Neutral-Atom Quantum Computing: Technology, Industry and System Architecture

Oct 22, 2026 · David F. Hakala

This is a partial snapshot of a field that is moving quickly: what neutral-atom quantum computers are, who is building them, how the hardware is organized, how atoms act as qubits and gates, and who supplies the components. It combines two earlier notes, *Equipment and Operation (Reformatted)* and *Equipment and Operation A*, with company facts updated to October 2026. Material the sources stated without support is flagged in the review notes at the end.

1. Overview: neutral-atom quantum computing

A neutral-atom quantum computer stores each qubit in a single atom, usually rubidium, cesium, strontium or ytterbium, held in vacuum by a tightly focused laser beam called an **optical tweezer**. Hundreds to thousands of tweezers, written by a spatial light modulator (SLM) or acousto-optic deflectors (AODs), form a programmable array a few hundred micrometers across.

The qubit is a pair of long-lived internal states of the atom, for example two hyperfine ground states of ^{87}Rb or two nuclear-spin states of ^{171}Yb. Lasers or microwaves rotate the qubit. To entangle two atoms, a laser briefly excites them to a **Rydberg state**, a highly excited state with an enormous electron orbit, in which neighboring atoms interact strongly. A camera reads the result by imaging which atoms fluoresce.

Benefits

- **Identical qubits.** Every atom of an isotope is the same, so there is no fabrication spread to calibrate out, unlike superconducting circuits.
- **Scale.** Arrays of about 1,000 atoms are now routine, and research systems have held a few thousand to around ten thousand.
- **Long coherence.** Ground-state qubits keep their phase for seconds; nuclear-spin qubits for tens of seconds or more.
- **Reconfigurable connectivity.** Atoms can be moved during a computation, so any qubit can be brought next to any other without chains of SWAP gates.
- **Parallel operations.** One global laser pulse can apply the same gate to many atom pairs at once.
- **Mostly room-temperature hardware.** The atoms are laser-cooled to microkelvin, but the apparatus itself needs no dilution refrigerator. Some newer systems add cryogenic vacuum to extend atom lifetime.
- **Error-correction friendly.** Movable atoms suit efficient codes such as qLDPC codes, and some atomic errors can be converted into detectable “erasures”.

Challenges

- **Slow cycle time.** Loading, imaging and rearranging take milliseconds to a second per run, much slower than superconducting systems.
- **Atom loss.** Atoms are lost to background-gas collisions and imaging, so large machines need continuous reloading.

- **Two-qubit gate fidelity.** The best Rydberg CZ gates now reach about 99.85%, or 99.94% when runs that lost an atom are discarded. That is near the roughly 0.1% error level efficient error correction needs, but holding it across large arrays and long circuits is still the challenge.
- **Deep circuits.** Running many error-correction cycles in sequence is still a demonstrated-in-part capability; Google names it "an outstanding challenge for neutral atoms".
- **Laser complexity.** A system needs many stabilized lasers, from UV to infrared, with sub-kHz to sub-Hz linewidths for some functions.
- **Readout.** Fluorescence imaging is slow and can disturb neighbouring atoms, which makes mid-circuit measurement harder.
- **Optics at scale.** Trap uniformity, SLM power handling and flicker, and addressing many sites individually become harder as arrays grow.

	Strengths	Weaknesses
Qubit quality	Identical atoms, long coherence	Best two-qubit gates about 99.85%, not yet across whole arrays
Scale	Thousands of sites per vacuum cell	Atom loss and reloading at scale
Connectivity	Atoms movable mid-circuit	Moving atoms takes time
Speed	Highly parallel global gates	Millisecond cycle times
Engineering	No dilution refrigerator	Many precision lasers, optics and control loops

2. Major companies

All the full-stack builders trap single atoms in optical tweezers and use Rydberg states for entangling gates. They differ in atom species, array geometry, how they move or replace atoms, and business focus.

Company	Atom	Architecture	Status (October 2026)
QuEra Computing (Boston; Harvard/MIT spin-out, 2018)	^{87}Rb	Reconfigurable 2D arrays with coherent shuttling	Aquila 256-atom analog system on Amazon Braket; Gemini 260-atom gate-model system (>99.9% one-qubit, >99.2% two-qubit fidelity); system installed at AIST, Japan (2025); 3,000-qubit continuously reloaded array published in *Nature* (2025) (Wikipedia)
Atom Computing (USA)	^{171}Yb	Large 2D arrays, nuclear-spin qubits	1,180-qubit system on a 35 × 35 grid; co-developing fault-tolerant machines with Microsoft (Atom Computing)
Pasqal (France; NASDAQ: PSQL)	^{87}Rb	Arbitrary 2D geometries, analog and digital modes	Defect-free 1,024-atom registers; redesigned cryogenic platform with about 40× longer atom lifetime; QPUs in Saudi Arabia and Italy; public via SPAC in August 2026 (Quantum Computing Report)

Company	Atom	Architecture	Status (October 2026)
Infleqtion (Louisville; NYSE: INFQ; formerly ColdQuanta)	Cs	Sqale gate-model computer, plus clocks and sensors	Targeting 30 logical qubits in 2026; FY2026 revenue guidance about $43M (Quantum Computing Report)
planqc (Garching; MPQ spin-off, 2022)	Sr and alkali atoms in optical lattices	Digital neutral-atom computers for HPC centers	100-qubit DiNAQC demonstrator for DLR (spring 2027); 1,000-qubit MAQCS for LRZ (end of 2027) (Max Planck Society)
Google Quantum AI	not stated	New neutral-atom hardware team in Boulder, led by Adam Kaufman	Announced March 2026 alongside its superconducting program; continues working with QuEra (Google)
Microsoft	—	Software, error-correcting codes and cloud	Partner to Atom Computing; Azure Quantum
Amazon Web Services	—	Cloud access	QuEra Aquila on Amazon Braket
Oratomic (USA), Q-Factor (Israel), Atom Quantum Labs, OpenQuantum	various	Emerging companies	Named in the source; not verified here

The source article organizes each company's approach as *creation* (how the array is made), *maintenance* (how it's kept full and stable) and *manipulation* (how qubits are operated). That structure is kept below.

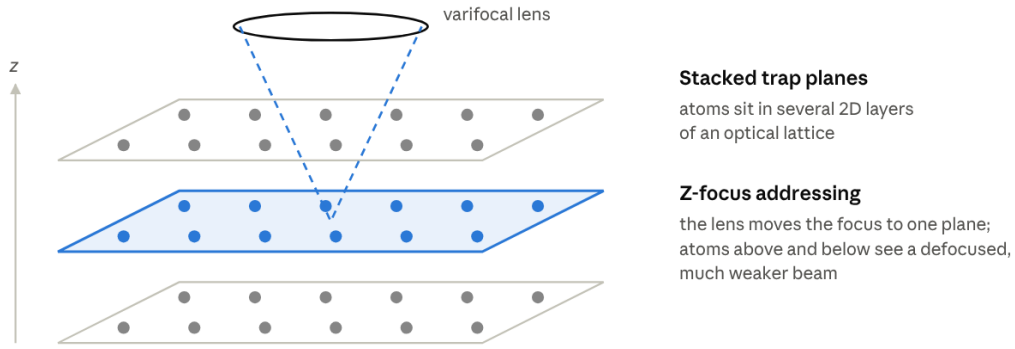
QuEra: reconfigurable 2D arrays

- **Creation.** A trap laser feeds an LCoS SLM, which splits it into an arbitrary 2D grid of hundreds of tweezers.
- **Maintenance.** Loading leaves about half the traps empty. AODs make a second, moving set of tweezers; a camera images the grid, an algorithm finds the gaps, and the moving tweezers slide atoms into a defect-free array (Section 7).
- **Manipulation.** AOD tweezers move atoms mid-computation (coherent shuttling), so distant qubits can interact without SWAP gates. A global Rydberg pulse at 420 + 1013 nm then entangles every pair that has been brought together (Section 6).

Atom Computing: nuclear-spin qubits at large scale

- **Qubit.** ^{171}Yb stores the qubit in its nuclear spin, separate from the valence electrons, so stray magnetic fields barely affect it. A “magic-wavelength” trap holds the atom without shifting the qubit frequency.
- **Manipulation.** Narrow-linewidth lasers drive single-qubit rotations; ultraviolet light briefly excites the electrons to a Rydberg state for entangling gates.
- **Scale.** The published 1,180-qubit system is a 35 × 35 grid. The source describes a 3D optical lattice with plane-by-plane “fill-and-discard” loading; that description is kept in Section 7 and the diagram below, but it is not supported by Atom Computing's published system (review notes).

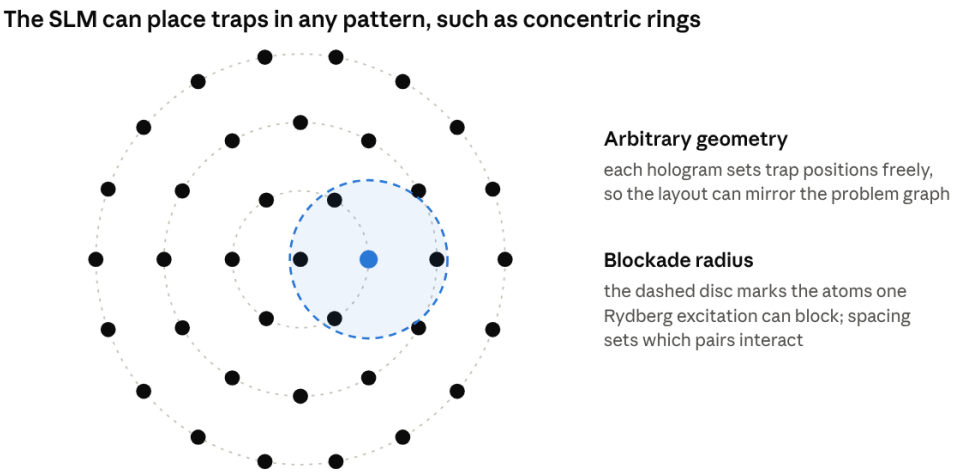
A varifocal lens picks out one plane of a 3D trap stack



3D lattice planes and Z-focus addressing · as described in the source (see review notes)

Pasqal: custom geometries, analog and digital modes

- **Creation.** High-resolution SLMs write trap patterns that match the problem, for example a material's crystal lattice for a simulation, or rings and frustrated triangles for optimization.
- **Maintenance.** Low background pressure and, in its newer platform, cryogenic vacuum extend atom lifetime; QPUs are integrated with HPC centers.
- **Manipulation.** An analog mode drives the whole array at once as a quantum simulator; a digital mode uses tightly focused beams on selected atoms for site-specific gates.



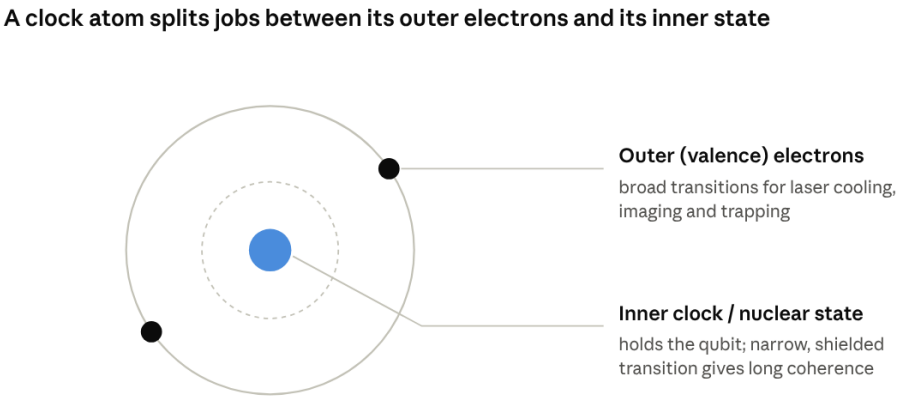
arbitrary trap geometry · concentric rings

Infleqtion: cesium gate-model systems

Infleqtion builds the Sqale gate-model computer from cesium atoms and also sells atomic clocks (Tiqker) and quantum sensors, plus the Superstaq software platform. Its 2026 target is 30 logical qubits, with an Illinois system planned for 2027 designed to pass 50 logical qubits.

planqc: clock atoms and optical lattices

- **Creation.** Atoms in high-precision optical lattices with uniform trap depth across thousands of sites.
- **Maintenance.** Alkaline-earth atoms like strontium allow two-stage cooling: a broad transition cools quickly, then a narrow transition cools atoms toward their lowest vibrational state, reducing gate errors from motion.
- **Manipulation.** Trapping and computing are separated: the trap light holds the atom while separate, localized laser beams address the qubit.



clock-atom qubit · schematic, not to scale

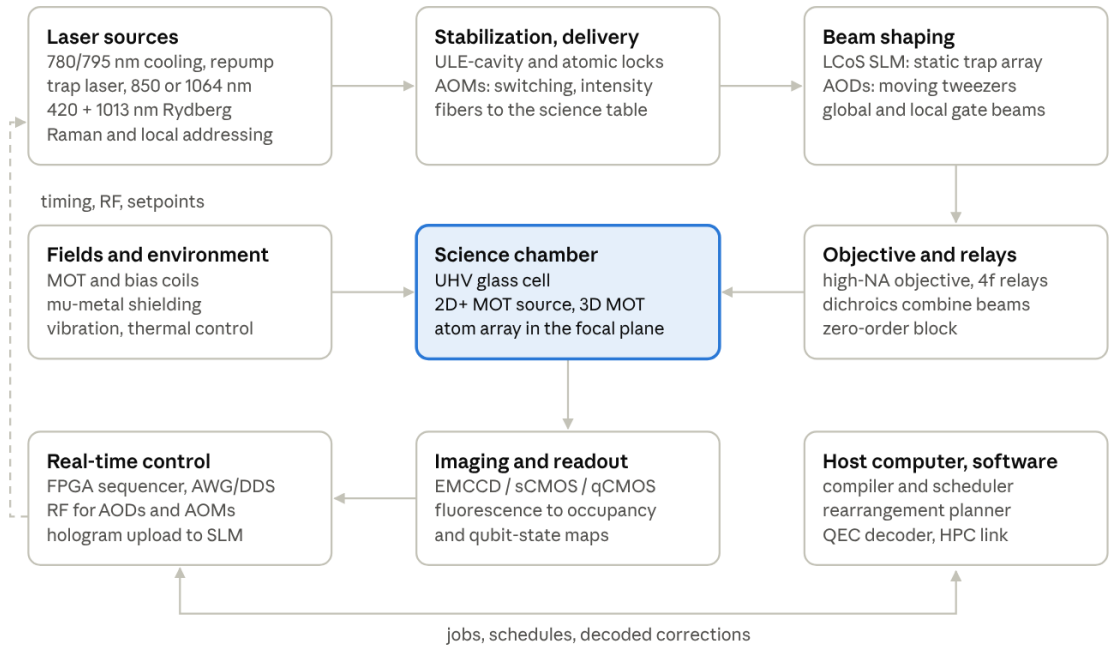
Big-tech programs

- **Google Quantum AI** added a neutral-atom hardware program in March 2026, alongside superconducting qubits, built on three pillars: error correction, modeling and simulation, and experimental hardware. It names deep circuits with many error-correction cycles as the key open challenge.
- **Microsoft** supplies error-correcting codes, qubit-virtualization software and Azure Quantum for Atom Computing's hardware.
- **Amazon Web Services** gives public cloud access to QuEra's Aquila through Amazon Braket.
- **Nvidia** works with QuEra and Pasqal on CUDA-Q integration for hybrid quantum–classical computing.

3. System architecture

At block level, a neutral-atom computer is three chains that meet at the atoms: **light** is generated, stabilized and shaped into traps and gate beams; the **atoms** are produced, cooled and held in a vacuum cell; and **control** sequences everything and reads out the result.

Light, atoms and control: the main blocks of a neutral-atom computer



system block diagram · light flows left to right and down to the atoms; control runs everything

What each block does

Laser sources. Four families of lasers do different jobs:

- *Cooling, repump and imaging* (780 nm for Rb): narrow-linewidth diode lasers, often tapered amplifiers for power, tuned near the D2 line. They cool atoms in the MOT and molasses and

make them fluoresce for imaging. A 795 nm D1 laser adds gray-molasses cooling and optical pumping.

- *Trap laser* (850 nm or 1064 nm for Rb): a high-power source, typically a fiber amplifier giving tens of watts, far detuned from the atomic lines so it holds atoms without scattering photons. It feeds both the SLM and the AODs.
- *Rydberg lasers* (420 nm + 1013 nm for Rb, about 302 nm for Yb): drive the two-photon transition to the Rydberg state for entangling gates. They need low phase noise, often sub-kHz linewidth.
- *Raman and addressing lasers* (for example 795 nm or a far-detuned beam): rotate single qubits, or shift selected atoms so a global pulse acts on only some of them.

Stabilization and delivery. Frequency locks hold each laser on its atomic line: saturated-absorption spectroscopy for cooling lasers, and ultra-low-expansion (ULE) glass reference cavities with Pound–Drever–Hall locking for Rydberg and clock lasers. Acousto-optic modulators (AOMs) switch beams in nanoseconds to microseconds, set detunings, and stabilize intensity in a photodiode feedback loop. Polarization-maintaining fibers carry the light from a separate laser table to the science table, which keeps amplifier heat and vibration away from the atoms.

Beam shaping. The trap light splits into two paths:

- *LCoS SLM:* displays a phase hologram (Gerchberg–Saxton or weighted GS) that turns one beam into a static array of hundreds to thousands of tweezers, in any 2D or 3D pattern. Its phase accuracy, flicker and power handling set trap-depth uniformity.
- *AODs:* a crossed X–Y pair driven by multi-tone RF makes a smaller set of moving tweezers. Sweeping a tone drags its tweezer, and the atom in it, across the array.
- *Gate beams:* the Rydberg lasers are shaped into a global sheet covering the array, and addressing beams are steered to single sites by their own AODs or a second SLM.

Objective and relays. 4f relay lenses image the SLM and AODs onto the back aperture of a high numerical-aperture objective (NA 0.5–0.7) just outside the vacuum cell. The objective focuses each beam to a spot about 1 µm wide. Dichroic mirrors combine the trap, AOD and imaging paths, and a small block in a Fourier plane removes the SLM's undiffracted zero-order spot.

Science chamber. An ultra-high-vacuum glass cell, typically below 10^{-11} mbar, holds the atoms. A separate source chamber with a dispenser or oven and a $2D^{+}$ MOT sends a slow atomic beam through a differential-pumping tube into a 3D MOT at the focal plane, where the tweezers load from the cold cloud. Ion and getter pumps keep the vacuum; some new systems are cryogenic to extend atom lifetime.

Fields and environment. Anti-Helmholtz coils make the MOT gradient; three-axis bias coils set a quantization field and cancel stray fields. Mu-metal shields and active field compensation stop magnetic noise from shifting the qubit frequency. The whole apparatus sits on a vibration-isolated optical table in a temperature-stabilized enclosure.

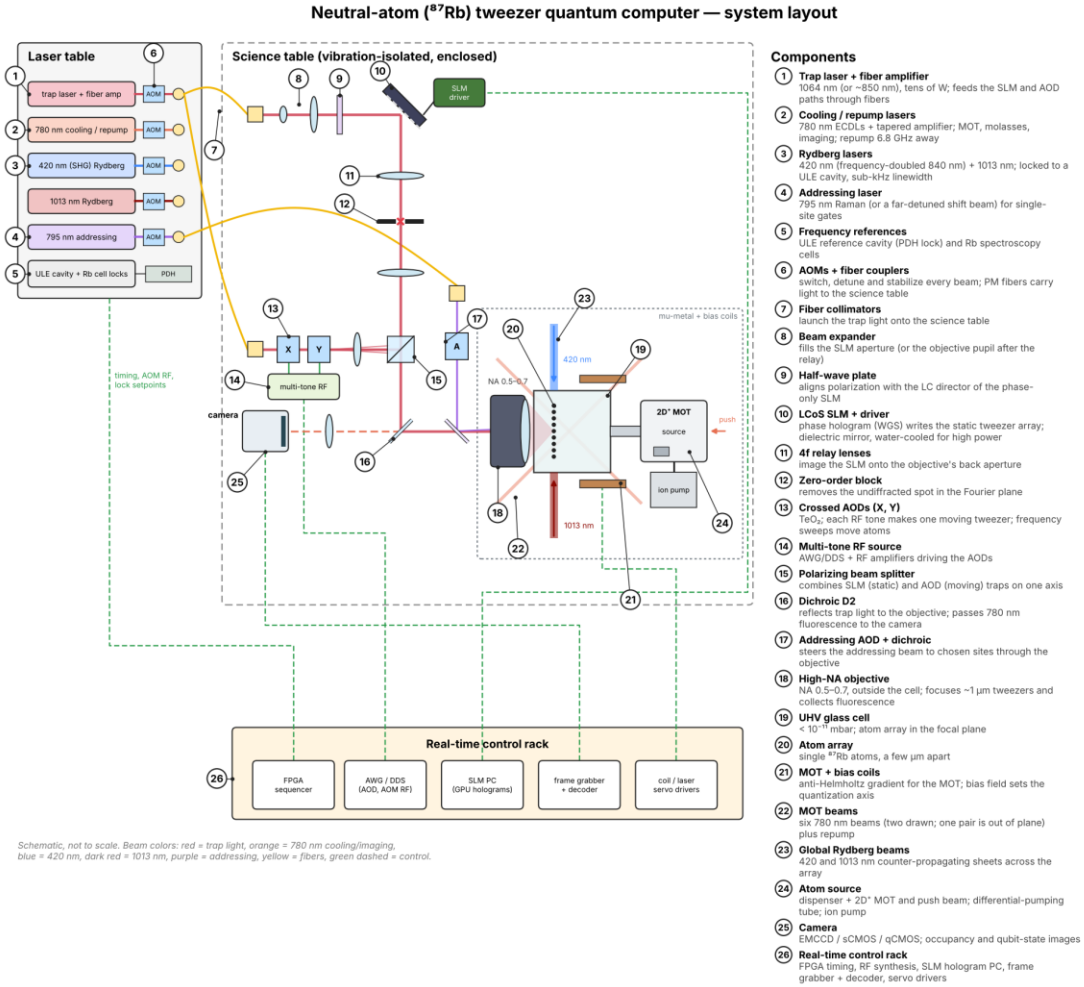
Imaging and readout. A camera images fluorescence through the same objective: EMCCD, sCMOS or qCMOS sensors detect a few hundred photons per atom in milliseconds. One image gives the occupancy map used for sorting; a state-selective image (after pushing out one qubit state, or shelving it in a dark state) gives the measurement result.

Real-time control. An FPGA-based sequencer times every laser pulse, RF tone and camera trigger to nanosecond precision. Arbitrary waveform generators or direct digital synthesizers make the multi-tone RF for the AODs and the shaped pulses for the AOMs. The controller uploads holograms to the SLM and closes fast loops for intensity, position drift and magnetic field.

Host computer and software. Compiles a quantum circuit into a schedule of atom moves and laser pulses, computes holograms and rearrangement paths (often on a GPU), runs the real-time quantum error-correction decoder, and connects the machine to users and to HPC centers.

4. System layout

The drawing below places the blocks of Section 3 on two optical tables, in the same style as the LCoS fringe-projection setup drawings. It shows a ^{87}Rb system with a 1064 nm trap laser; the numbers match the key on the right.



How light reaches the atoms.

1. **Laser table (1–6).** Every laser passes an AOM for switching and frequency control, then a polarization-maintaining fiber. The Rydberg lasers are locked to a ULE cavity, the cooling lasers to rubidium spectroscopy.

2. **Static traps (7–12).** The trap light is expanded, its polarization set along the LC director with a half-wave plate, and reflected off the LCoS SLM. A 4f relay images the SLM onto the objective's back aperture, with a block at the Fourier plane removing the undiffracted zero order.
3. **Moving traps (13–15).** A second fiber feeds the crossed AODs. Each RF tone makes one tweezer, and a polarizing beam splitter combines these moving traps with the SLM's static ones.
4. **Into the objective (16–18).** A dichroic sends the trap light into the high-NA objective, and a second dichroic adds the AOD-steered addressing beam. The objective focuses everything to micron spots in the atom plane.
5. **The cell (19–24).** The atoms sit in the focal plane of the glass cell, fed from the $2D^+$ MOT source through a differential-pumping tube. MOT beams, coils and the counter-propagating 420 nm and 1013 nm Rydberg sheets all cross there, inside a magnetic shield.
6. **Readout and control (25–26).** Fluorescence returns through the objective and the dichroic to the camera. The control rack times every beam, synthesizes the RF, uploads holograms and decodes images.

5. Atoms as qubits

Which atoms

Companies use either **alkali atoms** (Rb, Cs: one valence electron) or **alkaline-earth and alkaline-earth-like atoms** (Sr, Yb: two valence electrons).

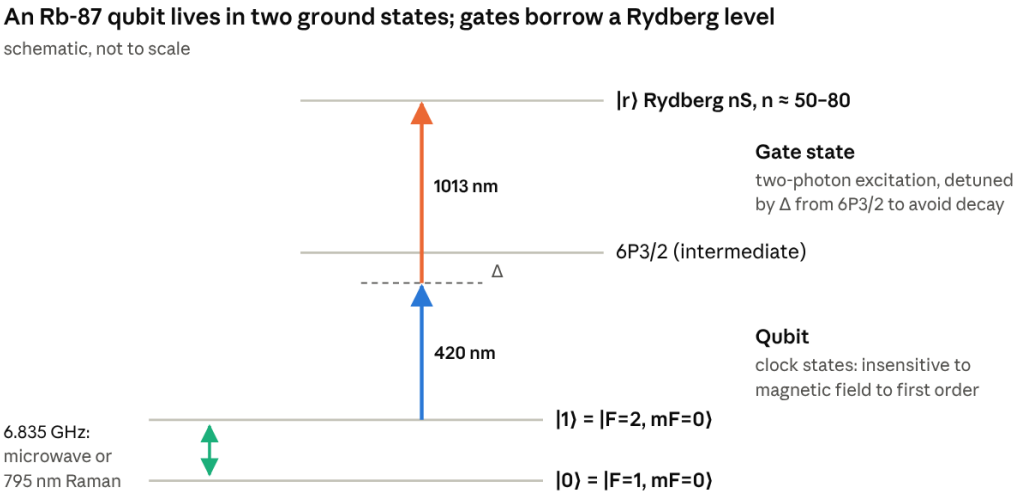
Company	Atom	Category	Qubit encoding
QuEra	^{87}Rb	Alkali	Ground hyperfine states; Rydberg state used only during gates (analog mode drives ground–Rydberg directly)
Pasqal	^{87}Rb	Alkali	Ground hyperfine states (digital); ground–Rydberg (analog)
Infleqtion	Cs	Alkali	Ground hyperfine states
Atom Computing	^{171}Yb	Alkaline-earth-like	Nuclear spin (I = ½) of the ground state
planqc	Sr (and alkali atoms)	Alkaline-earth	Clock-state or fine-structure qubits in Sr; ^{87}Sr also allows a nuclear-spin qubit

How an atom stores a qubit

A good qubit is a pair of states that lives a long time, is hard for the environment to disturb, and can be driven precisely.

- **Hyperfine clock qubit (Rb, Cs).** The two ground hyperfine levels with m_F = 0, for ^{87}Rb $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$, are split by 6.835 GHz (9.19 GHz for Cs). Their energy difference doesn't change to first order with magnetic field, so they dephase slowly. Microwaves or a pair of Raman laser beams drive transitions between them.

- **Nuclear-spin qubit (^{171}Yb, ^{87}Sr).** In a two-electron atom the ground state has no electronic spin, so the qubit lives in the nucleus, which barely couples to stray fields. Coherence of tens of seconds or more is possible, and magic-wavelength traps leave the qubit frequency unchanged.
- **Optical clock and fine-structure qubits (Sr, Yb).** One state is the ground state and the other a metastable excited state, such as the Sr $^1S_0$–$^3P_0$ clock transition at 698 nm. These need sub-Hz lasers but give very clean single-qubit control. The Munich work behind planqc uses a fine-structure qubit in metastable ^{88}Sr.



Rb-87 qubit and Rydberg ladder

For ^{87}Rb, the qubit sits in the two ground hyperfine states and is rotated by microwaves or 795 nm Raman beams. For a gate, $|1\rangle$ is briefly excited to a Rydberg state through two photons, 420 nm then 1013 nm, detuned from the intermediate $6P_{3/2}$ level so it is barely populated. The effective two-photon Rabi frequency is

$$\Omega_{eff} = \frac{\Omega_{420}\,\Omega_{1013}}{2\Delta}$$

Trade-offs between atom families

	Alkali (Rb, Cs)	Alkaline-earth (Sr, Yb)
Lasers	Commercially mature, reliable, relatively inexpensive	Complex, highly stable UV and deep-blue lasers; harder to engineer and scale
Coherence	Seconds with dynamical decoupling; shorter than nuclear spins	Nuclear spins well isolated; tens of seconds or more
Magnetic noise	Electronic spin more sensitive to	Nuclear-spin qubit nearly insensitive

	Alkali (Rb, Cs)	Alkaline-earth (Sr, Yb)
	stray fields	
Gates	Single valence electron couples strongly to Rydberg lasers: fast, high-fidelity entangling gates	Gates via nuclear spin or narrow clock transitions are typically slower
Preparation and readout	Simple optical cycling	Readout from the nuclear spin needs extra optical steps
Other	Electrons can leak into unwanted sublevels	Metastable states allow erasure conversion and shelving during readout

Laser wavelengths by species

Control splits into trapping and cooling (high power, broad linewidth) and quantum logic (narrow linewidth, low noise). Values are as given in the source article.

Company	Atom	Function	Wavelength	Typical linewidth
QuEra, Pasqal	^{87}Rb	Laser cooling (D2 line)	780 nm	~1 MHz
QuEra, Pasqal	^{87}Rb	Optical tweezers	850 nm (1064 nm also used)	Broadband (> 10 GHz)
QuEra, Pasqal	^{87}Rb	Ground-state qubit control (D1)	795 nm	< 100 kHz
QuEra, Pasqal	^{87}Rb	Rydberg excitation (two-photon)	420 nm + 1013 nm	< 10 kHz
Atom Computing	^{171}Yb	Laser cooling (broad line)	399 nm	~30 MHz (natural linewidth)
Atom Computing	^{171}Yb	Magic-wavelength trapping	759 nm	Broadband
Atom Computing	^{171}Yb	Qubit clock rotation	578 nm	< 1 Hz
Atom Computing	^{171}Yb	Rydberg transition	302 nm (UV)	< 5 kHz
planqc	^{87}Sr	Main cooling (blue line)	461 nm	~32 MHz (natural linewidth)
planqc	^{87}Sr	Narrow cooling (red line)	689 nm	~7 kHz
planqc	^{87}Sr	Magic-wavelength trapping	813.4 nm	Broadband
planqc	^{87}Sr	Qubit clock gate	698 nm	< 1 Hz

- **Why sub-hertz linewidths.** Clock transitions in Sr and Yb are “forbidden”, with extremely narrow natural linewidths. Even small laser phase noise makes the laser drift out of step with the atom, and the qubit rotation fails (Physics World).
- **Magic wavelength.** Trap light normally shifts atomic levels (the AC Stark shift). At a magic wavelength, 759 nm for Yb or 813.4 nm for Sr, both qubit levels shift by the same amount, so the atom is held firmly without disturbing the stored qubit (ADS).

6. Gates: what happens physically

A quantum computer needs single-qubit rotations and one entangling two-qubit gate; together they can build any circuit. Neutral atoms use **rotations driven by microwaves or lasers** and a **controlled-Z (CZ) gate driven by Rydberg blockade**.

Single-qubit gates

A resonant drive between $|0\rangle$ and $|1\rangle$ makes the atom oscillate between them (Rabi oscillation). A pulse of area $\theta = \Omega t$ rotates the qubit by θ about an axis set by the drive's phase: a π pulse flips the qubit, and a $\pi/2$ pulse makes an equal superposition, the atomic version of a Hadamard.

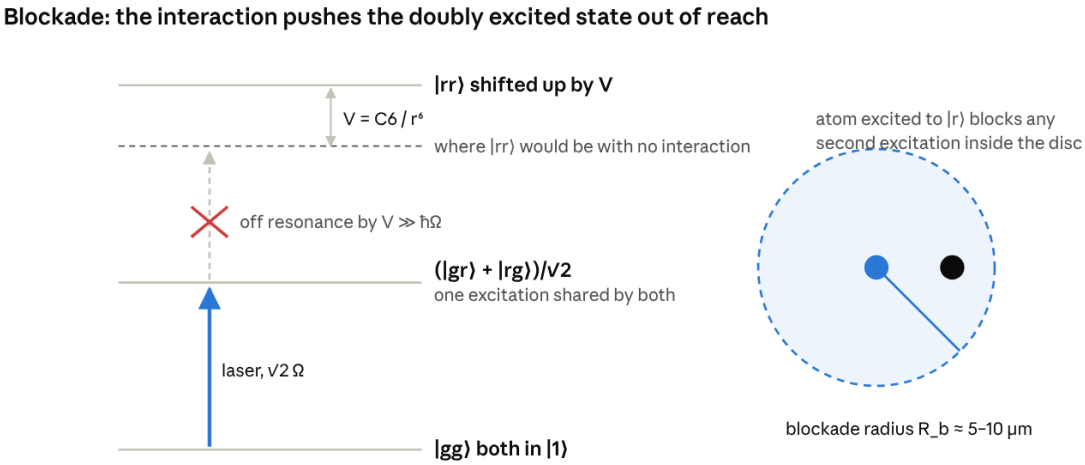
- **Global rotations.** A microwave field or a broad Raman beam pair rotates every atom in the array at once. Fidelities above 99.9% are typical.
- **Local rotations.** A tightly focused beam, steered by an AOD or a second SLM, light-shifts selected atoms. That gives them a Z rotation, or moves them out of resonance so a following global pulse skips them. Focused Raman beams can also drive individual X or Y rotations.
- **Errors.** Mainly intensity and pointing noise of the beams, and crosstalk onto neighbouring sites.

Two-qubit gates: the Rydberg blockade

Exciting an atom to a Rydberg state ($n \approx 50$–80) makes its electron orbit hundreds of nanometers across, so two Rydberg atoms a few micrometers apart interact strongly through the van der Waals force:

$$V(r) = \frac{C_6}{r^6}, \qquad R_b = \left(\frac{C_6}{\hbar\Omega}\right)^{1/6}$$

If V is much larger than the laser coupling $\hbar\Omega$, both atoms can't be excited together: the doubly excited state is pushed out of resonance. Within the **blockade radius** R_b, typically 5–10 µm, one excitation blocks the other.



Rydberg blockade · two-atom energy levels and blockade radius, schematic

What happens physically during a CZ gate.

1. Atoms whose qubit is $|0\rangle$ are not coupled to the Rydberg laser and do nothing.
2. An atom in $|1\rangle$ with no partner in $|1\rangle$ is driven up to $|r\rangle$ and back, and picks up a phase.
3. If both atoms are in $|1\rangle$, the blockade allows only a shared single excitation, $(|gr\rangle + |rg\rangle)/\sqrt{2}$. It oscillates $\sqrt{2}$ times faster and returns with a different phase.
4. The pulse is designed so that after the atoms return to the ground state, $|11\rangle$ has picked up an extra phase of π relative to the other states. That is a CZ gate, equivalent to a CNOT up to single-qubit rotations.

Pulse protocols.

- *Jaksch protocol:* π pulse on the control atom, 2π on the target, π on the control. Simple, but it needs individual addressing.
- *Levine–Pichler protocol:* two global pulses of equal length with a phase jump between them, so both atoms are driven together with no local addressing. With $\Delta/\Omega = 0.377$, $\Omega\tau = 4.293$ per pulse and a phase jump of 3.902 rad, it gives an ideal CZ.
- *Time-optimal pulses:* a smoothly varying laser phase replaces the jump. In our earlier simulation it completes the gate in about 263 ns at $\Omega = 2\pi \times 4.6$ MHz and spends less time in the decaying Rydberg state.

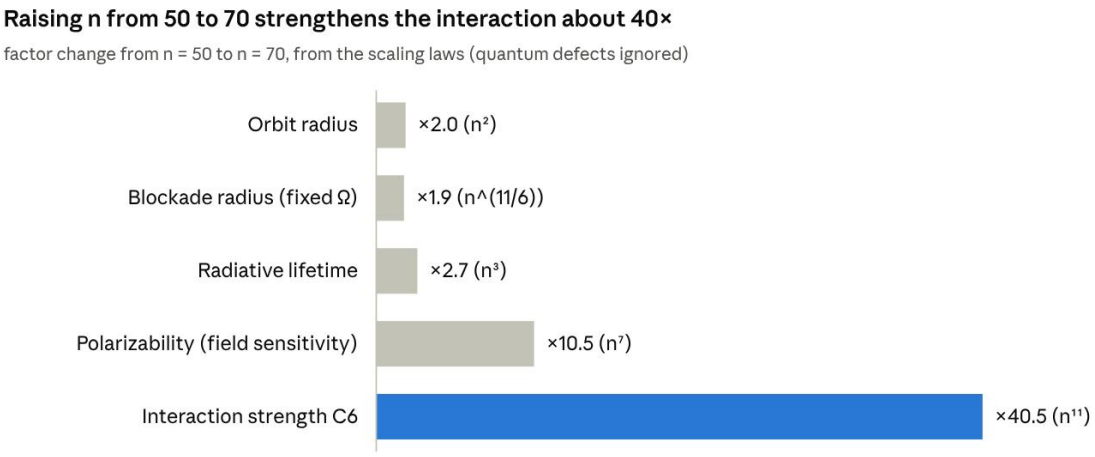
Fidelity. QuEra quotes better than 99.2% for Gemini two-qubit gates. The best published neutral-atom CZ gate, from the Harvard–MIT–QuEra group in April 2026, reaches 99.854% raw and 99.941% when runs with a lost atom are discarded (Evered et al., arXiv 2604.25987). The limits are Rydberg-state decay, laser phase and intensity noise, Doppler shifts from atom motion, and leakage to other Rydberg levels. The separate note *Improving Rydberg Gate Fidelity* covers the remedies.

How the 99.85% gate was achieved (Evered et al., 2026):

- **Atoms and levels.** ^{87}Rb hyperfine clock qubits, excited to the $53S_{1/2}$ Rydberg state through $6P_{3/2}$ with 420 nm and 1015 nm light, detuned $2\pi \times 7.8$ GHz from the intermediate level.
- **Fast, smooth pulse.** An optimal-control pulse with smooth amplitude and phase at a peak Rabi frequency of about $2\pi \times 17$ MHz. Going faster cuts the time spent in the decaying Rydberg state, and the pulse keeps population out of the intermediate level to limit scattering.
- **Side effects managed.** The high laser intensity causes large light shifts, so the Rydberg beam powers and positions are tightly stabilized. A stronger magnetic field suppresses coupling to the neighbouring Rydberg sublevel.
- **Fast calibration.** Up to 442 atoms are arranged in zones, with 17 gate sites and a large reservoir that refills lost atoms. State-selective readout reveals which atoms were lost, so calibration runs at 20–30 cycles per second and takes under 40 minutes from scratch.
- **Stability.** Fidelity stayed at about 99.85% for over 10 hours without recalibration. The gate is sensitive to changes of about 0.1% in Rydberg beam power, and remaining drifts come from beam power and pointing.
- **Error budget.** Atom loss is about 60% of the remaining error (0.087%) and can be detected, which helps error correction. Most of the rest is Rydberg-state decay and scattering from the intermediate level. The authors project 99.9–99.95% raw fidelity with better coherence, less coupling to the neighbouring Rydberg level, and more laser power.
- **In circuits.** Gates on atoms moved by coherent rearrangement measured 99.843% raw and 99.956% after discarding runs with lost atoms, so moving atoms between gates did not cost fidelity.

Choosing the Rydberg level

Higher n makes the interaction much stronger, but also makes the atom far more sensitive to stray electric fields. The choice is a balance.



computed from the standard Rydberg scaling laws in n

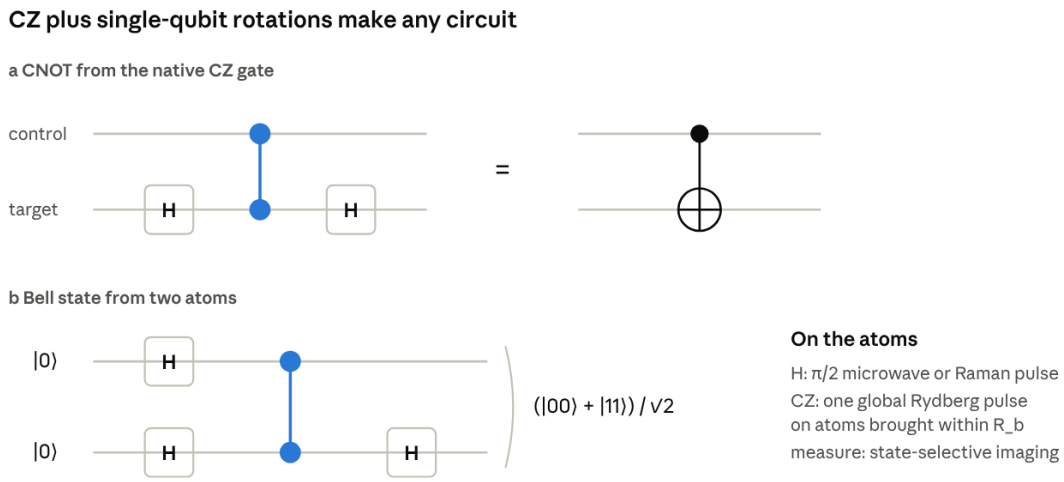
The interaction grows as n^{11} while the lifetime grows only as n^3, so stronger blockade comes with a much larger sensitivity to stray electric fields (n^7). That is why most systems settle on n ≈ 50–80.

Multi-qubit gates and parallelism

- **Parallel gates.** One global Rydberg pulse applies CZ to every pair of atoms brought within R_b of each other, while pairs farther apart are unaffected. Moving atoms with the AODs between pulses chooses which pairs interact.
- **Native multi-qubit gates.** Three atoms all inside one blockade radius give a native CCZ (Toffoli-class) gate in a single pulse.
- **Erasure conversion.** In alkaline-earth atoms, many gate errors leave the atom in a state that can be detected without disturbing the others. Knowing *where* an error happened makes it much cheaper to correct.

Circuits

The native gate set is global and local single-qubit rotations plus CZ. Every other gate is compiled from these. For example, CNOT is a CZ with Hadamards on the target, and a Bell pair needs only two Hadamards, one CZ and one more Hadamard.



native gates and two small circuits

In practice the compiler groups gates into layers. Each layer is a set of atom moves, then a global Rydberg pulse that entangles all pairs in that layer, then local single-qubit rotations. Error-correction circuits repeat such layers, with ancilla atoms measured mid-circuit.

7. Operating cycle

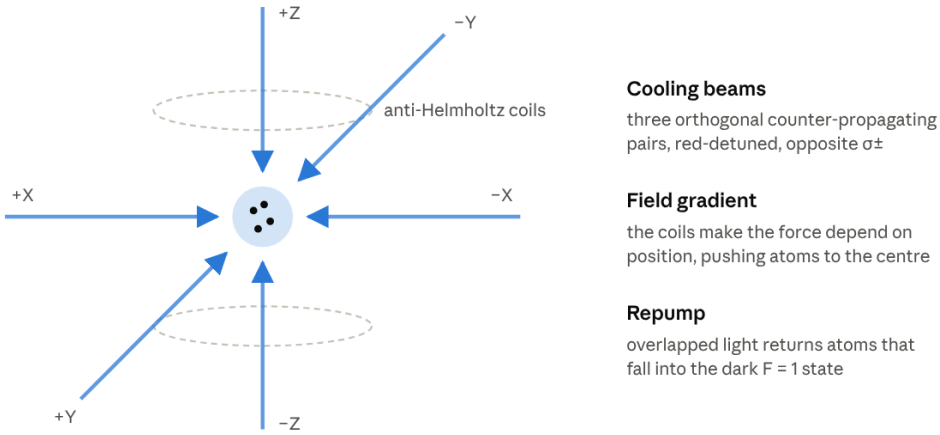
One run of the machine follows the same steps on every platform: load atoms, image them, fill the gaps, run the circuit, read out. A full cycle takes about 0.3–1 s, and loading is usually the longest part.

Loading

Rubidium vapor from a dispenser is cooled in a 2D⁺ MOT and pushed into the science cell, where a 3D MOT captures it at about 100–200 µK. Sub-Doppler molasses cools it to about 5–20 µK, and the

tweezers, already on, catch atoms from this cold cloud. Each tweezer is small enough that a second atom causes a light-assisted collision that ejects the pair (collisional blockade), so each site ends up with 0 or 1 atom, at about 50–60% occupancy, or about 80–90% with blue-detuned “enhanced loading”. The separate note *Rb-87 Atom Source, Laser Cooling and Tweezer Loading* covers this in detail.

A MOT uses three crossed beam pairs plus a magnetic gradient

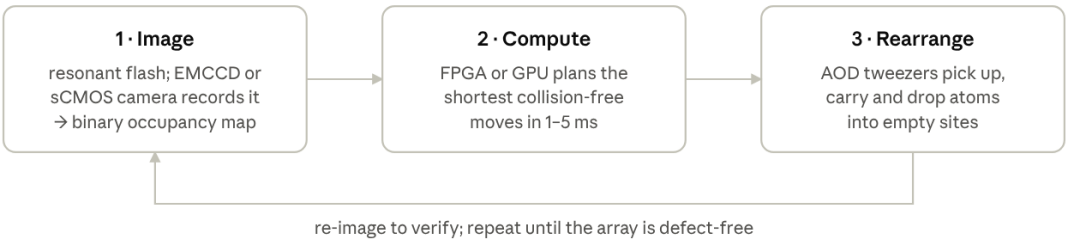


six-beam magneto-optical trap

Imaging and defect-free sorting

Random loading leaves about half the traps empty. A closed loop of imaging, planning and rearrangement turns that half-empty grid into a defect-free array in about 20–50 ms, after which the sorting light switches off and the gate sequence starts.

Sorting closes the loop in about 20–50 ms



sorting loop

1. **Imaging.** An EMCCD or sCMOS camera outside the vacuum cell records fluorescence through the high-NA objective during a brief resonant pulse. Filled traps light up and empty ones stay dark; the image becomes a binary occupancy matrix.
2. **Planning.** An FPGA, or a CPU/GPU coupled to the camera, computes the moves within 1–5 ms, before background-gas collisions cause more losses.

3. **Rearrangement.** The static SLM grid stays on while AOD tweezers pick atoms out of surplus sites, carry them across the array without disturbing their quantum state, and drop them into vacancies.

Algorithm	How it works	Main advantage
Heuristic / greedy matching	Scans row by row, matching the nearest atom to the nearest vacancy	Very fast to compute; suits simple grids
Network-flow optimization	Treats the array as a graph and finds global minimum-cost paths for groups of atoms	Minimizes total travel distance, reducing losses
Compression	Sweeps scattered atoms toward one side to form a dense block	Reliable for packing arbitrary shapes

Company	Sorting	Hardware	Distinctive approach
QuEra	2D moving grid	Two-axis AOD controlling rows and columns	Moves whole lines of atoms at once, fixing hundreds of defects in a few sweeps
Pasqal	Arbitrary geometries	High-precision 2D AOD	Rearranges atoms into non-grid layouts such as rings, spirals or frustrated triangles
Atom Computing (per the source)	3D lattice selection	Multi-plane focusing optics	"Fill-and-discard": ejects unneeded atoms instead of moving them between planes

How AODs move atoms

Acousto-optic deflectors split one laser beam into an array of tweezers and move individual atoms (arXiv 2510.19816, Optica OPN).

Array generation by multi-tone RF. A crossed X–Y AOD pair is driven by many simultaneous tones f_1, f_2, …, f_n. The deflection angle is proportional to the acoustic frequency, so each tone makes one focal spot (RP Photonics):

$$\theta \approx \frac{\lambda f}{v}$$

where λ is the optical wavelength, f the acoustic frequency and v the acoustic velocity in the crystal.

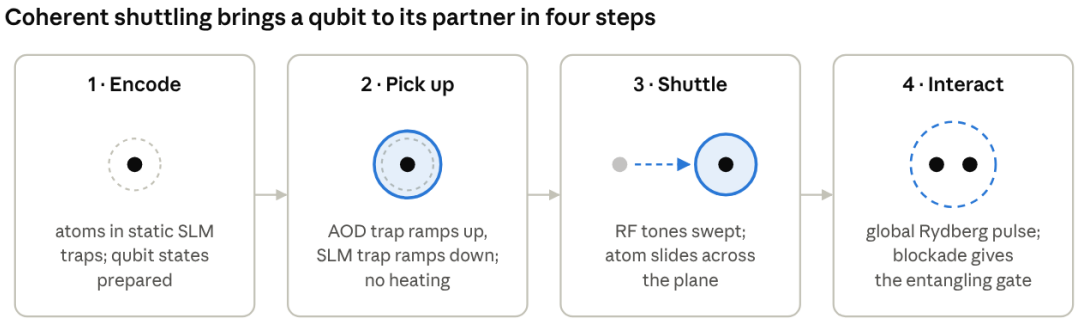
Shuttling by frequency chirping. The controller sweeps the RF tone of a trap, f(t). The acoustic phase front in the TeO_2 crystal changes, steering the beam and sliding the tweezer smoothly. The RF amplitude follows a calibrated look-up table that compensates the AOD's frequency-dependent efficiency, keeping trap depth constant (Optica preprint).

Avoiding acoustic lensing. During a fast chirp, a range of acoustic frequencies spans the crystal aperture and acts as an unwanted cylindrical lens, blurring the focus and heating or dropping the atom. Two fixes (arXiv 2510.09398, DAMOP 2024):

- tailored chirp trajectories, for example cubic chirps, that smooth the phase distortion
- double-pass or multi-AOD "acousto-optic lens" layouts that cancel the focal shift and also allow fast 3D shuttling

Coherent shuttling during a computation (QuEra)

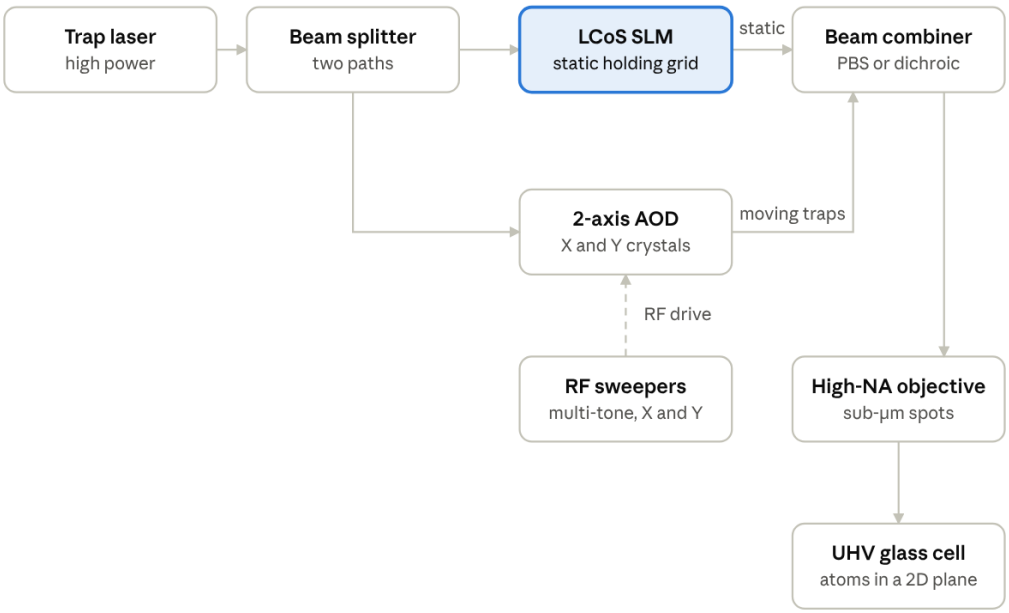
QuEra moves atoms during a calculation, so any qubit can interact directly with any other without SWAP gates.



coherent shuttling · four phases

1. **Initialization.** Atoms sit in the static SLM grid; single-qubit or local entangling operations are applied.
2. **Handshake and pick-up.** AOD tweezers turn on, overlapping selected static traps. The AOD trap is ramped up while the static trap ramps down, transferring the atom without heating it or losing coherence.
3. **Shuttling.** The RF tones driving the AOD crystals are swept, sliding the atoms across the plane.
4. **Gate and drop-off.** Moved atoms arrive next to their partners; a global Rydberg pulse entangles each pair. Atoms then return or move on.

One laser feeds both the static SLM traps and the moving AOD traps



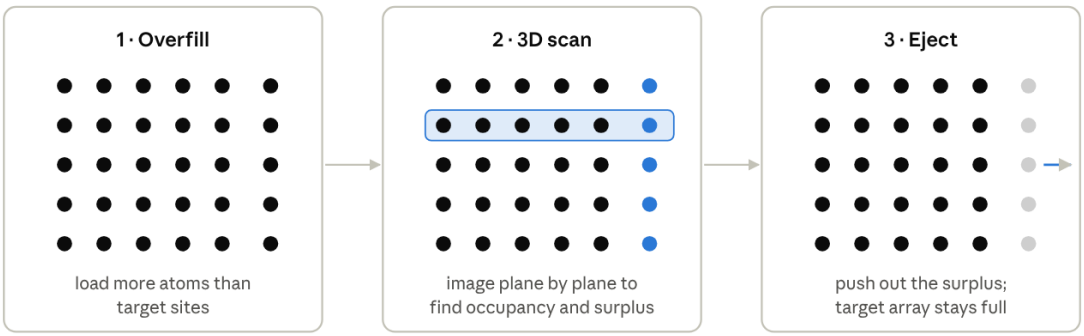
The combiner is added here for completeness; the source sketch shows the two paths meeting at the objective.

QuEra-style optical engine

Fill-and-discard in 3D (as described in the source)

The source describes Atom Computing stacking qubit layers in a 3D lattice, overfilling it and removing the surplus rather than routing atoms between planes. Atom Computing's published 1,180-qubit system is a 2D grid, so treat this as a described technique, not a confirmed product.

Fill and discard: overload the array, image it, eject the surplus



fill-and-discard loading · as described in the source (see review notes)

1. **3D stochastic loading.** A dense cloud loads a 3D lattice at roughly 50% occupancy across all layers.
2. **3D depth imaging.** The camera refocuses, or a multi-plane imaging system is used, to image planes Z_1, Z_2, Z_3 in quick succession; a controller builds a 3D occupancy map.

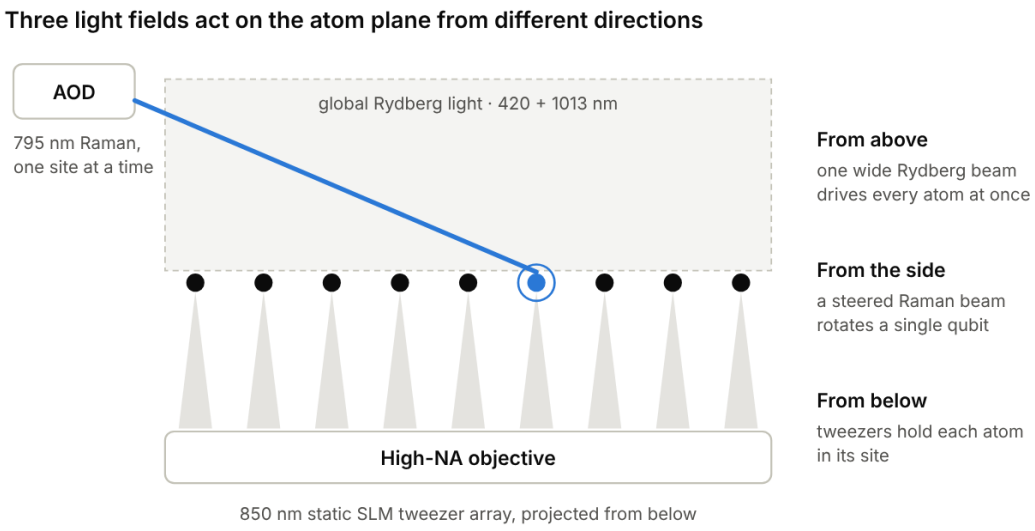
3. **Targeted ejection.** The controller picks a fully filled sub-block and fires focused pulses at the surplus atoms around it, leaving a defect-free 3D array.

The hardware for this would be three orthogonal, retro-reflected lattice beam pairs forming a 3D standing-wave “egg carton”, plus fast variable-focus lenses or electro-optic devices that move the addressing beam's focal plane along Z in microseconds without moving the objective.

Feature	2D shuttling (QuEra)	3D selection (as described)
Array dimension	Dense 2D plane	Layered 3D stack
Primary method	Atoms moved across the array mid-calculation	Surplus atoms ejected to leave a clean core
Core equipment	Two-axis AODs and RF sweepers	3D lattice beams and varifocal Z optics
Main benefit	Connectivity: layout can change mid-algorithm	Volumetric scaling: many qubits in a small footprint

Running the circuit and reading out

In 2D platforms the static tweezer array comes up through the objective, a global Rydberg beam covers the whole array, and an AOD-steered beam targets single qubits.



2D addressing geometry · side view

Readout. One qubit state is pushed out of its trap with resonant light (or shelved into a dark state), then the array is imaged: atoms still present were in the other state. Readout fidelity is typically above 99%.

Mid-circuit measurement and error correction. For error correction, ancilla atoms are read out while data atoms keep their states. Either the ancillas are moved to a separate readout zone, or the data atoms are shelved in states that don't scatter the imaging light. A real-time loop then runs, as described in the source:

1. Execute a layer of entangling gates.
2. Measure the ancilla qubits without disturbing the data qubits.
3. Digitize the camera image to bits.
4. Decode on an FPGA or ASIC, for example with minimum-weight perfect matching or union-find, to locate the error.
5. Feed forward: update a software “Pauli frame” or adjust the next laser pulses.

The decoding has to finish well within the qubits' coherence time, which is why decoders sit right next to the camera's digital output.

8. Component manufacturers and their roles

A neutral-atom computer is assembled from photonics, vacuum and electronics components that mostly come from established scientific-instrument suppliers. The tables below follow the block diagram of Section 3. The first two come from the source articles with their links; the rest list representative suppliers from general industry knowledge, so verify current product lines before relying on them.

LCoS spatial light modulators (static trap arrays)

The SLM displays the computed phase hologram that turns one laser beam into the static tweezer array, so its phase accuracy, flicker, pixel crosstalk and power handling directly set trap-depth uniformity. In neutral-atom platforms (Pasqal, QuEra and academic groups at Harvard, Caltech and Institut d'Optique) it produces static, arbitrary 2D or 3D arrays (LBTEK, arXiv 2608.05010, Columbia).

Manufacturer	Models	Why chosen for atom trapping	Source
Hamamatsu Photonics	X10468, X13138, X15213	Most widely adopted in atomic-physics labs; dielectric-mirror versions handle tens to hundreds of watts; low-phase-fluctuation modes prevent parametric heating of atoms	arXiv 2601.08906
Meadowlark Optics	1920 × 1200 analog series; 1024 × 1024 Ultra-High Power	High diffraction efficiency, water-cooled options, analog phase control that reduces flicker; also sold as an optical-tweezing kit	Meadowlark kit
HOLOEYE Photonics	PLUTO-2.1, LETO	High-reflectance dielectric coatings for ~780 nm, ~800–850 nm and 1064 nm; small pixel pitch for tight trap spacing	MeetOptics
Santec AOC	SLM-200, fast-switching SLM-210	10-bit phase control, high phase reproducibility, fast response for beam shaping	Santec SLM-200

Atom transport and imaging (from the source)

Subsystem	Role in the system	Representative suppliers
Acousto-optic deflectors	Make the moving tweezers; crossed X–Y pairs, TeO_2 crystals for visible and near-IR	AA Opto-Electronic, Gooch & Housego, IntraAction, Brimrose
RF signal generation	Multi-tone RF waveforms; each tone sets one tweezer position	Quantum Machines (OPX), Swabian Instruments, Keysight
RF power amplifiers	Raise the waveform to the watt level that drives the AOD transducer	Mini-Circuits, AA Opto-Electronic, G&H
High-NA objective	Focuses ~1 µm tweezers inside the cell and collects fluorescence	Navitar (Special Optics), Mitutoyo, Nikon
Camera	Images single-atom fluorescence for sorting and readout	Hamamatsu (ORCA sCMOS, qCMOS), Teledyne Princeton Instruments

Other major subsystems (representative suppliers, to verify)

Subsystem	Role in the system	Representative suppliers
Diode and tapered-amplifier lasers	Cooling, repump, imaging, D1 and Rydberg seed lasers	TOPTICA Photonics, MOGLabs, Vexlum, M Squared Lasers
High-power fiber lasers and amplifiers	Trap light at 1064 nm, tens of watts	NKT Photonics (part of Hamamatsu since 2024), Azurlight Systems, IPG Photonics, Precilasers
Frequency conversion	Frequency-doubled blue and UV light (420 nm, 461 nm, 399 nm, 302 nm)	TOPTICA, M Squared, Precilasers; PPKTP and LBO crystals from Raicol, Covesion
Reference cavities and combs	Sub-kHz to sub-Hz linewidths for Rydberg and clock lasers	Stable Laser Systems, Menlo Systems
AOMs and EOMs	Fast switching, frequency shifts, sidebands, intensity servos	Gooch & Housego, AA Opto-Electronic, Brimrose, Exail (iXblue), QUBIG, Jenoptik
Fibers and couplers	Polarization-maintaining delivery from laser table to atoms	Schäfter + Kirchhoff, Thorlabs, NKT Photonics (high-power fibers)
Dichroics, filters, coatings	Combine trap, gate and imaging wavelengths	Semrock (IDEX), Chroma, Layertec

Subsystem	Role in the system	Representative suppliers
Vacuum hardware	Glass cells, ion and getter pumps, alkali dispensers	Infleqtion (glass cells), Precision Glassblowing, Agilent (ion pumps), SAES Getters (NEG pumps, dispensers)
Cryostats (cryogenic platforms)	4 K environment for long atom lifetime, with optical access	Montana Instruments, attocube
Magnetic shielding and sensing	Mu-metal shields and fluxgates for field cancellation	Magnetic Shield Corp., Twinleaf, Bartington
Optical tables and isolation	Rigid, vibration-isolated platform for the science optics	Newport (MKS), TMC (Ametek), Thorlabs
Real-time control	Sequencing, AWG/DDS, feedback	Quantum Machines, Zurich Instruments, Qblox, Spectrum Instrumentation (AWG cards), Sinara/ARTIQ open hardware
Classical computing	Hologram calculation, decoding, hybrid workflows	NVIDIA (GPUs, CUDA-Q), FPGA vendors AMD (Xilinx) and Intel (Altera)

What the supply chain says about the technology

- **The optics are the system.** The qubits are free, but every function, from trapping and cooling to gates and readout, is performed by light. Laser power, linewidth, pointing stability and wavefront quality are the main cost and performance drivers.
- **The SLM is a key performance component.** Trap uniformity across thousands of sites, phase flicker that can heat atoms, and power handling at tens of watts all trace back to the LCoS device and its drive electronics. The device physics covered in the LCoS simulation work, such as alignment layers, pretilt and fringing fields, therefore matters at system level.
- **Hamamatsu sits across the optical chain.** It supplies LCoS SLMs, qCMOS and sCMOS cameras and, through NKT Photonics, high-power fiber lasers, so a single vendor can cover trap generation, light sources and detection.
- **Scaling pushes toward integration.** Pasqal's demonstration of trapping atoms with light from a photonic integrated circuit points to future systems with fewer free-space optics and smaller laser tables.

9. Supporting engineering

This section gathers the remaining material from *Equipment and Operation A*: how the short-wavelength and ultra-stable lasers are made, how the apparatus is protected from vibration, magnetic fields and heat, how cryogenic systems differ, the control loops that run in the background, how modules may be linked, and the hologram algorithm.

Making short wavelengths: frequency doubling

Few lasers emit directly at the short wavelengths atoms need, such as 302 nm for Yb Rydberg excitation or 461 nm for Sr cooling. Systems start from a stable infrared laser, amplify it, and double its frequency in a nonlinear crystal (second-harmonic generation, SHG).

1. **Master laser.** A narrow-linewidth near-IR laser, for example 922 nm for 461 nm output, or 604 nm doubled to 302 nm.
2. **Amplifier.** A tapered or fiber amplifier raises the power from milliwatts to watts.
3. **Nonlinear crystal.** In a temperature-controlled PPKTP or LBO crystal, often inside an enhancement cavity, two photons of wavelength λ combine into one of wavelength $\lambda/2$, with twice the frequency and energy.
4. **Output.** Clean blue or UV light is delivered to the experiment.

Sub-hertz clock lasers

The clock transitions used by Yb (578 nm) and Sr (698 nm) need lasers with sub-hertz linewidth. The laser is locked to an ultra-low-expansion (ULE) glass cavity in its own vacuum housing. Light bounces thousands of times between two very high-reflectivity mirrors, and a feedback loop (Pound–Drever–Hall) holds the laser on the cavity resonance, narrowing it from megahertz to a fraction of a hertz.

Vibration isolation

An optical tweezer is about 1 µm wide, so beam motion of a few hundred nanometers relative to the atoms heats them out of the trap or dephases the qubit. Systems use a layered mounting stack:

Layer	What it does
Vacuum cell, objective and optics, rigidly fixed together	A monolithic mount (stainless steel or Invar) makes the cell, lenses and camera move as one rigid body
Honeycomb-core optical table	Steel skins over a honeycomb core: very stiff, low mass, damps acoustic bending
Pneumatic and active isolation legs	Air springs filter floor noise above about 10 Hz; piezo or magnetic actuators actively cancel slow building motion from about 0.1 to 10 Hz
Laboratory floor	Concrete slab, ideally isolated from the building walls

Magnetic shielding

Stray magnetic fields shift atomic energy levels (the Zeeman effect) and dephase qubits, especially electronic-spin qubits in Rb and Cs. Systems combine passive and active shielding:

- **Mu-metal enclosures.** Nested boxes of high-permeability nickel–iron alloy draw external flux into the metal and route it around the cell.

- **Active compensation coils.** Three orthogonal coil pairs, driven by a feedback loop from fluxgate magnetometers, cancel low-frequency field changes from elevators, power lines and the Earth's field.

Thermal management

High-power beams deposit some heat in every modulator, crystal and mirror, and thermal expansion walks beams off the atoms.

- **Fiber decoupling.** Laser engines, amplifiers and power supplies sit on a separate bench, often in another room, and send light through polarization-maintaining fibers.
- **Active cooling of nearby parts.** AODs, EOMs and other parts that must sit near the atoms are held to about ±0.01 °C with thermoelectric coolers on water-cooled copper heat sinks fed by a vibration-isolated chiller.
- **Enclosures.** The science table sits in a rigid enclosure with gentle laminar, filtered airflow, holding the table to about ±0.1 °C.

Cryogenic systems

At room temperature, background gas collisions limit atom lifetime to tens of seconds. In a 4 K cryostat, residual gas freezes onto the walls and lifetimes become much longer; Pasqal reports about 40× longer lifetimes on its redesigned cryogenic platform. The challenges are optical access through nested, anti-reflection-coated fused-silica or sapphire windows that tolerate thermal contraction, and vibration from the pulse-tube cooler, which is decoupled with flexible copper thermal braids and separate damping structures.

Background control loops

Loop	Sensor	Controller	Actuator	Bandwidth
Magnetic-field cancellation	Fluxgate magnetometers on the shield	Digital PID	Three-axis compensation coils	DC to ~5 kHz
Laser power stabilization	Pick-off photodiodes (~1% of each beam)	Analog or digital servo	AOM RF power	~100 kHz
Position drift correction	sCMOS camera on reference spots	Computer vision on FPGA/CPU	Small offsets to the AOD RF frequencies	~10–50 Hz

Bandwidths are as stated in the source.

Linking modules with photons

A single vacuum module may reach about 1,000–10,000 atoms before optical access, laser power and window size limit it. To go further, modules can be linked through optical fibers:

1. **Atom–photon entanglement.** A laser pulse makes an “interface” atom emit a photon whose polarization is entangled with the atom's state.

2. **Collection.** A high-finesse optical cavity around the interface atom funnels the photon into a single-mode fiber.
3. **Bell-state measurement.** Photons from two modules meet on a beam splitter in front of single-photon detectors. A coincident detection heralds entanglement between the two remote atoms, which can then be used to teleport quantum states between modules.

The hologram algorithm

The SLM can only change the phase of the light, not its amplitude, so the hologram is found by iterating between the SLM plane and the focal plane with Fourier transforms: the Gerchberg–Saxton algorithm. Practical systems use refined versions:

Variant	Modification	Benefit
Weighted GS (WGS)	Per-spot weights updated each iteration from the computed spot brightness	Trap-depth uniformity below 1%
Phase-fixed WGS	Spot phases frozen after a few iterations, then only weights update	Faster, steadier convergence
GPU-accelerated WGS	FFTs run on graphics hardware	Holograms in milliseconds, so layouts can change on the fly

The source also lists an "offset GS" that moves ghost images outside the atom region. The separate document *Gerchberg–Saxton Holograms for Optical-Tweezer Arrays* covers the algorithms, bit depth and the spot-by-spot ("gratings and lenses") method in detail, with simulations.

10. Review notes and items to verify

Both source articles began as AI-chat output that cites mostly vendor pages, videos and forum threads. Company facts in Section 2 were checked against current sources in October 2026. The points below were corrected in this version, or still need checking.

Claim in the sources	Status in this version
Atom Computing uses a 3D optical lattice with fill-and-discard ejection	Not supported. The published 1,180-qubit system is a 35 × 35 (1,225-site) 2D array (Quantum Computing Report). Kept in Section 7 as a described technique, clearly labelled.
planqc stores the qubit in a nuclear spin, in ^{88}Sr / ^{87}Sr	Corrected. ^{88}Sr has no nuclear spin; the Munich work behind planqc uses a fine-structure qubit in metastable ^{88}Sr (arXiv 2401.11054). Only ^{87}Sr supports a nuclear-spin qubit. planqc also works with alkali atoms.
QuEra uses an "ultraviolet" global Rydberg pulse	Corrected. For ^{87}Rb the Rydberg lasers are 420 nm (blue) + 1013 nm (infrared); UV (about 302 nm)

Claim in the sources	Status in this version
	applies to Yb.
QuEra's qubits are “Rydberg states”	Corrected. Digital Rb processors store the qubit in ground hyperfine states and use the Rydberg state only during gates; analog mode is the exception.
Pasqal “HPC cryo-integration”	Now supported. An earlier review note called this an error, but Pasqal reports a redesigned cryogenic platform with about 40× longer atom lifetimes (Quantum Computing Report).
Infleqtion trades as NYSE: INFQ	Confirmed (Quantum Computing Report). Pasqal now also trades publicly, as NASDAQ: PSQL.
Oratomic, Q-Factor, Atom Quantum Labs, OpenQuantum	Not verified; kept as named in the source.
WGS gives “>99% identical” trap depths; “offset GS”	Plausible in simulation (our model reaches 0.3–0.5% non-uniformity at 8 bits), but hardware results depend on SLM flicker, calibration and feedback. “Offset GS” was not checked.
Laser linewidths, control-loop bandwidths, temperature stabilities, decoder time budget	As stated in the source; typical-looking but not checked.
Photonic links “instantly teleport” states	Reworded. A heralded Bell-state measurement creates remote entanglement, which can then be used for teleportation.
Best published two-qubit gate fidelity about 99.5%	Corrected: Evered et al., arXiv 2604.25987 (April 2026), report 99.854(4)% raw and 99.941(3)% loss-postselected CZ fidelity.
Representative suppliers in Section 8 (beyond the source tables)	From general industry knowledge; verify current product lines.