

# Rydberg Blockade, Optical Tweezers, and the Long Road to Fault Tolerance: What Neutral Atom Quantum Computing Actually Means for Drug Discovery, Health Data Encryption, and Medical Sensing

JUL 13, 2026



Share



## Part I Podcast free on Spotify.



Thoughts on Healthcare Markets & Technology



### **Part I: Rydberg Blockade, Optical Tweezers, and the Long Road to Fault Tolerance: What Neutral Atom Quantum Computing Actually Means for Drug Discovery, Health Data Encryption, and Medical Sensing**

Neutral atom quantum computers are raising billions. Here is what health systems and pharma boards are not being told about what they can actually do...

▶ Listen now

9 minutes ago · Thoughts on Healthcare



## Part II Podcast episode for paid subscribers only. Also available on Spotify.



Thoughts on Healthcare Markets & Technology



### **Part II: Rydberg Blockade, Optical Tweezers, and the Long Road to Fault Tolerance: What Neutral Atom Quantum Computing Actually Means for Drug**

## Discovery, Health Data Encryption, and Medical Sensing

Neutral atom quantum computers are raising billions. Here is what health systems and pharma boards are not being told about what they can actually do...

▶ Listen now

7 minutes ago · Thoughts on Healthcare

To listen to paid episodes in Apple or Spotify, link settings on those platforms (instructions inside the Podcast).

## Abstract

- Neutral atom quantum computers trap ionized strontium atoms in focused laser beams and add an electron to a huge orbital (principal quantum number 40). The atom swells to a couple of microns across. The quantum jumps by roughly 11 orders of magnitude. The quantum blockade, is the entire gate mechanism.
- The platform's selling points: no fab required, thousands of identical qubits in a single vacuum chamber, and physical atom movement that maps beautifully onto the high-rate qLDPC error correcting codes that everyone now agrees are the path to fault tolerance. Its weaknesses: slow clock speed, atom loss, and a fidelity record now pushing toward four nines while trapped ions sit a bit higher.
- The healthcare relevance splits cleanly into three buckets, and most of what gets pitched to hospital systems and pharma boards falls into none of them.
- Bucket one (real, slow): quantum chemistry for strongly correlated systems, mostly metalloenzymes and transition metal catalysis. Resource estimates for something like the nitrogenase active site or a cytochrome P450 oxidation still land in the neighborhood of a million physical qubits and days of runtime. That is a 2030s problem at best, and error correction overhead, not qubit count, is the variable that moves the date.
- Bucket two (real, boring, urgent now): post-quantum cryptography migration for protected health information and genomic data. Harvest now, decrypt later is a live threat model for data with a 100 year identifiability half-life. NIST finalized the first set of post-quantum cryptography standards in 2024. Health IT has not started.



## Discover more from Thoughts on Healthcare Markets & Technology

Expert analysis of healthcare and life sciences markets, technology, investment, entrepreneurship policy, and AI — for investors, entrepreneurs,...

Over 4,000 subscribers

Enter your email...

Subscribe

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).

Already have an account? [Sign in](#)

- Bucket three (real, nearest term, most investable): the sensing spinoffs. Optically pumped magnetometers, Rydberg RF electrometry, and cold atom clocks are already becoming medtech. Wearable MEG and magnetocardiography are the first places a patient will ever touch this physics, and the bottleneck there is reimbursement, not decoherence.
- Bucket zero (fake): quantum machine learning on claims data, quantum optimization of formularies and provider networks, “quantum AI” anything. The data loading problem kills it. Details below.

Thoughts on Healthcare Markets &  
Technology is a reader-supported  
publication. To receive new posts and  
support my work, consider becoming a  
free or paid subscriber.

## Table of Contents

The minus sign that runs the whole show

Why atoms that ignore each other turned out to be worth a billion dollars

Optical tweezers, coin flips, and the atom conveyor belt

Gates, and the slow climb from horrible to three nines

The error correction tax, and why it decides everything healthcare cares about

The drug discovery pitch, and the small part of it that is true

The input problem, or why quantum claims analytics is a category error

Harvest now, decrypt later, and the most boring important thing in health IT

The cold atom tech that will touch a patient first

Twenty companies, one talent pool, and how to read the tape

A scoreboard worth keeping

## The minus sign that runs the whole show

Start with the weirdest fact in the whole business, because it is also the load-bearing one. Take a spin one-half object, rotate it a full 360 degrees, and it does not come back

to where it started. It picks up a minus sign. Go around twice, 720 degrees, and only then are things back to normal. Feynman wrote a whole charming little book trying to make this intuitive and the honest summary of that book is that it has something to do with particle exchange statistics and the difference between bosons and fermions, and also that nobody is fully satisfied.

That minus sign is not a curiosity. It is the mechanism. Every entangling gate in a neutral atom quantum computer is, at bottom, a scheme for making one atom pick up that minus sign only if the other atom is in a particular state. Conditional phase. That is all a controlled-Z gate is, and a controlled-Z plus single qubit rotations gets a CNOT, and CNOT plus single qubit gates gets universal computation, and universal computation plus about eight more years of engineering gets a machine that might tell a chemist something about an iron-sulfur cluster.

The original recipe, from a theory paper in 2000 that kicked off the field, was called  $\pi$ ,  $2\pi$ ,  $\pi$ . Encode the qubit in two stable ground states of the atom, call them zero and one. Tune a laser so it couples the one state, and only the one state, to a very high-lying excited orbital. Hit atom A with a  $\pi$  pulse. If A was in one, it goes up to the Rydberg state. Then hit atom B with a full  $2\pi$  pulse, up and back down. If A did not go up, B completes its round trip and collects the minus sign. If A did go up, A's presence up there blocks B from being excited at all, B does nothing, and no minus sign. Then bring A back down. The phase of the whole wavefunction now depends on the state of atom A. Conditional dynamics. Entanglement. Done.

The blocking part is called the Rydberg blockade, named by analogy to Coulomb blockade in quantum dots, where you cannot cram another charge onto a small island because the potential will not let you. Same idea, different mechanism. Two atoms cannot both sit in the same high-lying state at the same energy because they interact so strongly with each other that the second excitation is pushed off resonance. Some people call it dipole blockade. Both names are fine and physicists will fight about it anyway.

## **Why atoms that ignore each other turned out to be worth a billion dollars**

Here is the structural advantage, and it is worth understanding because it is the entire investment thesis.

Trapped ions were first to everything. First CNOT, first entanglement, and still the best fidelity in the business at roughly four nines. Great. But ions are charged, and charges talk to each other whether you want them to or not. That is how you entangle

them, so it is a feature, until you have a lot of them, at which point the conversation gets loud. Dozens of normal modes of collective motion, all coupling, all needing control. Scaling is genuinely hard.

Neutral atoms are neutral. In their ground state they essentially ignore each other. Put two of them five microns apart and the interaction strength, expressed in frequency units, is millihertz. Effectively zero. You can pack thousands of them in a small volume and they sit there like polite strangers on a train, which makes for an excellent quantum memory and an absolutely useless computer, because computation requires them to talk.

The Rydberg state is the switch. Promote the outer electron to principal quantum number  $n$  and the atom's physical size scales like  $n$  squared times the Bohr radius. At  $n$  equal to 100 that is ten thousand Bohr radii, a couple of microns across, which is a genuinely enormous atom. And the interaction strength between two such atoms scales like  $n$  to the eleventh or twelfth power depending on which states are involved. So that millihertz interaction becomes hundreds of megahertz. Eleven to twelve orders of magnitude, turned on and off with a laser pulse.

The figure of merit is the product of the interaction strength and the Rydberg state lifetime, because you need the atoms to talk loudly and you need them to survive long enough to finish the conversation. Lifetime scales as  $n$  cubed, which is counterintuitive (higher energy state, longer life) but falls out of density of states and Fermi's golden rule. Multiply the two scalings and the figure of merit goes as  $n$  to the fourteenth. Fourteenth power. There are not many places in physics where the knob is that sensitive.

Naturally the reaction is to crank  $n$  to the moon. People have done Rydberg experiments up at  $n$  equals 1500, atoms twenty microns wide, absolutely absurd objects. It does not help. At that point the neighboring energy levels are so densely packed that the blockade stops working, because there is always some other level close enough to absorb the second excitation. The sweet spot ends up somewhere between 50 and a little over 100. The universe, as usual, gives with one hand.

## **Optical tweezers, coin flips, and the atom conveyor belt**

The gate scheme was proposed in 2000. Getting a single atom to sit still where you want it took the field most of another decade, and the details matter because they determine the failure modes that any healthcare-relevant computation has to survive.

Single atoms get held in far-off-resonance traps, which are just very tightly focused laser beams, now universally called optical tweezers. The physics of choosing the wavelength is a nice piece of engineering logic. Trap depth scales inversely with how far detuned the trapping light is from an atomic resonance. So you want to be close to resonance for a deep trap. But close to resonance you also scatter photons, and every scattered photon kicks the atom and heats it. Scattering rate scales inversely with the square of the detuning. So going further out costs one factor of detuning in trap depth and buys back two factors in scattering suppression. Net, you win by going far away. Typical operating detunings run 10 to 100 terahertz. That is very far off resonance, which is why it is in the name.

Loading those traps is a coin flip, and the reason is a lovely bit of collision physics. Overlap a tweezer with a cloud of cold atoms and a few atoms wander in. If two of them collide, they typically release enough energy that both are ejected. So atoms leave the trap two at a time. Start with an odd number, you end up with exactly one. Start with an even number, you end up with zero. Roughly fifty-fifty. There are tricks with additional laser wavelengths that engineer the collisional potential so atoms leave one at a time, which pushes loading above ninety percent, but the naive number is a half.

A half is a disaster at scale. Two traps, both loaded, twenty-five percent of the time. One hundred traps, all loaded, one over two to the hundredth. That is not a small number, that is a never number. The fix, which arrived in 2016 from three groups more or less simultaneously, is embarrassingly simple in hindsight and it is the single development that lit the fuse under this entire field. Load your traps at random. Take a picture. Now you know which sites are filled. Then physically pick up the atoms with a movable tweezer and rearrange them into a defect-free array. Sort the atoms. That is it. That is the trick.

And once you can move atoms around, you can do much more than sort them. You can move them mid-circuit, which turns out to be the platform's superpower. More on that shortly.

The current frontier is loss. The vacuum is not perfect, atoms leave, and if a Rydberg excitation fails to come back down the atom gets kicked out of the trap. A long computation will bleed atoms. So the machines now being built continuously reload the array while maintaining coherence in the atoms already there. A conveyor belt. Reservoir over here, cold atoms streaming in, tweezers pulling them into the vacancies while the computation carries on. Groups in Munich, at Atom Computing, and in the Lukin lab have all demonstrated versions. Nobody has to load perfectly anymore. They just have to keep the array topped up, like a bar restocking during service.

# Gates, and the slow climb from horrible to three nines

The first Rydberg-mediated CNOT, done at Wisconsin around 2009 to 2010, had a fidelity that its own author cheerfully describes as horrible. It produced entanglement, which had never been done this way, and being first is worth a great deal in science, but the numbers were bad. The atoms were at more than 100 microkelvin, which is cold by any human standard and blazing hot by cold atom standards. Modern experiments run at 5 microkelvin or below. Hot atoms move, moving atoms Doppler shift off resonance, and moving atoms wander around inside a tightly focused addressing beam so the effective pulse area changes shot to shot. The lasers were diode lasers with too much phase noise. The intermediate state chosen for the two-photon excitation was not the right one. And the  $\pi$ , two- $\pi$ ,  $\pi$  protocol itself is fragile, because it parks one atom up in the Rydberg state, exposed to decoherence, while it does the slower two- $\pi$  rotation on the other.

Everything since has been a grind against those errors. Laser stabilization borrowed from the atomic clock community, ultrastable reference cavities, linewidths down near 100 hertz, and then the later realization that linewidth is not enough and phase noise in the servo bumps matters just as much. Better intermediate states. And better gate protocols. The Levine-Pichler gate drives both atoms at once with a single beam, does a phase jump halfway through, and drives again. The time-optimal gate, worked out by a student of Guido Pupillo, does away with the phase jump entirely: one continuous pulse, constant amplitude, with a time-varying phase engineered so the differential phase comes out to  $\pi$ . Smoother, faster, less time parked in the Rydberg state, which means less scattering, which means higher fidelity.

State-of-the-art two-qubit Rydberg gate fidelities now reach the mid-four-nines range and have been demonstrated in multiple atomic species, so this is not a rubidium-specific parlor trick. Cesium, ytterbium, and strontium all work. A strontium experiment at Caltech using single-photon excitation from a metastable state set an earlier neutral-atom gate fidelity benchmark, but more recent rubidium platforms using optimized smooth-amplitude pulses have now exceeded those fidelities, reaching 99.854(4)% raw fidelity and 99.941(3)% with loss postselection.

The ceiling is what should interest anyone modeling timelines. Gate infidelity scales as some constant divided by the product of interaction strength and Rydberg lifetime. Theory from twenty years ago guessed the constant could not go below about two, but the bound was loose and non-constructive. Work last year tightened it to roughly two and a half and showed how to actually reach it. And a very recent result from Madison suggests a slightly different gate class can push the constant down to about one and a

half. Plug in realistic numbers for interaction strength and lifetime and that implies achievable gate fidelities around five nines, an infidelity of ten to the minus five. The experimental record is a hundred times worse than that. Which means the physics has a hundred times of headroom left and all of it is technical: laser noise, atom temperature, stray electric fields. That gap is the most important number in the field.

## **The error correction tax, and why it decides everything healthcare cares about**

Nobody is going to compute a drug binding energy on physical qubits. Not at three nines, not at five nines. Every serious application requires logical qubits, which means error correction, which means a redundancy tax, and the size of that tax is what separates “pharma gets a useful machine in 2032” from “pharma gets a useful machine in 2045.”

Under the standard surface code, with physical error rates around a tenth of a percent, the exchange rate is roughly a thousand physical qubits per logical qubit. That is where all the million-qubit resource estimates come from. It is also why every quantum computing roadmap has a cliff in it.

The alternative that has taken over the field’s attention is high-rate quantum LDPC codes, of which IBM’s bivariate bicycle code is the famous example. These encode many logical qubits into a modest number of physical ones, cutting the overhead by an order of magnitude or more. The catch is that the error syndrome measurements are non-local. You have to connect qubits that are not neighbors. On a superconducting chip that means multi-layer wiring and a genuine fabrication nightmare, and IBM is doing it anyway, which tells you how valuable the codes are.

On neutral atoms, non-local connectivity is free. You pick up the atom and move it. This is the platform’s structural argument and it is a strong one. Physical transport inside the circuit, introduced by the Lukin group, lets any qubit talk to any other qubit. The cost is time: move too fast and the atom heats, so you have to recool, and recooling is slow. Motion is measured in hundreds of microseconds to milliseconds while the Rydberg gate itself is hundreds of nanoseconds. The motion, not the physics, is the clock.

Which is why the two-species work matters more than it sounds like it should. Error correction requires ancilla qubits that get measured constantly, and measuring an atom means scattering resonant light off it, and scattered resonant light hits the neighboring data qubits and wrecks them. The standard fixes are to shelve the data

qubits in states that do not see the light, or to physically haul the ancillas somewhere else, measure them, and haul them back. Both cost operations and time. But if the data qubits are rubidium and the ancillas are cesium, the cesium readout light is so far detuned from any rubidium transition that the rubidium simply does not notice. Measure in place. No shelving, no transport. And the two species still entangle with each other perfectly well through Rydberg states, at fidelity comparable to same-species gates, because you can always find levels where a rubidium Rydberg atom and a cesium Rydberg atom interact strongly. Saffman's group proposed this about a decade ago and has now demonstrated it with a team at Infleqion, using it for real syndrome extraction.

Strip the motion out of the error correction cycle and the machine's effective clock speed goes up by a large factor, because most of what a fault-tolerant quantum computer does, minute to minute, is not compute. It corrects errors. Everything else is a rounding error on the duty cycle.

## **The drug discovery pitch, and the small part of it that is true**

Now the part the audience actually cares about.

The genuine quantum chemistry case is narrow and it is real. Classical methods handle most molecules fine. Density functional theory is cheap and wrong in known ways. Coupled cluster with perturbative triples is the gold standard and scales as the seventh power of system size, which means it dies fast. The problems where classical methods break are strongly correlated electronic systems: transition metal centers, iron-sulfur clusters, excited state photochemistry. In biology that means metalloenzymes. The nitrogenase FeMo cofactor is the standing example, and the cytochrome P450 family is the one with obvious pharma relevance, since P450s are responsible for the metabolism of the majority of marketed small molecules and predicting sites of oxidation from first principles remains genuinely hard.

The resource estimates for these have been coming down steadily and are still brutal. The state of the art, after aggressive algorithmic work on tensor factorizations and qubitization, lands somewhere around a million physical qubits and days of runtime for a FeMoco-class or P450-class calculation at chemical accuracy, assuming surface code overheads. High-rate codes could cut that meaningfully. That is a real number and a real path, and it is also not a 2027 number.

Meanwhile the thing that annoys physicists is that most of drug design is not an electronic structure problem. Binding affinity is dominated by conformational

sampling and solvation entropy, which is a statistical mechanics problem, and the industry already throws free energy perturbation and, increasingly, machine-learned interatomic potentials at it with results that are improving fast. The structure prediction problem got eaten alive by deep learning. Nobody in 2019 predicted that. Anyone building a quantum drug discovery thesis has to answer the question of why the classical ML curve, which is compounding annually, will still leave a gap for a quantum machine that arrives in the 2030s. The honest answer is: for a specific class of metalloenzyme and photochemistry problems, probably yes. For most of medicinal chemistry, probably no.

The pharma partnerships that get press releases should be read accordingly. A large pharma putting a few million dollars and two FTEs into a quantum collaboration is buying an option and a recruiting brochure. That is a rational thing to buy. It is not a signal that the technology is close.

## **The input problem, or why quantum claims analytics is a category error**

This section exists because the pitch deck will show up eventually, if it has not already.

Quantum computers are not fast classical computers. They are exquisitely specific machines that gain advantage on problems with particular algebraic structure. Two hard constraints kill nearly every healthcare data application that gets proposed.

First, data loading. To operate on a classical dataset, a quantum computer has to encode it into a quantum state. There is no known efficient way to load a large, unstructured classical dataset into a quantum register. The theoretical construct that would do it, quantum random access memory, requires hardware nobody knows how to build and its resource requirements typically swamp whatever speedup it was supposed to enable. A payer's claims warehouse is terabytes of unstructured classical bits. Loading it into a quantum state is not a hard engineering problem, it is closer to a category error.

Second, the speedup class. The famous exponential speedups apply to a short list: factoring, discrete logarithms, simulating quantum systems, and a few structured linear algebra problems with heavy fine print. Most optimization and search problems get at best a quadratic speedup, and there is a strong argument, made carefully by the Google team, that quadratic speedups are effectively useless for the foreseeable future because the constant factor overhead of error correction, plus the slow physical clock speed of quantum hardware, eats the advantage for any problem size a human being

will ever actually pose. A neutral atom gate is microseconds. A classical logic gate is sub-nanosecond. Starting a race a million times behind and hoping the square root saves you is not a strategy.

And the quantum machine learning literature took a serious beating when a series of dequantization results showed that several proposed exponential speedups had efficient classical analogues all along. The famous quantum recommendation systems algorithm got classically dequantized by an undergraduate.

So: quantum optimization of a PBM formulary, quantum provider network adequacy, quantum prior authorization triage, quantum risk adjustment. All of these are mixed integer programs that commercial solvers have gotten roughly a thousand times faster at over the last twenty-five years, and none of them have the structure a quantum computer exploits. If a vendor pitches these, the correct question is which specific speedup they are claiming and against which classical baseline, and the correct expectation is that the answer will be vapor.

## **Harvest now, decrypt later, and the most boring important thing in health IT**

The one place where quantum computing is already a health IT operations problem, today, involves no quantum computer at all.

Shor's algorithm breaks RSA and elliptic curve cryptography. That is not controversial and it is not a maybe. The only question is when a machine large enough exists, and the estimates cluster in the range of a few million physical qubits, meaning the same fault-tolerance milestone that drug discovery needs. Cryptography may actually arrive first, since the problem is more structured and less demanding on precision than chemistry.

The threat model that matters is not the day the machine turns on. It is today. An adversary with storage can capture encrypted traffic now and decrypt it in fifteen years. For most data that is worthless. For health data it is not. A genome is identifying for the life of the patient and partially identifying for the lives of their siblings, children, and grandchildren. A behavioral health record, an HIV status, a genetic predisposition, a substance use treatment episode: none of these get less sensitive in 2040. Health data has the longest secrecy half-life of almost any commercial data class, which makes it the single most attractive target for harvest-now-decrypt-later collection.

NIST finalized the first post-quantum standards in 2024: a lattice-based key encapsulation mechanism and two signature schemes. Federal guidance points toward deprecating classical public key algorithms around 2030 and disallowing them around 2035. Meanwhile the health data ecosystem runs on TLS-secured FHIR APIs, decades-old HL7 v2 interfaces over VPNs, a national interoperability framework whose participants have wildly varying security maturity, and hospital networks that still have devices running operating systems that reached end of life during the first Trump administration. Crypto agility, the ability to swap out an algorithm without rewriting the application, is essentially absent.

Nobody gets promoted for a successful cryptographic migration. It is a decade of unglamorous inventory work that ends with everything working exactly as it did before. It is also the highest-expected-value quantum-related activity available to any covered entity right now, and the reason it will not happen on time is that no regulator has made it a condition of anything.

## **The cold atom tech that will touch a patient first**

Here is the part that gets underweighted, and it is where the near-term money probably is.

The physics that makes a neutral atom quantum computer work also makes exceptionally good sensors, and the companies in this space understood that from the beginning. ColdQuanta, founded in 2007 by Dana Anderson, bootstrapped on family money, SBIR contracts, and selling components to other cold atom labs, was never purely a computing company. It was a cold atom technology company that sold clocks, sensors, and quantum computing as three legs of the same stool, and it later renamed itself Infleqion, a decision its own scientific advisors seem lukewarm about.

The medical application is magnetometry. Optically pumped magnetometers use exactly the atomic physics described above, alkali vapor, laser interrogation, precise state control, to measure magnetic fields at femtotesla sensitivity without cryogenics. That kills the liquid helium requirement that has kept magnetoencephalography confined to a few dozen research centers worldwide, each with a multi-million-dollar shielded room and a fixed helmet that fits an adult and nobody else. OPM sensors can be worn. They can sit on a child's head. They can move with the subject. MEG already has established CPT codes and established clinical utility in presurgical epilepsy localization and eloquent cortex mapping, so the reimbursement pathway is not hypothetical, it exists and is simply underused because the installed base is tiny. Companies in the UK and the US are commercializing wearable OPM-MEG now.

The cardiac analogue is magnetocardiography, measuring the heart's magnetic field rather than its electrical projection onto the skin. The clinical pitch is rapid rule-out of ischemia in the emergency department without biomarkers or radiation, an enormous volume indication where the current pathway is a serial troponin protocol and a long boarding time. Whether it works well enough is an open clinical question, not a physics question.

And Rydberg atoms themselves make superb radio frequency sensors, because those giant electron orbitals are absurdly polarizable and therefore absurdly sensitive to electric fields. That same electric field sensitivity, which is a nuisance in the quantum computer and forces careful control of the electrode environment, is a feature in an electrometer. Self-calibrating, SI-traceable RF measurement over a huge bandwidth. The medical uses are less obvious but the imaging and device-communication implications are not nothing.

So the honest sequencing looks like this. A patient will encounter this family of physics through a sensor long before they encounter it through a molecule that a quantum computer helped design. And the gating factor for the sensor is a payer coverage decision, which is a problem this audience knows how to solve.

## **Twenty companies, one talent pool, and how to read the tape**

There are now something like twenty companies worldwide building neutral atom quantum computers. Four of them started at essentially the same moment around 2018. Google, which spent a decade all-in on superconducting circuits, has spun up a neutral atom effort, which is either a vote of interest or a vote of fear and probably both.

The barrier to entry looks low because there is no fab. No cleanroom, no lithography, no yield problem. That is misleading. The lasers are expensive, the vacuum systems are expensive, and the real constraint, as the people actually building these things will tell you, is not dollars at all. It is people. There is a small pool of physicists who can actually make one of these machines work, the universities are not producing them fast enough, and twenty companies are bidding for the same postdocs. Great for the postdocs. Not an efficient allocation of capital.

For anyone underwriting this sector, that has a specific implication: the scarce asset is a functioning team with a decade of laser stabilization and vacuum systems scar tissue, not a patent portfolio. Diligence should look like biotech diligence on a

platform company, which is to say it should be mostly about the people and the specific technical problem they have already solved that nobody else has.

The public market comparables are not helpful; as of mid-2026, the sector's listed names still trade on expectations rather than meaningful operating revenue.

Valuations in listed quantum names have detached from revenue by multiples that make 2021 biotech look sober. Revenue in this sector is overwhelmingly government contracts, research systems sold to national labs, and cloud access fees that amount to marketing budgets from large enterprises buying optionality.

## **A scoreboard worth keeping**

Qubit count is the metric the press reports and it is nearly meaningless. A six thousand atom array is a beautiful piece of work and it computes nothing.

The numbers that actually predict when a healthcare-relevant computation happens are four. First, two-qubit gate fidelity, currently in the high-three- to mid-four-nines range on this platform, with a theoretical ceiling near five nines and roughly an order of magnitude technical gap in between. Second, logical qubit count and logical error rate, where the interesting threshold is a logical error rate below one in a million, since that is roughly what a chemistry algorithm with billions of gate operations requires. Third, the error correction overhead ratio, where the move from surface codes at a thousand to one down to high-rate LDPC codes at something closer to ten or twenty to one is the single biggest lever anyone has, and where the neutral atom platform's ability to physically move qubits is a genuine structural advantage over anything on a chip. Fourth, the clock, which for neutral atoms is currently set by atom transport and recoiling, and which the two-species in-place syndrome measurement work is aimed squarely at.

Watch those four. Ignore the qubit count press releases. And keep a healthy respect for the timeline, which one of the field's founders describes as five years plus or minus infinity, and which the same founder, twenty-five years into building the thing, is still unwilling to pin down more precisely than "longer than two or three years, and it is definitely going to happen."

For a sector that routinely underwrites ten-year clinical development timelines with a ninety percent failure rate, that should sound familiar. The difference is that in this case the physics is not in doubt. Only the engineering. And the engineering, as anyone who has tried to get two hospital systems onto the same interface standard will appreciate, is usually the hard part

---

## Subscribe to Thoughts on Healthcare Markets & Technology

By Thoughts On Healthcare Markets & Technology · Hundreds of paid subscribers

Expert analysis of healthcare and life sciences markets, technology, investment, entrepreneurship, policy, and AI — for investors, entrepreneurs, hospital and insurance executives, and physicians navigating the business of healthcare.

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).



1 Like

← Previous

### Discussion about this post

Comments

Restacks



Write a comment...