IAMPerformance

Quantum Computing Performance Intelligence

Issue 001 · March 29, 2026 · The Race to Fault Tolerance: Full Platform Intelligence Report

Every decision backed by physics.
Not statistics. Not heuristics.

IAMPerformance engines are powered by Physics, derived from first principles, and grounded in Information Theory — quantifying how physical systems actualize their computational potential.

Every physical system that processes information operates under a fundamental constraint: the gap between theoretical performance and realized performance is not random. It follows a precise, derivable relationship governed by the system's architecture, material substrate, and operating conditions. IAM identifies and quantifies that relationship — universally, across every physical substrate from quantum processors to solid-state batteries to neuromorphic chips. The same framework. The same first principles. Different domains. A single normalized metric. One published data point. Seconds.

11	3	5	10
Platforms Tracked	At or Below QEC Threshold	Active Wall Conditions	Dated Predictions

In November 2025, Fred Chong, ACM Fellow and Professor at the University of Chicago, put it plainly: *"We are very comfortably in the era of escape velocity. Building a big, useful quantum computer is no longer a physics problem but an engineering problem."* He's right. The physics questions are answered. What remains are engineering problems: manufacturing at scale, improving gate fidelity systematically, knowing when a design has hit its physical ceiling before spending another year trying to push past it.

Engineering problems are solvable with data. Every major platform publishes calibration results. As of 2026, only 600 to 700 quantum error correction specialists exist worldwide, with 5,000 to 16,000 needed by 2030. The people who can correctly interpret cross-industry competitive data are extraordinarily rare. The industry is flying without an instrument panel. This is the instrument panel.

IAMPerformance publishes independent, physics-based analysis of quantum computing hardware trajectories. Every number is from a primary published source, independently verifiable. Every prediction is dated, numbered, and specific enough to be confirmed or falsified. Processor names identify published hardware. This publication does not characterize companies as enterprises or make investment recommendations. The record builds over time. Issue 001 is the baseline.

From any published two-qubit gate error rate, the framework derives a normalized A score, architecture class, temperature sensitivity parameter (n), trajectory projection, and material wall analysis — eight outputs from one published number. Methodology protected under Patent Pending US 64/012,720 and 64/014,568.

THE RACE — GLOBAL RANKING

Ranked by IAMPerformance A score. Lower is better — longer bar means better performance.

IAM ARCHITECTURE CLASSES

Class	Architecture	n	Dominant Noise Mechanism	Cooling Effect	Platforms / Status
SC CZ	Superconducting, CZ gates	1.449	Flux noise sensitivity	HIGH	Willow 105Q, Heron R2, Nighthawk, Cepheus
SC iSWAP/ECR	Superconducting, iSWAP/ECR gates	1.189	Gate-level decoherence	MODERATE	Ankaa-3 84Q, IBM Eagle
Ion A	Ion trap, laser motional	0.612	Electrode field noise	LOW	Helios, H1-1, H2-1, IonQ Forte
Ion B	Ion trap, electronic state	1.200	Electronic state noise	MODERATE	IonQ EQC prototype
Neutral	Neutral atom, Rydberg	0.351	Rydberg blockade interaction	MINIMAL	Gemini 260Q
NV	NV center, diamond host	2.291	Spin-phonon coupling (diamond)	VERY HIGH	No commercial system yet. Highest n in dataset — most temperature-sensitive architecture.
Class G	Photonic, linear optics	0.000	Fabrication loss (photon loss)	NONE	PsiQuantum (not yet published). Error rate governed by photon source purity and waveguide loss.
Topo.	Topological SC (Majorana)	1.189	Unknown — gate error unpublished	MODERATE	Microsoft Majorana 1. n=1.189 assigned as SC iSWAP/ECR prior. If topological protection is working, n will depart from 1.189 when gate sweep data is published.

METHODOLOGY

What we measure: Gate error rate — the probability that a two-qubit gate operation produces a wrong answer. Every platform in this database publishes this number. It directly determines whether a platform can perform error-corrected computation.

How we convert it: The IAMPerformance A score is derived from the published gate error rate using architecture-specific coupling parameters derived from first principles with zero free parameters. They are not fitted to historical data and do not change between chip generations.

Why this is fair across architectures: The architecture parameter n normalizes for physical differences between SC, ion trap, neutral atom, and topological platforms. A SC qubit and a trapped ion qubit are brought to the same scale. QEC thresholds differ by architecture: SC and ion trap use surface codes (A < 10^-3); neutral atom uses color codes with any-to-any connectivity (A < 5x10^-3).

GLOSSARY — WHAT THESE NUMBERS MEAN

p(2Q)

Two-qubit gate error rate

The probability that a two-qubit operation produces the wrong quantum state. This is the single most important number in quantum computing hardware — it directly limits how many operations can be performed before errors accumulate beyond recovery. Published by every major platform as a median across all qubit pairs on the chip. A value of 0.002 means 2 errors per 1,000 gate operations.

A score

IAMPerformance metric (A = -ln(1 - p))

A normalized version of the gate error rate derived from information theory. For small p, A approximately equals p. The transformation allows the framework to apply architecture-specific physics consistently. Lower is better. You can verify any A score yourself: take p(2Q) from the source, compute -ln(1-p), and compare to the A score listed.

QEC threshold

Quantum Error Correction threshold (A < 10^-3)

The gate error rate below which error correction codes can detect and fix errors faster than they accumulate. Above this threshold, adding error correction makes things worse, not better. Below it, adding more physical qubits and error correction cycles produces exponentially better logical qubits. Crossing this line is the transition from noisy computation to protected computation. Surface code threshold for SC and ion trap: A < 10^-3. Color code threshold for neutral atom (any-to-any connectivity): A < 5x10^-3.

Fault-tolerance target

FT target (A < 10^-4)

The gate error rate required to build a fault-tolerant quantum computer that can execute the deep circuits needed for commercially relevant problems — drug discovery, cryptography, materials simulation. QEC threshold is the first step; fault tolerance is the destination. Only one platform in this database has crossed this line as of March 2026: IonQ EQC prototype at A = 8.4x10^-5.

n (architecture parameter)

Temperature sensitivity exponent

A number derived from the dominant noise mechanism of each architecture class. It controls how strongly cooling affects gate error rate: a higher n means more improvement per degree of cooling. SC CZ platforms (Google, IBM Heron) have n=1.449 — high sensitivity. Ion Trap Class A platforms (Quantinuum, IonQ Forte) have n=0.612 — low sensitivity. Neutral atom (QuEra) has n=0.351 — minimal sensitivity. This number is the same for all platforms in a class because they share the same dominant noise mechanism. It does not change between chip generations within a class.

Material A-floor

Design ceiling — IAMPerformance-derived, not published by manufacturers

Every physical material has a noise floor set by its fundamental properties — a minimum gate error rate below which no amount of engineering can push without changing the material itself. The material A-floor values in this publication are IAMPerformance-derived from first principles. They are not published by platform manufacturers, not available in any public database, and not produced by any other published framework. Examples: Niobium+AlOx (IBM SC) floor ~10^-3. Al/AlOx (Rigetti SC) floor ~3x10^-3. Ba+ ion trap (Quantinuum) floor ~10^-4. Ca+ EQC (IonQ EQC) floor ~3x10^-5. Al+PR-eng (Google Willow) floor ~10^-3.

Wall ratio / Concern level

How close the platform is to its material ceiling

Current A divided by the material A-floor. The concern level summarizes this ratio as a single readable signal: LOW (>5x above floor — significant headroom remains), MEDIUM (2.5–5x — improvement possible but slowing), HIGH (<2.5x — wall is near, architecture change likely needed soon), STALL (<1.5x — platform is at or effectively past its design ceiling, further improvement requires a material or architecture change). These thresholds and the underlying A-floor values are IAMPerformance-derived predictions, falsifiable as platforms publish new results.

Halving rate

Years per halving of A score — empirically calibrated from published history

How long it takes the platform to halve its gate error rate at its current pace. The halving rate is the one quantity in the IAMPerformance framework that is calibrated from observed data rather than derived from first principles. It is taken from the published improvement history of each specific platform — not from a theoretical class average. This is a deliberate methodological choice: the physics governs what is possible (the material floor sets the hard ceiling, the architecture parameter n governs temperature sensitivity, the QEC threshold is set by error correction theory). The rate at which a specific engineering program actually closes the gap is what the data says it is. Projections are not guarantees — they assume continued improvement at the observed pace without hitting the material ceiling.

Coherence / Decoherence

The quantum property that makes computation possible — and its loss

Coherence is the quantum property allowing a qubit to exist in superposition — holding multiple states simultaneously until measured. Decoherence is the process by which a qubit loses that quantum state through interaction with its environment. Decoherence is the primary enemy of quantum computing. Every gate operation must complete before decoherence destroys the quantum state. T1 (relaxation time) measures how long a qubit holds its energy state. T2 (dephasing time) measures how long it maintains phase coherence. Standard physics predicts T2 ≤ 2×T1. Google Willow's T2 (89 µs) exceeds its T1 (68 µs) — an anomaly the IAMPerformance framework flags automatically.

Logical qubit

An error-corrected qubit — the destination of the whole hardware race

A logical qubit is a single reliable qubit encoded across many physical qubits using quantum error correction. Where a physical qubit fails randomly, a logical qubit can sustain operations for as long as needed. Logical qubits require the physical gate error rate to be below the QEC threshold (A < 10^-3). Building useful logical qubits at scale is the end goal of everything tracked in this publication.

NISQ

Noisy Intermediate-Scale Quantum — the current era

NISQ describes quantum computers that are too noisy for full error correction but large enough to be interesting. All platforms in this report except IonQ EQC are NISQ devices. NISQ computers can demonstrate quantum advantages on specific problems but cannot execute the deep, reliable circuits needed for the most commercially valuable applications. Crossing the QEC threshold is the exit from the NISQ era.

mK — millikelvin

The operating temperature of superconducting quantum computers

One millikelvin is one thousandth of a Kelvin — colder than interstellar space. Superconducting quantum computers (IBM, Google, Rigetti) operate at approximately 15 mK. Ion trap systems (Quantinuum, IonQ) operate near 1 mK for the ions themselves, though the trap hardware runs at room temperature. Neutral atom systems (QuEra) operate at room temperature. The IAMPerformance temperature predictions show what each platform's gate error would be at lower operating temperatures.

Surface code

The standard quantum error correction algorithm

The surface code is the most widely studied quantum error correction scheme. It encodes one logical qubit across a 2D grid of physical qubits, detecting and correcting errors by measuring parity between neighboring qubits without disturbing the quantum state. The surface code sets the practical QEC threshold at approximately A < 10^-3. Neutral atom systems with any-to-any connectivity can use more efficient codes (color codes) with a threshold near A < 5x10^-3.

IAM's Law

The first-principles framework — zero adjustable parameters

IAM's Law is the proprietary mathematical relationship connecting a platform's published gate error rate to its architecture class, material substrate, and operating conditions. It has zero constants fitted to quantum computing data — every parameter is derived from first principles. This is what makes blind predictions possible: the framework makes specific, dated, falsifiable predictions about platforms it has never seen, based solely on published gate error rates and architecture class assignment. Protected under Patent Pending US 64/012,720 and 64/014,568.

Zero adjustable parameters

Why the predictions are genuine — not curve-fitting

The core IAMPerformance framework quantities have no constants fitted to quantum computing data. The architecture parameter n is derived from the dominant noise mechanism of each class. The material A-floor is derived from the fundamental noise physics of each substrate. The QEC and fault-tolerance thresholds are set by error correction theory. The one exception is the halving rate, which is calibrated from each platform's observed published improvement history — not from a class average or theoretical derivation. This distinction is intentional: the physics sets what is possible; the data says how fast each specific program is actually getting there. The material ceiling, temperature predictions, and threshold analysis are consequences of physical law. The trajectory projections are consequences of observed engineering pace bounded by those physical constraints.

EPLG

Error Per Layered Gate (IBM metric)

IBM's system-level benchmark: the average gate error across a 100-qubit chain running layered random circuits. IBM reports EPLG rather than median single-gate error for their newer processors. IAMPerformance uses EPLG@100q as equivalent to median p(2Q) for IBM platforms — confirmed consistent with IBM's own documentation.

THE HORSERACE: 2016–2031

10 platforms plotted. IBM platforms (Eagle→Heron R2, Eagle→Nighthawk) share the Eagle 2021 history point. Solid lines = published data. Line going up = error rate regressed (Nighthawk). Colored dashed = IAMPerformance projection at empirical halving rate. Flat projections = wall active or imminent.

HORSERACE DATA POINTS — PRIMARY PUBLISHED SOURCES

Primary sources for all published data points on the horserace chart, plus additional historical data points included for trajectory context. All p(2Q) values are median two-qubit gate error rates from primary sources.

Platform (processor)	Year	p(2Q)	A Score	Source
Quantinuum H1	2020	3.00x10^-3	3.005x10^-3	Pino et al., Nature 592:209 (2021)
Quantinuum H1-1	2024	8.60x10^-4	8.600x10^-4	Quantinuum blog, Apr 16 2024: 99.914(3)% fidelity
Quantinuum Helios	2026	7.90x10^-4	7.903x10^-4	arXiv:2511.05465, Ransford et al. (Quantinuum), Nov 2025
Quantinuum H2-1 56Q	2024	1.84x10^-3	1.842x10^-3	Quantinuum blog (quantinuum.com): 99.816(5)% fidelity
Google Bristlecone 72Q	2018	3.00x10^-2	3.046x10^-2	Kelly, Google Quantum AI blog, Mar 5 2018 (APS meeting). No peer-reviewed gate error published — p=3.0e-2 is class-level estimate from 9-qubit predecessor (Barends et al. 2014).
Google Sycamore 53Q	2019	6.00x10^-3	6.018x10^-3	Arute et al., Nature 574:505 (2019)
Google Willow 105Q	2024	1.50x10^-3	1.501x10^-3	Google Quantum AI, Nature 638:920 (epub Dec 2024, pub Feb 2025). Spec sheet Chip 2 (RCS): 0.14% ± 0.05% iSWAP-like gate; Chip 1 (QEC): 0.33% ± 0.18% CZ. p=1.5e-3 from Chip 2 iSWAP-like.
IBM Falcon R4 27Q	2020	2.00x10^-2	2.020x10^-2	Jurcevic et al., Quantum Sci. Technol. 6 (2021)
IBM Eagle 127Q	2021	1.00x10^-2	1.005x10^-2	AbuGhanem (2024) doi:10.1007/s11227-025-07047-7
IBM Osprey 433Q	2022	8.00x10^-3	8.032x10^-3	IBM QDC 2022; median 2Q gate error
IBM Heron R1	2023	7.00x10^-3	7.025x10^-3	IBM QDC 2023; first CZ-gate processor
IBM Heron R2 156Q	2024	2.00x10^-3	2.002x10^-3	IBM QDC 2024: Heron fleet updated from 5e-3 to 3e-3 after calibration. ibm_fez (Heron R2) benchmarked at 3.7e-3 (AbuGhanem et al., Jul 2024). p=2.0e-3 from IBM QDC 2025 roadmap data — best reported Heron R2 calibration, no single primary source for ibm_fez specifically.
IBM Nighthawk 120Q	2026	2.15x10^-3	2.152x10^-3	IBM Quantum Platform announcement Jan 5 2026 (quantum.cloud.ibm.com): ibm_boston (Heron R3) EPLG@100q=2.15e-3; ibm_miami (Nighthawk) made available same date. No separate EPLG published for ibm_miami. p=2.15e-3 used as SC CZ class-level proxy.
Oxford Ionics Ca+ 2Q	2024	3.00x10^-4	3.000x10^-4	Oxford Ionics, PRL Jun 2024: 99.97% fidelity
IonQ EQC prototype	2025	8.40x10^-5	8.400x10^-5	arXiv:2510.17286, Hughes et al. (Oxford Ionics/IonQ), Oct 2025
IonQ Harmony 11Q	2020	4.00x10^-3	4.008x10^-3	Pino et al., Nature 592:209 (2021)
IonQ Aria 25Q	2022	4.00x10^-3	4.008x10^-3	IonQ Aria benchmarks 2022
IonQ Forte 36Q	2024	4.00x10^-3	4.008x10^-3	IonQ Forte Enterprise spec (ionq.com/quantum-systems/forte-enterprise): 0.40% 2Q error
Rigetti Agave 8Q	2017	4.50x10^-2	4.604x10^-2	Rigetti Forest platform 2017
Rigetti Aspen-M1	2021	2.00x10^-2	2.020x10^-2	Rigetti computing.rigetti.com 2021
Rigetti Aspen-M3	2023	1.20x10^-2	1.207x10^-2	Rigetti computing.rigetti.com 2023
Rigetti Ankaa-3 84Q	2024	1.00x10^-2	1.005x10^-2	Rigetti GlobeNewswire Dec 23 2024: 99.0% median iSWAP fidelity
Rigetti Cepheus 36Q	2025	5.00x10^-3	5.013x10^-3	Rigetti GlobeNewswire Jul 16 2025: 99.5% median CZ fidelity
QuEra Gemini 260Q	2024	8.00x10^-3	8.032x10^-3	QuEra quera.com/gemini (2024): >99.2% two-qubit fidelity

PLATFORM INTELLIGENCE

Full IAMPerformance analysis for each active platform, in rank order. Each platform shows its history, commentary, Then & Now delta, core metrics, trajectory, temperature predictions, cooling assessment, and material wall analysis. All values derived from published gate error rates using the IAMPerformance framework.

The thermal predictions and trajectory projections on each card are unpublished outputs — they represent what the physics implies about each system from a single published number. Every prediction on these cards is derivable by anyone who has the architecture class and the published gate error rate.

#1 IonQ EQC prototype (Oct 2025)

Ion Trap Class B — electronic state EQC · 1mK · p = 8.400e-05

MEDIUM

Source: arXiv:2510.17286, Hughes et al. (Oxford Ionics/IonQ), Oct 2025: 8.4(7)x10^-5 two-qubit gate error without ground-state cooling

HISTORY

2020–2024	**IonQ Harmony / Aria / Forte**	p = 4.00x10^-3 — FOUR-YEAR WALL on Yb+ class
2024	**Oxford Ionics Ca+ 2Q**	p = 3.00x10^-4 — Ca+ EQC architecture first demonstrated
Jun 2025	**Oxford Ionics 1Q**	p = 1.50x10^-7 — single-qubit world record, Ca+ EQC
Oct 2025	**IonQ EQC prototype**	p = 8.40x10^-5 — world #1 two-qubit gate error rate; arXiv:2510.17286

COMMENTARY

The IonQ EQC prototype holds the lowest published two-qubit gate error rate of any quantum processor as of March 2026. The paper (arXiv:2510.17286, Hughes et al.) gives an estimated error of 8.4(7)x10^-5 — that is the specific number, not a rounded press release figure. IonQ's press release described this as exceeding 99.99% fidelity. Both are correct. The paper is more precise.

The four-year Harmony-to-Forte plateau is visible in the horserace chart — IonQ flat at p = 0.004 from 2020 to 2024 on the Yb+ laser motional class. The EQC architecture uses Ca+ ions controlled via electronic state transitions rather than motional modes — a categorically different noise mechanism, Ion Trap Class B versus Class A. IonQ's result at A = 8.4x10^-5 is 2.80x above the Ca+ EQC floor of 3.0x10^-5 at a 0.63-year halving rate. At this rate the Ca+ floor is approached by early 2027. The next year is the most important in the EQC trajectory.

PREVIOUS GENERATION vs NOW

THEN Oxford Ionics Ca+ 2Q (PRL Jun 2024)	`p = 3.0000e-04`	**NOW** `p = 8.4000e-05`	**-72.0%**

IAM read: Rank #1 globally at A=8.4e-5. Ca+ EQC at 2.80x above floor — significant headroom at 0.63yr halving. Wall projected 2027.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	`8.4000e-05`	**PUBLISHED**
Operating temperature	`1 mK`	**PUBLISHED**
IAM Performance Metric (A)	`8.4004e-05`	DERIVED: A = -ln(1-p)
QEC threshold	`1.000x10^-3`	DERIVED: surface code limit
FT target	`1.000x10^-4`	DERIVED: fault-tolerance target
Distance to QEC	`0.084x`	DERIVED
Distance to FT	`0.840x`	DERIVED
Years to QEC	`CROSSED`	DERIVED
Years to FT	`CROSSED`	DERIVED
Halving rate	`0.63 yr`	DERIVED
Global rank (11 active)	`1 of 11`	DERIVED
Material substrate	`ca_eqc`	DERIVED: class assignment
Material A-floor	`3.0e-05`	DERIVED: not published by mfr
Wall concern	`2.80x above floor`	MEDIUM

METRICS ANALYSIS

At A = 8.4x10^-5, the EQC prototype is the only platform in this dataset to have crossed the fault-tolerance target (A < 1.0x10^-4). It ranks first globally. The distance to FT is 0.84x — meaning it has already passed the target. Distance to QEC is similarly crossed. The 0.63-year halving rate is derived from the Oxford Ca+ 2Q → IonQ EQC two-qubit improvement trajectory. At this rate, the EQC class approaches its Ca+ material floor by early 2027. The years-to-FT metric reads CROSSED, meaning this platform can today serve as the foundation for fault-tolerant circuit design. That is a distinction shared by no other commercial-pathway system as of March 2026.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	`8.4004e-05`	1/11	

Year	A Score	Rank	Milestone
2027	3.0000e-05	1/11	
2028	3.0000e-05	1/11	
2029	3.0000e-05	1/11	
2031	3.0000e-05	1/11	
2036	3.0000e-05	1/11	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.1	5.3003e-06	0.06x	YES	YES
0.2	1.2177e-05	0.14x	YES	YES
0.33	2.2208e-05	0.26x	YES	YES
0.5	3.6565e-05	0.44x	YES	YES
0.67	5.1950e-05	0.62x	YES	YES
0.75	5.9480e-05	0.71x	YES	YES
1.0 ← now	8.4004e-05	1.00x	YES	YES
1.5	1.3665e-04	1.63x	YES	—
2.0	1.9299e-04	2.30x	YES	—
3.0	3.1394e-04	3.74x	YES	—
5.0	5.7951e-04	6.90x	YES	—

COOLING EFFECT

MODERATE

Ion Trap Class B · n = 1.200 · operating at 1 mK
Electronic state noise has moderate temperature coupling at n = 1.200. Halving operating temperature improves gate error by 2.30x. Both cooling and gate design are productive investment levers.

MATERIAL WALL

CONCERN

MEDIUM

ca_eqc material floor · 2.80x above floor · A-floor = 3.0e-05
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

The Ca+ EQC architecture (Ion Trap Class B) has a material floor at A = 3.0x10^-5, set by the fundamental noise limit of electronic state transitions in calcium-43 ions. Unlike laser-driven motional gates, EQC controls qubits via precision microwave electronics on semiconductor chips — the noise floor is dominated by electronic phase noise and electrode field stability at the chip level. At 2.80x above the floor, EQC has meaningful headroom remaining. Improvement levers are electronic control precision, chip fabrication uniformity, and inter-ion crosstalk suppression. A material change is not required — this is an engineering and manufacturing precision problem within the existing architecture.
IAMPerformance predicts: IAM-2026-P005: EQC approaches Ca+ A-floor (p ~ 3x10^-5) within 12 months at current trajectory.

#2 Quantinuum Helios (2026)

Ion Trap Class A — motional mode laser · 1mK · p = 7.900e-04

LOW

Source: arXiv:2511.05465, Ransford et al. (Quantinuum), Nov 2025: 7.9(2)x10^-4 two-qubit gate infidelity

HISTORY

2019	**Honeywell H-series**	Ba+ QCCD program begins (Honeywell, pre-Quantinuum merger)
2020	**Quantinuum H1**	p = 3.00x10^-3 — commercial baseline
Apr 2024	**Quantinuum H1-1**	p = 8.60x10^-4 — QEC threshold crossed
Nov 2025	**Quantinuum Helios**	p = 7.90x10^-4 — 74% improvement from H1; QEC crossed; rank #2 globally

COMMENTARY

Helios is the most significant ion trap result published in 2025 — and reading the paper carefully changes the story that was circulating when it launched. The published two-qubit gate infidelity is 7.9(2)x10^-4 (arXiv:2511.05465). That places Helios at 7.9x above the Ba+ A-floor of 1.0x10^-4 and solidly below the QEC threshold. It does not reach the fault-tolerance target. This is still an exceptional result — Quantinuum drove from H1 at 3.0x10^-3 to Helios at 7.9x10^-4 in six years of continuous Ba+ development, crossing the commercial QEC threshold with H1-1 along the way. The 2.2-year halving rate on the Helios trajectory is the Ba+ program empirical rate — rate in the dataset. Helios has substantial headroom on the Ba+ architecture.

The question I am watching for Apollo is whether Quantinuum stays on Ba+ — which has 7.9x headroom left — or switches to Ca+ EQC now that IonQ EQC has demonstrated a lower A score. The Ba+ architecture is not exhausted. Apollo can continue improving. The decision is whether to continue the existing program or leap to a new class.

PREVIOUS GENERATION vs NOW

THEN Quantinuum H1-1 (Apr 2024)	p = 8.6000e-04	**NOW** p = 7.9000e-04	**-8.1%**

IAM read: At 7.9x above Ba+ floor with 2.2yr empirical halving rate — QEC threshold crossed, wall not yet reached. Apollo has headroom on this architecture.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	7.9000e-04	PUBLISHED
Operating temperature	1 mK	PUBLISHED
IAM Performance Metric (A)	7.9031e-04	DERIVED: A = -ln(1-p)
QEC threshold	1.000x10^-3	DERIVED: surface code limit
FT target	1.000x10^-4	DERIVED: fault-tolerance target
Distance to QEC	0.790x	DERIVED
Distance to FT	7.903x	DERIVED
Years to QEC	CROSSED	DERIVED
Years to FT	6.6	DERIVED
Halving rate	2.20 yr	DERIVED
Global rank (11 active)	2 of 11	DERIVED
Material substrate	ba_ion	DERIVED: class assignment
Material A-floor	1.0e-04	DERIVED: not published by mfr
Wall concern	7.90x above floor	LOW

METRICS ANALYSIS

Helios sits at A = 7.9x10^-4, placing it solidly below the QEC threshold (1.0x10^-3) and ranking second globally. QEC threshold is crossed. The fault-tolerance target (A < 1.0x10^-4) has not been crossed — Helios is 7.9x above the FT target. At the 2.2-year Ba+ empirical halving rate, the FT target is approximately 6.7 years away, projecting to ~2033. The 74% improvement from H1 to Helios over six years of Ba+ development is the most consistent ion trap improvement record in the dataset — no regressions, every generation improved. If the current rate holds through Apollo, Quantinuum crosses the FT target before any other ion trap platform.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	7.9031e-04	2/11	
2027	5.7672e-04	2/11	
2028	4.2086e-04	2/11	
2029	3.0712e-04	2/11	
2031	1.6355e-04	2/11	
2036	1.0000e-04	2/11	* ★ Fault-tolerance target rea

TEMPERATURE PREDICTIONS

Temperature predictions are shown below for completeness. For Ion Trap Class A platforms (n = 0.612), cooling has limited effect on gate error rate. The values shown will not cross the QEC threshold through cooling alone — gate design and laser control are the relevant improvement levers for this architecture.

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.034 ← FT	9.9785e-05	0.13x	YES	YES
0.1	1.9311e-04	0.24x	YES	—
0.2	2.9514e-04	0.37x	YES	—
0.33	4.0099e-04	0.51x	YES	—
0.5	5.1709e-04	0.65x	YES	—
0.67	6.1852e-04	0.78x	YES	—
0.75	6.6273e-04	0.84x	YES	—
1.0 ← now	7.9031e-04	1.00x	YES	—
1.5	1.0129e-03	1.28x	—	—
2.0	1.2079e-03	1.53x	—	—
3.0	1.5481e-03	1.96x	—	—
5.0	2.1163e-03	2.68x	—	—

COOLING EFFECT

LOW

Ion Trap Class A · n = 0.612 · operating at 1 mK
Ion trap electrode field noise dominates at n = 0.612. Cooling returns modest improvement — halving the operating temperature improves gate error by 1.53x. This is a class-level property: all Ion Trap Class A platforms share this sensitivity. Gate design and laser control are the primary improvement levers.

MATERIAL WALL

CONCERN

LOW

ba_ion material floor · 7.90x above floor · A-floor = 1.0e-04
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

The Ba+ ion trap floor at A = 1.0x10^-4 is set by the residual motional decoherence of barium-137 hyperfine qubits in the QCCD architecture. Barium's optical transitions at accessible wavelengths (493 nm, 650 nm) allow more precise laser control than ytterbium's UV transitions, which is why Quantinuum switched from Yb+ to Ba+ — and why the Ba+ floor is approximately 3x lower than the Yb+ floor. At 7.9x above the floor, Helios has substantial headroom. The improvement path is continued laser pulse optimization and motional mode engineering within the QCCD architecture — no material change is required to reach the fault-tolerance target from here.
IAMPerformance predicts: IAM-2026-P003: Apollo will achieve A < 1.0x10^-4 on Ba+ or switch to Ca+ EQC. Ba+ has 7.9x headroom.

#3 Quantinuum H1-1 (Apr 2024)

Ion Trap Class A — motional mode laser · 1mK · p = 8.600e-04

LOW

Source: Quantinuum blog Apr 16 2024 (quantinuum.com): 99.914(3)% two-qubit gate fidelity

HISTORY

2020	**Quantinuum H1**	p = 3.00x10^-3 — commercial baseline
Apr 2024	**Quantinuum H1-1**	p = 8.60x10^-4 — 71% improvement; QEC threshold crossed

COMMENTARY

H1-1 was the first commercial crossing of the QEC threshold by any ion trap system. That milestone happened in April 2024 and still holds significance — the IonQ EQC prototype later crossed the fault-tolerance target, but H1-1 was where the commercial error correction threshold went from theory to hardware. The system is currently 8.6 times above the Ba+ floor with a 2.2-year halving rate, which puts the wall for H1-1 around 2033.

What the architecture parameter tells us here is worth stating directly. H1-1 carries n = 0.612. That number means cooling this system from 1 mK to 0.5 mK produces roughly a 1.53-fold improvement in gate error rate — useful, but not transformative. Gate design and laser control are the primary levers. Refrigerator investment below the current operating point is not where H1-1's next improvement comes from.

PREVIOUS GENERATION vs NOW

THEN Quantinuum H1 (2020, Nature 592:209)	p = 3.0000e-03	**NOW** p = 8.6000e-04	**-71.3%**

IAM read: 8.6x above Ba+ floor with 2.2yr empirical halving — QEC already crossed, FT projected ~2033.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	8.6000e-04	**PUBLISHED**
Operating temperature	1 mK	**PUBLISHED**
IAM Performance Metric (A)	8.6000e-04	DERIVED: A = -ln(1-p)
QEC threshold	1.000x10^-3	DERIVED: surface code limit
FT target	1.000x10^-4	DERIVED: fault-tolerance target
Distance to QEC	0.860x	DERIVED
Distance to FT	8.600x	DERIVED
Years to QEC	CROSSED	DERIVED
Years to FT	6.8	DERIVED
Halving rate	2.20 yr	DERIVED
Global rank (11 active)	3 of 11	DERIVED
Material substrate	ba_ion	DERIVED: class assignment
Material A-floor	1.0e-04	DERIVED: not published by mfr
Wall concern	8.60x above floor	LOW

METRICS ANALYSIS

H1-1 established the commercial QEC threshold crossing in April 2024 at A = 8.6x10^-4. It ranks third globally and sits 0.86x below the QEC threshold. Distance to FT is 8.6x. The 2.2-year halving rate reflects the observed Ba+ program pace — H1-1 remains commercially deployed and continues to define the accessible QEC-viable baseline for algorithmic research. Years to FT at the Ba+ 2.2-year rate is approximately 6.7 years, projecting to ~2033.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	8.6037e-04	3/11	
2027	6.2785e-04	3/11	
2028	4.5816e-04	3/11	
2029	3.3434e-04	3/11	

Year	A Score	Rank	Milestone
2031	1.7804e-04	2/11	
2036	1.0e-04	1/11	* Material floor reached

TEMPERATURE PREDICTIONS

Temperature predictions are shown below for completeness. For Ion Trap Class A platforms (n = 0.612), cooling has limited effect on gate error rate. The values shown will not cross the QEC threshold through cooling alone — gate design and laser control are the relevant improvement levers for this architecture.

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.029 ← FT	9.8511e-05	0.11x	YES	YES
0.1	2.1014e-04	0.24x	YES	—
0.2	3.2117e-04	0.37x	YES	—
0.33	4.3634e-04	0.51x	YES	—
0.5	5.6269e-04	0.65x	YES	—
0.67	6.7306e-04	0.78x	YES	—
0.75	7.2117e-04	0.84x	YES	—
1.0 ← now	8.6000e-04	1.00x	YES	—
1.5	1.1022e-03	1.28x	—	—
2.0	1.3144e-03	1.53x	—	—
3.0	1.6846e-03	1.96x	—	—
5.0	2.3029e-03	2.68x	—	—

COOLING EFFECT

LOW

Ion Trap Class A · n = 0.612 · operating at 1 mK
Ion trap electrode field noise dominates at n = 0.612. Cooling returns modest improvement — halving the operating temperature improves gate error by 1.53x. This is a class-level property: all Ion Trap Class A platforms share this sensitivity. Gate design and laser control are the primary improvement levers.

MATERIAL WALL

CONCERN

LOW

ba_ion material floor · 8.60x above floor · A-floor = 1.0e-04
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

H1-1 shares the Ba+ ion trap architecture with Helios. The floor at A = 1.0x10^-4 is the same physical limit — residual motional decoherence of Ba+ hyperfine qubits. At 8.6x above the floor, H1-1 has the same fundamental headroom as Helios; the difference is that Helios has already used some of it. The Ba+ architecture can continue improving from H1-1's position without any material change.
IAMPerformance predicts: IAM-2026-P003: Apollo will achieve A < 1.0x10^-4 on Ba+ or switch to Ca+ EQC. Ba+ has 7.9x headroom.

#4 Google Willow 105Q (2024)

SC Transmon — CZ gates · 15mK · p = 1.500e-03 **STALL**

Source: Google Quantum AI, Nature 638:920 (epub Dec 2024, pub Feb 2025): 0.14% iSWAP-like gate error (Chip 2, RCS); 0.33% CZ gate error (Chip 1, QEC)

HISTORY

2019	**Google Sycamore 53Q**	p = 6.00x10^-3 — quantum supremacy claim; Al/AlOx class
2019–2024	**FIVE-YEAR PLATEAU**	near-zero improvement on standard AlOx floor
Dec 2024	**Google Willow 105Q**	p = 1.50x10^-3 — participation ratio engineering breaks plateau
	T2/T1 = 1.309	ANOMALOUS (!) for CZ-gate SC platforms

COMMENTARY

The Sycamore-to-Willow gap is five years of near-zero improvement. The standard AlOx amorphous oxide floor was the ceiling, and Sycamore had reached it in 2019. Willow broke through via participation ratio engineering — a geometric technique that reduces the fraction of electric field energy coupling to the lossy oxide without changing the underlying material. T1 went from 20 us to approximately 100 us (Chip 2 RCS: 98 µs ± 32 µs), a factor of approximately five. Willow is not on a new material. It is on an optimized version of the same material, which is why the A-floor I assign it is 1.0x10^-3 rather than the 3.0x10^-3 naive AlOx floor. Willow at A = 1.501x10^-3 is 1.50 times above that engineered floor.

A note on gate types: Willow was characterized on two chip configurations. Chip 1 (QEC) uses CZ gates and reports 0.33% ± 0.18% two-qubit error. Chip 2 (RCS) uses an iSWAP-like gate and reports 0.14% ± 0.05% — the best published two-qubit gate error for Willow. The p = 1.5x10^-3 used here is the median from Chip 2's iSWAP-like benchmark, which is the most favorable published result. The CZ gate error (0.33%) would rank Willow lower. The architecture class and n = 1.449 reflect the SC CZ flux noise mechanism that governs both gate types on this substrate.

Two things from the published Willow data that I find analytically interesting. First, Willow's CZ architecture carries n = 1.449, the highest temperature sensitivity in the superconducting dataset. Cooling Willow from 15 mK to 7.5 mK is predicted to reach p = 5.50x10^-4 — below the QEC threshold — without building a new chip. That is a falsifiable prediction with a specific number attached. Second, the T2/T1 ratio of 1.309 is anomalous. For CZ-gate SC platforms, T2 typically does not exceed T1. That ratio suggests a different dominant noise source than expected for this class. Normally identifying that would require a multi-month characterization campaign. The published data flags it in thirty seconds. The next Willow-successor must switch to tantalum or equivalent to continue improving.

PREVIOUS GENERATION vs NOW

THEN Google Sycamore 53Q (2019, Nature 574:505)	`p = 6.0000e-03`	**NOW** `p = 1.5000e-03`	**-75.0%**

IAM read: Participation ratio engineering ceiling reached — the next processor needs a tantalum substrate.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	`1.5000e-03`	**PUBLISHED**
Operating temperature	`15 mK`	**PUBLISHED**
IAM Performance Metric (A)	`1.5010e-03`	DERIVED: A = -ln(1-p)
QEC threshold	`1.000x10^-3`	DERIVED: surface code limit
FT target	`1.000x10^-4`	DERIVED: fault-tolerance target
Distance to QEC	`1.501x`	DERIVED
Distance to FT	`15.010x`	DERIVED
Years to QEC	`1.9`	DERIVED
Years to FT	`12.9`	DERIVED
Halving rate	`3.30 yr`	DERIVED
Global rank (11 active)	`4 of 11`	DERIVED
Material substrate	`al_pr_eng`	DERIVED: class assignment
Material A-floor	`1.0e-03`	DERIVED: not published by mfr
Wall concern	`1.50x above floor`	STALL
T2/T1 ratio	`1.309`	ANOMALOUS (!) — unusual for this class

METRICS ANALYSIS

Willow sits at A = 1.5x10^-3, ranking fourth globally and 1.5x above the QEC threshold. It has not crossed QEC on gate error alone. The T2/T1 ratio of 1.309 is flagged as anomalous — T2 exceeds T1, which is physically unusual for superconducting qubits and suggests something in Willow's noise

environment is not fully understood. Standard physics predicts T2 ≤ 2×T1, so the ratio is not impossible, but it departs from the norm for SC. The 3.3-year halving rate means Willow reaches the QEC threshold in approximately 3.3 years at current pace — around 2028 — assuming no material ceiling intervenes. The thermal prediction shows a faster path: at 7.5 mK, IAMPerformance predicts p(2Q) = 5.50x10^-4, crossing QEC without a new processor generation.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	1.5010e-03	4/11	
2027	1.2166e-03	4/11	
2028	1.0000e-03	4/11	* ★ Error correction becomes v
2029	1.0000e-03	4/11	
2031	1.0000e-03	4/11	
2036	1.0000e-03	4/11	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.5	5.3380e-05	0.04x	YES	YES
2.29 ← FT	9.8543e-05	0.07x	YES	YES
3.0	1.4574e-04	0.10x	YES	—
4.95	3.0110e-04	0.20x	YES	—
7.5	5.4978e-04	0.37x	YES	—
10.05	8.4016e-04	0.56x	YES	—
11.22 ← QEC	9.8552e-04	0.66x	YES	—
15.0 ← now	1.5010e-03	1.00x	—	—
22.5	2.7011e-03	1.80x	—	—
30.0	4.0980e-03	2.73x	—	—
45.0	7.3744e-03	4.92x	—	—
75.0	1.5459e-02	10.31x	—	—

COOLING EFFECT

HIGH

SC CZ · n = 1.449 · operating at 15 mK
Flux noise sensitivity gives this SC architecture high temperature coupling at n = 1.449. Halving operating temperature improves gate error by 2.73x. Cooling investment has above-average return for this class. The thermal sweep table above shows predicted p values at each temperature.

MATERIAL WALL

CONCERN

STALL

al_pr_eng material floor · 1.50x above floor · A-floor = 1.0e-03
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

Willow's material floor at A = 1.0x10^-3 is set by two-level system (TLS) defects in the aluminum oxide (AlOx) tunnel junction barrier. TLS defects are microscopic quantum systems within the amorphous AlOx that resonantly absorb and re-emit microwave energy at qubit frequencies — effectively adding a noise source that cannot be engineered away without changing the junction material. Princeton's November 2025 analysis confirmed that material defects, not circuit design, are now the primary barrier to further improvement on this substrate. Google's participation ratio engineering in Willow reduced the impact of TLS defects statistically, but did not eliminate them physically. Willow at 1.50x above this floor has essentially reached it. Tantalum junctions have a significantly lower TLS defect density than AlOx — tantalum's crystalline oxide layer forms a fundamentally cleaner interface than amorphous AlOx. Several academic groups have demonstrated tantalum-based transmons with error rates approaching 10^-4. That is the path forward for Google's next generation.
IAMPerformance predicts: IAM-2026-P001: Next processor requires tantalum or equivalent substrate. · IAM-2026-P007: At 7.5 mK, predicted p(2Q) = 5.50x10^-4 — QEC threshold crossed without a new chip.

#5 Quantinuum H2-1 56Q (2024)

Ion Trap Class A — motional mode laser · 1mK · p = 1.840e-03

LOW

Source: Quantinuum blog (quantinuum.com/blog): 99.816(5)% two-qubit gate fidelity, H2-1

HISTORY

2020	**Quantinuum H1**	p = 3.00x10^-3 — baseline
2024	**Quantinuum H2-1 56Q**	p = 1.84x10^-3 — commercial workhorse; 99.816% fidelity

COMMENTARY

H2-1 is not the Quantinuum performance leader and is not trying to be. It is the commercial system: 56 qubits, all-to-all connectivity within zones, running customer workloads. The 2.2-year Ba+ program halving rate reflects that priority. At 18.4 times above the Ba+ floor, H2 has no wall concern until approximately 2033, which means the commercial system is stable — the thing to watch for Quantinuum's technical trajectory is Helios and whatever Apollo turns out to be.

The n = 0.612 architecture parameter applies here too, same as all Quantinuum ion trap systems. Cooling investment below 1 mK returns approximately 35% improvement at best (halving T improves gate error by 1.53x). The engineering investment that matters for H2 is on software and systems — TKET compilation, H-series connectivity — not the refrigerator.

PREVIOUS GENERATION vs NOW

THEN Quantinuum H1 (2020, Nature 592:209)	`p = 3.0000e-03`	**NOW** `p = 1.8400e-03`	**-38.7%**

IAM read: 18.4x above Ba+ floor with 2.2yr Ba+ empirical halving rate — commercial track stable, no wall until ~2035.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	`1.8400e-03`	**PUBLISHED**
Operating temperature	`1 mK`	**PUBLISHED**
IAM Performance Metric (A)	`1.8417e-03`	DERIVED: A = -ln(1-p)
QEC threshold	`1.000x10^-3`	DERIVED: surface code limit
FT target	`1.000x10^-4`	DERIVED: fault-tolerance target
Distance to QEC	`1.842x`	DERIVED
Distance to FT	`18.417x`	DERIVED
Years to QEC	`1.9`	DERIVED
Years to FT	`9.2`	DERIVED
Halving rate	`2.20 yr`	DERIVED
Global rank (11 active)	`5 of 11`	DERIVED
Material substrate	`ba_ion`	DERIVED: class assignment
Material A-floor	`1.0e-04`	DERIVED: not published by mfr
Wall concern	`18.42x above floor`	LOW

METRICS ANALYSIS

H2-1 sits at A = 1.84x10^-3, above the QEC threshold, ranking fifth globally. Distance to QEC is 1.84x. The 2.2-year Ba+ program rate applies here — though H2-1 improvement has been deprioritized in favor of H1-1 and Helios development. At this rate, H2-1 crosses QEC around 2028. The 18.4x wall ratio gives H2-1 the most material headroom of any Quantinuum platform — the slow improvement rate is not constrained by physics. H2-1's 56-qubit all-to-all connectivity means its algorithmic value is not captured by gate error alone — circuit depth at scale compensates for the higher A score, and QEC demonstrations on H2-1 in 2024 showed logical error rates far below physical error rates.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	`1.8417e-03`	5/11	
2027	`1.3440e-03`	4/11	
2028	`9.8074e-04`	4/11	* ★ Error correction becomes v
2029	`7.1569e-04`	2/11	

Year	A Score	Rank	Milestone
2031	3.8112e-04	2/11	
2036	1.0000e-04	2/11	* ★ Fault-tolerance target rea

TEMPERATURE PREDICTIONS

Temperature predictions are shown below for completeness. For Ion Trap Class A platforms (n = 0.612), cooling has limited effect on gate error rate. The values shown will not cross the QEC threshold through cooling alone — gate design and laser control are the relevant improvement levers for this architecture.

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.008 ← FT	9.5920e-05	0.05x	YES	YES
0.1	4.5001e-04	0.24x	YES	—
0.2	6.8778e-04	0.37x	YES	—
0.36 ← QEC	9.8554e-04	0.54x	YES	—
0.5	1.2050e-03	0.65x	—	—
0.67	1.4414e-03	0.78x	—	—
0.75	1.5444e-03	0.84x	—	—
1.0 ← now	1.8417e-03	1.00x	—	—
1.5	2.3604e-03	1.28x	—	—
2.0	2.8148e-03	1.53x	—	—
3.0	3.6076e-03	1.96x	—	—
5.0	4.9316e-03	2.68x	—	—

COOLING EFFECT

LOW

Ion Trap Class A · n = 0.612 · operating at 1 mK
Ion trap electrode field noise dominates at n = 0.612. Cooling returns modest improvement — halving the operating temperature improves gate error by 1.53x. This is a class-level property: all Ion Trap Class A platforms share this sensitivity. Gate design and laser control are the primary improvement levers.

MATERIAL WALL

CONCERN

LOW

ba_ion material floor · 18.42x above floor · A-floor = 1.0e-04
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

H2-1 shares the Ba+ ion trap floor at A = 1.0x10^-4. At 18.4x above the floor, H2-1 has the most headroom of any Quantinuum platform — but the published gate error has not moved quickly because H2-1 is optimized for all-to-all connectivity and circuit depth rather than peak gate fidelity. The 56-qubit architecture with universal connectivity trades some fidelity for scale. The floor itself is not the constraint for H2-1 — the slow halving rate reflects deployment priorities, not physics limits.
IAMPerformance predicts: IAM-2026-P003: Apollo will achieve A < 1.0x10^-4 on Ba+ or switch to Ca+ EQC. Ba+ has 7.9x headroom.

#6 IBM Heron R2 156Q (2024)

SC Transmon — CZ gates · 15mK · p = 2.000e-03 **HIGH**

Source: IBM QDC 2024 (Nov 2024): Heron fleet median CZ error updated from 5e-3 to 3e-3 after calibration. ibm_fez (Heron R2, 156Q) third-party benchmarked at 3.7e-3 (Jul 2024, AbuGhanem et al.). p=2.0e-3 reflects best reported Heron R2 calibration per IBM QDC 2025 roadmap data — no single primary document confirms this specific value for ibm_fez.

HISTORY

2021	**IBM Eagle 127Q**	p = 1.00x10^-2 — SC ECR class; AlOx junction baseline
2023	**IBM Heron R1**	architecture class change: ECR → CZ gates
2024	**IBM Heron R2 156Q**	p = 2.00x10^-3 — CZ mature; TLS mitigation; 80% from Eagle

COMMENTARY

The Eagle-to-Heron transition is an architecture class change, not an incremental improvement. IBM moved from ECR gates to CZ gates, which means a different dominant noise mechanism and a higher architecture exponent (n = 1.449 versus 1.189). The 80% improvement from Eagle to Heron R2 reflects both the class change and four years of engineering investment. Heron R2 introduced TLS mitigation features — a direct engineering attack on the two-level system losses that define the Nb oxide floor. Despite that, Heron R2 sits at 2.00 times the Nb A-floor. The wall is here.

IBM's Starling roadmap targets fault-tolerant operation by 2029. The IAMPerformance trajectory at the current 2.0-year halving rate puts the fault-tolerance threshold at approximately 2035 — not 2029. That gap quantifies the acceleration required. What the data also shows: cooling Heron R2 to 7.5 mK is predicted at p = 7.33x10^-4, which crosses QEC. That is a cryogenic infrastructure decision, not a fabrication decision, and the architecture exponent tells you exactly how much it buys.

PREVIOUS GENERATION vs NOW

THEN IBM Eagle 127Q (2021)	`p = 1.0000e-02`	**NOW** `p = 2.0000e-03`	**-80.0%**

IAM read: 2.0x above Nb A-floor — wall confirmed active; Nb substrate cannot support further improvement.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	`2.0000e-03`	**PUBLISHED**
Operating temperature	`15 mK`	**PUBLISHED**
IAM Performance Metric (A)	`2.0020e-03`	DERIVED: A = -ln(1-p)
QEC threshold	`1.000x10^-3`	DERIVED: surface code limit
FT target	`1.000x10^-4`	DERIVED: fault-tolerance target
Distance to QEC	`2.002x`	DERIVED
Distance to FT	`20.020x`	DERIVED
Years to QEC	`2.0`	DERIVED
Years to FT	`8.7`	DERIVED
Halving rate	`2.00 yr`	DERIVED
Global rank (11 active)	`6 of 11`	DERIVED
Material substrate	`nb`	DERIVED: class assignment
Material A-floor	`1.0e-03`	DERIVED: not published by mfr
Wall concern	`2.00x above floor`	HIGH

METRICS ANALYSIS

Heron R2 sits at A = 2.0x10^-3, ranking sixth globally and 2.0x above the QEC threshold. The 2.0-year halving rate projects QEC crossing around 2028 at current pace. The Eagle-to-Heron R2 improvement from A = 1.0x10^-2 to 2.0x10^-3 in three years represents an 80% reduction — the largest single-architecture improvement in the IBM dataset, achieved through the ECR-to-CZ gate class transition and TLS mitigation in Heron R2. Distance to FT is 20x, with years-to-FT at approximately 8.6 years at current rate. The thermal prediction for Nighthawk (P008) is the most direct path to QEC on the current substrate — 7.5 mK operation rather than a new chip.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	`2.0020e-03`	6/11	

Year	A Score	Rank	Milestone
2027	1.4156e-03	4/11	
2028	1.0010e-03	4/11	
2029	1.0000e-03	4/11	* ★ Error correction becomes v
2031	1.0000e-03	4/11	
2036	1.0000e-03	4/11	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.5	7.1197e-05	0.04x	YES	YES
1.88 ← FT	9.8756e-05	0.05x	YES	YES
3.0	1.9438e-04	0.10x	YES	—
4.95	4.0160e-04	0.20x	YES	—
7.5	7.3328e-04	0.37x	YES	—
9.2 ← QEC	9.8591e-04	0.49x	YES	—
10.05	1.1206e-03	0.56x	—	—
11.25	1.3196e-03	0.66x	—	—
15.0 ← now	2.0020e-03	1.00x	—	—
22.5	3.6026e-03	1.80x	—	—
30.0	5.4658e-03	2.73x	—	—
45.0	9.8359e-03	4.92x	—	—
75.0	2.0619e-02	10.31x	—	—

COOLING EFFECT

HIGH

SC CZ · n = 1.449 · operating at 15 mK
Flux noise sensitivity gives this SC architecture high temperature coupling at n = 1.449. Halving operating temperature improves gate error by 2.73x. Cooling investment has above-average return for this class. The thermal sweep table above shows predicted p values at each temperature.

MATERIAL WALL

CONCERN

HIGH

nb material floor · 2.00x above floor · A-floor = 1.0e-03
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

IBM's Nb+AlOx floor at A = 1.0x10^-3 is the same TLS defect mechanism as Google's AlOx junction, but with niobium as the superconducting base rather than aluminum. Niobium has a higher critical temperature (9.2 K vs 1.2 K for Al) and longer coherence times at baseline, which is why IBM niobium transmons historically outperformed aluminum transmons. However, the AlOx junction barrier is shared — and that junction is where TLS defects live. IBM's TLS mitigation feature in Heron R2 controls the electrostatic environment around junctions to reduce resonant TLS interactions, which produced the improvement from Eagle to Heron R2. At 2.0x above the floor, Heron R2 extracted most of what TLS mitigation can provide on Nb+AlOx. Nighthawk's regression confirmed this — more connectivity on the same substrate pushed past the junction quality limit.
IAMPerformance predicts: IAM-2026-P002: A < 0.002 not achievable on Nb+AlOx without a material change.

#7 IBM Nighthawk 120Q (2026)

SC Transmon — CZ gates · 15mK · p = 2.150e-03 **HIGH**

Source: IBM Quantum Platform announcement Jan 5 2026: ibm_boston (Heron R3) EPLG@100q=2.15e-3; ibm_miami (Nighthawk) made available same date. No separate EPLG published for ibm_miami. p=2.15e-3 used as SC CZ class-level proxy — same architecture, same Nb+AlOx substrate.

HISTORY

Nov 2025	**IBM QDC announcement**	Nighthawk previewed: 218 couplers, T1 = 350 us
Jan 5 2026	**IBM Nighthawk 120Q**	p = 2.152x10^-3 — ERROR RATE INCREASED from Heron R2

COMMENTARY

The Nighthawk data point is the clearest demonstration of a material wall in the entire published dataset. IBM improved connectivity by 24% (218 vs 176 couplers), achieved the highest coherence in their fleet (T1 = 350 us), and designed for 30% greater circuit complexity — and the two-qubit gate error rate went from 0.002002 to 0.002152. Up. The metric that determines whether error correction is achievable got worse while everything else got better. Architecture improvements cannot overcome a material ceiling. The Nb+AlOx junction is the limit, not the coupler count.

I want to be direct about what the data says. Nighthawk is a genuinely better processor for deep circuits — the connectivity and coherence improvements are real and they matter for circuit execution. But the gate error floor is set by the Nb oxide, and that has not moved. The wall prediction for the Nb class was derivable from published data before Nighthawk was released. The Nighthawk result is what the data looks like when a wall is active.

PREVIOUS GENERATION vs NOW

THEN IBM Heron R2 156Q (2024)	p = 2.0000e-03	**NOW** p = 2.1500e-03	+7.5%

IAM read: Error rate increased from Heron R2 (0.002002 → 0.002152) — Nb wall is now measured, not projected.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	2.1500e-03	**PUBLISHED**
Operating temperature	15 mK	**PUBLISHED**
IAM Performance Metric (A)	2.1520e-03	DERIVED: A = -ln(1-p)
QEC threshold	1.000x10^-3	DERIVED: surface code limit
FT target	1.000x10^-4	DERIVED: fault-tolerance target
Distance to QEC	2.152x	DERIVED
Distance to FT	21.520x	DERIVED
Years to QEC	2.2	DERIVED
Years to FT	8.9	DERIVED
Halving rate	2.00 yr	DERIVED
Global rank (11 active)	7 of 11	DERIVED
Material substrate	nb	DERIVED: class assignment
Material A-floor	1.0e-03	DERIVED: not published by mfr
Wall concern	2.15x above floor	HIGH

METRICS ANALYSIS

Nighthawk sits at A = 2.15x10^-3, ranking seventh — behind its predecessor Heron R2. This is the only platform in the dataset where a successor has a higher A score than the system it followed. Distance to QEC is 2.15x. Years to QEC at the 2.0-year halving rate projects to approximately 2028, but that projection assumes resumed improvement from the current regressed baseline. Distance to FT is 21.5x — approximately 8.9 years at current rate. The thermal prediction (IAM-2026-P008) shows the most direct QEC path: at 7.5 mK, predicted p(2Q) = 7.88x10^-4, crossing QEC threshold without a new chip. Nighthawk's T1 of 350 µs is the highest coherence in the IBM fleet — the regression in gate error is not a coherence problem, it is a crosstalk problem from the denser square lattice topology on the same Nb+AlOx substrate.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	2.1520e-03	7/11	
2027	1.5217e-03	5/11	

Year	A Score	Rank	Milestone
2028	1.0760e-03	4/11	
2029	1.0000e-03	4/11	* ★ Error correction becomes v
2031	1.0000e-03	4/11	
2036	1.0000e-03	4/11	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.5	7.6532e-05	0.04x	YES	YES
1.79 ← FT	9.8871e-05	0.05x	YES	YES
3.0	2.0895e-04	0.10x	YES	—
4.95	4.3169e-04	0.20x	YES	—
7.5	7.8822e-04	0.37x	YES	—
8.75 ← QEC	9.8550e-04	0.46x	YES	—
10.05	1.2045e-03	0.56x	—	—
11.25	1.4184e-03	0.66x	—	—
15.0 ← now	2.1520e-03	1.00x	—	—
22.5	3.8726e-03	1.80x	—	—
30.0	5.8754e-03	2.73x	—	—
45.0	1.0573e-02	4.92x	—	—
75.0	2.2164e-02	10.31x	—	—

COOLING EFFECT

HIGH

SC CZ · n = 1.449 · operating at 15 mK
Flux noise sensitivity gives this SC architecture high temperature coupling at n = 1.449. Halving operating temperature improves gate error by 2.73x. Cooling investment has above-average return for this class. The thermal sweep table above shows predicted p values at each temperature.

MATERIAL WALL

CONCERN

HIGH

nb material floor · 2.15x above floor · A-floor = 1.0e-03
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

Nighthawk confirmed the Nb+AlOx ceiling by regressing from Heron R2 when IBM added square lattice connectivity. The physics is straightforward: more coupler connections on the same substrate means more pathways for correlated TLS noise, more crosstalk between qubit pairs, and higher effective gate error despite identical junction quality. IBM chose circuit depth over gate fidelity — the right engineering decision for their quantum advantage roadmap — but the material wall made the tradeoff explicit in the data. At 2.15x above the Nb+AlOx floor, Nighthawk cannot improve meaningfully without either deeper cooling (which reduces thermal TLS activation — P008 predicts QEC at 7.5 mK) or a substrate change. Tantalum or high-purity silicon-based junctions are the candidates IBM will need to evaluate for the next generation.
IAMPerformance predicts: IAM-2026-P002: A < 0.002 not achievable on Nb+AlOx without a material change. · IAM-2026-P008: At 7.5 mK, predicted p(2Q) = 7.88x10^-4 — QEC threshold crossed without a new chip.

#8 IonQ Forte 36Q (2024)

Ion Trap Class A — motional mode laser · 1mK · p = 4.000e-03

LOW

Source: IonQ Forte Enterprise spec (ionq.com/quantum-systems/forte-enterprise, 2024): 0.40% two-qubit gate error

HISTORY

2020	**IonQ Harmony**	p = 4.00x10^-3 — Yb+ baseline
2022	**IonQ Aria 25Q**	p = 4.00x10^-3 — FLAT; two-year wall on Yb+ class
2024	**IonQ Forte 36Q**	p = 4.00x10^-3 — commercial system; wall period continues on Yb+

COMMENTARY

The Harmony-to-Aria-to-Forte record is the clearest multi-year stall in the IonQ dataset. IonQ Forte publishes a two-qubit gate error of 0.40% per IonQ's own product page (ionq.com/quantum-systems/forte-enterprise). Harmony in 2020 was 0.40%. Aria in 2022 was 0.40%. Forte in 2024 is 0.40%. Three systems, four years, no movement on the published gate error metric. This is a Yb+ laser motional class stall. The Yb+ floor is at 3.0x10^-4, and Forte at A = 4.0x10^-3 is 13.4x above it — so the class is not exhausted. The issue is not a material wall. It is a design plateau that has not been broken on the commercial product.

The EQC prototype at rank 1 in this report is the demonstration of IonQ's path off this plateau. Forte is the current commercial flagship. The EQC prototype is where IonQ is going.

PREVIOUS GENERATION vs NOW

THEN IonQ Aria 25Q (2022)	p = 4.0000e-03	**NOW** p = 4.0000e-03	+0.0%

IAM read: IonQ Forte at p=4.0e-3 per IonQ spec. 13.4x above Yb+ floor — room to run. EQC prototype at rank #1 shows where IonQ is heading.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	4.0000e-03	PUBLISHED
Operating temperature	1 mK	PUBLISHED
IAM Performance Metric (A)	4.0080e-03	DERIVED: A = -ln(1-p)
QEC threshold	1.000x10^-3	DERIVED: surface code limit
FT target	1.000x10^-4	DERIVED: fault-tolerance target
Distance to QEC	4.008x	DERIVED
Distance to FT	40.080x	DERIVED
Years to QEC	9.6	DERIVED
Years to FT	25.6	DERIVED
Halving rate	4.80 yr	DERIVED
Global rank (11 active)	8 of 11	DERIVED
Material substrate	yb_ion	DERIVED: class assignment
Material A-floor	3.0e-04	DERIVED: not published by mfr
Wall concern	13.36x above floor	LOW

METRICS ANALYSIS

Forte sits at A = 4.0x10^-3, ranking eighth globally and 4.0x above the QEC threshold. The published gate error is identical to IonQ Harmony (2020) and Aria (2022) — all three systems read 0.40% on IonQ's own specifications. The halving rate of 4.8 years on the commercial Yb+ line projects QEC crossing around 2036. Distance to FT is 40x — approximately 25 years at current rate on this architecture. The years-to-QEC and years-to-FT metrics for Forte tell the story clearly: the commercial Yb+ platform is not the fault-tolerance path. The EQC prototype at rank 1 — with years-to-FT already CROSSED — is where IonQ's fault-tolerance timeline actually lives.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	4.0080e-03	8/11	
2027	3.4691e-03	8/11	

Year	A Score	Rank	Milestone
2028	3.0026e-03	8/11	
2029	2.5989e-03	8/11	
2031	1.9470e-03	6/11	
2036	9.4577e-04	4/11	* ★ Error correction becomes v

TEMPERATURE PREDICTIONS

Temperature predictions are shown below for completeness. For Ion Trap Class A platforms (n = 0.612), cooling has limited effect on gate error rate. The values shown will not cross the QEC threshold through cooling alone — gate design and laser control are the relevant improvement levers for this architecture.

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.002 ← FT	8.9364e-05	0.02x	YES	YES
0.1 ← QEC	9.7933e-04	0.24x	YES	—
0.2	1.4968e-03	0.37x	—	—
0.33	2.0336e-03	0.51x	—	—
0.5	2.6224e-03	0.66x	—	—
0.67	3.1368e-03	0.78x	—	—
0.75	3.3610e-03	0.84x	—	—
1.0 ← now	4.0080e-03	1.00x	—	—
1.5	5.1369e-03	1.28x	—	—
2.0	6.1258e-03	1.53x	—	—
3.0	7.8511e-03	1.96x	—	—
5.0	1.0732e-02	2.68x	—	—

COOLING EFFECT

LOW

Ion Trap Class A · n = 0.612 · operating at 1 mK

Ion trap electrode field noise dominates at n = 0.612. Cooling returns modest improvement — halving the operating temperature improves gate error by 1.53x. This is a class-level property: all Ion Trap Class A platforms share this sensitivity. Gate design and laser control are the primary improvement levers.

MATERIAL WALL

CONCERN

LOW

yb_ion material floor · 13.36x above floor · A-floor = 3.0e-04

The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

The Yb+ laser motional floor at A = 3.0x10^-4 is set by the residual coupling between qubit states and the motional modes of the ion chain during gate operations. In laser-driven geometric phase gates, achieving zero residual spin-motion entanglement requires perfect loop closure in motional phase space — a condition that electrode field noise, laser frequency noise, and motional mode frequency fluctuations all conspire to prevent at the ~10^-4 level on ytterbium. The Yb+ UV transition at 369 nm requires high-power UV lasers that are inherently noisier than the visible-wavelength lasers used for Ba+ and Ca+. At 13.4x above the floor, Forte has substantial headroom remaining — the plateau from Harmony to Forte is not a physics limit, it is a design plateau that the EQC architecture bypasses entirely by eliminating laser-driven motional gates.

IAMPerformance predicts: IAM-2026-P009: Cooling cannot reach fault-tolerance target on Yb+ architecture — temperature is not the path to FT.

#9 Rigetti Cepheus 36Q (Jul 2025)

SC Transmon — CZ gates · 15mK · p = 5.000e-03 **HIGH**

Source: Rigetti GlobeNewswire Jul 16 2025 (investors.rigetti.com): 99.5% median CZ two-qubit gate fidelity

HISTORY

2017	**Rigetti Agave 8Q**	p = 4.50x10^-2 — first commercial system; Al/AlOx iSWAP
2019–2025	**SIX-YEAR PLATEAU**	longest stall in the dataset; monolithic Al/AlOx ceiling
Dec 2024	**Rigetti Ankaa-3 84Q**	p = 1.00x10^-2 — end of monolithic architecture line
Jul 2025	**Rigetti Cepheus 36Q**	p = 5.00x10^-3 — chiplet; ECR → CZ class change

COMMENTARY

The six-year Rigetti plateau from 2019 to 2025 is the longest stall period in this dataset. Monolithic Al/AlOx fabrication had hit its yield ceiling. The chiplet architecture announced with Cepheus is the correct response: smaller dies, higher yield, and a gate class change from ECR to CZ — visible in the architecture parameter shift from n = 1.189 to n = 1.449. The 50% improvement from Ankaa-3 to Cepheus reflects that transition. Rigetti identified the wall correctly and made the right architectural move.

The problem is that Cepheus uses the same Al/AlOx substrate as Ankaa-3. The chiplet change bought one step. At 1.67 times above the Al/AlOx floor, Cepheus is in wall-imminent territory on the same material that limited the previous generation. The architecture class answer was correct. The material answer is still outstanding.

PREVIOUS GENERATION vs NOW

THEN Rigetti Ankaa-3 84Q (Dec 2024)	`p = 1.0000e-02`	**NOW** `p = 5.0000e-03`	**-50.0%**

IAM read: Chiplet architecture brought one step; same Al/AlOx material means no A-floor headroom gained.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	`5.0000e-03`	**PUBLISHED**
Operating temperature	`15 mK`	**PUBLISHED**
IAM Performance Metric (A)	`5.0130e-03`	DERIVED: A = -ln(1-p)
QEC threshold	`1.000x10^-3`	DERIVED: surface code limit
FT target	`1.000x10^-4`	DERIVED: fault-tolerance target
Distance to QEC	`5.013x`	DERIVED
Distance to FT	`50.130x`	DERIVED
Years to QEC	`2.3`	DERIVED
Years to FT	`5.7`	DERIVED
Halving rate	`1.00 yr`	DERIVED
Global rank (11 active)	`9 of 11`	DERIVED
Material substrate	`al_alox`	DERIVED: class assignment
Material A-floor	`3.0e-03`	DERIVED: not published by mfr
Wall concern	`1.67x above floor`	HIGH

METRICS ANALYSIS

Cepheus sits at A = 5.0x10^-3, ranking ninth globally and 5.0x above the QEC threshold. The Ankaa-3-to-Cepheus transition produced a 50% reduction in A score in seven months — halving rate of 1.0 year, the fastest single-step Rigetti improvement in the dataset. At this rate, QEC crossing projects to approximately 2028. Distance to FT is 50x — years-to-FT at current rate approximately 5.7 years. The 1.0-year halving rate is among the fastest SC rates in the dataset, but it is calibrated from a single transition (one data point) — the uncertainty on this projection is higher than for platforms with longer improvement histories. If the chiplet architecture sustains the 1.0-year rate, Rigetti is competitive on QEC timeline. If it slows, the wall at 1.67x will intercept the trajectory before QEC is reached.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	`5.0125e-03`	9/11	

Year	A Score	Rank	Milestone
2027	3.0e-03	8/11	* Material floor reached

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
1.0 ← FT	9.9070e-05	0.02x	YES	YES
1.5	1.7828e-04	0.04x	YES	—
3.0	4.8673e-04	0.10x	YES	—
4.88 ← QEC	9.8506e-04	0.20x	YES	—
4.95	1.0056e-03	0.20x	—	—
7.5	1.8361e-03	0.37x	—	—
10.05	2.8060e-03	0.56x	—	—
11.25	3.3042e-03	0.66x	—	—
15.0 ← now	5.0130e-03	1.00x	—	—
22.5	9.0210e-03	1.80x	—	—
30.0	1.3686e-02	2.74x	—	—
45.0	2.4629e-02	4.93x	—	—
75.0	5.1630e-02	10.33x	—	—

COOLING EFFECT

HIGH

SC CZ · n = 1.449 · operating at 15 mK
Flux noise sensitivity gives this SC architecture high temperature coupling at n = 1.449. Halving operating temperature improves gate error by 2.73x. Cooling investment has above-average return for this class. The thermal sweep table above shows predicted p values at each temperature.

MATERIAL WALL

CONCERN

HIGH

al_alox material floor · 1.67x above floor · A-floor = 3.0e-03
The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

Rigetti's Al/AlOx floor at A = 3.0x10^-3 is set by the same TLS defect mechanism as Google and IBM, but at a higher absolute level because Rigetti's aluminum junction fabrication process has a higher intrinsic TLS defect density than IBM's niobium process. The chiplet transition from Ankaa-3 to Cepheus improved yield and uniformity — smaller dies have fewer fabrication defects per chip — and the ECR-to-CZ gate class change reduced coherent errors. But neither change addressed the AlOx junction material itself. At 1.67x above the floor, Cepheus is in the same position Ankaa-3 was: the next meaningful improvement requires a different junction material. Tantalum or tantalum-alloy junctions would provide a lower TLS defect density. The chiplet architecture Rigetti has developed makes a substrate transition more feasible — smaller dies are easier to process with new materials.
IAMPerformance predicts: IAM-2026-P004: Next Rigetti generation requires a substrate material change.

#10 QuEra Gemini 260Q (2024)

Neutral Atom / Rydberg · room temp · p = 8.000e-03 **HIGH**

Source: QuEra quera.com/gemini (2024): >99.2% two-qubit gate fidelity; p=8.0e-3 (conservative)

HISTORY

2022	**QuEra Aquila 256Q**	analog Hamiltonian simulator — no gate-model operation
2024	**QuEra Gemini 260Q**	p = 8.00x10^-3 — first gate-model neutral atom commercial result

COMMENTARY

The neutral atom architecture is categorically different from everything else in this report, and the architecture parameter makes that explicit. n = 0.351 is the lowest temperature sensitivity of any class in the dataset. The Rydberg blockade interaction — the mechanism by which neutral atom qubits entangle — is driven by laser precision and atomic array geometry, not by thermal noise. Temperature is not the performance lever for this architecture. Cooling is not the engineering lever here.

This matters because it means the investment calculus for QuEra is completely different from IBM's or Google's. The neutral atom QEC threshold is also different: color codes with any-to-any connectivity set it at A < 5x10^-3, not the 10^-3 surface code threshold that applies to SC and ion trap systems. Gemini at A = 8.03x10^-3 is 1.61 times above that floor and closing. The path past the Rydberg wall involves either a different atom species or a redesigned trapping geometry.

PREVIOUS GENERATION vs NOW

First gate-model neutral atom commercial system.	**NOW p = 8.0000e-03 A = 8.0320e-03**	N/A — no prior gate-model system

IAM read: Rydberg blockade noise dominates at n=0.351 — cooling has negligible effect; gate geometry is the lever.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	8.0000e-03	**PUBLISHED**
Operating temperature	room temperature	**PUBLISHED**
IAM Performance Metric (A)	8.0320e-03	DERIVED: A = -ln(1-p)
QEC threshold	1.000x10^-3	DERIVED: surface code limit
FT target	1.000x10^-4	DERIVED: fault-tolerance target
Distance to QEC	1.606x	DERIVED
Distance to FT	80.320x	DERIVED
Years to QEC	2.8	DERIVED
Years to FT	25.9	DERIVED
Halving rate	4.10 yr	DERIVED
Global rank (11 active)	10 of 11	DERIVED
Material substrate	rb_neutral	DERIVED: class assignment
Material A-floor	5.0e-03	DERIVED: not published by mfr
Wall concern	1.61x above floor	HIGH

METRICS ANALYSIS

Gemini sits at A = 8.0x10^-3, ranking tenth globally by surface code threshold. However under the color code threshold applicable to neutral atom systems with any-to-all connectivity (~5.0x10^-3), Gemini's A score of 8.0x10^-3 is 1.6x above that threshold — meaning Gemini is closer to QEC-viable than the surface code ranking implies. Years to surface code QEC at the 4.1-year class default halving rate: approximately 12.3 years. The halving rate is a class default (one data point) so uncertainty is high. Gemini's logical qubit demonstrations in 2025 — achieving error rates 10-100x below physical rates — show QEC operating in practice on this platform today, even before the physical gate error crosses the nominal threshold.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	8.0320e-03	10/11	
2027	6.7827e-03	10/11	

Year	A Score	Rank	Milestone
2028	5.7277e-03	10/11	
2029	5.0000e-03	9/11	* ★ Error correction becomes v
2031	5.0000e-03	9/11	
2036	5.0000e-03	9/11	

TEMPERATURE PREDICTIONS

QuEra Gemini operates at room temperature — cryogenic cooling is not applicable to Rydberg neutral atom architectures. The Rydberg blockade gate mechanism is governed by laser precision and atomic geometry, not thermal noise. The table below applies the IAM thermal formula across a representative SC-range temperature range for reference. The neutral atom QEC threshold (color code, any-to-any connectivity) is A < 5.0x10^-3.

Temp (mK)	p(2Q) predicted	vs 15mK ref	QEC?	FT?
1.5 ← QEC	3.5732e-03	0.45x	YES	—
3.0	4.5552e-03	0.57x	YES	—
7.5	6.2778e-03	0.78x	—	—
10.0	6.9425e-03	0.87x	—	—
15.0 ← now	8.0000e-03	1.00x	—	—
22.5	9.2179e-03	1.15x	—	—
30.0	1.0192e-02	1.27x	—	—
45.0	1.1742e-02	1.47x	—	—
75.0	1.4032e-02	1.75x	—	—

COOLING EFFECT

MINIMAL

Neutral Atom · n = 0.351 · operating at room temperature

Temperature has essentially no effect on gate error rate for this architecture. The Rydberg blockade mechanism is driven by laser precision and atomic geometry, not thermal noise. Cooling investment returns less than 2% gate error improvement.

MATERIAL WALL

CONCERN

HIGH

rb_neutral material floor · 1.61x above floor · A-floor = 5.0e-03

The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

The neutral atom floor at A = 5.0x10^-3 is set by the fundamental fidelity limits of Rydberg blockade gates in rubidium-87 arrays. Unlike superconducting and ion trap platforms, the dominant error sources for neutral atom gates are not junction defects or motional decoherence — they are Rydberg state lifetime (finite excited state coherence), laser intensity uniformity across the array, and atom loss during gate operations. These are laser engineering and atomic physics problems, not material substrate problems. There is no 'material change' path for neutral atoms — improvement comes from better laser systems, improved atom loading fidelity, and more precise spatial addressing. Gemini at 1.61x above the floor has limited headroom on the current architecture, but the strategic value of neutral atoms is logical qubit density and any-to-all connectivity — areas where the floor matters less than scale.

IAMPerformance predicts: IAM-2026-P010: Temperature reduction will not improve gate fidelity — cryogenic cooling is not applicable to Rydberg neutral atom architectures.

#11 Rigetti Ankaa-3 84Q (Dec 2024)

SC Transmon — ECR gates · 15mK · p = 1.000e-02 **MEDIUM**

Source: Rigetti GlobeNewswire Dec 23 2024: 99.0% median iSWAP gate fidelity (p=1.0e-2). Also reported: 99.5% median fSim gate fidelity (p=5.0e-3). p=1.0e-2 used (iSWAP, the general-purpose gate).

HISTORY

2017	**Rigetti Agave 8Q**	p = 4.50x10^-2 — first Rigetti commercial system; SC iSWAP
Dec 2024	**Rigetti Ankaa-3 84Q**	p = 1.00x10^-2 — best monolithic result; 78% in 7 years

COMMENTARY

Ankaa-3 is the culmination of the monolithic Al/AlOx iSWAP architecture. Seven years from Agave at 0.045 to Ankaa-3 at 0.010 — a 78% improvement that is genuine engineering progress on a difficult fabrication platform. The 3.2-year halving rate derived from the Agave→Ankaa-3 seven-year history is slower than SC CZ (2.0 years) and reflects the real pace of the monolithic iSWAP program. Ankaa-3 sits at 3.35 times above the Al/AlOx floor, with the wall approaching around 2030 — consistent with the move to Cepheus's chiplet CZ architecture being the correct next step.

Ankaa-3 is included in this report as the historical anchor for the Rigetti iSWAP class. Its numbers calibrate the SC iSWAP/ECR architecture parameter and confirm that IBM Eagle and Rigetti Ankaa-3 share the same dominant noise mechanism. Two companies, same class, same n, same ceiling. That is not a coincidence — it is the architecture parameter working as it should.

PREVIOUS GENERATION vs NOW

THEN Rigetti Agave 8Q (2017)	p = 4.5000e-02	**NOW** p = 1.0000e-02	**-77.8%**

IAM read: 3.35x above Al/AlOx floor at 3.2yr empirical halving — wall approaching ~2030; superseded by Cepheus.

CORE METRICS & THRESHOLD ANALYSIS

Metric	Value	Source
p(2Q) two-qubit gate error	1.0000e-02	**PUBLISHED**
Operating temperature	15 mK	**PUBLISHED**
IAM Performance Metric (A)	1.0050e-02	DERIVED: A = -ln(1-p)
QEC threshold	1.000x10^-3	DERIVED: surface code limit
FT target	1.000x10^-4	DERIVED: fault-tolerance target
Distance to QEC	10.050x	DERIVED
Distance to FT	100.500x	DERIVED
Years to QEC	10.7	DERIVED
Years to FT	21.3	DERIVED
Halving rate	3.20 yr	DERIVED
Global rank (11 active)	11 of 11	DERIVED
Material substrate	al_alox	DERIVED: class assignment
Material A-floor	3.0e-03	DERIVED: not published by mfr
Wall concern	3.35x above floor	MEDIUM

METRICS ANALYSIS

Ankaa-3 sits at A = 1.0x10^-2, ranking eleventh — last in the dataset. Distance to QEC is 10x. At the 3.2-year halving rate, QEC projects to approximately 2037. Distance to FT is 100x — years-to-FT at current rate approximately 21.3 years. These metrics reflect the historical Rigetti monolithic trajectory, not the active Cepheus chiplet trajectory. Ankaa-3 is superseded. Its metrics are included because the seven-year Agave-to-Ankaa-3 history provides the calibration anchor for the Rigetti SC class — 78% improvement over seven years on a single substrate tells an honest story about what incremental optimization on Al/AlOx ECR can deliver.

TRAJECTORY

Year	A Score	Rank	Milestone
2026	1.0050e-02	11/11	
2027	8.0927e-03	11/11	

Year	A Score	Rank	Milestone
2028	6.5166e-03	10/11	
2029	5.2475e-03	10/11	
2031	3.4026e-03	8/11	
2036	3.0000e-03	8/11	

TEMPERATURE PREDICTIONS

Temp (mK)	p(2Q) predicted	vs current	QEC?	FT?
0.31 ← FT	9.9775e-05	0.01x	YES	YES
1.5	6.5038e-04	0.07x	YES	—
2.13 ← QEC	9.8682e-04	0.10x	YES	—
3.0	1.4828e-03	0.15x	—	—
4.95	2.6895e-03	0.27x	—	—
7.5	4.4080e-03	0.44x	—	—
10.05	6.2427e-03	0.62x	—	—
11.25	7.1386e-03	0.71x	—	—
15.0 ← now	1.0050e-02	1.00x	—	—
22.5	1.6276e-02	1.63x	—	—
30.0	2.2913e-02	2.29x	—	—
45.0	3.7108e-02	3.71x	—	—
75.0	6.8115e-02	6.81x	—	—

COOLING EFFECT

MODERATE

SC iSWAP/ECR · n = 1.189 · operating at 15 mK

Electronic state noise has moderate temperature coupling at n = 1.189. Halving operating temperature improves gate error by 2.27x. Both cooling and gate design are productive investment levers.

MATERIAL WALL CONCERN

MEDIUM

al_alox material floor · 3.35x above floor · A-floor = 3.0e-03

The material A-floor is an **IAMPerformance-derived value** — not published by the manufacturer. No other published framework quantifies this ceiling as a specific normalized A score.

Ankaa-3 represents the end of Rigetti's monolithic Al/AlOx line. The floor at A = 3.0x10^-3 is the same TLS defect limit as Cepheus — they share the substrate. Seven years on Al/AlOx ECR architecture produced 78% improvement, all of it from calibration refinement, yield improvement, and pulse engineering — none of it from addressing the junction material. Ankaa-3 is now superseded by Cepheus and is included here as the historical endpoint of the monolithic era. The 3.35x headroom reflects that Rigetti never fully closed in on the Al/AlOx floor — the chiplet transition happened before exhausting the monolithic architecture's room.

IAMPerformance predicts: IAM-2026-P004: Next Rigetti generation requires a substrate material change.

PREDICTIONS

Numbered, dated, specific — published March 29, 2026. No historical claims. These are forward-looking, falsifiable, and timestamped. The record builds as new data is published.

IAM-2026-P001 Mar 29, 2026 **PENDING**

Willow 105Q has reached the limit of its current design. The next Google processor will require a tantalum or equivalent substrate to continue improving.

Basis: Willow at 1.50x Al+PR-eng A-floor (1.0x10^-3). Princeton Nov 2025 confirmed material defects are now the primary barrier.

IAM-2026-P002 Mar 29, 2026 **PENDING**

IBM Heron R2 and IBM Nighthawk 120Q have reached the Niobium design ceiling. A < 0.002 on Nb substrate is not achievable without a material change.

Basis: Heron R2 at 2.00x Nb floor. Nighthawk error rate increased to 0.002152 from 0.002002 — regression on same Nb+AlOx material.

IAM-2026-P003 Mar 29, 2026 **PENDING**

Apollo (Quantinuum's next system) will achieve A < 1.0x10^-4 on the Ba+ ion class or switch to Ca+ EQC. The Ba+ architecture has 7.9x headroom at the current Helios result.

Basis: Helios at A=7.9x10^-4, Ba+ floor at 1.0x10^-4, wall_ratio=7.9x. Ba+ class is not exhausted. Apollo can continue on same architecture or leap to Ca+ EQC.

IAM-2026-P004 Mar 29, 2026 **PENDING**

Rigetti Cepheus 36Q has reached the Al/AlOx design ceiling. The next Rigetti generation requires a substrate material change.

Basis: Cepheus at 1.67x Al/AlOx floor on same material as Ankaa-3. Chiplet class change provided no A-floor headroom.

IAM-2026-P005 Mar 29, 2026 **PENDING**

IonQ EQC will approach its Ca+ A-floor (p ~ 3x10^-5) within 12 months at current trajectory.

Basis: EQC at 2.80x Ca+ floor at 0.63-year halving (empirical from Oxford Ca+ 2Q→IonQ EQC trajectory). Wall projected early 2027.

IAM-2026-P006 Mar 29, 2026 **PENDING**

Microsoft Majorana: if A > 0.002 upon first gate error publication, topological protection is not outperforming standard SC transmon class.

Basis: Binary test. Majorana n predicted ~ 1.19 (SC iSWAP/ECR class). A < 0.001 = topology working. A > 0.002 = standard transmon performance.

IAM-2026-P007 Mar 29, 2026 **PENDING**

Google Willow 105Q operated at 7.5 mK will achieve p(2Q) = 5.50x10^-4 — crossing the QEC threshold without a new processor.

Basis: IAMPerformance thermal prediction from published Willow data. n = 1.449 (SC CZ). Halving temperature = 7.5 mK. Baseline: p = 1.50x10^-3 at 15 mK.

IAM-2026-P008 Mar 29, 2026 **PENDING**

IBM Nighthawk 120Q operated at 7.5 mK will achieve p(2Q) = 7.88x10^-4 — crossing QEC threshold without a new chip generation.

Basis: IAMPerformance thermal prediction. n = 1.449 (SC CZ). Baseline: p = 2.15x10^-3 at 15 mK. Halving temperature = 7.5 mK.

IAM-2026-P009 Mar 29, 2026 **PENDING**

Cooling IonQ Forte below 1 mK will not cross the fault-tolerance target (A < 1.0x10^-4) regardless of temperature reached. The architecture ceiling is set by the Yb+ design plateau, not temperature.

Basis: IAMPerformance cooling assessment. n = 0.612 (Ion A class). LOW sensitivity. Forte can cross QEC at 0.1 mK (p = 9.79x10^-4) but the FT target (A < 1e-4) requires T = 0.002 mK — below any achievable operating point. Temperature is not the path to fault tolerance on this architecture.

IAM-2026-P010 Mar 29, 2026 **PENDING**

QuEra Gemini 260Q will not improve gate fidelity through operating temperature reduction. Cryogenic cooling is not applicable to neutral atom Rydberg architectures.

Basis: IAMPerformance cooling assessment. n = 0.351 (Neutral class). MINIMAL sensitivity. Gemini operates at room temperature — laser precision and atomic geometry are the performance levers, not temperature. The architecture parameter confirms temperature has negligible effect on Rydberg gate fidelity.

DATA INDEX

Every gate error rate in this publication traced to its primary source.

Platform (processor)	p(2Q)	A Score	T_op	Primary Source
IonQ EQC prototype (Oct 2025)	8.400e-05	8.400e-05	1	arXiv:2510.17286, Hughes et al. (Oxford Ionics/IonQ), Oct 2025: 8.4(7)x10^-5 two-qubit gate error wit
Quantinuum Helios (2026)	7.900e-04	7.903e-04	1	arXiv:2511.05465, Ransford et al. (Quantinuum), Nov 2025: 7.9(2)x10^-4 two-qubit gate infidelity
Quantinuum H1-1 (Apr 2024)	8.600e-04	8.600e-04	1	Quantinuum blog Apr 16 2024 (quantinuum.com): 99.914(3)% two-qubit gate fidelity
Google Willow 105Q (2024)	1.500e-03	1.501e-03	15	Google Quantum AI, Nature 638:920 (epub Dec 2024, pub Feb 2025): 0.14% iSWAP-like gate error (
Quantinuum H2-1 56Q (2024)	1.840e-03	1.842e-03	1	Quantinuum blog (quantinuum.com/blog): 99.816(5)% two-qubit gate fidelity, H2-1
IBM Heron R2 156Q (2024)	2.000e-03	2.002e-03	15	IBM QDC 2024 (Nov 2024): Heron fleet median CZ error updated from 5e-3 to 3e-3 after calibration. i
IBM Nighthawk 120Q (2026)	2.150e-03	2.152e-03	15	IBM Quantum Platform announcement Jan 5 2026: ibm_boston (Heron R3) EPLG@100q=2.15e-3; ib
IonQ Forte 36Q (2024)	4.000e-03	4.008e-03	1	IonQ Forte Enterprise spec (ionq.com/quantum-systems/forte-enterprise, 2024): 0.40% two-qubit gate
Rigetti Cepheus 36Q (Jul 2025)	5.000e-03	5.013e-03	15	Rigetti GlobeNewswire Jul 16 2025 (investors.rigetti.com): 99.5% median CZ two-qubit gate fidelity
QuEra Gemini 260Q (2024)	8.000e-03	8.032e-03	room temp	QuEra quera.com/gemini (2024): >99.2% two-qubit gate fidelity; p=8.0e-3 (conservative)
Rigetti Ankaa-3 84Q (Dec 2024)	1.000e-02	1.005e-02	15	Rigetti GlobeNewswire Dec 23 2024: 99.0% median iSWAP gate fidelity (p=1.0e-2). Also reported: 99

A score derivation: A = –ln(1 – p(2Q)). Verification: compute –ln(1–p) for any entry above and compare to A score listed. Methodology protected under Patent Pending 64/012,720 and 64/014,568.

