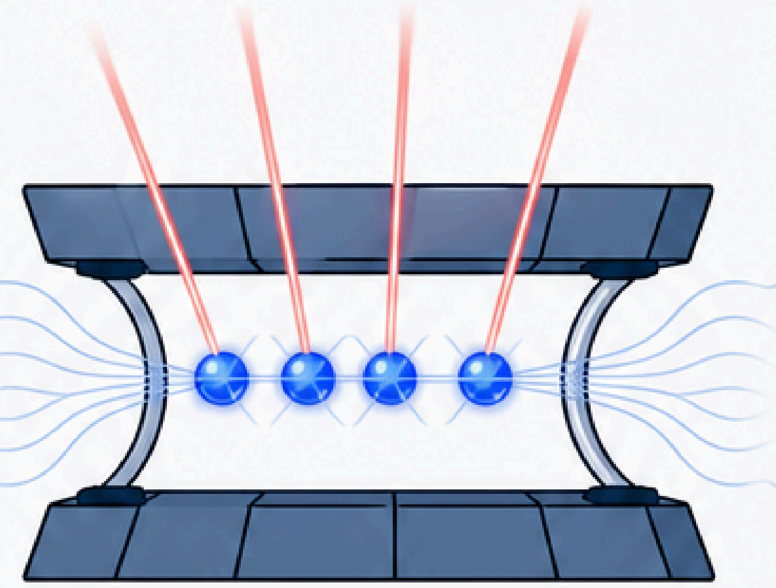

How it works:

Individual ions — charged atoms, typically ytterbium or barium — are suspended in a vacuum using electro-magnetic fields. Laser beams cool them to near stillness and perform gate operations by carefully nudging the ions' energy states. Because the ions are identical by nature (every ytterbium atom is exactly the same), you start with a higher-quality qubit than anything you'd manufacture on a chip.

Who's doing it:

IonQ and Quantinuum (a Honeywell spin-out) are the two dominant companies. Both are publicly or semi-publicly listed and investable.

- IonQ has published two-qubit gate fidelity results approaching 99.9% or above in recent demonstrations; verify current figures and the specific technologies involved from IonQ's official disclosures and published research.
- Quantinuum's H-series processors have similarly reported leading fidelity numbers; confirm current benchmarks from Quantinuum's published H-series documentation.

The strengths:

Trapped-ion systems hold world records for gate fidelity. That matters for error correction: the higher the native fidelity, the fewer physical qubits you need to build each error-corrected logical qubit. Quantinuum has published logical qubit demonstrations using its H-series hardware; refer to their published research for verified logical qubit counts and the specific error-correction codes employed.

Coherence times are orders of magnitude longer than superconducting systems, measured in seconds rather than microseconds.

All-to-all connectivity is another practical advantage: any ion in the trap can interact with any other, which simplifies algorithm design considerably.

The limitations:

Gate speed is the trade-off. Laser-based operations are slow compared to microwave-driven superconducting gates — the difference is roughly a factor of 1,000 to 10,000 in raw speed.

The laser control infrastructure required to address individual ions precisely is complex and hard to miniaturize.

Scaling qubit count beyond a few hundred ions in a single trap requires either larger traps (which introduce noise) or modular networking approaches where multiple traps communicate with each other, which adds latency.

The investor-relevant question for trapped-ion companies:

What's the scaling strategy beyond a few hundred qubits, and how does ion-trap networking affect the effective computational speed?

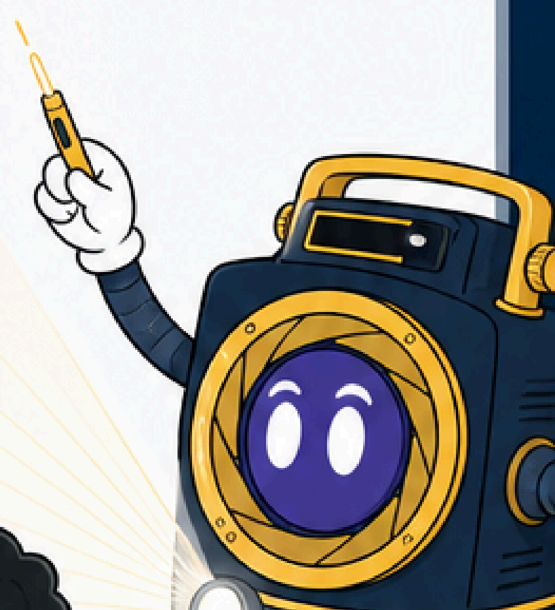

Neutral-atom qubits: reconfigurable, scalable, and catching up fast

How it works: Neutral atoms (uncharged, unlike trapped ions) are loaded into an array of individual optical tweezers — highly focused laser beams that hold each atom in place like a microscopic forceps. The atoms can be rearranged mid-computation. Gates are executed using a phenomenon called Rydberg excitation, where a laser briefly promotes an atom to a high-energy state that interacts strongly with its neighbors.

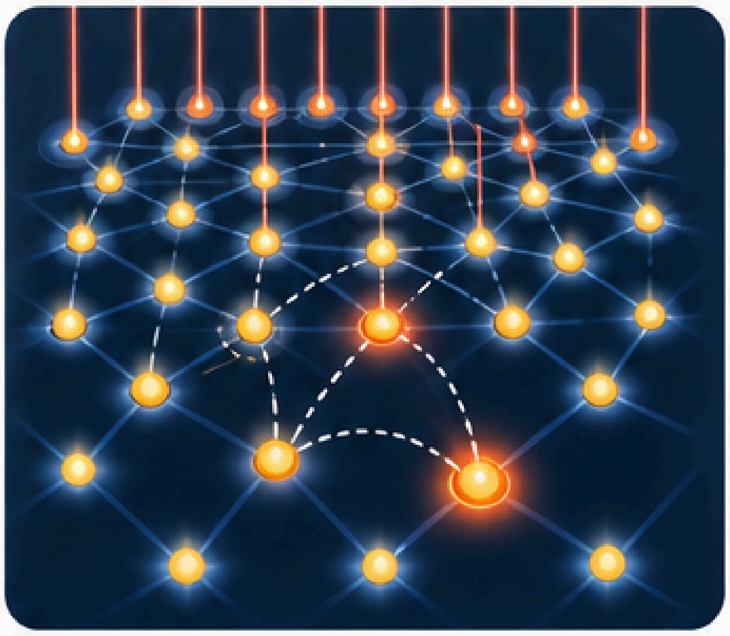

Who's doing it: QuEra (Harvard spin-out), Atom Computing, Pasqal, and Infleqtion are the primary players. Atom Computing's Phoenix system has been documented at approximately **1,180** atomic sites. QuEra published a logical qubit demonstration in 2024; verify the current state of their logical qubit results from their published research for the most up-to-date confirmed figures.

The strengths:

- **Qubit count scales well:** published research has noted that neutral-atom systems show promise in achieving high physical qubit counts relative to other architectures.
- **Reconfigurability is a practical advantage.** Because optical tweezers can be repositioned, the qubit connectivity graph can be dynamically restructured between computation steps. Two atoms can be brought next to each other regardless of where they started in the array — a flexibility fixed-chip architectures don't have.
- **No deep cryogenics required:** the atoms, once trapped, slow to a point that corresponds to a very low effective temperature, but the surrounding apparatus doesn't need to be cooled to millikelvin levels. That reduces operational cost and infrastructure complexity.
- **Mid-circuit measurement has been demonstrated:** the ability to measure some qubits partway through a computation and act on the result without disturbing the rest. This is a prerequisite for many fault-tolerant error correction protocols and is harder to implement cleanly on fixed-chip architectures.

The limitations:

- Gate fidelity is currently lower than trapped ions — in the 99.5%–99.73% range — which means more physical qubits are required per logical qubit.
- Gate operations are also slower than superconducting systems by a factor of roughly 100 to 1,000.
- Loading atoms into tweezers is a stochastic process, and atom loss during computation is a practical challenge that requires active correction.

The investor-relevant question for neutral-atom companies:

Are fidelity numbers improving fast enough to close the error-correction overhead gap with trapped ions, and how does the qubit-loss problem affect circuit depth in practice?

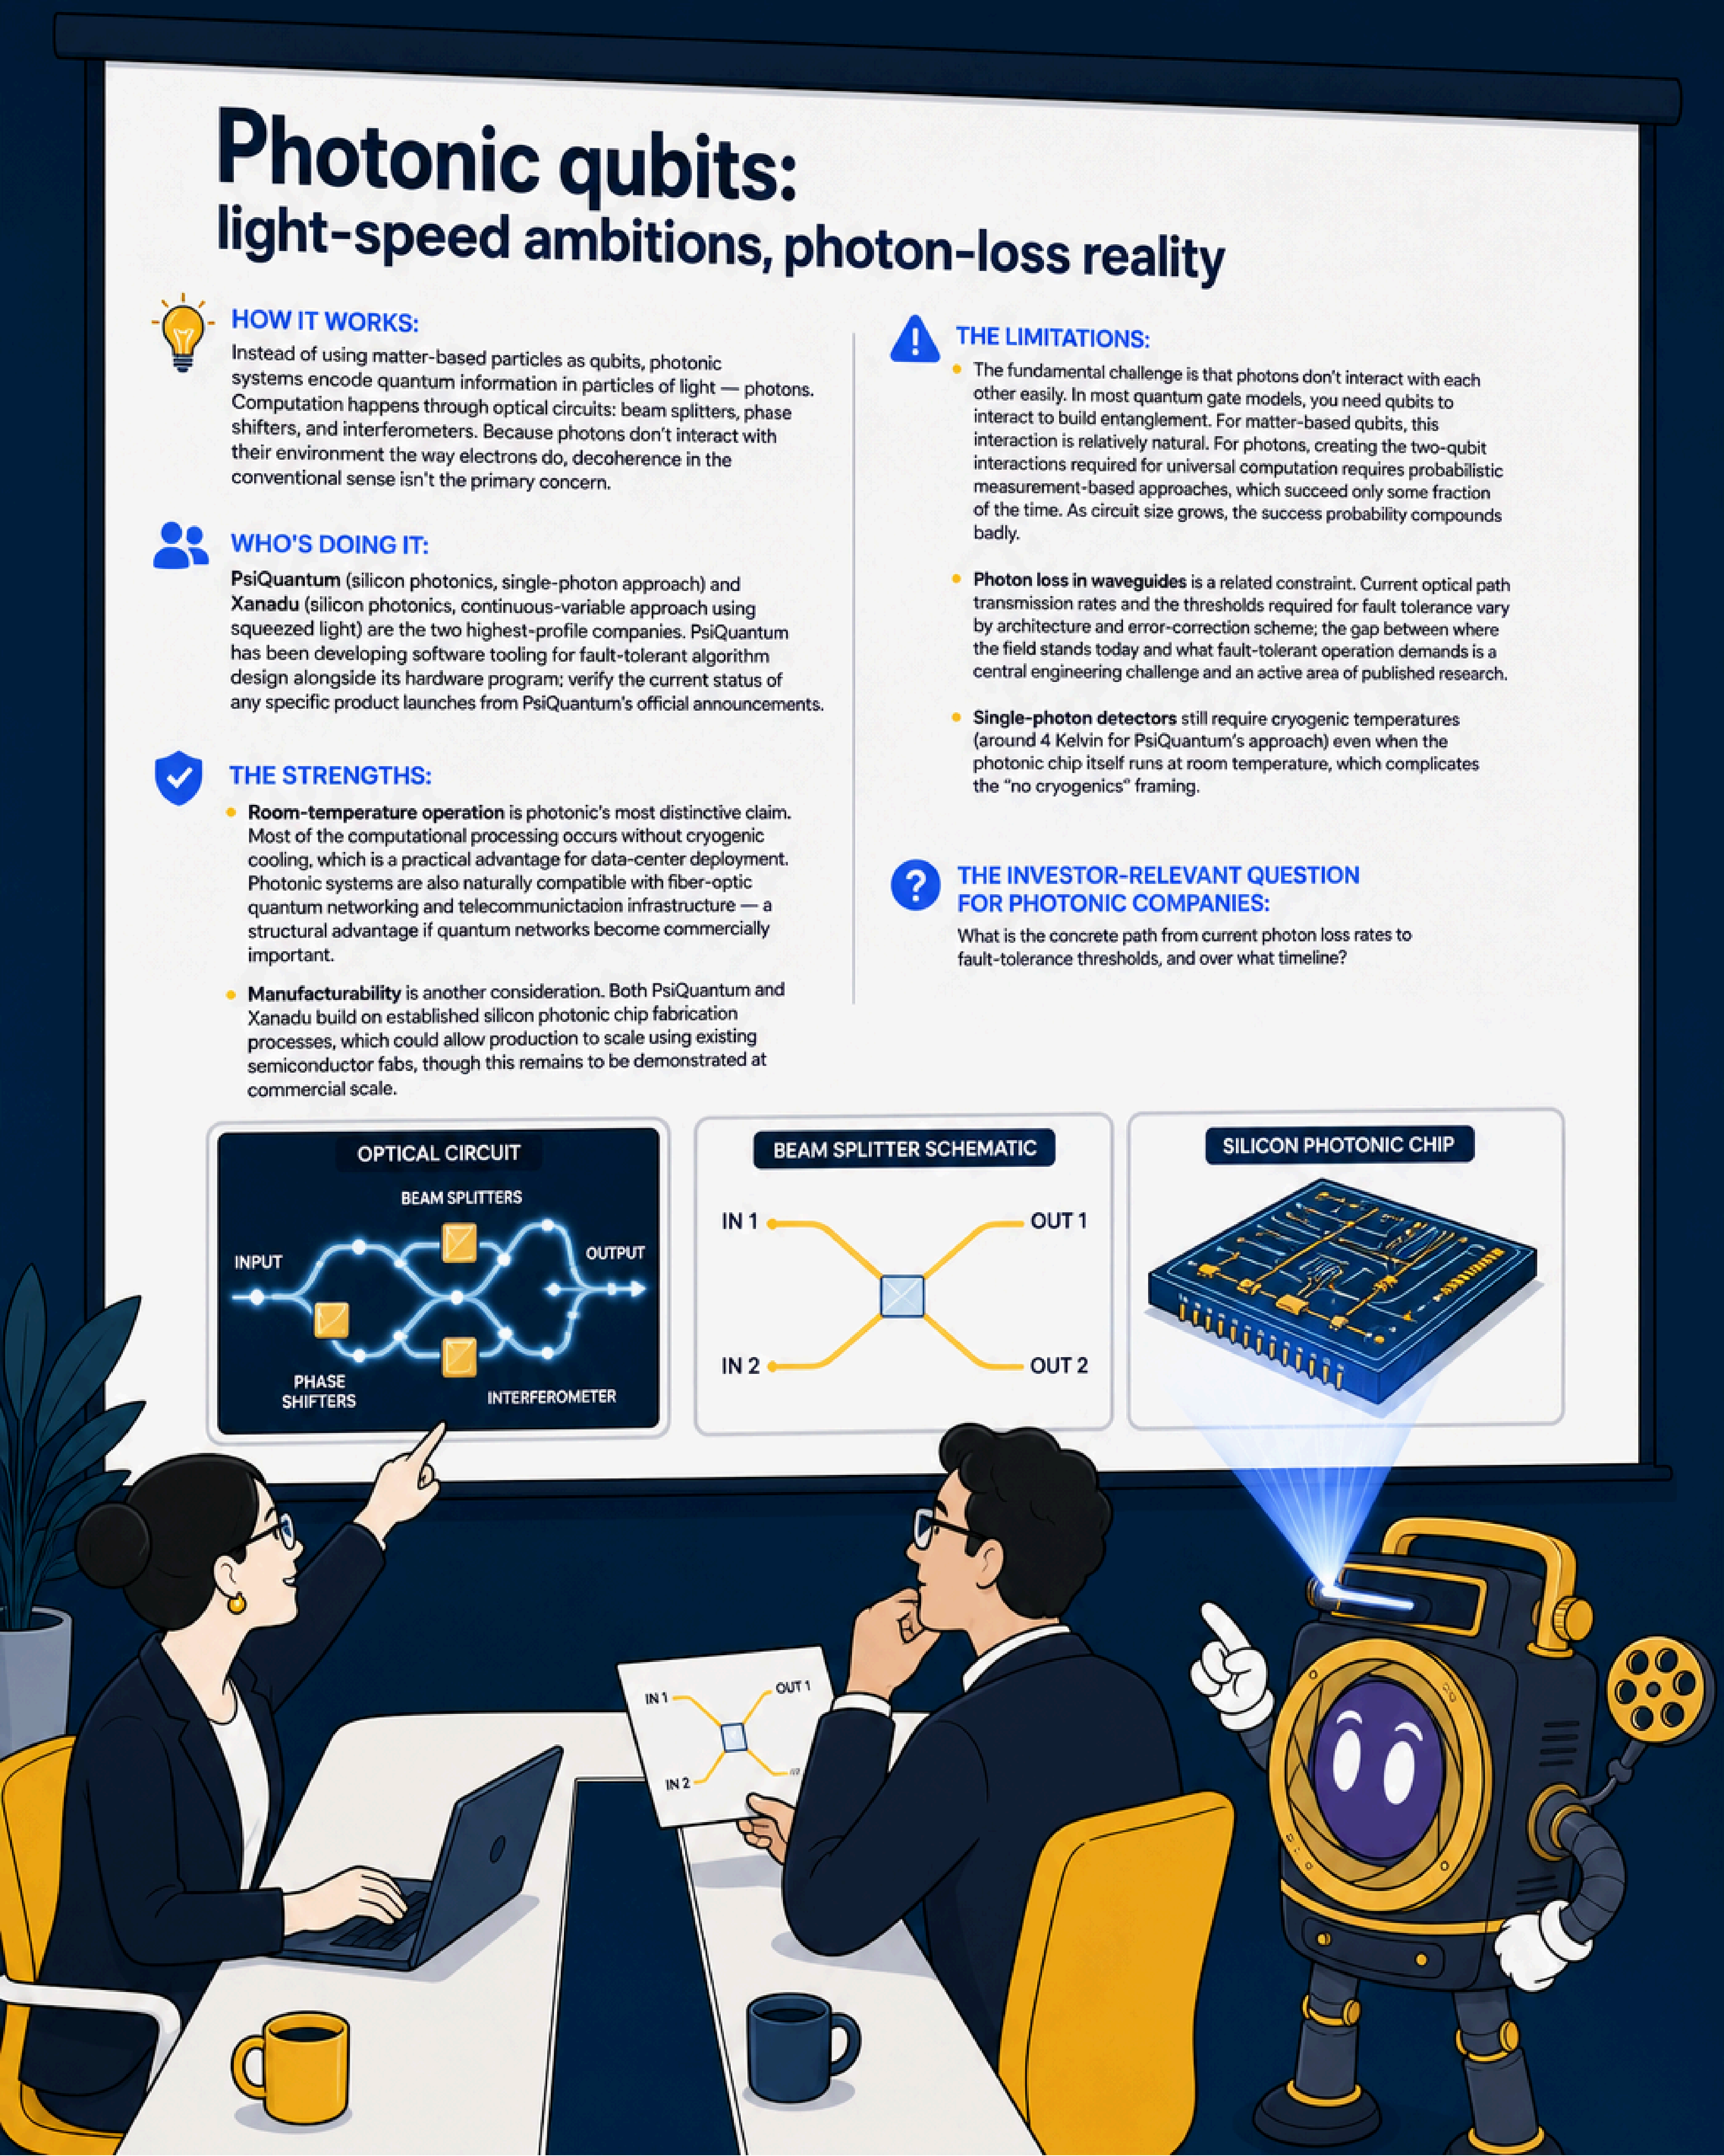

Photonic qubits:
light-speed ambitions, photon-loss reality
HOW IT WORKS:
Instead of using matter-based particles as qubits, photonic systems encode quantum information in particles of light — photons. Computation happens through optical circuits: beam splitters, phase shifters, and interferometers. Because photons don't interact with their environment the way electrons do, decoherence in the conventional sense isn't the primary concern.
WHO'S DOING IT:
PsiQuantum (silicon photonics, single-photon approach) and Xanadu (silicon photonics, continuous-variable approach using squeezed light) are the two highest-profile companies. PsiQuantum has been developing software tooling for fault-tolerant algorithm design alongside its hardware program; verify the current status of any specific product launches from PsiQuantum's official announcements.
THE STRENGTHS:
Room-temperature operation is photonic's most distinctive claim. Most of the computational processing occurs without cryogenic cooling, which is a practical advantage for data-center deployment. Photonic systems are also naturally compatible with fiber-optic quantum networking and telecommunictaoion infrastructure — a structural advantage if quantum networks become commercially important.
Manufacturability is another consideration. Both PsiQuantum and Xanadu build on established silicon photonic chip fabrication processes, which could allow production to scale using existing semiconductor fabs, though this remains to be demonstrated at commercial scale.
THE LIMITATIONS:
The fundamental challenge is that photons don't interact with each other easily. In most quantum gate models, you need qubits to interact to build entanglement. For matter-based qubits, this interaction is relatively natural. For photons, creating the two-qubit interactions required for universal computation requires probabilistic measurement-based approaches, which succeed only some fraction of the time. As circuit size grows, the success probability compounds badly.
Photon loss in waveguides is a related constraint. Current optical path transmission rates and the thresholds required for fault tolerance vary by architecture and error-correction scheme; the gap between where the field stands today and what fault-tolerant operation demands is a central engineering challenge and an active area of published research.
Single-photon detectors still require cryogenic temperatures (around 4 Kelvin for PsiQuantum's approach) even when the photonic chip itself runs at room temperature, which complicates the "no cryogenics" framing.
THE INVESTOR-RELEVANT QUESTION FOR PHOTONIC COMPANIES:
What is the concrete path from current photon loss rates to fault-tolerance thresholds, and over what timeline?
OPTICAL CIRCUIT
BEAM SPLITTERS
INPUT
OUTPUT
PHASE SHIFTERS
INTERFEROMETER
BEAM SPLITTER SCHEMATIC
IN 1
OUT 1
IN 2
OUT 2
SILICON PHOTONIC CHIP
IN 1
OUT 1
IN 2

The comparative framework: what matters for your investment decision

Here's how the four architectures sit across the dimensions that actually drive commercial timelines.

Dimension	Superconducting	Trapped-ion	Neutral-atom	Photonic
Coherence time	Microseconds	Seconds	Seconds	Not decoherence-limited (photon loss instead)
Gate fidelity (2-qubit)	~99%–99.85% (varies by system; verify current benchmarks)	Among the highest demonstrated; verify current H-series and IonQ figures	99.5%–99.73%	Architecture-dependent; loss-limited
Gate speed	Nanoseconds (fastest)	Milliseconds (slowest)	~100x slower than SC	Fast clock speed potential
Qubit count (2026)	100s of physical qubits across production systems; IBM Condor reached 1,121 as a research demonstration (not a production system)	Tens of physical qubits in leading systems; logical qubit demonstrations published by Quantinuum	Hundreds to ~1,180+ physical sites (Atom Computing Phoenix); logical qubit results published by QuEra	Pre-commercial (large-scale); software stack progressing
Operating environment	Millikelvin cryogenics	Laser/vacuum lab	Laser/vacuum; no deep cryogenics	Room-temperature processing; detectors at ~4K
Connectivity	Fixed chip topology	All-to-all	Dynamically reconfigurable	Flexible optical routing
Fault-[illegible] timeline	Targeting 2029 (Google)	Targeting 2027–2030 (Quantinuum)	Active logical qubit demonstrations underway (QuEra, 2024–2025)	[illegible]

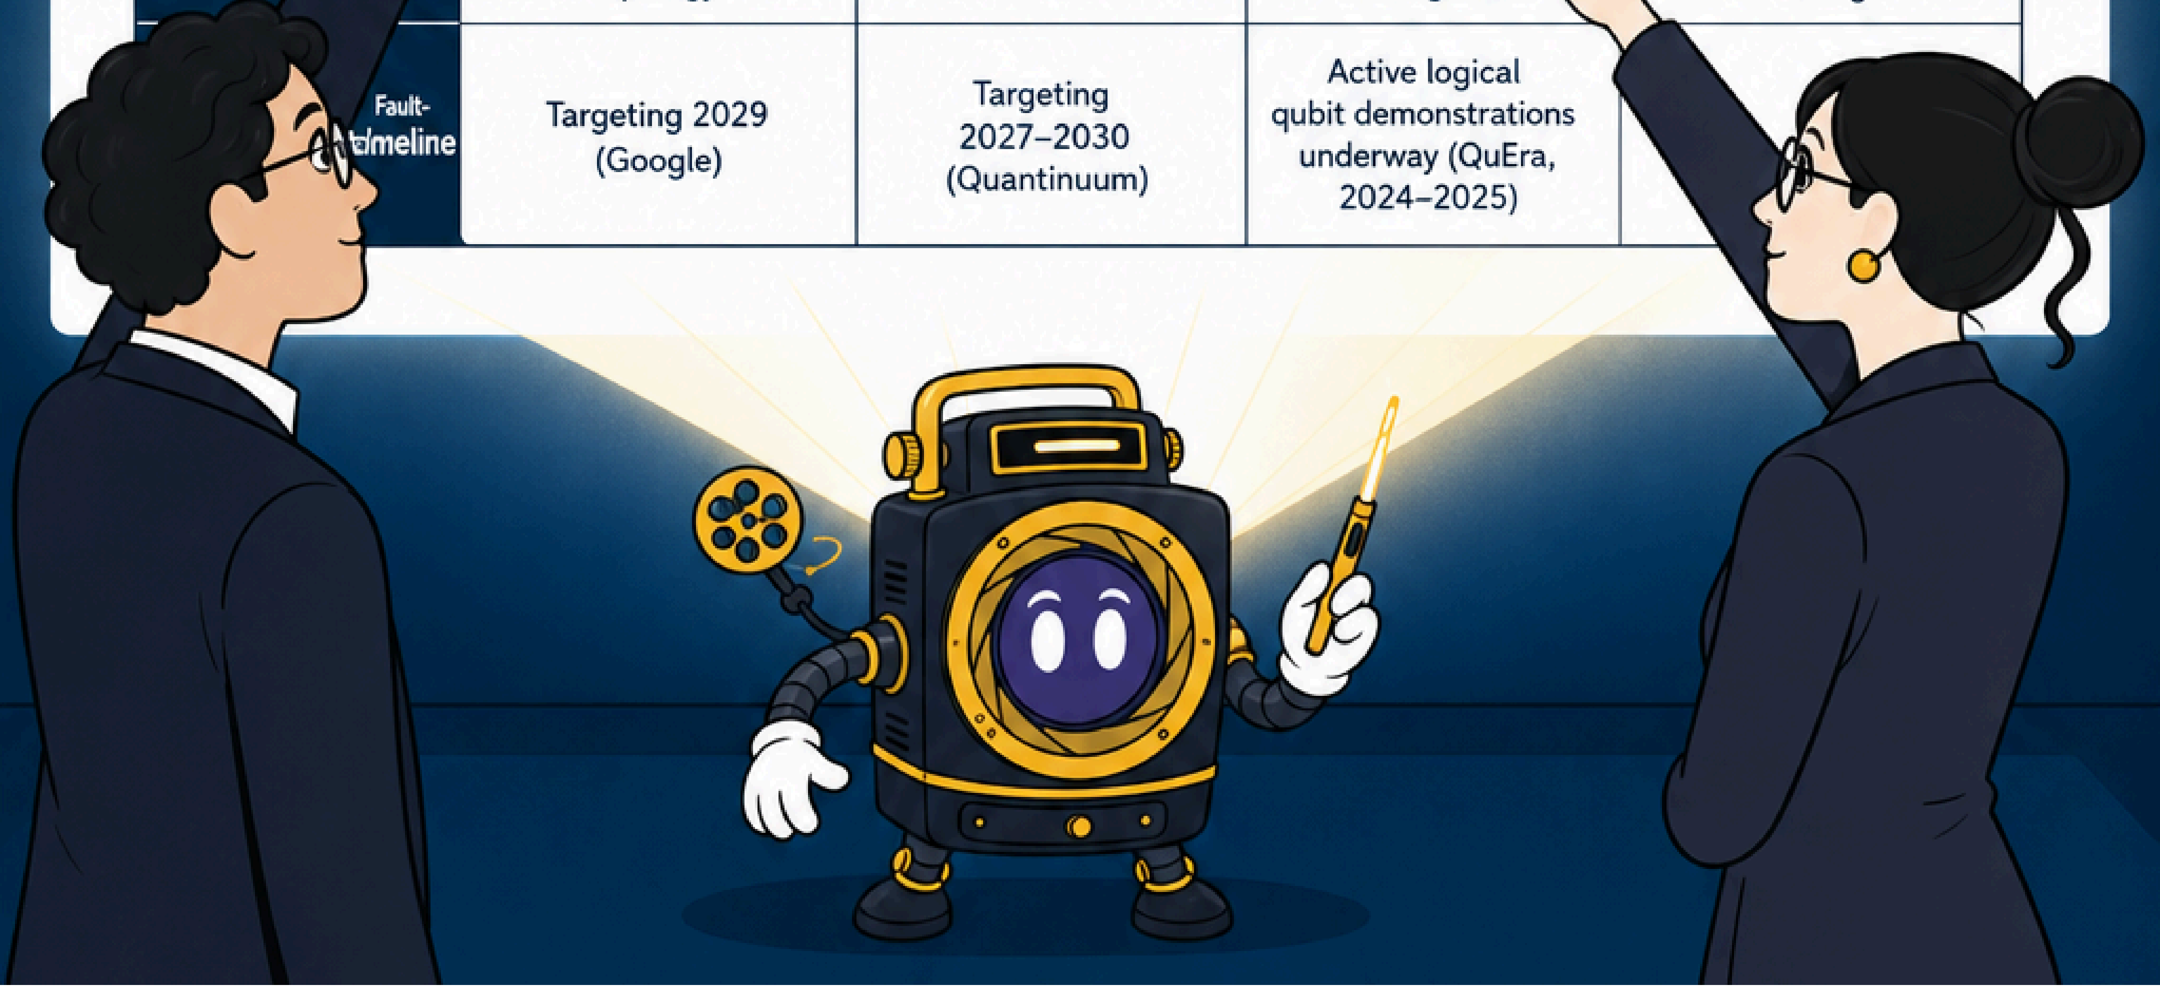

What this means in plain language
SUPERCONDUCTING
Most mature architecture with the broadest ecosystem.
Fighting coherence time and crosstalk at scale.
IBM and
tooling than anyone else.
This matters for near-term utility.
TRAPPED-ION
Offers the cleanest qubits.
Best-demonstrated error correction ratios.
Quantinuum's logical-qubit results are among the most impressive on a per-physical-qubit basis.
The scaling and speed constraints are real...
...but the fidelity numbers are compelling.
NEUTRAL ATOMS
Making fast progress on logical qubit count.
Structural advantages in reconfigurability and operating complexity.
Among the architectures producing the highest verified logical qubit numbers.
The fidelity gap with trapped ions is real but narrowing.
PHOTONIC
The long-range bet.
Compatibility with networking infrastructure and the room-temperature manufacturing story are differentiated.
But the loss rate problem and the probabilistic gate challenge are not minor engineering obstacles — they require fundamental breakthroughs.
PsiQuantum and Xanadu are building the software infrastructure in parallel, which is sensible...
...but investors should be clear-
about where the hardware stands today.

Three investor principles worth writing down
1
QUBIT COUNT IS NOT THE METRIC
≠
1,000 PHYSICAL QUBITS
48 LOGICAL QUBITS
Logical qubits — error-corrected, fault-tolerant qubits — are what algorithms actually run on.
Press releases often conflate the two.
2
ARCHITECTURE CHOICE IS A LONG-TERM COMMITMENT
TRAPPED IONS
NEUTRAL ATOMS
SUPERCONDUCTING QUBITS
Each approach has different scaling economics, different talent requirements, and different hardware supply chains.
A company that has built expertise in trapped-ion systems is not pivoting to neutral atoms next year.
Evaluate the team's depth in their chosen modality, not just the qubit count headline.
3
THE "QUANTUM UTILITY" ERA IS HERE NOW
TODAY
3–7 YEARS
The field is no longer purely speculative — several systems are producing results that classical computers find genuinely difficult.
But broad commercial fault-tolerance is still 3–7 years out across all architectures.
Investments made today are bets on which engineering approaches solve the remaining problems first, not bets on deployed commercial revenue in the next 12 months.
INVESTOR PRINCIPLES
①
②
③

Explaining quantum complexity is exactly the problem we solve

If you're an IR lead, a founder preparing for a Series A, or a CMO trying to explain a quantum product to enterprise buyers, the challenge isn't understanding the technology. It's compressing it into a narrative that lands in **60 seconds** without losing the technical truth.

That's what we do at Infrairis.
We work with deep tech companies across ANZ to turn complex products into clear, professional explainer videos in 2–3 weeks — directed by people who have shipped tech themselves.

→ If your quantum pitch, investor update, or product explainer isn't landing the way it should, **start a conversation with us at startups.infrairis.com.**

infrairis

https://startups.infrairis.com